CAUSALITY STUDY FORTNIGHT

8-19 September 2008

Centre for Reasoning, University of Kent

www.kent.ac.uk/secl/philosophy/jw/2008/Csf/

THE BOOK OF ABSTRACTS

The 'Causality Study Fortnight' is organised by Federica Russo and Jon Williamson as part of a project funded by the British Academy.

Thanks to British Academy, British Society for Philosophy of Science, Mind Association, KIASH (University of Kent), SECL (University of Kent) for financial support.

INTRODUCTORY LECTURES

8-9 September 2008

Philosophy of causality

Julian Reiss

The aim of this unit is to provide an overview and discussion of the major contenders for a universal theory of causation. In particular, we will study:

Regularities, Ceteris Paribus Laws, Capacities

The modern philosophical discussion of causation begins with David Hume's notorious analysis of the causal relation as no more than event-regularity: A causes B, where A and B are event-types if and only if A is constantly conjoint with B, spatio-temporally contiguous and such that A precedes B in time. We will review the main difficulties with this analysis and discuss John Mackie's state of the art version of it.

Mackie's account of causes as insufficient but non-redundant parts of unnecessary but sufficient or INUS conditions is based partly on John Stuart Mill's work on the topic. But there is another interpretation of Mill's contribution, one which does not understand causal relations as regularities: Nancy Cartwright's account of causal capacities. A capacity claim differs from a regularity claim in that concerns not outcomes of situations but rather what the different factors contribute to a situation. Cartwright understands these causal contributions realistically but there are alternative metaphysics, which we will also briefly study.

Counterfactuals

Another response to the various difficulties Hume's regularity theory was facing was David Lewis's idea that effects (now understood as token events) counterfactually depend on their causes. That is, roughly, A causes B if and only if had A not been the case, B would not have followed. Over thirty years of contributions to this strand of the literature have produced an enormous wealth of insights and highly informative (and entertaining) examples and thought experiments – but thus far to solution to the problem of redundant causation: whenever two factors compete in bringing about an effect and one is successful while the other is somehow held back, there is causation but no counterfactual dependence. We will discuss the most recent attempts to fix this problem as well as some more radical responses that have at least partially abandoned this theory.

Probabilistic Accounts

Causes are rarely universally followed by their effects. Especially in the more complex sciences where sets of causal conditions are seldom sufficient for the effect of interest we rather expect the cause to increase the probability of the effect. Clearly, not every probability-raising factor is also a cause: drops in barometer readings increase the chance of the occurrence of a storm, but it is a common factor – decreasing atmospheric pressure – that causes both. At best, therefore, a cause raises the probability of its effect given

certain other factors. But what are these factors? Any event that occurs before the putative cause? Any statistically relevant factor? Any causally relevant factor? Moreover, conditioning on relevant variables, do we require that a cause raise the probability of its effect for every value of every variable? Or at least one? Or on average? We will review the major contributions answering these questions.

Mechanistic/Process Accounts
Another idea is that causes are somehow physically connected to their effects. Perhaps smoking increases the probability of lung cancer. However, we would expect that it does so via some physiological mechanism that transports the causal message from cause to the effect, for instance via carcinogenic chemicals in tobacco smoke and mutations in lung tissue. What is the nature of these processes or mechanisms from cause to effect? There has been an explosion in the literature on causal mechanisms in recent years, and we will look at the most important contributions and see whether they are able to distinguish between genuinely causal relations and other types of relations.

Agency/Interventionist Accounts
Causes bring about their effects, they 'make things happen'. This leads one to think that one can exploit causal relations for change: if A causes B, one should be able to manipulate A in order to affect B. There are different ideas though as to what should count as an appropriate manipulation of the cause-variable A, whether it is always possible to manipulate variables in the right way and, if not, how informative these accounts in fact are. Again, we will review the main ideas in this area and draw some conclusions for an appropriate notion of cause.

Although the unit is organised around theories of causation, issues concerning the interpretation of probability will show up time and again. Most of the theories come in a deterministic as well as an indeterministic version (the probabilistic account being the indeterministic version of the regularity theory), and not all interpretations of probability sit equally well with all indeterministic theories of causation. We will discuss these issues as they come up in the context of each theory.

Causal modeling and causal discovery

Kevin Korb

This is an introduction to the principles and ideas behind causal modeling and causal discovery. Causal modeling began with the work of Sewall Wright on linear models, which introduced methods that were forerunners of Bayesian network modeling techniques and causal discovery. Wright's methods provide an easy introduction to structural equation models (SEM), which are widely used in the social sciences. Bayesian networks build upon these ideas and apply them to non-linear probabilistic models, for a much richer range of application. Recent research has focused on how to automate the learning of causal Bayesian networks. I will review this work and discuss some of the difficulties that arise for causal discovery algorithms.

Topics:

I. Path modeling
  Sewall Wright's path analysis method
  Simon-Blalock method for parameterization
  Structural equation models

II. Reichenbach's Common Cause Principle

III. Bayesian networks
  Factorization
  The Markov condition
  Propagation
  Dynamic Bayesian networks
  Decision networks

IV. Causal Bayesian networks
  Manipulation and causality
  The causal Markov condition
  Causal power theory

V. Causal discovery algorithms
  Constraint-based discovery
  Metric discovery
  Parameterization
  Problems:
    Markov equivalence
    Unfaithful networks
    Latent variables
    Variable identification
    Evaluation of discovery algorithms

VI. Knowledge engineering
  Elicitation of structure and probabilities
  Combining expert priors with discovery
  Validation

What is causal decision theory, and do we need it?

Jim Joyce

The distinction between evidential and causal decision theories concerns, not what agents should do, but whether one can adequately characterize the rationale for actions without making explicit reference to agent's beliefs about what their choices are likely to cause. Causal decision theorists maintain that there is no avoiding causality: to rationalize an agent's behavior in a given decision situation we must know what she believes about the effects of her actions, and to do this we must advert to beliefs about propositions with explicitly causal content or to modes of belief revision that are subject to explicitly causal constraints. In contrast, evidential theorists claim that, for the purposes of decision

theory, one can characterize all the relevant information by appealing only the agent's ordinary conditional subjective probabilities for non-causal propositions.

This tutorial will provide a sophisticated introduction to causal decision theory, arguing for its advantages over the evidential theory. Various mathematical formulations of causal decision theory will be presented, thought the main focus will be on philosophical issues, rather than formal ones. Special attention will be paid to questions of representing causal knowledge in decision-making contexts. In particular, models of causal knowledge based on unconditional beliefs about causal or counterfactual conditionals will be compared to those based on causally sensitive belief updating mechanisms and those based on Bayesian causal networks. We will also discuss the concept of ratifiability, and relate it to a "reflection" principle for desires. It will be shown that, despite its roots in evidentialism, ratifiability is most profitably understood in causal terms. We will also discuss an important set of issues concerning the role of act probabilities in decision making. Some philosophers, notably Wolfgang Spohn and Isaac Levi, have argued that it is incoherent to for an agent to assign subjective probabilities to their own acts (when they actively deliberating about these acts). It will be suggested, to the contrary, that a complete account of rational deliberation requires act probabilities. These probabilities must, however, be carefully interpreted. It will turn out that, for rational agents who are actively engaged in deliberation, they both reflect the agent's best estimates of what she is likely to do, and also serve to weigh the strength of her reasons for or against various courses of actions in a way different from expected utilities.

Useful Reading

Arntzenius, Frank (2008) "No Regrets, or: Edith Piaf Revamps Decision Theory," Erkenntnis 68: 277-297

Bradley, Richard (2000). "Conditionals and the Logic of Decision,"' Philosophy of Science (Proceedings) 67.

Egan, Andy (2007) "Some Counterexamples to Causal Decision Theory", Philosophical 116: 93-114.

Gibbard, Allan and William Harper (1978) "Counterfactuals and Two Kinds of Expected Utility," in Foundations and Applications of Decision Theory, edited by C. Hooker, J. Leach, and E. McClennen, pp. 125-62. Dordrecht: Reidel.

Jeffrey, Richard (2004) Subjective Probability: The Real Thing, Chapter 6 on line at http://www.princeton.edu/~bayesway/Book*.pdf

Joyce, James M. (1999) The Foundations of Causal Decision Theory. Cambridge, UK: Cambridge University Press.

------ (2002) "Levi on Causal Decision Theory and the Possibility of Predicting One's Own Actions," Philosophical Studies 110: 69-102.

----- (2007) "Are Newcomb Problems Really Decisions?" Synthese, 156[3]: 537-562.

Causal models in evidential reasoning

David Lagnado

How do people integrate a mixed and complex body of evidence? For example, when judging the guilt of a suspect, how do jurors assess and combine different items of

evidence such as witness testimonies, alibis, confessions and forensic evidence? Do they conform to appropriate normative principles, or are they systematically biased?

This tutorial will introduce students to the main psychological theories of evidential reasoning: belief adjustment, the story model and coherence-based models. It will discuss the evidence for and against these models, and compare them to a normative Bayesian standard. Particular attention will be given to the important role of causal models in human reasoning about evidence.

We will also discuss recent experimental studies that present problems for these theories. In particular, we will consider (i) what happens when evidence is discredited; (ii) the distinction between witness and alibi evidence; (iii) the effect of presenting evidence in different orders.

On the basis of these findings, novel approaches to evidential reasoning will be discussed. The possibility of generalizing these beyond the domain of legal reasoning will also be considered.

CAPITS 2008

10-12 September 2008

Invited speakers

Nancy Cartwright - Causal claims: from science to practice

Considerable effort in pure science is invested in establishing causal claims. This is true across a great variety of different scientific disciplines using a great variety of different scientific methods. Once 'established' we expect these claims to be reliable, repeatable and to continue to be borne out outside the precincts of pure science: they are among the reliable products that pure science puts on the storeroom shelf to be taken down and put to use at other times and places by other people with other purposes. How does this work? And what kinds of facts must a causal claim be reporting to allow it to work in this way? In particular, what roles will probability play, and where? This are the questions to be addressed in this paper.

Damien Fennell - Identifying causes in econometrics

For much of its history, econometrics has avoided use of the term 'cause' in describing its methods and goals. This is surprising given the fact that the work of the founders of econometrics, such as Frisch, Haavelmo, Koopmans and Simon, explicitly attempted to construct methods for identifying and measuring causes from economic data. The resulting orthodox econometric theory, however, in face of positivistic interdictions on causality, came to be couched purely in mathematical and statistical terms, with its causal content hidden. Recently the situation has relaxed, however, and different views of causality in econometrics have been expounded and debated. The work of Clive Granger, Kevin Hoover and Stephen LeRoy are good examples of this work. Continuing this revival of causality in econometrics, this paper focuses on the important identification conditions in econometrics. In econometric textbooks these are presented as necessary conditions for identifying 'structural' parameters from data. These conditions are central to the econometrician's aim of identifying the structure of the economy from non-experimental data. Historically these conditions have a close connection with concepts of causal structure, as seen in Herbert Simon's highly influential (1952) paper on the subject. In this paper, I argue that, despite their mathematical presentation, the identification conditions only make sense if structural equations are read in a substantive, non-mathematical way. I then build on Simon's work to develop a causal interpretation of the conventional identification conditions. The end result is one which brings out the genius of the econometric method, that of building experiments (ex post) from data generated in non-experimental settings. However, it also brings out just how strong and metaphysical the identification conditions are, and thus shows some important limits of econometric approach to causal inference.

James Joyce - Causally unstable acts and the role of regret in decision theory

We are accustomed to encountering mixed actions in game theory, where players often wind up believing that others will perform a variety of "pure acts" with positive probability. In single-agent decision theory, however, mixed acts are only invoked in "Buridan's Ass" problems to break ties. One exception arises in decision problems in which each act is causally unstable in the sense that deciding to perform it gives the agent compelling reason for thinking that it will have less desirable consequences than some alternative. Andy Egan has suggested that decision problems with causally unstable acts provide counterexamples to causal decision theory. As Frank Arntzenius has argued, however, Egan's cases fail as counterexamples to a sophisticated version of causal decision theory that incorporates a plausible model of rational deliberation. As Arntzenius shows, this "deliberational causal decision theory" satisfies the principle of Weak Desire Reflection (WDR), something that he takes to be a requirement for any reasonable decision theory. WDR says, roughly, that one should not prefer one option to another when one also expects that acquiring more evidence will lead one to reverse this preference, and thereby to regret one's choice. While I agree with Arntzenius's critique of Egan, I will take issue with two of aspects of his overall position. First, while Arntzenius sees himself as proposing an alternative to causal decision theory, I will show that it is nothing more than ordinary causal decision theory combined with a sound epistemology of causal judgments. Second, I will argue that the status of WDP is more complicated that Arntzenius suggests. The principle does apply straightforwardly the mixed actions that result from the Egan cases. However, these are situation in which the agent assigns a non-extreme probability ($0 < p < 1$) to more than one of her acts. When this uncertainty is resolved (as a result of her action) WDR seems to be violated. I shall explain why this is not a problem. Causal decision theorists must either allow for violations of WDR or must seek to preserve it by adopting a very discriminating notion of what it is to acquire evidence. One can see the rationale for taking one or the other of these courses by coming to understand how an agent's probabilities for her own pure actions in these mixed states reflect the value of her acts in a way different from causal expected utility. This new way of thinking about act probabilities helps to explain how mixed actions can naturally arise for rational agents in single-agent decision theory.

Kevin Korb - A new causal power theory

The causal power of C over E is (roughly) the degree to which changes in C cause changes in E. A formal measure of causal power would be very useful, as an aid to understanding and modeling with complex stochastic systems. Previous attempts to measure causal power, such as those of Good (1961), Cheng (1997), and Glymour (2001), while useful, suffer from one fundamental flaw: they only give sensible results when applied to very restricted types of causal system, all of which exhibit causal transitivity.

Causal Bayesian networks, however, are not in general transitive. We develop an information-theoretic alternative, causal information, which applies to any kind of causal Bayesian network.

Causal information is based upon three ideas First, we assume that the system can be represented causally as a Bayesian network. Second, we use hypothetical interventions to select the causal from the non-causal paths connecting C to E. Third, we use a variation

on the information-theoretic statistic mutual information to summarize the total causal influence of C on E. Our measure gives sensible results for a much wider variety of complex stochastic systems than previous attempts and promises to simplify the interpretation and application of Bayesian networks.

David Lagnado - Thinking about evidence

Are there general principles for how people update their beliefs in the face of uncertain evidence? How do these relate to formal theories of evidence integration? In particular, how do jurors, judges or investigators draw conclusions from large bodies of interrelated legal evidence? On the basis of empirical studies this paper will argue that people construct qualitative causal networks to interpret evidence and draw inferences, and that they use heuristic rather than fully Bayesian computations. For the most part this enables them to make reasonable and flexible inferences, but it can also lead to systematic biases. Moreover, people are sensitive to the order in which information is received. This conflicts with the dictates of Bayesian updating, but can reflect an adaptive response to memory and processing demands. In addition, certain forms of evidence (e.g., alibi testimony) introduce a social dimension into the analysis of evidence. For example, the level of detail given in an alibi statement is potentially a cue to its credibility, but this depends on the judge's model of the alibi provider (and the alibi provider's model of the judge). Thus alibi evidence needs to be understood in terms of strategic interactions between actors in the legal context. Overall, some of these psychological factors can be incorporated into current formal models whereas others require novel extensions. These findings also have potential implications for the way evidence is presented in court.

Michel Mouchart - Causal explanation: mechanisms and recursive decomposition

(Joint work with Federica Russo)

Explanation is still matter of debate in the scientific and philosophical communities (Psillos 2002, Woodward 2003, Halpern & Pearl 2005, Craver 2006). In spite of a large disagreement on various aspects of explanation, a consensus seems to concern the fact that there are various types of explanation that might fit different contexts or different purposes. This paper focuses on causal explanation, with special attention to the social science domain. Notably, we investigate the features of explanation in the context of a more quantitatively-oriented causal analysis, also known as 'causal modelling'. We argue that the social sciences, in the attempt to understanding social phenomena, look for explanations having the following features. (i) Explanation is 'causal', that is we look for cause-effect relations or causal mechanisms. (ii) Explanation is relative and partial, that is relative to the specific conceptual framework, dependent on available empirical and theoretical information, and involving an implicit stopping rule in order to avoid an otherwise ad infinitum chain of 'explaining the explanatory'. (iii) An explanation is given by decomposing a complex causal mechanism into a sequence of (more elementary) explanatory mechanisms.

The paper is organised in two parts. In the first part we discuss the general features of causal explanation. Building up on previous philosophical accounts, we stress that causal explanations in the social sciences is an appropriate answer to a specific why-question (van Fraassen 1980) but that we also have to 'put the cause into because' (Salmon 1984).

This is not enough though, as when it comes to the study of social phenomena, we need to uncover the mechanism underlying the phenomenon. We then reinforce the view that causal modelling is and ought to be the modelling of mechanisms (Russo 2008). Moreover, explanations need to be tailored to the specific background context and conceptual framework. In other words, the social sciences primarily seek to provide specific explanations for the problem under investigation, and then to generalise upon that.

The second part of the paper shows that the philosophical considerations above have 'statistical counterparts' in the formal apparatus of causal models. We examine the formal structure of quantitative causal models (e.g., structural equation models, covariance structure models) and show that a causal model provides an explanation which is relative to the available data by constructing a vector of variables deemed to be relevant and available. But the explanation is also partial, because the statistical model includes a random component thus delineating the frontier between what we can and cannot explain. Finally, a complex mechanism, modelled through multivariate probability distributions, may be recursively decomposed into a sequence of marginal and conditional distributions, each one being associated to a (simpler) explanatory mechanism (Mouchart, Russo, Wunsch 2008).

Stathis Psillos - Mechanisms: from explanation to causation and back again

Mechanisms are back in vogue, but our understanding of them has changed. A mechanism, nowadays, is virtually any relatively stable arrangement of entities such that, by engaging in certain interactions, a function is performed or an effect is brought about. To call a structure a mechanism is simply to describe it in a certain way—focusing on its internal composition and the interactions or processes through which an effect is brought about and not so much on the function it performs. When we think about mechanisms—even in this relatively attenuated sense—there are two issues we need to consider. The first is broadly epistemic and has to do with the understanding yielded by identifying and knowing the internal workings of mechanisms; the second is broadly metaphysical and has to do with the ontic status of mechanisms as building blocks of nature. These two issues can be brought together under certain assumptions. But they are distinct—the first treats mechanisms as (useful) vehicles of explanation, while the second treats mechanisms as the means through which causation operates—the cement of the universe. Current views about mechanisms tend to run together these two issues. There is a tendency to draw conclusions about the nature of causation from the explanatory role attributed to mechanisms. But a) causation need not be mechanistic; and b) mechanisms can be explanatorily useful even if causation is not mechanistic.
In this paper, I discuss some of the current conceptions of mechanism (notably those by Machamer, Darden and Craver, Machamer and Bogen, Glennan, and Woodward), aiming to analyse their role in causation and explanation. I draw attention to how different conceptions of mechanisms ascribe different roles to them and deal with the two issues above in significantly divergent ways.

Miklos Redei - Open problems and recent results on causal completeness of probabilistic theories

A probabilistic theory is called causally complete if it provides explanation of all the correlations it predicts. The causal explanation can be of two sorts: in terms of a causal connection between the correlated entities or in terms of a so-called common cause (or common cause system) of the correlation. The paper defines causal closedness explicitly both in classical Kolmogorovian probability theories and in non-classical (quantum) probability theories -- in the latter ones the Boolean algebra representing the random events is replaced by a more general orthocomplemented lattice and the classical probability measure by a more general countably additive bounded map -- and reviews the recent results and open problems about causal completeness. Special attention will be paid to the status of causal completeness of (local, algebraic, relativistic) quantum field theory, where the spacelike correlations predicted by the theory are supposed to be explained by properly localized common causes.

Julian Reiss - Causation: an evidentialist perspective

Evidence for causal claims comes in a variety of forms in the social sciences. The most important of these are evidence of what would have been, evidence of regularities, evidence of certain statistical relations, evidence of connecting mechanisms and evidence of invariant relationships. Social scientists often use evidence from more than one source in order to confirm a single causal hypothesis and sometimes even demand a plurality of evidence in order for a causal hypothesis to be regarded as established. The overall aim of this paper is to provide an analysis of this state of affairs and draw some methodological conclusions. A number of philosophers have recently offered pluralistic perspectives on causation. Here I distinguish epistemic, conceptual and metaphysical versions of pluralism and consider some of the arguments in favour of these. Ignoring metaphysical issues here, I end up with a form of evidential monism but conceptual pluralism about causation in the social sciences.

Paolo Vineis - The nature of medical theories and conditionalized realism

What is the nature of an observational medical theory, such as "tobacco smoking causes lung cancer" ? It can hardly be claimed that such theories represent universal laws of nature, comparable to the laws of thermodynamics or molecular genetics. In the meantime, they cannot be dismissed as simple empirical generalizations. We tend to believe that the statement "smoking causes lung cancer" is clearly related to some natural phenomenon. The feeling that it reflects something more than an empirical generalization does not mean that we are ready to accept that such statement is comparable to laws describing basic natural phenomena like the genetic code.

According to Schaffner (1), biology is characterized by "middle range" theories, i.e. laws that are intermediate between the simple observation of empirical regularities and universal statements about nature. I believe that the three models I have proposed are middle range theories of cancer. Middle range theories have the peculiarity of being strongly based on mutual reinforcement between different types of evidence, at different levels of reality and including some reference to basic laws of nature. The two main features of middle range theories are their being temporal models (i.e. they refer to

phenomena that undergo a process, like carcinogenesis) and their being "overlapping interlevel models" (i.e. they serve to connect different levels of reality).

Let us consider the relationship between tobacco and cancer. Even after the publication of persuasive evidence linking lung cancer to tobacco smoking, some investigators questioned whether the epidemiologic evidence incriminated smoking as a cause of cancer in humans. In particular, R.A. Fisher, an eminent statistician of last century, claimed that the early epidemiologic observations could not be interpreted as a proof of cause-effect relationship, arguing that one could not rule out that a genetic factor both increased the propensity to smoke and the risk of lung cancer. A key criticism was that exact knowledge of the mechanisms of tobacco carcinogenesis was necessary to establish a cause-effect relationship. Such criticism was at the root of skepticism towards epidemiological evidence and its applications in Public Health.

In fact, in addition to the (redundant) epidemiological observations linking tobacco to lung cancer in humans, we have now several types of evidence at different levels. Tobacco smoke contains many mutagenic and carcinogenic substances. Both tobacco smoke and extracts induced tumors in experimental animals. A general trend in molecular studies is the increasing evidence that point mutations in tumour suppressor genes (i.e., p53) and oncogenes (i.e., ras) may be specific both for the type of tumour and for the critical environmental exposure; this is true also for tobacco. In addition, as we have seen, smoking leads to strong cell selection, induces epigenetic events, and also causes chromosome instability.

Furthermore, Fisher's hypothesis that genetic predisposition both induces smoking habits and increases the risk of lung cancer has been refuted on the basis of twin studies

To admit that smoking causes lung cancer one need not be either a realist or an empiricist, to refer to a long-lasting debate in medicine. The realist postulates that empirical observations do refer to some reality in the external world (independently of theoretical models); the empiricist strictly sticks to observable entities, avoiding any judgement about the essence of reality. For example, realists in medicine tend to believe that basic biochemical or molecular mechanisms explain the effectiveness of therapies; while empiricists strongly advocate empirical evidence coming from randomized controlled trials. Wide areas of observational medicine, and particularly epidemiology, clearly belong to the empiricist field. As a third alternative, Schaffner proposes a "conditionalized realism". This means that a "middle range" theory is held to be true if two conditions are met: 1) that also "auxiliary hypotheses" are true; 2) that no valid alternative explanation can be put forward. The second condition is well known to epidemiologists, since it corresponds to the concept of "confounding". The first condition is also easily understandable : examples of auxiliary hypotheses are that the design of a particular study did not introduce bias; that the evidence collected from animal experiments can be extrapolated to humans; that tobacco-related mutations in specific genes (oncogenes) actually are relevant to the carcinogenic process, etc.

Which type of message does the "tobacco and cancer" example convey ? First, we believe that smoking causes cancer not only on the basis of empirical observations in humans (which can be limited by their non-experimental nature, but are in fact overwhelmingly convincing in the case of tobacco), but also because we have independent proof referring to different levels of reality. Such proof includes reference to some of our most profound beliefs concerning nature, such as the crucial role played by DNA damage in carcinogenesis. Therefore, prior beliefs in nature are essential in the interpretation of empirical observations.

Secondly, as in other fields of science, also in observational medicine the truth of a theory is conditionalized on auxiliary hypotheses and the lack of alternative explanations. This conditionalized nature of biologic realism (Schaffner) is an example of the interplay between direct evidence and interpretation, in that even an experiment - such as a randomized trial - will be interpretable only in the context of background knowledge concerning auxiliary hypotheses (although a randomized experimental trial needs less auxiliary hypotheses than observational medicine).

References

Schaffner KF. Discovery and Explanation in Biology and Medicine (Science and Its Conceptual Foundations series). University Of Chicago Press, 1994

Contributed speakers

Rani Anjum & Johan Arnt Myrstad - The paradoxes of conditional probability

In his paper 'What conditional probability cannot be' Alan Hajek argues that the intuitive probability of 'A given B' is not derived from or analysable in terms of the technically defined conditional probability 'P(A/B)' as a ratio of 'P(A&B)' and 'P(B)':

$$\text{RATIO} \qquad P(A/B) = P(A\&B) / P(B) \qquad (P(B) > 0)$$

Hajek presents us with several cases in which the probability P(A given B) is 'basic and immediate', while the probability P(A/B) is undefined (basically because the probabilities P(A&B) or P(B) are undefined). He concludes that the notion of conditional probability is basic and primitive, and cannot be analysed in terms of unconditional probabilities. On the contrary, he says, unconditional probabilities are always derived from conditional probabilities. For instance, the probability of a coin landing heads is conditional upon assumptions about the coin being fair, being tossed, etc.

We agree with Hajek on all these points. However, while Hajek does not regard RATIO as an adequate analysis of the probability 'P(A given B)', he seems willing to regard it as a constraint over rational probability assignments to 'P(B given A)', so that "whenever P(A&B) and P(B) are sharply defined and P(B) is non-zero, then the probability of A given B is constrained to be their ratio" (p. 314).

This is clearly too optimistic. The divergence between the intuitive understanding of conditional probability and the RATIO goes even deeper than Hajek reveals. In this paper we present three paradoxes of the ratio analysis of conditional probability, all of which

applies even when P(B) and P(A&B) are defined and P(B) > 0. This shows that the technical notion of conditional probability gets its plausibility and alleged applicability basically from our intuitive understanding of the probability of P(C given A).

What we conclude is that our intuitive notion of conditional probability cannot easily be replaced when analysing and evaluating the distributions of absolute probabilities and indicative, subjunctive, or counterfactual conditionals about probable outcomes. Thus, we point to the view that causality and probabilistic causation can be understood best by relegating RATIO and material conditional interpretations from the central positions they occupy in recent studies of these matters. Instead we should study how our intuitive notion of conditional probability is intimately involved with our conceptions of causality and probabilistic causation – also as these are articulated through the technical notions.

Marianne Belis – Causality and subjective probability

Human beings have evolved the ability of transforming the incomplete and changing data received from the environment into an inner qualitative measure of uncertainty, which later becomes a formal subjective probability, which nowadays measures the concept of belief. From the outset, the meaning of probability was either purely subjective in a deterministic universe or partially objective, when genuine randomness was recognised in the world. This led to confusion in which subjective probability was defined by the same rigorous axioms as the objective variety, or else simply ignored. Nowadays, the advent of belief functions creates a new mathematical framework for human uncertainty, liberated from the axioms of objective probability.

The growing interest in subjective probability is due to the development of applied sciences, like engineering, medicine, economics, and geology, which raise the problem of deciding under uncertainty when frequencies are unavailable. One of the chief virtues of subjective probability is its ability to take into account the probability of the singular event, which frequentists reject in spite of its widespread use in every day life.

.

The main effort in approaching subjective probability has been devoted to abstract, mathematical models and less to the act of judgement by which given evidence is transformed into a degree of belief. The prevalence of the mathematical formalism led to the startling conclusion that any subjective values are correct so long as they respect the coherence conditions. Artificial intelligence, in its quest to simulate thought processes, is highly interested in clarifying the act of judgement, especially since its goal-oriented applications require efficiency and not only axiomatic rigor.

The paper analyses of this act of judgement, taking into account the various kinds of information available to the observer, their relevance in a given context and their combination in order to assess a numerical measure of probability. Among the various sources of information used in belief formation, I consider that the causal connection plays a crucial role.

Causality and probability are both models of change closely related with what is observed in the world. Born at different historical times and for different purposes, these concepts

progressed along distinct paths. With hindsight, the scientific community now looks for a unified model to which each of them makes its contribution. The practical value of the concept of causality, broadly used in science and in every day life, prevails over the difficulties to define it exhaustively. I consider that causal connections established by scientific means are a sound source of information in building subjective probabilities.

The theory of decision under uncertainty relies on probabilities, but these are supposed to be deduced from frequencies when such are available. In the vast majority of cases when frequencies are unknown or the event is singular, probabilities are subjective. When used in decision situations subjective probabilities have to fulfil the coherence axioms, but they have mainly to lead to correct decisions by maximising expected utility. The more correct their evaluation, the more correct the estimates of the expected values of the utilities and the higher the probability of a good decision. The paper emphasises the connection which exists between causality and decision theory through subjective probability.

Lorenzo Casini - Fetzer's solution of the reference class problem

Can we attribute a probability to the occurrence of a single event? According to Hájek, the well-known "reference class problem" (RCP) besets all theories of probability that are 'genuinely informative and that plausibly constrain our inductive reasoning and decisions'—"informativeness requirement" (IR). Depending on the way the event is classified, the probability, whether objective or subjective, changes accordingly. I adopt Hájek's distinction between a metaphysical (MRCP) and an epistemological problem (ERCP). To solve MRCP requires answering the question: What is the probability of a single event? i.e.: Is there a "relevant" reference class? To solve ERCP, instead, requires answering: Given the available evidence, what numerical value is to be attributed to this probability? i.e.: What is the relationship between reference class and probability? Whilst I agree that the ERCP has no definite solution, I contend that Fetzer's propensity interpretation does solve MRCP, and that this is not possible if we take, as subjectivists do, every event as unrepeatable, depending on an infinite list of conditions, and yet want to fulfil IR.

First, I illustrate how RCP arises in a frequency interpretation à la Reichenbach-von Mises. Insofar as probability is operationally defined as the limit of a frequency ratio, the narrower the reference class gets the lesser statistical analyses of variance under ordinal selection and absence of after-effects are informative. Hence, this interpretation is in principle unable to overcome MRCP. I argue that Fetzer's interpretation of probability in terms of propensity, applied to an intensional language, solves RCP, and meets IR, therefore classifies as a genuine solution of MRCP. Probability is the dispositional tendency for an experimental set-up to produce a particular result on its single trial, whereas its numerical value indicates this tendency's strength, hypothesised and tested from—but not defined in terms of—frequencies. Single-case probability statements are provided with a lawlike and counterfactual value, by specifying the complete list of relevant causes on which the event depends. I defend the testability claim against epistemological and methodological objections, by appealing to Fetzer's ontological distinction between "kinds of things" and "things of kinds".

Secondly, I explain why RCP plagues also Pearl's counterfactual analysis within causal Bayesian networks (CBNs), arguing that we cannot consistently take probabilities as subjective and give causality an objective value. Pearl aims at fulfilling IR, by letting finite frequencies suggest causal dependencies. However, causal assumptions rest on confidence, not on objective and testable probabilities. When dealing with the single case in counterfactual analyses, RCP is avoided only on pain of assuming unknown factors besides known ones—to meet the graph-completeness condition, and their mutual independence—for the network to be causally interpretable. This does guarantee that, given the subjective probability of non-occurrence of a single (unrepeatable) event, intervening on its causes would have been sufficient for the same event to occur. Yet this assumption is in principle untestable—IR is violated, therefore metaphysical. I show how Fetzer's solution permits to answer counterfactual queries with CBNs by solving the completeness problem without requiring untestable assumptions.

Tom Claassen - Zooming in on probability and causality

At present, several views on the deceptively straightforward concept of causality co-exist quite peacefully. Despite their fundamental differences, probabilistic, mechanistic, counterfactual, agent, epistemic and even pluralistic theories are all quite apt at capturing and explaining much of the expected characteristics of causality. Perhaps the only aspect they all agree on is that causality is rooted in the real, physical world. Surprisingly, few have looked in detail at this physical level to find out what happens to the interpretations once we 'zoom in' on a system by following its dynamical behaviour through phase-space using the equations of motion. Our aim is to understand probability and causality at this level.

We investigate a model of a simple dynamical system and try to establish where the probabilities and causal nodes actually reside in the parameter space, i.e. we try to derive what (if anything) the causal structure actually is. We compare results with the characteristic problems and features commonly associated with probability/causality, and study how the causal structure affects the distribution measured in an experiment. Our approach offers a fresh look on the concept of probability and its close link to causality. It provides us with a first glimpse of what an 'underlying causal structure' might actually look like. Intuitively: probability relates to proportions of trajectories in phase-space that lead to a certain outcome, causality relates to changes in these proportions with respect to changes in system parameters; causal nodes represent areas where these changes are effectively induced.

In this phase-space view causality takes on a much less 'tangible' role than usual: at the 'lowest level' the neat structural equations and chains of basic causal mechanisms we use in our models simply do not exist. It becomes clear that the observed random variables we typically represent in graphical models need not correspond to nodes in the underlying causal structure: there is an implicit mapping of variables in going from the causal to the probabilistic level. This mapping may introduce 'non-causal' nodes in the graphical model, leading to an apparent breakdown of transitivity in certain cases. We study how and why they emerge and how they might be recognised. We explain why

causal relations can depend on the level of detail in measurements and see in what cases apparent probabilistic dependencies may arise without a causal link.

Fortunately, in causal analysis for other systems and situations, there is no need to keep going back to the physical equations of motion: our results apply to the nature of causal relations and the way they manifest themselves in experiments. Causality is a single concept, characterised by manipulability through physical links. It shows itself and may be analysed through both probabilities (or counterfactuals) and 'mechanisms', as long as we recognise that the nodes and the relations that we infer represent but a mapping of an unknown underlying structure. And this may or may not impact any causal conclusion or interpretation we aim to make.

Brendan Clark – A parallel literature: causation in medicine

While causation is used and discussed extensively in medicine, the medical and philosophical literatures on the subject rarely intersect. However, there are examples of publications which display uncanny parallels to those in the other discipline. Rothman, 1976, put forward a notion of causation that became known as the 'sufficient/component cause model' [Greenland, 1995]. In summary, any factor which can be considered causal is at least a necessary component of a group of factors which, acting together, are sufficient to evoke a cause. These groups may have interchangeable components, and a component may be necessary if it is part of all possible groups of cause – Rothman gives the example of 'having an appendix' as a necessary component of all causal groups for appendicitis.

This causal scheme has been influential in medicine, and it deals rather nicely with many problems of applied causation when compared to earlier causal schemes, such as the Koch-Henle postulates, which were based in the 'aetiological standpoint' – that is, a position where disease is defined in terms of its causative entity, rendering the causal pathogen a universal, necessary and sometimes sufficient cause for the disease. Rothman's model is useful not only in conceiving of the causes of multifactorial disease, such as cancer, but also in dealing with the modern concept of host factors in disease. To illustrate, while the aetiological standpoint would have the bacteria Mycobacterium tuberculosis as a universal, necessary and sufficient cause of the disease tuberculosis, the sufficient/component model leaves the presence of Mycobacterium tuberculosis as an insufficient but necessary part of the cause, requiring other host factors – vaccination status, immune function, appropriate exposure etc. – to achieve sufficiency. This appears to me a much better reflection of causation in clinical medicine.

A feature of the model that has not been previously discussed in the literature is its resemblance to Mackie's INUS (insufficient and non-redundant parts of unnecessary but sufficient causes) conditions. [Mackie, 1974: 62]. Dealing with the "plurality of causes" [Mackie, 1974: 61], Mackie claims that causes are instantiated by the collective action of logical conjunctions of necessary but insufficient constituent factors. The resemblance and temporal proximity of these two models of cause is intriguing, suggesting either some unacknowledged influence of Mackie's ideas on Rothman or, more interestingly, some sort of convergence between disciplines.

In this paper, I therefore propose to introduce both of these schemes of causation and compare them. I will then look at some of the problems faced by Mackie's INUS conditions (for instance, issues of causal priority) and transcribe them into a critique of Rothman's component causes. Finally, I'd like to discuss more generally the implications for philosophy of this parallel literature on causation.

Flavio D'Abramo - Final, efficient and complex causes in biology

In this paper I offer an analysis of the connection between efficient and ultimate causes in biology. As it is well known, the first approach to biology has been historical in character. Authors such as William Paley, Darwin and Lamarck were seeking ultimate, or final causes, i.e. the past conditions that ultimately originated the existing organisms.

I suggest that this approach was indeed correct. More precisely, I will argue that historical methods are necessary to biology. This claim must be understood with some qualification, however. For most historical methods are goal oriented: the past phenomena are used to explain present effects. The study of many structural dynamics between living things and their environment like those faced up by Waddington in his epigenetics biology show that organisms possess a number of features that have no particular function and that evolution can proceed through direct mutations. This kind of causation, I suggest, may be seen as complementary not only to the finalistic one, but also to the neo-Darwinian concept of evolution in which selection acts on random mutations.

I will focus on some aspects of D'Arcy Thompson's work in which he has given up final explanation to recover the physical forces acting on the living objects. Only after that this forces act on biological substance will arise a structure, a biological form, and then a function. And only after that structures accomplish functions then natural selection can "chose" adapt organisms.

D'Arcy Thompson defines the evolution as a system that evolves in space and he describes this "evolution" through a geometrical-mathematical approach. In his On Growth and Form, D'Arcy Thompson criticizes the idea of progressive evolution. In fact, the very idea of evolution involves reference to an ordered and continuous series of events.

Such a progressive connotation characterizes much of the biological views. The law of recapitulation, summed up by Ernst Haeckel, is the best example of a finalistic criterion. Even if this law is false, it is still used in the present biological explanation.

I want suggest that in order to analyse form and function this particular finalistic-progressive criterion has to be put aside. I suggest that if we avoid the use of finalistic concepts in history, we can then characterize some biological processes also as a causal sequence of phenomena that origins structures without any immediate function or function that change in time or arose as novelty.

I will also mention some classical problems that spring out suspending the progressive, ordered and uniformitarian view of evolution, such as the lack of classical prediction, that

arise when history is considered without such a constitutive concept of finality: evolutive past processes are not always the same of the present one. I argue that this particular historical method unrelated to classical philosophical dualistic view rooted in the dichotomy "final/efficient causes" is today necessary for biology.

Isabelle Drouet - Probabilistic theories and levels of causality

In spite of the debate on their relationship, generic causality and singular causality are nowadays commonly distinguished. A noticeable difference between these two levels of causality is that probabilistic theories naturally fit generic causality, whereas singular causality does not admit any uncontroversial probabilistic analysis. Thus, an appealing position seems to be as follows: generic causality can be given a probabilistic analysis but singular causality cannot. This position was defended in particular by E. Sober.
The aim of my paper is to examine how this position articulates the avalaible analyses of the relationship between levels of causality. More precisely, I first identify two main families of analyses of the relationship between generic and singular causality – neo-Humean analyses on the one hand and generalization analyses on the other hand – and then I argue that neither is straightforwardly compatible with the idea that only generic causality (and not singular causality) can be given a probabilistic analysis. Finally, I explore the consequences of this claim.

Anton Froeyman – Concepts of causality in history

Over the last few years, it has become increasingly clear that a monistic approach to the philosophy of causality will always encounter considerable difficulties. As a result, more and more attention has been going to a pluralistic characterization of causality. Instead of competing with each other, the different philosophical theories of causality (counterfactual theories, process theories, manipulability theories, probabilistic theories, …) are standing side by side. In such a pluralistic approach, the specific concepts of causality which are used in the various special sciences are of great importance. In this paper, I will present the different methods and forms of explanation which are and have been used in different areas of historiography, and I will try to trace the different concepts of causality and causal explanation involved in these methods. This is interesting because history is a multi-paradigmatic discipline with a large variety of areas, methods and forms of explanation. It focuses both on the individual (in the so-called "history from below") and on general structures (as in Marxist or structuralist approaches). It can, for example, be in search of the causes of both the industrial revolution in Europe and of those of the worldview of an 16- century miller. Furthermore, it applies to a wide variety of methods, ranging from formal and statistical methods (in cliometrics) to more subjectivist interpretative methods (in the traditional hermeneutic approach). I will try to show in what way the plurality of philosophical concepts of causality is reflected in the plurality of methods and research areas in historiography, by tracing the various dominant philosophical accounts of causality in these different domains and methods. Furthermore, I will take a short look at so-called non-causal explanations in history (in the hermeneutic and postmodern tradition) and try to ascertain whether these non-causal forms of explanation really are non-causal, or whether there isn't a special case of causal explanation involved after all.

Luke Glynn - Groundwork for a probabilistic analysis of causation: some foundational issues in the philosophy of probability

The starting point in the development of probabilistic analyses of causation is usually the naïve intuition that causes, in some relevant sense, raise the probability of their effects. A natural way to cash this out is in terms of an inequality between two conditional probabilities. One might say that c raises the probability of e in the sense relevant to causation iff inequality (*) holds (where capital '*C*' and '*E*' are, respectively, the propositions 'c occurs' and 'e occurs'):

$P(E|C) > P(E|{\sim}C)$ (*)

Probabilistic analyses that appeal to such inequalities between conditional probabilities are given by Good (1961a, b; 1962), Reichenbach (1971), Suppes (1970) and Kvart (2004).

In a well-known passage, Lewis (1986b) raises the following problem for probabilistic analyses that use conditional probabilities:

"Conditional probabilities, as standardly understood, are quotients. They go undefined if the denominator is zero. If we want to say, using conditional probabilities, that c raises the probability of e, we will need probabilities conditional on the non-occurrence of c …. But there is no guarantee that this conditional probability will be defined. What if the probability that c occurs … is one? … For that matter, what if we want to apply our probabilistic analysis of causation to a deterministic world in which all probabilities (at all times) are extreme: one for all events that do occur, zero for all that don't? The requisite conditional probabilities will go undefined, and the theory will fall silent." (p. 178)

In the paper, I draw upon some recent developments in the philosophy of probability in order to address this apparent problem.

Hájek (2003a, b; 2007) has recently argued that the standard (Kolmogorov, 1933) axiomatization of the probability calculus is flawed. The difficulties concern precisely that principle which creates problems for probabilistic analyses: the principle that conditional probability is to be analysed as a quotient of unconditional probabilities.

Hájek argues compellingly that probability is a fundamentally relative or conditional notion and that we should therefore abandon the standard axiomatization in favour of one (such as that given by Popper, 1972) which reverses the traditional direction of analysis by taking conditional probabilities as primitive and analyzing unconditional probabilities in terms of them. I argue that, if we do so, then we can allow that the conditional probabilities required by probabilistic analyses are well-defined after all, even in cases where the unconditional probability of the cause is 1 (or 0).

It might still be objected that the required conditional probabilities will be undefined in deterministic worlds. Whilst one might allow that there are well-defined probabilities conditional upon probability zero events, one might nevertheless doubt that there are well-defined probabilities conditional upon events that are physically impossible

(assuming that, like most mathematicians and philosophers of probability, one doesn't equate zero probability with physical impossibility).

Since causation and determinism do not seem to be incompatible, someone who wants to defend a probabilistic analysis that uses conditional probabilities ought to show that the required conditional probabilities are well-defined even in deterministic worlds. This is a burden that I seek to discharge by arguing that there exist probabilistic special scientific laws even in deterministic worlds. Because the special scientific laws are probabilistic, they yield non-trivial (objective) probabilities for the events that they govern. Hence, I argue, there is no problem about well-defined probabilities conditional upon the absence of these events, even in deterministic worlds.

The conclusion is that probabilistic analyses of causation that use conditional probabilities can, after all, accommodate causation by probability 1 (and probability 0) events and causation in deterministic worlds. There is therefore no fundamental objection, along the lines suggested by Lewis, to attempts to develop such an analysis.

Dawn Holmes - Why making Bayesian networks Bayesian makes sense.

It is well-known that Bayesian networks are so-called because of their use of Bayes' theorem for probabilistic inference. However, since Bayesian networks commonly use frequentist probabilities, is this sense they are not Bayesian. In this paper it is argued that Bayesian networks that are truly Bayesian, in other words those whose prior distribution is derived from subjective probabilities, themselves based on all and only the available information, have certain desirable properties and strengths over and above those based solely on the frequentist approach to probability. It is demonstrated, through an example that these specially constructed graphical models may be used in otherwise intractable situations where data is unavailable or scarce and decisions need to be made. To this end, an algorithm is described for constructing a prior distribution based on subjective probabilities.

Samantha Kleinberg and Bud Mishra - multiple testing of causal hypotheses

A primary problem in causal inference is the following: From a set of time course data, such as that generated by gene expression microarrays, is it possible to infer all significant causal relationships between the elements described by this data? In prior work [2], we have proposed a framework that combines notions of causality in philosophy, with the algorithmic approaches built on model checking and statistical techniques for multiple hypotheses testing. The causal relationships can be then described in terms of temporal logic formulas, reframing the problem in terms of model checking. The logic used, PCTL, allows description of both the time between cause and effect and the probability of this relationship being observed. Borrowing from philosophy, we define prima facie causes in terms of probability raising, and then determine genuine causality by computing the average difference a prima facie cause makes to the occurrence of its effect, given each of the other prima facie causes of that effect. However, it faces many interesting issues confronted in statistical theories of hypotheses testing, namely, given these causal formulas with their associated probabilities and our

average computed differences, instead of choosing an arbitrary threshold, how do we decide which are "significant"? To address this problem rigorously, we use the concepts of multiple hypothesis testing (treating each causal relationship as a hypothesis), and false discovery control. In particular, we apply the empirical Bayesian formulation proposed by Efron in [1]. This method uses an empirical rather than theoretical null, which has been shown to be better equipped for cases where the test statistics are dependent - as may be true in the case of complex causal structures.

References

[1] B. Efron. Large-Scale Simultaneous Hypothesis Testing: The Choice of a Null Hypothesis. Journal of the American Statistical Association, 99(465):96–105, 2004.

[2] S. Kleinberg and B. Mishra. Inferring causation in time course data with temporal logic. Under review, 2008.

Jan Lemeire - When are Graphical Causal Models not Good Models?

The theory of graphical causal models is based on that of Bayesian networks. A Bayesian network offers a dense description of a probability distribution by a Directed Acyclic Graph (DAG) and a set of Conditional Probability Distributions (CPDs). A causal model gives a causal interpretation to the edges of the DAG. The causal interpretation is defined by the Causal Markov Condition and interventions. The ability to predict the effect of interventions relates to modularity and the decomposition of the system into independent components, which are described by the CPDs of the Bayesian network. These components correspond to the mechanisms of the system. We argue that all aspects of causal model theory, such as faithfulness and the possibility to learn causal models from observations, are valid if the causal model is the minimal model. We mean minimal in the general sense, i.e. that there is no other model, possibly based on another modeling framework, that describes the system simpler. Non-minimality results in compressibility of the system's description by a Bayesian network. In this paper we examine for several systems, many of which appearing in literature, whether minimality, faithfulness and the modularity assumption holds. Violations of faithfulness come from non-minimality. The modularity assumption is still valid when the non-minimality of the description is due to the compressibility of an individual CPD. Modularity becomes, however, implausible if the concatenation of the descriptions of the CPDs is compressible. This suggests that either there is a kind of meta-mechanism governing some of the mechanisms or either a single mechanism responsible for setting the state of multiple variables.

Bert Leuridan and Erik Weber - The IARC, mechanistic evidence and the precautionary principle

The IARC, the International Agency for Research on Cancer, is a division of the World Health Organisation that seeks to identify the causes of human cancer [1]. More specifically, it evaluates the evidence on the carcinogenicity of a wide range of agents (chemicals, etc.).

In this paper, we will briefly present the procedures and criteria used by the IARC, show that they are at odds with contemporary philosophical theories of causation (more specifically regarding the role of mechanistic evidence in establishing causal claims) and explore the ways in which these procedures and criteria could be defended by invoking some version of the precautionary principle.

How is the carcinogenic risk of exposures assessed by the IARC? The available evidence is divided into three different groups: *epidemiological studies*, *animal experiments* and information about *mechanisms*. The assessment consists of three phases. In the first phase, each study is evaluated separately. In the second phase, an assessment is made of the strength of the evidence for each of the three groups of evidence. For example, it may be concluded that there is 'sufficient', or 'limited', … epidemiological evidence of carcinogenicity. Likewise, it may be concluded that there is 'strong', or 'moderate', or 'weak' mechanistic evidence of carcinogenicity. In the third phase, the evidence of the different groups is combined and the agent's carcinogenicity is assessed, ranging from "Group 1: The agent is carcinogenic to humans" to "Group 4: The agent is probably not carcinogenic to humans".

In the second part of the paper we will show that these procedures and criteria are at odds with contemporary theories of causation in the following way: in the third phase, they systematically downplay the weight of mechanistic evidence. For example, if there is 'sufficient' epidemiological evidence of carcinogenicity, the agent is placed in group 1 (irrespective of any mechanistic evidence pro or contra carcinogenicity). Mechanistic evidence is only used to guide extrapolation from animal experiments to humans in case epidemiological evidence is less than 'sufficient'. However, mechanistic evidence can have at least two other important roles that are neglected in the IARC procedures. Firstly, it may help to solve the problem of confounders that threatens non-experimental studies. In the absence of mechanistic background knowledge, it can be argued, no 'sufficient' epidemiological evidence may be obtained (cf. [2] for a similar argument in the context of the social sciences). Secondly, it may help to assess the temporal stability of correlations and causal relations.

Domenico Mancuso - Why is statistical relevance undesirable?

After its formal development in the early thirties, Kolmogorov's theory of probability has soon come to be appreciated as a powerful tool for constructing non-deterministic models of causation.

Clearly, the main setback of such a theory lies in the exceptions that any probability distribution – no matter how uneven – necessarily admits: supposing event E has 80% chance of obtaining in the light of an alleged cause C, the occurrence of non-E would still be perfectly *justified* – not less than the opposite event – and hence, C would qualify as the "cause" of both outcomes.

The principle of statistical relevance, introduced by I.J. Good and Patrick Suppes and extensively defended by Wesley Salmon, seems to respond precisely to such difficulties: in order for an event C to be counted as a cause of E, it must not make the occurrence of E more probable than its non-occurrence, but more probable *than it would be* in the absence of C.

Yet, as Good himself has pointed out, it is not difficult to find counterexamples to the principle – that is, situations where an event that we would intuitively regard as a cause of E turns out to *lower* the chance of its occurrence.

Since the sixties, a rich literature has flourished on these counterexamples, and several proposals have been worked out on how to cope with them without abandoning the notion of statistical relevance. The key idea appears to be that of progressively *refining* the principle, by ruling out those *prima facie* causes which are only apparently relevant (Suppes divided them into "spurious" and "indirect" causes), and on the other hand, by taking into account an extra set of conditions (a "conceptual framework", in Suppes' words) that would make the cause appear *positively* relevant to the effect; this process, as Salmon remarks, could go on indefinitely, since causal processes are *continuous*, whereas our descriptions are always *discrete* approximations.

After discussing some counterexamples and their purported solutions, I will explain why I think the conceptual framework strategy cannot be successful. First of all, it appears dangerously *ad hoc*: why should we add further preliminary information only so long as the statistical relevance is negative, and stop once it becomes positive? Secondly, and most important, the ideal account of an event – one where the conceptual framework is complete, and processes are continuous – would no longer constitute a *probabilistic* model, but a *deterministic* one. If statistical relevance were to hold systematically, it would be impossible for things to happen "the hard way" – nobody could win a large sum by betting on the weaker horse, or resist the temptation of a cigarette while trying to quit smoking; at the very best, such events would only *appear* to be against the odds, because of our incapacity to see the deterministic mechanism that produced them. Therefore, amending statistical relevance in order to account for all exceptions is not simply unfeasible; it is not even *desirable* – at least, if probabilistic causation is to preserve its specificity.

Phyllis McKay and Jon Williamson - Mechanisms and explanation

Mechanisms have become much-discussed in the current philosophy literature, to begin to match the long-enduring key place of explanation using mechanisms in the sciences. Yet there is still no consensus on the best way to characterise mechanisms. In this paper, we will start with one of the few things everyone – both philosophers and scientists – agrees on about mechanisms: mechanisms explain. We investigate what constraints this imposes on our metaphysics of mechanisms.

We disentangle two senses of explanation and argue that the deeper sense in which mechanisms explain the phenomena they produce requires mechanisms to be both *real* and *local*. We argue that reality and locality require a broadly dispositionalist metaphysics for mechanisms, and argue against the characterisations of mechanisms given by Psillos, Woodward and Glennan on these grounds.

We raise the widely agreed requirement of the *stability* of mechanisms as a possible problem for a broadly dispositional metaphysics, but argue that a broadly dispositionalist metaphysics has the resources to deal with this too.

We finish by raising some remaining problems for the reality, locality and stability of mechanisms. We suggest solutions for some of them, so identifying the more serious worries.

Alessio Moneta - Probabilistic evidence, causal knowledge and econometric models

According to an established tradition in econometrics and in the philosophy of science, in order to learn causal relationships in a non-experimental setting one needs both some background of causal knowledge and an appropriate statistical model of the data. Hoover (2007) has recently argued that these two requirements should be carefully distinguished and that the establishment of a probability model does not require causal presuppositions. In this paper I elaborate some theses about the kind of evidence used to draw causal relations in econometric models which qualify Hoover's statement.

First, this statement is true in the sense that, for the sake of establishing a probability model, one does not need to know in which directions causal relations operate (if present). For example, knowing that some causal directions are not possible (e.g. cycles) is extremely useful, and in some cases necessary, to infer causal structures from evidence of statistical associations, but this knowledge is not needed in order to build a probabilistic model, and in particular in order to test conditional independence among the variables. Furthermore, this type of knowledge does not give any contribution to a better specification of the probabilistic model.

The situation is quite different when we consider background knowledge about the "mode" in which causal relations operate. A crucial mode in which causal relations operate is determined by the functional form which describes these relations. Knowing the functional form, I shall argue, is extremely relevant to test features of the probabilistic model. For instance, knowing that the functional form of the causal relations within a group of variables is linear permits the researcher to use straightforward statistics to test conditional independence relations among them. Another important causal presupposition is knowledge about the level of aggregation in which causal relations operate. For instance, knowing that causal relations operate at an individual level and that the observed variables are aggregate (e.g. sums or averages) of individual variables brings some important consequences for the specification of the probabilistic model. I will argue that the role played by these two presuppositions, which regard the mode in which causal operations operate, has often been neglected in the econometric literature.

My final thesis is that these presuppositions are not per se necessary to build a probabilistic model, although they give in general a significant contribution to better specify it. Referring to the literature on nonparametric statistics, I will show how, in case of lack of knowledge about the functional form, one can build, under certain constraints, probabilistic models allowing causal inference.

Amit Pundik - Could there be any epistemic reason to restrict the use of statistical evidence in court?

In a variation on Jonathan Cohen's gate-crasher paradox, Humpty, an enthusiastic football fan, is sued for gate-crashing a local football match. He was seen by Alice sneaking through the fence. Alice's visual identification ability is tested and shown to be accurate in nine times out of ten. This case seems straightforward and finding Humpty liable in private law seems intuitively right. In another football match, Hatter, another enthusiastic football fan, is also sued for gate-crashing. This time, the only evidence available is that just one hundred tickets had been sold whilst yet a thousand spectators were counted. Hatter being sued for gate-crashing is based on the fact that nine out of ten spectators gate-crashed. Surprisingly, this case is relatively straightforward too: finding Hatter liable based on this statistical evidence alone seems intuitively wrong. It is surprising because in both cases, at least prima facie, the risk of error is similar. In both cases, there is a probability of one in ten that finding liability would be mistaken.

Numerous accounts have suggested in the literature to explain the difference between the cases and to justify some restrictions (either inadmissibility or insufficiency) on the use of statistical evidence in court. A dominant direction is the idea that the statistical evidence in the second case lacks a certain quality which makes it epistemologically unwarranted to establish the fact this evidence is brought to prove. Amongst the qualities suggested are luck, appropriate causal connection (Thomson 1986, Wright 1988), weight (Cohen 1977), case-specificity (Stein 2005), and relevance (Cohen 1981 and Stein 2005). The epistemic accounts are thus attempts to identify a quality that non-statistical evidence (e.g. eyewitnesses, confessions, the individual's medical records) has but statistical evidence lacks. These attempts are grouped as "the epistemic accounts" because they share the view that making inference from the statistical evidence to the particular case epistemologically deficient.

To the existing debate around these accounts (e.g. Schoeman 1987, Pundik 2006), this paper adds three generic and inherent problems from which any epistemic account suffers. First, using statistical evidence improves accuracy and thus it is epistemologically unwarranted to restrict the use of such information. Second, an epistemic account should apply similarly in any discipline. However, some types of statistical evidence are only objectionable when used in court. Last, evidential rules require moral/political justification (Stein 2005). Yet it is unclear how an epistemic account can ever provide such a moral/political justification. Based on these three generic problems, it is concluded that the epistemic direction lacks potential. If there is any reason to restrict the use of statistical evidence in court, it cannot be epistemic.

Eric Raidl - The 'early' method of arbitrary functions: A logical frame for a causal theory of probability

Assigning prior probabilities is un-doubtfully one of the main foundational problems in the application of probability theory to science and particularly in statistics. Probability theory transforms probability assignments on certain events into probability assignments on related events, but says nothing about how to choose initial assignments. The 'principle of indifference', assigning equal probability in absence of sufficient reason, is

known as the classical solution to this problem. A general trouble with this or similar assignments (e.g. Bayes' uniform distribution) is arbitrariness, or lack of justification: Why prefer one probability assignment over all others?

One of the modern attempts made to justify prior probabilities is the 'method of arbitrary functions' (named after Poincaré, but originating with Von Kries). Philosophically the approach considers chance as resulting from physical instability and provides a frame for a 'causal theory of probability' (Nagel), neutral regarding the question of indeterminism in the laws of nature. It thus presents an alternative to a quantum-mechanical (e.g. propensity) analysis of chance phenomena. Mathematically the approach relies on a theorem which establishes equi-probability of result types on the basis of two non-trivial hypotheses. The first assumes the existence of a continuous probability distribution over a parameter space representing initial conditions, the second supposes that this space is partitioned into (locally approximately) equal, alternatively two-coloured intervals, where colours represent (unions of intervals of initial conditions leading to) types of possible results.

The second hypothesis is unanimously thought to be justified by physical symmetry of the chance setup. I will show that there are however two possible and actually defended interpretations of the first hypothesis and subsequently of the theorem: an a priori and an a posteriori reading. Early attempts (Kries, Smoluchovski, the early Reichenbach), viewed the theorem as modified principle of indifference. Same in conditional form, it replaces the dubious epistemic condition of equi-possibility by continuity of distribution, which is to be justified a priori. The recent interpretation (Plato, Strevens), on the contrary, considers continuity as general statistical information, obtainable by enumerative inductive inference. The first view thereby turns the theorem into a principle of a priori reasoning applicable under certain conditions of empirical adequacy, whereas the second considers it as providing a method of a posteriori (or inductive) inference.
The recent literature (which mostly relies on the branch of literature inspired by Poincaré, e.g. Hostinsky, Hopf) at best mentions, if not ignores early attempts. Thus a strong interpretation of the approach is left aside, as well as a general frame for a causal theory of probability which, in addition of being physical, is logical.

I propose to focus on Reichenbach early texts (1919-1920), since those are most ignored, and extract his mathematical results left aside by the literature, a general formulation of the conditions of empirical application of the theorem and his a priori justification of the continuity hypothesis. Secondly I will compare Kries' and Reichenbach's a priori interpretation with the recent a posteriori interpretation and finally ask how each might be integrated into a causal theory of probability.

Jan-Willem Romeijn & Jon Williamson - Intervention, underdetermination, and theory generation

It has often been argued that the problem of underdetermination is a logical curiosity rather than a real life possibility. Indeed, in the natural sciences it is not easy to find empirically equivalent rivals to our best candidate theories. However, in the social sciences underdetermination is fairly widespread, in particular where these sciences

employ statistical modelling. Within a given experimental setup or population study, it may very well happen that the theory allows for distinctions between hypotheses that do not correspond to a difference in the likelihood function of the hypotheses. The problem of underdetermination then appears as the problem that the data do not single out a unique best hypothesis: the best fitting hypothesis might have a number of equally well fitting rivals. The standard response to this is to look for theoretical criteria that force a choice between the rivals, such as simplicity or explanatory force. The problem of underdetermination is then resolved by an appeal to theoretical considerations.

This paper investigates a different response to underdetermination in statistical modelling. It is that the resolution of underdetermination may also be driven by further empirical criteria, specifically by intervention data. Relative to some experimental setup, two statistical hypotheses may have exactly the same likelihood functions and thus perform equally well on the observation data. However, they need not be empirically equivalent altogether. We can consider specific changes to the experimental setup, or interventions for short, such that the likelihood functions for the hypotheses come apart. This paper is concerned with the use of such interventions. We will further show that if the data resulting from the intervention dismiss all the rival hypotheses, they suggest an extension of the statistical model to accommodate them. So whenever the intervention leaves us empty-handed, it directs us towards a revision of the underlying conceptual scheme. The upshot of all this is that aspects of scientific method that are typically associated with theoretical considerations, such as resolving underdetermination and generating new theory, are here seen to be driven by empirical fact.

The paper is set up in the following way. First we describe the indeterminacy of estimations in factor analysis, a statistical modelling tool from the social sciences. We show that this indeterminacy is an instantiation of the problem of underdetermination and that accordingly, it is usually resolved by imposing further theoretical criteria. We then argue that the estimation problem of factor analysis is essentially identical to estimating the parameters in a Bayesian network with hidden nodes. We proceed to show that Bayesian networks, and thus models in factor analysis, allow for incorporating intervention data in statistical modelling, and we argue that in specific cases, intervention data can be used to resolve the indeterminacy of the estimations. When resolving underdetermination, interventions may therefore take the place of theoretical criteria. We further argue that in all the other cases, the intervention data guide us towards adding a hidden node of a latent variable in the statistical model. In other words, interventions may guide theory generation.

Federica Russo – Measuring variations

Causal analysis in the social sciences takes advantage of a variety of methods and of a multi-fold source of information and evidence. Within this pluralistic methodology and source of information, I argue, there is a single rationale of testing regimenting causal reasoning in causal modelling, which is based on the notion of *variation*. I present this position by first drawing a distinction between a definition, a criterion, and a rationale of causality, and then by offering methodological arguments. In particular, I focus on the interpretation of structural equations and suggest that invariance and regularity are

constraints to impose on variations in order to guarantee their causal interpretation. I finally discuss three possible threats to the variation thesis, namely (i) that invariance is more basic, (ii) that regularity is more basic, and (iii) that no variations can be found and tested in highly homogenous populations.

Niels van Miltenburg - Confidence in causality

It is a pity that Georg Henrik von Wright is only mentioned as a forefather of the more recent Agency approach [Menzies and Price 1993] and Interventionist approaches [Woodward 2001, et. al.], in contemporary writings on the manipulability theory of causation. That vonWright's theory is only seen as a precursor of today's manipulability accounts seems to be due to some ambiguousness in von Wright's work. Two interpretations are possible:
`Strong' experimentalism: `p causes q' means that we can produce q by bringing about p.
`Weak' experimentalism: We distinguish between causal relations and accidental regularities by using the notion of action, which is conceptually prior to that of causation.
The contemporary proponents of manipulability theory are only taking strong experimentalism into account. This is possibly because they have an empiristic conception of action. If you think, as Menzies and Price do, that action is some- thing that we (only) experience, weak experimentalism becomes too weak. The thesis that we distinguish between cause and effect on the basis of the notion of action, is not really substantial if we still have to root our notion of action in observational terms.

We should note that in his earlier writings on the subject von Wright is more ambiguous than in his later writings. In *Explanation and Understanding* [Von Wright 1971] weak and strong experimentalism are conflated but in his `Reply to my Critics' [Von Wright 1989], he is a only defending weak experimentalism.

VonWright's weak experimentalism, like Wittgenstein's or Anscomb's views on action, isn't an empiristic theory and it is offering answers to questions that differ from the questions addressed by contemporary manipulability theorists. It focuses on the epistemic access that we have to causal processes happening in the world around us. Hume already found himself confronted with the following question: \Why, when we come across some regularly succeeding events, we don't see them as just regular succession, but we see them as causally linked?" Von Wright offers an answer. He argues that there is a conceptual link between ‘causation’ and ‘action’. This link is grounded in the counterfactual element that resides in both causation and action so that, when we are confident that we can act, we are confident in the fact that causal processes are happening in the world in which we live. What is this counterfactual element in action? Certain changes in nature would not have occurred had we not produced them. When we act we ‘know’ that this is the case because otherwise we would not have had the confidence that we can act. This ‘know’, I want to argue, should be viewed as something closely related to what is called `practical knowledge' by Miss Anscombe [1957]. I find it odd that these ideas of von Wright are not at all discussed in the contemporary debate on causality because both the conceptual dependence of ‘causation’ on ‘action’ and the notion of confidence are key to understanding our human understanding of causality in the

experimental natural sciences. When scientists experiment in their laboratories they are not just observing what hap pens. They are interfering with the world, they are acting.

Jon Williamson & Federica Russo - Generic vs. single-case causal knowledge. The case of autopsy

An interesting aspect of the levels of causation is how generic knowledge of causal relations can or should inform knowledge of single-case causal relations and viceversa. On the one hand, generic causal knowledge is gained from a number of relevant and appropriate instantiated single-case causal relations. On the other hand, in many cases a single-case causal relation is established based on available generic knowledge. We argue that the inference from one level to another is not simply one way, but instead requires a constant interaction between the two levels. To exemplify, we discuss the case of autopies, where single-case causal knowledge about death is usually sought using established generic medical knowledge. However, results of autopsies of similar deaths, which are single-case, are also used to contribute to generic medical knowledge.

Karen Zwier - The causal Markov condition: should you choose to accept it?

The Causal Markov Condition (CMC) is an axiom specifying a relationship between a causal graph and the probability distribution over its vertices. In recent years, there has been much debate over this axiom. In this paper, I classify the criticisms against the CMC into two groups. The first type of criticism is metaphysical: How do we know that causal relationships, in reality, always exhibit the precise statistical relationships specified by the CMC? Could there not be multiple types of causation that each exhibit different statistical relationships? Specifically, could there not be a form of causation in which two distinct effects of the same cause are correlated, even when all common causes are given? The second type of criticism involves the application of the CMC: Why should we continue to use the axiom when there are so many known “counterexamples"?

In regard to the first type of criticism, I do not attempt to argue on a metaphysical basis that causal processes must be independent. I do argue, however, that use of the CMC is reasonable given certain aims. If the object of a particular scientific study is to discover and isolate the independent causal relationships in a system, the CMC will prove ideal. I give several arguments to motivate its use. In regard to the second type of criticism, I confirm the accusation that CMC is not applicable in every case. I show, however, that many cases that have been cited as "counterexamples" of the CMC are actually instances of misapplication or misunderstanding. Furthermore, I argue that the strength of the CMC lies precisely in the ability to discern appropriate situations for its use.

ADVANCED RESEARCH SEMINARS

15-19 September 2008

Causality and the mind (Julia Tanney)

The problem of mental causation is well-known in contemporary philosophy of mind as it constitutes one seemingly undeniable element of the mind-body problem. Mental causation forces us to see mental events as real: but a commitment to physicalism forces us to see them as dependent, if not reducible, to their alleged physical correlates. Various, ultimately unsuccessful, versions of reductive and non-reductive physicalism have emerged as ways of accommodating these conflicting intuitions.

Seminar 1

The first seminar will engage with some up-to-date contemporary literature on reductive and non-reductive physicalism.
Reading:
Louise Antony (2007), Everybody Has Got It: A Defence of Non-Reductive Materialism, in B. L. McLaughlin and J. Cohen (eds.), Contemporary Debates in Philosophy of Mind, Oxford: Blackwell

Jaegwon Kim (2007) Causation and Mental Causation, in B. L. McLaughlin and J. Cohen (eds.), Contemporary Debates in Philosophy of Mind, Oxford: Blackwell

Seminar 2

The second seminar will suggest a difficulty with Davidson's arguments for construing reasons as causes.
Reading
Donald Davidson (1980) "Actions, Reasons, and Causes", in Essays on Actions and Events (Clarendon Press).
http://www.oxfordscholarship.com/oso/public/content/philosophy/9780199246274/toc.html

Julia Tanney (1995) "Why Reasons May Not Be Causes", Mind & Language, vol. 10, nos. ½ pp 103-126. (See http://www.kent.ac.uk/secl/philosophy/staff/tanney.html)

Seminar 3

The third seminar will explore a different way of understanding the explanatory power of mental predicates.

[Reading to be circulated]

Causality in the history of philosophy (Ken Westphal)

Let's make that 'Living History': The three seminars I shall lead are issues-oriented even though historically based in neglected though philosophically important features of Kant's, Hegel's and Newton's views. (Yes, Hegel actually understood Newton and Newtonian gravitation very, very well, rumours to the contrary not withstanding.) Though these three seminars may be followed independently, they are thematically linked by issues in semantics of cognitive reference.

Seminar 1

Are causal theories of the mind legitimate?

This seminar considers whether we can make legitimate causal judgments within the philosophy of mind and psychology more generally. Kant's account of causal judgment (which stands independently of his infamous transcendental idealism) provides considerable reasons to show that can*not* make such causal judgments legitimately (that is, justifiedly). Some of these reasons are presented in the assigned reading for this seminar. Further such reasons will be presented for discussion in the seminar, drawing on Kant's powerful semantics of cognitive reference (which also stands independently of his transcendental idealism). The aim of this s

Reading: Westphal, Kenneth R., 'Kant's Critique of Determinism in Empirical Psychology'. In: H. Robinson, ed., *Proceedings of the 8th International Kant Congress* (Milwaukee: Marquette University Press, 1995), II.1: 357–70. (PDF)

Seminar 2

Interpreting and Misinterpreting Newton's Definition of Accelerative Force.

Despite certain cautions and precautions Newton took, he was a realist about gravitational force and this realism was (and remains) justified by his methods, theory and data. Yet Newton's theory has been understood by empiricist philosophers of science from Hume to Bas van Fraassen in non-realist terms. One key focal point of this important issue is this question: To what extent do non-realist interpretations of Newtonian gravitational theory require or rest on interpreting the formula, '$f = ma$' (and likewise *Principia*, Definitions 6–8), as defining *force*, rather than as defining the *quantity* of force (and likewise quantity of matter, quantity of acceleration), as Newton expressly states?

Readings: Excerpts from Max Jammer, *Concepts of Force* and from some forthcoming research by K. R. Westphal. These will be supplied in PDF format.

Seminar 3

Newton's Realism about Gravitation and his Rule Four of Philosophy.

Recent scholarship, especially that of William Harper, has shown that Newton was significantly more sophisticated about scientific method than contemporary philosophers of science, that his standards of theoretical adequacy justified his realism about gravitational force and that these standards *do* apply to the shift from Newtonian to Einsteinian mechanics. This seminar shall focus on one key feature of Newton's method: his Rule Four of Philosophy, which dismisses mere logical possibilities as illegitimate, pseudo-scientific, merely nominal 'hypotheses'. This rule has significant implications for empiricist, non- or anti-realist interpretations of Newton's gravitational theory.

Reading: William Harper, 'Newton's Argument for Universal Gravitation', in I. B. Cohen & G. Smith, eds., *The Cambridge Companion to Newton* (Cambridge: Cambridge University Press, 2002), pp. 174–201.

Levels of causality and the interpretation of probability (Federica Russo & Jon Williamson)

This series of seminars is devoted to the levels of causality and the interpretation of probability. There are various ways of laying dawn the problem. One is that causation is a relation between tokens ("*my* taking an aspirin relieves *my* headache") or between types ("taking aspirin relieves headaches"). Typically, relations at both levels are expressed in probabilistic terms. Questions concern the relations between the two levels and how probability is or ought to be interpreted. Another facet of the problem is that causal modellers aim to draw causal conclusions either from individual-level data or from population-level data. Questions then concern whether and to what extent multilevel modelling succeeds in accounting for different levels of causation and for the interrelations between them.

Seminar 1

The first seminar has two goals. First, I present and discuss metaphysical accounts of the levels of causation (e.g., Eells'); I will try to show that this is a fallacious way of reasoning about the levels of causation and suggest a better *epistemological* understanding of the levels of causation. Second, I will investigate how to connect the levels and discuss a possible reformulation of Sober's 'Connecting Principle'. Attendees will be required to read the paper: Sober, E. (1986). Causal factors, causal influence, causal explanation. *Proceedings of Aristotelian Society*, 60, 97-136.
http://www.jstor.org.www.sia.ucl.ac.be:888/stable/pdfplus/4106899.pdf

Seminar 2

The second seminar provides a brief introduction to multilevel modelling and to the so-called atomistic and ecological fallacies. I will point out that the philosophers' problem of the levels of causation slightly differs from the (social) scientists' problem. Finally, I will argue that, within the context of multilevel modelling, the problem of the levels of causation is better understood and dealt with if reformulated as a problem of levels of analysis. Attendees will be require to read the paper: Courgeau, D. (2003) (Ed.). Introduction. *Methodology and epistemology of multilevel modelling. Approaches from different social sciences*. Dordrecht: Kluwer.

http://www.courgeau.com/articles/Intro%20Methodos.pdf

Seminar 3

The third seminar will be a reading group around the paper: Máire Ní Bhrolcháin and Tim Dyson (2007) 'Causation in demography: issues and illustrations', *Population and development review*, 33(1), 1-36.

Abstract. We address the issue of causation in demography, focusing. particularly on change through time in aggregate demographic phenomena. We regard the intervention framework associated with the Neyman-Rubin-Holland model of causal estimation as too rarely applicable to be useful as a general solution to causal inference in demography. Discussion centres on a series of examples of demographic change ranging from mechanical demographic processes, where causation is known to occur, through to relatively straightforward cases where causation is not, or only rarely, contested to more complex demographic phenomena. We develop and propose ten criteria, or guidelines, intended to assist in the practical evaluation of causation in the context of demographic change, analogous to those in use in epidemiology.

Machine learning and causality (David Corfield & Alex Freitas)

Abstract:

This seminar series will be based on the Causality Challenge #1:
Causation and Prediction, which is part of the Pascal challenges and of the WCCI (World Congress on Computational Intelligence) 2008 competition program. This challenge involves feature selection in the challenging supervised machine learning (classification) scenario of "dataset shift". In this scenario the training and test sets (used, respectively, to build a classification model and evaluate its generalization ability, or predictive accuracy) are not necessarily identically distributed, which makes the discovery of both causal and predictive models considerably harder, by comparison with the conventional machine learning scenario where the training and test sets are assumed to be identically distributed. In this seminar series we will discuss the definition of this causality challenge, the relevant background concepts and methods, and the results of this causality competition.

Seminar 1

For the first seminar, there is an on-line GoogleTech video entitled "Challenges in Causality": http://www.youtube.com/watch?v=eL7gO-rUlDY
We will ask the delegates to listen to this video (no need to read papers) before the first one-hour slot. So, the contents of the first seminar will be: Discussion of the above video on Challenges in Causality, in the context of dataset shift in machine learning, where training and test sets are not necessarily identically distributed.

Seminar 2

For the second seminar, we will ask the delegates to read, before the seminar, the following tutorial/technical report: I. Guyon, C. Aliferis, A. Elisseeff. Causal Feature Selection. Technical Report. March 2, 2007. 40 pages. http://www.clopinet.com/isabelle/Papers/causalFS.pdf
So, the contents of the second seminar will be: Discussion on the above tutorial, containing the relevant background concepts and methods for a deeper understanding of the causality challenge.

Seminar 3

The third seminar will be a discussion about the results of the Causality Challenge that will be held in the WCCI 2008 conference in Hong Kong this early June (the above video is precisely about this challenge). There should be a paper discussing those results but we will not have the details of that until middle June at least. So, for now what we can say is that the contents of the third seminar will be:
Discussion of the results of the Causality Challenge presented at the WCCI 2008 conference, based on a paper to be indicated later.

Mechanisms and causality (Phyllis McKay & Jon Williamson)

Abstract

As the importance of mechanisms to science is recognized, a proper understanding of mechanisms is increasingly important to the metaphysics and philosophy of science. There are various characterisations of mechanisms available in the literature, of which the most influential is Machamer, Darden and Craver's. In this seminar series, we will become familiar with their views, and the controversy they have generated over what a mechanism is. We use the final seminar to consider how thinking about mechanisms might illuminate older debates in the philosophy of science – in this case the debate about the unity of science.

Seminar 1

Reading group around:
Machamer, Darden and Craver: 'Thinking about Mechanisms in *Philosophy of Science* Vol. 67, No. 1. (Mar., 2000), pp. 1-25.

This paper is the *locus classicus* of Machamer, Darden and Craver's controversial and influential claim: 'Mechanisms are entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions.' (p3.) Their work spurs much modern work on mechanisms. We will discuss this paper, and issues arising from it in the debate about what a mechanism is.

Seminar 2

Reading group around:

James Tabery ‘Synthesizing activities and interactions in the concept of a mechanism’ in *Philosophy of Science* 71 (January 2004) pp1-15.

Tabery examines the disagreement over MDC’s controversial introduction of the notion of an *activity* in understanding mechanisms, and attempts to reconcile their position with the position of Stuart Glennan, an opponent who argues that the notion of a mechanism depends in a fundamental way on laws of nature.

Seminar 3

Reading group around:
Craver: ‘Beyond reduction: mechanisms, multifield integration and the unity of neuroscience’ in *Studies in the History and Philosophy of the Biological and Biomedical Sciences* 36 (2005) 373-395.

Craver argues that a common search for mechanisms can help us understand the increasing unity of different fields in a way that owes nothing to reduction – and owes some of its persuasive success to *maintaining* a diversity of methods.