Causality and Probability in the Sciences

Canterbury 14th-16th June 2006

www.kent.ac.uk/secl/philosophy/jw/2006/capits.htm

Conference booklet

Organizers:
Federica Russo and Jon Williamson*

June 12, 2006

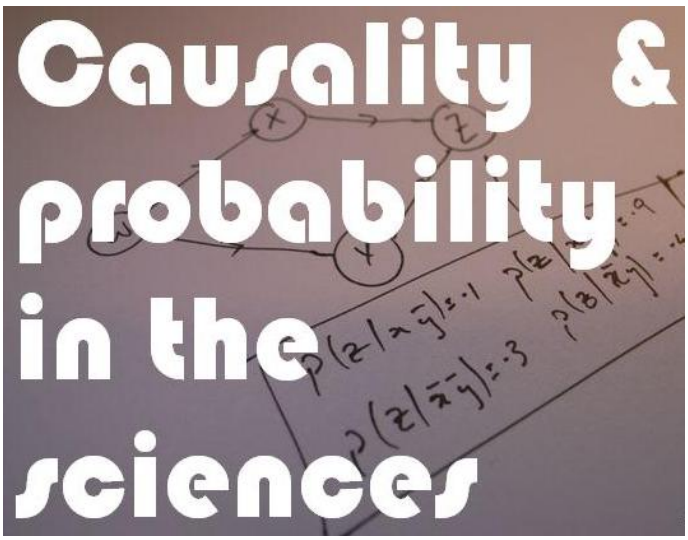

*We gratefully acknowledge the British Academy, the British Society for Philosophy of Science the Mind association and the Kent Institute for Advanced Study in the Humanities for financial support.

Programme

Wednesday 14th June
13.30 Registration
13.50 Welcome
14.00-15.00 Invited speaker
Nancy Cartwright *Causality: what can be the good of it?*
15.00-16.00
Robert Northcott *Natural-born determinists: causation in a probabilistic world*
Monika Koeppl *Where causal dualism comes from*
Coffee
16.30-18.00
Bert Leuridan *Galton's blinding glasses. Modern statistics hiding causal structure in early theories of inheritance*
Pawel Kawalec *Causality and explanation. A procedural approach to causal inference*
Marianne Belis *The relation between causality and probability*
Thursday 15th June
09.00-10.00 Invited speaker
Margherita Benzi *Contexts for causal models*
10.00-11.00
Phyllis Mckay *What do mechanisms do?*
Franoise Longy *The case of the probability of dysfunction*
Coffee
11.30-13.00
Nuala Sheehan *Causal inference using instrumental variables in an epidemiological application*
Federica Russo & Jon Williamson *Interpreting probability in causal models for cancer*
Aviezer Tucker *A naturalized account of the inference of common causes*
Lunch
14.00-16.00
Isabelle Drouet *Causal inference: what is new with Bayes nets?*
Alex Freitas & Ken McGarry *Integrating Bayesian networks and Simpson's paradox in data mining*
Andrea L'Episcopo *Causality and axiomatic probability calculus*
Jan Lemeire *Causal models as minimal descriptions of multivariate systems*
Coffee
16.30-17.00 Sam Maes *Causal inference in graphical models with latent variables. From theory to practice*
17.00-18.00 Invited speaker
Kevin Korb *Informative interventions*
19.30 Conference Dinner
Friday 16th June
09.00-10.00 Invited speaker
Mauricio Suárez *Causal inference in quantum mechanics: a reassessment*

10.00-11.00
Friedel Weinert *A conditional view of causality*
Erik Weber *Conceptual tools for causal analysis in the social sciences*
Coffee
11.30-13.00
Damien Fennell *Causality, mechanisms and modularity: structural models in econometrics*
Julian Reiss *Time series, invariance, and the principle of common cause*
Alessio Moneta *Mediating between causes and probability: the use of graphical models in econometrics*
Lunch
14.00-14.30
Amit Pundik *The causal relata in the legal context*
14.30-15.30 Invited speaker
Philip Dawid *Causality: counterfactual and hypothetical*
Coffee

Titles and Abstracts

Nancy Cartwright
Causality: what can be the good of it?

We don't have the abstract yet

Robert Northcott
Natural-born determinists: causation in a probabilistic world

The central idea of probabilistic causality is that a cause is something that raises the probability of its effect. It is clear immediately that two kinds of counterexample are possible, threatening respectively this criterion's necessity and sufficiency: first, when a cause lowers the probability of its effect; and second, when the probability of an effect is raised by a non-cause.

I propose an account that tracks successfully our causal judgments in a range of problem cases of both these kinds - but at an interesting price. That price is that our causal judgment must be taken to be deterministic. That is, even though the world itself may be indeterministic, according to our judgment causation isn't. Instead, we deem something a cause in all-or-nothing fashion.

I argue that in probabilistic cases we judge only in hindsight - deeming something a cause depending only on whether its chancy effect actually did subsequently occur. So far as our judgment is concerned, causation is therefore not an intrinsic property, since it depends on temporally distant events.

Probabilistic causality has come to seem indispensable to scientific practice. A key virtue of this paper's view is that it turns out to be readily reconcilable with that. To see why, we need to track carefully the connection between type and token cases, epistemic constraints, and ex ante and ex post probabilities. In particular, on my view probabilistic causal claims are characteristic of type and ex ante token cases, and in both those sorts of cases the source of indeterminism is merely epistemic. That is, claims about apparently probabilistic causes in fact reflect merely epistemic uncertainty regarding a putative cause that would itself be certain.

Among other things, this account creates a dilemma for a certain brand of realist about probabilistic causation. In opposition to Hume, many have argued for causal realism on the grounds that we can perceive causation as well as we can perceive anything else, denying that we should privilege metaphysically a Humean base of occurrent events. But on the view of this paper, the only causes we perceive are non-probabilistic. Therefore privileging perception over Humean theoretical considerations buys causal realism only at the price of disallowing probabilistic causal realism.

Finally, this deterministic approach seems to leave open only two possibilities. Either, first, that while indeterministic processes may well exist in nature, nevertheless causation remains strictly a deterministic affair. Or else, second,

that causal judgment and causation itself diverge radically in probabilistic environments, and hence that typical philosophical examples - which trade on our causal judgments - are useless as test cases.

Monika Koeppl
Where causal dualism comes from

In view of abundant philosophical accounts of causation, causal pluralism has become a fashionable position. Besides a number of definitional issues, the question is, however, whether causal pluralism is a consistent position in itself.

I concentrate on one particularly common version, causal dualism, which has recently been defended by Hall (2004). He proposes to distinguish between causes as being connected to their effects and causes as making a difference to their effects.

I consider three ways of interpreting causal dualism. One option is to take it as a mere psychological statement about our way of thinking about causation. A second option is to take it as a statement about two kinds of causal relationships in the world. Although both options yield a consistent claim, they are founded on strong assumptions.

My main thesis is that there is a third way of explaining causal dualism. The two senses represent different methodological approaches towards an analysis of causation and are based on conflicting epistemological convictions. As both approaches cannot be defended at once, causal dualism is inconsistent.

Bert Leuridan
Galton's blinding glasses. Modern statistics hiding causal structure in early theories of inheritance

Nowadays it is commonplace to assign probability theory and statis- tical inference a central place in the philosophy of causality. Especially in sciences like economy, biology or psychology, techniques like structural equa- tion modeling provide a promising tool for uncovering causal relations. It is also commonplace that one should not be overly enthusiastic. Causal infer- ence is a tricky business that is heavily dependent on the quality of the data. Most importantly, possibly confounding factors should be incorporated to the largest extent. So even if statistical inference is the royal road to causal inference, it takes good climbing boots to walk it.

I want to focus here on a related problem. Theory can play a decisive role in the interpretation of statistical relations. Depending on one's causal presuppositions, one can succeed or fail to recognize causal patterns in a set of data. Having a good theory is of central importance for the working scientist. So even if statistical inference is the royal road to causal inference, it takes good climbing boots to walk it and a good map to follow the right direction.

How can the working scientist be sure that he has a good theory? A circular answer is that he should make sure that his theory ts the 'well established'

empirical/statistical relations. To show that this circularity is not merely a philosopher's witticism but poses real problems in scientic practice, I will present a case study from biology.

In the second half of the 19th century, when words like 'gene' or 'genetics' did not yet exist and Mendel's theory had sunk into oblivion for decennia, Francis Galton developed his own, very inuential theory of inheritance. Most typical in his work was the use of modern statistical techniques (e.g. linear regression), some of which he himself helped to develop. This contrasts sharply with Mendel's work. Unacquainted with these techniques, Mendel inferred in a rather intuitive way from particulate observations to general regularities. So shouldn't we expect Galton to have found the most 'true' regularities? Of course, every one knows that Mendel is still considered as the founding father of genetics, while Galton's name is now only associated with dubious disciplines such as phrenology and eugenics.

Pawel Kawalec

Causality and explanation. A procedural approach to causal inference

I draw upon the distinction outlined in (Menzies and Price 1993, 196) in setting apart criterial from philosophical issues pertaining to causal inference. My paper is exclusively concerned with the former in the context of causal discovery in non-experimental research. Detailed (historical) studies of the paradigm cases of causal discovery in non-experimental research (J. Snows work on the causes of cholera epidemics in the 19th c. London and R. Dolls, B. Hills and E. Wynders research on the cause of lung cancer in the 1950s) and methodological discussions thereof (esp. D. Cox, J. Goldthorpe, D. Freedman, M. Susser, N. Wermuth) motivate three conditions of adequacy of a causal criterion in non-experimental research: 1) domain-oriented; 2) theoretically informed and 3) unit invariant. The idea of unit invariance roughly: a causal criterion for phenomena in the set V is satisfied independently of which level of analysis is picked up for V is elaborated in terms of mereological operations (Lewis 1991) on the studied phenomena.

I continue with the characterization of the procedural criterion of causal dependence in non-experimental research that satisfies the three of conditions of adequacy. Taking into account the limitations of the graphical models of causal inference (Cartwright 2005; Rubin 2004) I follow J. Pearl, C. Glymour, P. Spirtes, R. Scheines, S. Lauritzen and others in elaborating the procedural criterion on the assumption of the global Markov condition for a DAG on V (Lauritzen 2001). If X is a manipulable variable, then it sets out the minimal intervention field for two variables X and Y in V. What really matters, however, in nonexperimental research is an expansion of the minimal intervention field, which is a Markov blanket modulo the intervention data for X and Y given the unit invariance. Thus, on the procedural criterion, roughly speaking, X and Y are causally related iff the local Markov condition holds in the expanded inter-

vention field for X and Y. The reference to intervention data in the procedural criterion (almost) uniquely indicates which level of analysis should be picked up in a given research to test whether the unit invariance condition holds. I shortly exemplify the application of the procedural criterion using D. Rubins (2004) example and point some of its advantages in causal discovery against the performance of an RCM model and a TETRAD model.

Marianne Belis
The relation between causality and probability

Causality and probability are two models which account for the happening of phenomena. Probability has many formal definitions, not always in accord with one another. The definition of causality and the relation which must exist between causality and probability are still subjects of debate.

Many domains of scientific research and practical activity require statistical data and so benefit from using frequentist probability and probabilistic causality. Other domains and especially everyday life confront singular events which require an objective definition of probability and a specification of its connection with causality.

In this paper I propose an interpretation of probability which considers the singular case to be fundamental for discovering the relation between causality and probability. Cartwright's capacities and Salmon's interactions in the frame of a process are steps toward this goal.

I consider that an objective probability exists in the case of the unique uncertain event and that this probability can be defined in the frame of a causal process. If the process is repeatable, then the probability's value comes out in the long run in the same way in which a signal embedded in noise is recovered through multiple transmissions. Accordingly, repetitive or mass phenomena are beneficial only because they are a sound, and sometimes the only, way to assess the value of objective probability, whose evaluation in the single case is generally difficult. Mass phenomena are not a condition of existence for objective probability.

The analysis of the structure of a causal process leads to a formal definition of the concept of propensity which, as Popper suggested, is very close to the definition of probability. In the paper, I make explicit the relation between the causal structure of a process and the propensity of the changes that the process generates. As a result, probability is derived from causal considerations via propensity. This causal interpretation of probability has the following advantages:

- It illuminates the deep phenomenological relation between causality and probability, providing an objective causal definition of probability for the single case.
- It unifies the various definitions of probability with respect to the basic definition of the "single case", explaining why in the long run, the value obtained by repetition is the same as that of the single case.

- It enables an intuitive evaluation of the subjective probability of a singular event, as required in decision problems under uncertainty and in some non-statistical sciences.

Some practical examples are given in the paper.

Margherita Benzi
Contexts for causal models

I will distinguish between two kinds of philosophical approaches to causal modeling: i) approaches based on the idea of a big omnicomprehensive causal network underlying local networks, and ii) approaches which see causal relations as essentially relative to the context of inquiry. I will analyse how the two approaches can be related to the choice of the interpretation of causality and probability and I will explore the plausibility of the contextual approach.

Phyllis Mckay
What do mechanisms do?

What do mechanisms do? In recent years there has been an explosion in literature on mechanism, in philosophy, but also more importantly across the special sciences. Many scientists, in explaining phenomena as diverse as chemical synthesis, economic growth, rational choices and memory, see what they are doing as looking for causal mechanisms responsible for the phenomena in question. In the social sciences, mechanisms have been seized on as a way of making social science more scientific by providing a midway between producing(implausible) scientific laws, and just telling stories.

Despite this broad interest in mechanisms, there is no consensus about what a mechanism is, with work done on the question producing a startling variety of different candidate definitions. In this paper I investigate how different disciplines see the search for mechanisms, which problems thinking about mechanisms is supposed to solve, and how the underlying conceptions of a mechanism can be demarcated (including: idealised vs. concrete, modular vs. non-modular).

Francoise Longy
The case of the probability of dysfunction

I intend to investigate whether we are not missing a distinction to account for scientifically relevant probabilities. I will explore this question considering the probability of dysfunction of a determinate artefact. I will show that relative to a particular artefact there are two different probabilities of having the capacity to do F in some definite circumstances, that are causally grounded and which are equally required for scientific explanations. The first one, rooted in the physical nature of the artefact, is the one which will be usually interpreted as a propensity, insofar as one accepts of propensities. The second probability is

rooted in the fact that an artefact belongs to what some authors like Ruth Millikan, call a real kind. It is causally grounded in the features of the production processes of manufactured artefacts. These two probabilities, I will argue, are equally objective: there is no ground that would justify to interpret the one in epistemic or subjective terms but not the other. One could certainly interpret these two probabilities objectively in term of frequencies, but it would obliterate their difference with other objective probabilities referring to frequencies in population that have no such causal grounding and not the same value for scientific explanations. For these reasons, a realist interpretation in terms of propensity may appear better, but than one is confronted with something looking paradoxical : an item having two different propensities of doing the same thing in exactly the same circumstances. So I will conclude, that we might need a new distinction for objective probabilities.

Nuala Sheehan & Vanessa Didelez

Causal inference using instrumental variables in an epidemiological application

In epidemiological research, the causal effect of a potentially modifiable phenotype or exposure on a particular outcome or disease is often of public health interest. Randomised controlled trials to investigate this effect are not always possible and inferences from observational data can be distorted in the presence of confounders affecting both phenotype and disease. However, it can often be argued that a known gene, with an indirect effect on the disease via its effect on the phenotype, can be reasonably assumed not to be itself associated with any confounding factors - a phenomenon called *Mendelian randomisation* (Davey Smith & Ebrahim 2003, Katan 2004). It is well known in the economics and causal literature that these properties define an instrumental variable and permit estimation of the causal effect, despite the confounding, under certain model restrictions.

We present a formal framework for causal inference in epidemiological applications where the causal effect is defined as the effect of an intervention. We focus in particular on the strength of the assumptions required to test for and estimate the causal effect of an intermediate phenotype on a disease. We demonstrate how these core assumptions can be suitably represented and hence veri

ed visually using directed acyclic graphs. In particular, this framework allows us to address limitations of the Mendelian randomisation technique that have often been overlooked in the medical literature (Didelez & Sheehan 2005). We stress the importance of having a formal language for causal inference in these applications. Without such a framework, it is difficult to disentangle association concepts from causal concepts and mistakes are more easily made.

References Davey Smith, G. & Ebrahim, S. (2003), Mendelian randomization: can genetic epidemiology contribute to understanding environmental de-

terminants of disease?, International Journal of Epidemiology 32, 1-22.

Didelez, V. & Sheehan, N. A. (2005), Mendelian randomisation and instrumental variables: what can and what can't be done, Technical Report 05-02, Department of Health Sciences, University of Leicester. (http://www.homepages.ucl.ac.uk/ucakvdi/vlon.html).

Katan, M. (2004), Commentary: Mendelian randomization, 18 years on, International Journal of Epidemiology 33, 10-11.

Federica Russo & Jon Williamson

Interpreting probability in causal models for cancer

Whilst it might seem uncontroversial that the health sciences search for causes that is, for causes of disease and for effective treatmentsthe causal perspective is less obvious in the social sciences, perhaps because it is apparently harder to glean general laws in the social sciences than in other sciences. Thus the search for causes in the social sciences is often perceived to be a vain enterprise and it is often thought that social studies merely describe the phenomena.

On the other hand an explicit causal perspective can already be found in pioneering works of Adolphe Quetelet and Emile Durkheim in demography and sociology respectively, and the social sciences have taken a significant step in quantitative causal analysis by following Sewall Wrights path analysis, which was first applied in population genetics. Subsequent developments of path analysise.g. structural models, covariance structure models and multilevel analysishave the merit of making the concept of cause operational by introducing causal relations into the framework of statistical modelling. However, these developments in causal modelling leave a number of conceptual issues unanswered: for instance the question of how probability should be interpreted in probabilistic causal models.

In the philosophy of probability many interpretations have been proposed and crucial objections raised. For instance, it has been argued that the frequency interpretation does not make sense in the single-case, that subjectivist accounts lead to a too strong arbitrariness in probability assignments, and that logical interpretations, though suited to gambling situations, are of scarce applicability in science.

In this paper we raise the problem of the interpretation of probability within a specific context: causal models in cancer epidemiology. This is motivated by the thought that competing interpretations are not right or wrong, but that they are better or worse suited to particular contexts and the demands we make of them. To this end, we first introduce causal analysis in the social and health sciences and then present the case of cancer epidemiology in some detail. We pay particular attention to explaining different possible meanings of probabilistic statements in this context and the importance of choosing one interpretation of probability over another. We then argue that any satisfactory interpretation of probability should satisfy five desiderata; this narrows down the choice to the frequency interpretation twinned with an empirically-based subjective interpretation or an objective Bayesian interpretation. We go on

to argue that the probabilities in causal models in cancer epidemiology should be given a frequency-cum-objective-Bayesian interpretation; the main reason for this choice is the need to cope with two different types of probabilistic inference, population-level and individual-level.

Aviezer Tucker
A naturalized account of the inference of common causes

This is an analysis of how do biologists, historians, comparative linguists, and textual critics infer common causes. The inference that there was some common cause is distinct of and prior to the inference of the properties of a particular common cause. The inference of common cause types is common in science and does not require a particular philosophical analysis. The inference of common cause tokens is special to the sciences that attempt to reconstruct aspects of the past. Some common cause token is inferred from a variational group whose members share information preserving properties, through the comparison of the likelihoods of the variational group given some common cause and given separate causes. The choice of one of the two alternatives is made usually by proving the negligible likelihood of the variational group given the alternative. Three consecutive comparisons of likelihoods determine particular common cause hypotheses: between some common cause token and separate causes; between five kinds of causal nets involving some common cause token, and between concrete common cause hypotheses that specify their properties.

I examine critically the tradition of inference of common cause from Reichenbach through Salmon to Hitchcock. I propose that Sober's Bayesian approach may form a better foundation for solving the problem of inference of common causes, though it requires some adjustment to fit the paradigmatic cases of inference of common cause in science.

Isabelle Drouet
Causal inference: what is new with Bayes nets?

Since the beginning of the 1990s, several causal inference algorithms based on Bayesian networks have been introduced. They have been extensively discussed in numerous papers, ranging from the most enthusiastic support to the harshest criticism. For a few years now, the debate has focused on the assumptions conveyed by these algorithms concerning the relation between causality and probability - in particular on the now famous "Causal Markov Condition". This debate is not over, but it now seems possible to assert that these assumptions are true for at least some causal systems.

The present paper takes these systems as its subject matter, and aims to provide a non polemical assessment of the "net contribution" of Bayes nets algorithms to the resolution of the causal inference problem in those particular cases. This assessment is made up of three analyses.

First I ponder whether and how Bayes nets algorithms differ from more traditional causal inference techniques. The focus is on causal inference methods

resorting to path analysis. Indeed they share with Bayes nets algorithms two of their essential features: they do not require experimental data in order to output causal information (which explains the fact that they are used mainly in social sciences) and they rely on the identification of conditional independencies in the discrete case and of vanishing partial correlations in the continuous one. I show that Bayes nets algorithms give better general results and highlight some of the conceptual and computational reasons for this superiority. The approach is mainly methodological, and I contend that most of the differences between Bayes nets algorithms and anterior methods can be understood in the light of the distinction between induction and hypothetico-deduction. To finish with, I some consider major objections that have been raised against Bayes nets algorithms and which do not concern the assumed relationship between causality and probability. I explain that they are not specific to these algorithms and already concern the rival methods I have considered.

The preceding analysis only deals with systems in which the Bayes nets assumptions for the connection between causal and probabilistic relations are true. In the second section of the paper, I come back to the arguments that have been given in favour of those assumptions. As already explained, my aim is not to discuss those arguments or to propose new ones, but rather to draw attention to their design. I identify two major kinds among the important arguments that have been proposed to support Bayes nets assumptions. According to the first kind of arguments, Bayes nets assumptions should not be viewed as substantive assertions about the systems under consideration, but as conventions dealing with the analysis grain and which can be satisfied for any system. The second kind of arguments consists of proofs of the assumptions for systems represented by models displaying certain properties. In both cases, the contention is that Bayes nets assumptions should not be rejected because they are sometimes true. I show that this contention is of no use in the inferential context we are interested in. The reason for this is roughly that so far we have no way to determine if Bayes nets assumptions are true of systems for which we may want to infer causes using Bayes nets algorithms.

The second section establishes that Bayes nets algorithms cannot be practically used to infer causes from observational statistical data. The conclusion seems to be that, despite the interesting characteristics outlined in the first section, Bayes nets algorithms do not contribute to the resolution of the causal inference problem. The third section qualifies this conclusion. I make clear that the criticisms developed in the second section deal with the sole idea of inferring causal relations from statistical data using Bayes nets algorithms exclusively. As a consequence, it must be possible to take the interesting characteristics of these algorithms into account in a slightly different methodological context. Finally I plead for a methodologically mixed approach to causal inference and explain how one can take advantage of the interesting characteristics of Bayes nets algorithms in this context.

Alex Freitas & Ken McGarry

Integrating Bayesian networks and Simpson's paradox in data mining

Data Mining consists of a set of principles and methods for discovering interesting patterns in real-world data sets [1]. The discovered patterns should be not only accurate, but also comprehensible and surprising (novel, unexpected) to the user. One kind of data mining method consists of constructing Bayesian Networks from data. Bayesian networks have the advantage of providing a graphical, easy-to-interpret representation of the structure of the important relationships in the data [2]. They also have the potential to represent causal relationships. Another kind of data mining involves the detection of Simpson's paradox, a surprising phenomenon in data analysis, as follows. Let the event C be the apparent "cause" of an event E, the "effect". Simpson's paradox occurs if the event C increases the probability of the event E in a given population Pop and, at the same time, decreases the probability of event E in every subpopulation of Pop [3]. Let Z and $\neg Z$ denote two complementary values of a confounding variable, representing complementary properties describing two subpopulations of Pop. Then, mathematically Simpson's paradox occurs if the following 3 inequalities hold: $P(E|C) > P(E|\neg C)$, $P(E|C,Z) < P(E|\neg C,Z)$, $P(E|C,\neg Z) < P(E|\neg C,\neg Z)$, where $P(X|Y)$ is the probability of X given Y. Although Simpson's paradox is well-known by statisticians, it is usually very surprising to data mining users, who typically have no statistical training. This makes the automatic detection of Simpson's paradox one of the few objective, data-driven methods explicitly designed for discovering patterns that are very likely to be considered truly surprising according to a user's subjective evaluation [4]. This work proposes to integrate a Bayesian Network's ability to represent comprehensible and potentially causal patterns with the detection of Simpson's paradox, a data mining method explicitly designed to discover surprising patterns. The basic ideas of this integration are as follows.

In general, Bayesian networks are Independence-maps (I-maps) of the true probability distribution [3]. This means that every independence between variables represented in the network corresponds to an actual independence in the true probability distribution, but the converse is not true, i.e., dependences between variables represented in the network are not guaranteed to correspond to actual dependences in the true probability distribution. Therefore, it is important to develop new methods to identify spurious dependences, i.e., dependences that are represented in a Bayesian network but do not correspond to actual dependences in the true probability model. A major type of spurious dependence is the apparent dependence between events C and E (apparent cause and effect) when in reality this dependence is due to a confounding event Z [3]. This kind of spurious dependence can be discovered by detecting occurrences of Simpson's paradox, as follows. An occurrence of Simpson's paradox involving a triple of events C, E, Z - with the aforementioned meanings - is evidence that we should not believe that C is a cause of E, which suggests that dependences of the form $C \rightarrow E$ should not be represented in the Bayesian network. To implement this

idea we can run a Simpson's paradox detection method - see [5] - as a kind of "preprocessing step" for the Bayesian network construction algorithm, producing a list of occurrences of the paradox found in the data. This list of paradox occurrences can then be used to modify the network construction algorithms' procedures for generating candidate networks, in such a way that the algorithms will be biased against generating candidate networks representing dependences that are considered apparent (rather than true) according to the list of paradox occurrences. This will effectively prune down the number of candidate networks to be considered, and will also have the useful side effect that any detected occurrence of Simpson's paradox can be returned to the user, as an instance of a surprising patterns discovered in the data, in the spirit of data mining.

Note that it is possible that the data to be mined does not contain any instance of Simpson's paradox. If this is the case, then the just-proposed method will not help to prune down the number of candidate networks. However, even in this case the application of the method can be considered useful, because, if no instance of Simpson's paradox was detected in the data, we would have an increased degree of confidence that the dependences represented in the network are true (rather than spurious) dependences, since the candidate dependences represented in the network would have passed an additional test - i.e., no confounding variable related to the dependence was detected. This additional test complements (not replaces) conventional methods for evaluating Bayesian networks.

References

[1] I.H. Witten and E. Frank. Data Mining: Practical Machine Learning Tools. 2nd Ed. Morgan Kaufmann, 2005.

[2] J. Pearl. Probabilistic Reasoning in Intelligent Systems. Morgan Kaufmann, 1988.

[3] J. Pearl. Causality: models, reasoning and inference. Cambridge, 2000.

[4] K. McGarry. A survey of interestingness measures for knowledge discovery. Knowledge Engineering Review J., 20(1), 39-61. 2005.

[5] C.C. Fabris and A.A. Freitas. Discovering surprising instances of Simpson's paradox in hierarchical multidimensional data. Intl. Journal on Data Warehousing and Mining 2(1), 26-48, Mar 2006.

Andrea L'Episcopo
Causality and axiomatic probability calculus

The aim of this paper is to highlight the epistemological distinction between the concept of causality and axiomatic theories, with particular regard to probability calculus, and between the intuitive and the physical notion of causality. Axiomatic theories are coherent structures with no reference to semantics and probability calculus is not an exception. The intuitive notion of causality, on the other hand, is just a semantical matter: saying that A causes B is equivalent to say that there is some kind of relation which links the two variables and this relation is causal. But what does it mean, for two variables, to be causally related? We can search for an answer by a conceptual analysis, or by an empirical

analysis, and the kind of answer will be very different. I will turn my attention to one of the most interesting theories of causation which has been proposed in the last years, Phil Dowes. Then, partly following, partly opposing Dowes theses, I will show how the seeming impossibility to reduce causality to probability is not simply a failure of current probabilistic theories. Quite to the contrary, this impossibility suggests a deep epistemological divide between Kolmogorovs axiomatization of probability calculus, a branch of measure theory, and the intuitive notion of causality. If causality is reduced to probability, one must renounce to many of the features which make causality desirable in scientific research. This is what has actually happened after Kolmogorovs axiomatization of probability: the notion of probability, once turned into a mathematical device, loses some of its desirable features, but it earns in clearness and, above all, in scientific applicability. In the first section of the work, I deal with some of the fundamental features of Kolmogorovs axiomatization of probability calculus, and with their repercussions on the intuitive notion of probability. In the second section, I briefly explore the intuitive and the physical notions of causality. In the third section, I review Dowes theses and list some evidence in support of my own theses 1) it is not possible to reduce the intuitive notion of causality to some axiomatic theory, i.e., probability calculus; 2) a partial reduction is not desirable, because it would lead to an abstract, less heuristic-valued notion of causality; 3) causality has to be seen as a fundamental heuristic device, which can help us in interpreting, but not in representing, contexts of interest which are presented in the conclusions.

Jan Lemeire

Causal models as minimal descriptions of multivariate systems

By applying the Minimum Description Length (MDL) approach for model selection, one should seek the model that describes the data by a code of minimal length. Learning is viewed as data compression that exploits the regularities found in the data. The theory of causal modeling can be interpreted by this approach. If existent and in the absence of other regularities, a faithful Bayesian Network offers the minimal description of a joint probability distribution. The Bayesian model consists of a Directed Acyclic Graph (DAG) and the Conditional Probability Distributions (CPD) of each node conditioned on its parents. The distribution is decomposed into atomic blocks which are able to explain all relational regularities found in the data. This reductionist approach motivates, but does not guarantee, the correctness of the causal interpretation that is attributed to the edges. Local structure of the CPDs allow further compression of the model. Particular regularities, however, generate conditional independencies that make the model unfaithful, such as deterministic relations or correspondence of the influences among different paths. These regularities should be incorporated into the model to insure minimality and the correctness of the inferences made with model. Regularities allowing compression of the

DAG or regularities among the CPDs challenge the minimality of the description, such that causality does not provide correct models. The assumptions that are made for the existence of faithful models and the correctness of the causal structure learning algorithms can be understood by this viewpoint.

Sam Maes

Causal inference in graphical models with latent variables. From theory to practice

Joint work with Stijn Meganck and Philippe Leray

This talk discusses causal inference in graphical mod- els that can handle latent variables without explicitly modelling them quantitatively. For such problem do- mains several approaches that originated in the ma- chine learning and AI community exist, such as semi- Markovian causal models or maximal ancestral graphs. Applying these techniques to a problem domain con- sists of several steps, typically: structure learning from observational and experimental data, parameter learn- ing, probabilistic inference, and, quantitative causal inference.

A problem is that each of the existing approaches only focus on one or a few of all the steps involved in the process of modelling a problem including latent vari- ables. The goal of this talk is to discuss the integral process from observational and experimental data unto different types of efficient inference. Semi-Markovian causal models (SMCMs) [TP02] are an approach developed by Tian and Pearl and they are specifically suited for performing quantitative causal inference in the presence of latent variables. However, at this time no efficient parametrisation of such models is provided and there are no methods for performing efficient probabilistic inference. Furthermore there are no techniques to learn these models from data, exper- iments or a combination of both.

Maximal ancestral graphs (MAGs) [RS02] are an ap- proach developed by Richardson and Spirtes and they are specifically suited for structure learning from ob- servational data. However, the techniques only learn up to Markov equivalence [ZS05] and provide no clues on which additional experiments to perform in order to obtain the fully oriented causal graph. See [MLM06] for that type of results for Bayesian networks without latent variables. Furthermore, as of yet no discrete parametrisation is provided for MAGs and no tech- niques for probabilistic and causal inference have been developed.

We have chosen to use SMCMs in our work, and given some simplifying assumptions we will combine existing techniques with newly developed methods to provide an integral approach that uses both observational data and experiments in order to learn fully oriented causal models with latent variables. More specifically, we will discuss what type of experiments to perform and how to interpret their results.

Furthermore we have developed an alternative repre- sentation for SMCMs together with a parametrisation for this representation, where the parameters can be earned from data with classical techniques. Finally, we will discuss techniques on how to perform proba- bilistic and quantitative causal inference

in these mod- els.

References

[MLM06] S. Meganck, P. Leray, and B. Manderick. Learning causal bayesian networks from ob- servations and experiments: A decision the- oretic approach. In Modeling Decisions in Artificial Intelligence, LNCS, 2006.

[RS02] T. Richardson and P. Spirtes. Ancestral graph markov models. Technical Report 375, Dept. of Statistics, University of Wash- ington, 2002.

[TP02] J. Tian and J. Pearl. On the identification of causal effects. Technical Report (R-290-L), UCLA C.S. Lab, March 2002.

[ZS05] J. Zhang and P. Spirtes. A characteriza- tion of markov equivalence classes for an- cestral graphical models. Technical Report 168, Dept. of Philosophy, Carnegie-Mellon University, 2005.

Kevin Korb
Informative interventions

Joint work with Erik Nyberg

Skepticism about causal discovery methods is fuelled in part by the fact that observational data alone can radically underdetermine the choice of causal model. Relying upon an additional general assumption such as faithfulness or simplicity is convenient but difficult to justify, since the true model can sometimes appear inferior to a faithful, simpler, but false model (cf. Cartwright, 2001). We develop a mathematical theory of causal interventions, extending earlier results of Korb & Nyberg (2006) and of Spirtes et al. (2000). Our results show that experimental data, together with plausible experimental assumptions, are sufficient to identify the true causal model in nonlinear cases as well as linear cases, provided that the interventions are chosen properly.

Mauricio Suárez
Causal inference in quantum mechanics: A reassessment

I review a few influential arguments against causal accounts of EPR correlations, due to Bas Van Fraassen, Huw Price and Daniel Hausman, and show them all to contain unwarranted assumptions. The Van Fraassen 1982 argument is an influential critique, following Reichenbachs Principle of the Common Cause, of causal models for EPR. I argue that its conditions on causal models are unduly restrictive. Price defends backwards-in-time causal accounts of the EPR correlations, but claims that the asymmetry of causation is a property of our agent-centered causal perspective on the macro-world, lacking at the quantum level. I argue that there is no real reason to suppose Prices perspectival thesis does not go through at the quantum level too. Finally Hausman has argued, following an intervention-based account of causation, that there is no causation in the quantum domain because interventions are inapplicable to the entangled pair of EPR particles. I contest the notion of intervention that is appropriate for quantum mechanics. My arguments lead to an openly neutralist and skeptical position concerning the status of causation in the quantum domain. The

status of causal claims concerning the quantum domain is highly sensitive on the general account of causation presupposed and its details so it follows that it is generally neither the case that causal models for EPR have been all refuted, nor that they are all impossible, nor that they all necessarily lack a fundamental feature of the ordinary concept of causation. Causality remains an optional perspective upon the quantum world, which we might or not want to adopt for entirely pragmatic reasons in analyzing quantum correlation phenomena.

Friedel Weinert
A conditional view of causality

The purpose of this paper is to investigate whether Mackie's *INUS*account of causation can be extended from its usual context of everyday causal relationships into the area of quantum mechanics; and whether it can shed light on Weber's notion of adequate causation in the social sciences. Mackie's *INUS* account will be interpreted more generally as a conditional model of causation. Such a model is concerned with the actual causal conditions, which obtain in particular situations. In its emphasis on actual parameters it differs from various counterfactual models of causation, which often presuppose knowledge of the lawful regularities of the physical world in order to impose constraints on the counterfactuals. This presupposition is striking in Woodward's recent interventionist model of causation.

In quantum mechanics Salmon's requirement of spatio-temporal continuity must fail. Dowe's conserved quantity account does not require spatio-temporal continuity but insistence on conserved quantities replaces a subset of causal conditions for the whole set. Using the example of the Franck-Hertz experiment, which established experimentally the discreteness of quantum energy levels in atomic systems, it will be shown that the conditional account in terms of a cluster of necessary and sufficient conditions delivers the correct answer. That is, the experimental outcome of the experiment clearly suggests that there is a conditional dependence of the consequent on the antecedent conditions; the conditional account shows that this experiment satisfies the conditions of causal relations. Salmon's account, by contrast, must give the wrong answer because it insists on spatio-temporal continuity, which is violated in the experiment. Nor is the Franck-Hertz experiment an example of deterministic causation. Reichenbach already established in the 1920s how a conditional model of causality satisfies the notion of probabilistic causality, which is required to interpret the Franck-Hertz experiment. A cause is to be analysed in terms of a set of antecedent conditions, an effect in terms of a set of consequent conditions. Given the antecedent conditions in the experiment, a certain probability distribution of the consequent conditions is predictable, and observed. The conditional dependence of the consequent conditions on the antecedent conditions is probabilistic in quantum mechanical experiments. But in the limit of the classical mechanics of the space-time trajectories of macro-particles, the conditional dependence becomes deterministic.

An interesting feature in the physical sciences is that the set of antecedent

conditions can be regarded as closed; e.g. the necessary and sufficient conditions can be clearly defined and any further conditions in the background information, B, can knowingly be excluded. In Mackie's terms, the normal running of things in the causal field can be held constant for the purpose of the laboratory experiment ($B = 0$). This epistemic situation contrasts sharply with the causal situation in the social sciences. If C_1, C_2, C_3 represent a cluster of antecedent causal conditions and E represents an effect in the social world, then generally $Pr_1(E|C_1, C_2, C_3; B \neq 0) > Pr_2(E|C_4, C_5, C_6; B \neq 0)$, where $Pr_1 > Pr_2$ captures Weber's notion of adequate causation. Weber's notion of causation can be specified in terms a cluster of necessary and sufficient conditions, which are regarded as the most adequate conditions to account for an effect, like the expected rise in crimes as a consequence of a rise in the young male population. Adequate causation in the social sciences differs from causation in the physical sciences in at least two ways: a) the background conditions, even in laboratory experiments with human subjects, cannot be set to zero ($B \neq 0$); b) while it is questionable whether lawlike regularities exist in the social sciences, trends clearly do exist in the social world. The existence of trends prevents us from speaking of probabilistic causation in the social sciences in a straight-forward sense. Due to these features it is doubtful whether the Salmon-Dowe account or Woodward's interventionist account can adequately capture the sense of adequate causation that seems to be a reality in the social sciences. A nice consequence of the conditional model of causality is its applicability to the social sciences. Although there are significant differences in the sets of antecedent and consequent conditions, depending on whether the conditional model discusses a physical or social event, the model is sufficiently general to be applied to both the natural and the social sciences.

Erik Weber
Conceptual tools for causal analysis in the social sciences

The ideas Wesley Salmon presented in his book *Scientific Explanation and the Causal Structure of the World* (1984) have been criticized and elaborated by many people. Thorough criticisms have been launched by e.g. Philip Kitcher (1989). Mixtures of criticism and elaboration have been offered by Phil Dowe (1992, 1995) and Christopher Read Hitchcock (1995). Last but not least, Salmon himself has published several articles in which his original model is improved (1994, 1997). If one looks at this literature, there is an obvious tendency: Salmons original definition of causal interaction (which uses counterfactuals) is replaced with counterfactual-free definitions. Dowe e.g. defines a causal interaction as an intersection of world lines which involves exchange of a conserved quantity.

This evolution has been very fruitful for the analysis of physical causation. But from the perspective of the social sciences, there is an obvious drawback: the concepts as they were developed are not applicable to the social sciences. The main reason is the reliance on the concept of conserved quantity which makes no sense in the social world.

The first aim of my paper is to elaborate the ideas of Salmon for the social sciences. Like Dowe and others have provided Salmon-inspired models of physical causation, I will develop Salmon-inspired models of social causation. In other words, I will develop concepts of causal interaction and causal process that are applicable to the social world and thus are useful tools for social scientists. My second aim is to develop two additional tools for causal analysis in the social sciences. Salmon assumes that causal interactions and causal processes are all we ever need for causal explanations. In his view, the conservative aspect of causation can be fully captured by the concept of causal process, while the innovative aspect can be fully covered by the concept of causal interaction. I think that the second part is problematic: I will argue that we also need the concepts of spontaneous evolution and permanent causal influence in order to capture the innovative aspect of causation. Of course, I will define these concepts in such a way that their application in the social sciences is straightforward. Here we only give the general idea. Spontaneous evolutions are triggered off by some interaction, but are self-determined: no interactions are needed to keep the evolution going, the constitution of the system itself is responsible for the evolution. Permanent causal influences are responsible for non-self-determined evolutions of systems.

References

Dowe P. (1992), W esley Salmon's Process Theory of Causality and the Conserved Quantity Theory', Philosophy of Science 59, pp. 195-216.

Dowe P. (1995), Causality and Conserved Quantities: a Reply to Salmon', Philosophy of Science 62, pp. 321-333.

Hitchcock C. R. (1995), Salm on on Explanatory Relevance', Philosophy of Science 62, pp. 304-320.

Kitcher P. (1989), Explanatory Unification and the Causal Structure of the World', in P. Kitcher & W.

Salmon (ed.), Scientific Explanation, Minneapolis (Minnesota), University of Minnesota Press, pp. 410-505.

Salmon W . (1984), Scientific Explanation and the Causal Structure of the W orld. Princeton, New Jersey: Princeton University Press.

Salmon W . (1994), Causality without Counterfactuals', Philosophy of Science 61, pp. 297-312.

Salmon W . (1997), Causality and Explanation: a Reply to Two Critiques, Philosophy of Science 64, pp. 461-477.

Damien Fennell

Causality, mechanisms and modularity: structural models in econometrics

Structural equation models (SEMs) play a central role in econometrics. Unlike standard structural modelling approaches, a distinctive aspect of the econometricians' approach is their widespread use of simultaneous SEMS to model the equilibrium relations of dynamic systems. In the more standard non-simultaneous

SEMs, a functional equation is attributed causal content by having the dependent variable denote the effect and the independent variables denote its direct causes. Semantics for such SEMs have been widely developed in the causal Bayes Net tradition, for instance, Spirtes, Glymour and Scheines (1993) and Pearl (2000). Unfortunately, these approaches are not straightforwardly adapted to the simultaneous equation models of econometrics. In contrast, the work of Herbert Simon (1952) does directly address simultaneous equation models but, unfortunately, is not as extensively developed as the Bayes-net semantics. The first part of the paper outlines a mechanistic interpretation of simultaneous SEMs based on Simon's semantics. The second part of the paper then fleshes out what properties these mechanisms must have and asks whether the brief discussions of mechanisms in Pearl (2000) and Woodward (2003) can be useful in developing the concept for non-simultaneous SEMs. In particular, it considers the claims of both Woodward and Pearl of a relationship between modularity and the individuation of mechanisms. The paper concludes that their treatment of mechanisms and modularity is inappropriate for simultaneous SEMs and argues instead for a weaker concept of modularity.

Julian Reiss

Time series, invariance, and the principle of common cause

This paper examines the peculiarities of causal inference in the context of time series, i.e., series of events or observations in which the time ordering matters. Time series have applications in fields as diverse as neurophysiology, climatology, epidemiology, astro- and geophysics and, of course, the social sciences. The aim of this paper is to assess the applicability of two approaches to causal inference that have dominated the literature in recent years, probabilistic and invariance accounts.

A central assumption of probabilistic accounts to causal inference is the Principle of the Common Cause (PCC), which states that if two variables are probabilistically dependent, then either one causes the other or both are joint effects of a common cause (which screens off the correlation)[1]. It is well known that so-called nonsense correlations constitute a problem for the PCC. Since there are a (possibly unlimited) number of different sources of nonsense correlation, it is best to define the term negatively: nonsense correlations are those correlations that do not arise on account of a causal relationship between variables. Common sources include: time series properties such as non-stationarity; logical, conceptual or mathematical connections; and nomological (but non-causal) connections between the variables.

Here I want to focus on specific properties of time series responsible for nonsense correlations. One counterexample to the PPC, which has been widely dis-

[1]There are obvious counterexamples to the phrase in parentheses in this formulation, e.g., when there are more than one common causes or when the causal structure is more complex. In this paper I won't deal with problems regarding the screening-off condition and can thus ignore these intricacies here.

cussed, concerns causally independent time series that monotonically increase: such time series are probabilistically dependent but-ex hypothesis-not causally connected. Against a number of recent papers (e.g., by Kevin Hoover and Daniel Steel, both published in the BJPS) I argue that time series properties such as this constitute a genuine problem for the PCC, which cannot easily be dismissed. First of all, not only monotonically increasing time series produce nonsense correlation but also other kinds of non-stationarity. Further, it can be proved that the same problem can arise even in stationary series of various kinds (albeit to a lesser extent). Third, and most importantly, it is not possible to get rid of the problem with "off-the-shelf" techniques such as differencing the series prior to analysis or tests for co-integration: differencing may result in a loss of long-run information; two series can be co-integrated and yet, not causally connected. Since the overwhelming majority of time series has properties capable of producing nonsense correlations and since they cannot be easily disposed of, the PCC is of little methodological value in the context of time-series analysis. In part responding to these difficulties, econometricians have developed the concept of "co-breaking" for causal inference (e.g., Kevin Hoover, David Hendry), which is an application of the idea that causal relations should be stable under interventions. One common source of non-stationarity consists in shifts in deterministic parameters such as the mean of a series or a trend. Two series are said to be co-breaking when a deterministic shift occurs in both series but a linear combination of the two series is stable. One can show that if the conditional model has both invariant parameters and invariant error variances across regimes, whereas the joint process varies across those regimes, then the reverse regression cannot have invariant parameters and thus determine causal direction.

While co-breaking is a sometimes useful method of inference for non-stationary series when deterministic shifts are frequent (as they are, e.g., in many macro economies), it is not a panacea. First, co-breaking is at best a test for causality, it does not constitute it. Series may be causally related and yet not be invariant under intervention. Second, co-breaking, too, can be spurious: when many variables are considered and/or sample sizes are small spurious co-breaking is almost certain.

The prospects for an inductive research strategy when inferences must be made on the basis of observational data alone thus look pretty grim. This paper thus strongly recommends the elaboration of alternatives that are not based on the two paradigms principle of the common cause and invariance under intervention.

Alessio Moneta

Mediating between causes and probability: the use of graphical models in econometrics

This paper examines the relationship between causes and probability in the particular context of econometrics and discusses a methodological framework

to address the problem of causal inference in econometrics. There are two conceptions of econometrics that have marked its history. The first one considers causes to be something that economic theory must provide and that statistical methods must measure. The second conception considers economic theory to be a not very reliable source of causal knowledge and opens the possibility of inferring causes from statistical properties of the data alone. The first conception was advocated by some exponents of the Cowles Commission during 1950s and is still largely accepted by econometricians who use structural models or calibration methods. The second conception was formalized by Granger's (1969) tests of causality, which are still very popular in nowadays econometrics.

I will argue that these conceptions can be interpreted as two opposite solutions to the problem of under-determination of theoretical causal relations by statistical data (usually named by econometricians the problem of identification). While the risk of the first approach is the commitment to an apriorist strategy, the second approach is hampered by difficulties which are typical of the probabilistic theories of causality, as many studies in the philosophy of science have shown. Econometrics offers a particular and clear example of the general problem of causal inference. I will argue that the general problem of causal inference can be solved only by delicately mediating between background knowledge and statistical properties of the data. The method for this careful handling is in large measure dependent upon the discipline considered.

With respect to macro-econometrics, graphical models, that is the methods for causal inference developed by Pearl (2000) and Spirtes, Glymour and Scheines (2000), can be very useful for this important task of mediating between probabilistic and causal knowledge. Indeed, graphical models permit to take into account the maximum amount of probabilistic information (partial correlations of all possible orders), which can be used to exclude false causal relations. Partial correlations, however, are never sufficient to isolate the unique true causal relations, except in very exceptional circumstances. Indeed, background knowledge has to be incorporated and this approach permits a very efficient use of background causal knowledge. This will be demonstrated by presenting some examples taken from time series models.

Amit Pundik
The causal relata in the legal context

Similarly to the natural sciences, the causal connection in law is usually been taken to hold between entities that exist in the actual world, outside of the courtroom. Whilst an extensive legal debate flourishes around the distinction between different types of causal connections, all sides to the legal debate seem to take the causal relata as actual events, compatible to how they are taken in the natural sciences.

Against this conception, I argue that causation in the legal context is incompatible with causation in the natural sciences. Whilst causation in the natural sciences holds between events (occurrences in the world), causation in the law holds between facts (propositions about the worlds). Hence, the causal connec-

tion in the legal context is between transcendental entities (propositions) that are produced for purpose of the legal proceedings rather than between immanent entities (events) that can be described in spatiotemporal terms, as the natural sciences regard them.

The structure of my argument is as follows. First, I argue that the concept of causation in the legal context can be investigated independently from the investigation about the nature of causation in the natural sciences. Although the questions about the nature of causation are the same, answers that fit the legal context are not necessarily consistent with the answers that fit the natural sciences. Thus, legal causation is committed neither to a general or unified concept of causation nor to the one of the natural sciences. On the other hand, I hold that a single concept of legal causation can be provided for the various legal fields, issues, and cases.

I then propose few alternative lines of argument for why law differs from the natural sciences on the question of the category of the causal relata. I suggest a general argument that recognises a fundamental difference between the scientific inquiry and the legal one. Legal fact-finding merely choose between two competing attempts to describe what happened in the actual world rather than make any proclamation about how the actual world is. As such, legal fact-finding is better understood as deciding between two alleged facts rather than between two actual events.

This general recognition clarifies the problem of synchronous over-determination in the legal context by shifting the focus from the metaphysical aspects to the epistemological ones. The role of the latter is more crucial in the legal context than it is in the natural sciences because courts cannot suspend judgement or defer decisions until enough information will be gathered. Hence, the problem for law is different than the one for the natural sciences, and these two problems should not be mixed. I then suggest two particular arguments. Following Jonathan Shaffer, I address two substantial questions that differentiate between the various theories about the category of the causal relata: immanence and individualisation. Taking position in these two questions can help in identifying the category of the causal relata. I show that contrary to the natural sciences, the legal context dictates certain positions that are consistent with facts as the category rather than events.

In regards to the first question, I argue that any legal concept of causation has to rely on facts because only facts are transcendental. Whilst the natural sciences take the causal relata to be immanent so they can be described in spatiotemporal terms, causal relata in law should be taken as abstract entities that cannot be described in these terms. I establish this conclusion through the following steps:

1. Immanent relata cannot account for omissions (or absences).

2. Omissions must be a possible basis for legal liability.

3. There must be some sort of connection between the omission and the harm.

4. This connection has to be causal.

5. Hence, the relata in the legal context must be transcendental.

Last, I turn to individualisation, and argue that whilst in the natural sciences different resolutions may fit, the legal context requires the most fine-grained option for relata. I then show that this argument from individualisation provides a solution to the problem of asynchronous over-determination in the legal context.

Philip Dawid
Causality: counterfactual and hypothetical

Modern statistical treatments of causality are based on the methodology of "potential responses", that was specifically constructed so as to be able to represent counterfactual reasoning. However the vast majority of applications do not require counterfactual reasoning, but rather hypothetical reasoning... and this can be handled far more simply and safely using standard probability and decision theory. Examples will be given.

Participants List

Marianne Belis
SUPINFO Ecole Supérieure d'Informatique (Paris)
Email: Marianne.BELIS@supinfo.com

Margherita Benzi
University of Genova
Email: benzi@unige.it

Nancy Cartwirght
London School of Economics
Email: N.L.Cartwright@lse.ac.uk

Phil Dawid
University College London
Email: dawid@stats.ucl.ac.uk

Isabelle Drouet
IHPST - Paris 1
Email: i.drouet@9online.fr

Damien Fennell
CPNSS - London School of Economics
Email: D.J.Fennell@lse.ac.uk

Alex Freitas
Computing Department, University of Kent
Email: A.A.Freitas@kent.ac.uk

Donald Gillies
Department of Science and Technology Studies
University College London
Email: donald.gillies@ucl.ac.uk

Laurence Goldstein
Philosphy, University of Kent
Email: L.Goldstein@kent.ac.uk

Colin Johnson
Computing Department, University of Kent
Email: C.G.Johnson@kent.ac.uk

Kiran Kalidindi Computing Department, University of Kent
Email: kk49@kent.ac.uk

Pawel Kawalec
Email: psk5+@pitt.edu

Bob Keim
Comuting Department, University of Kent
Email: R.G.Keim@kent.ac.uk

Monika Koeppl
University of Heidelberg
Email:
Monika.Koeppl@urz.uni-heidelberg.de

Kevin Korb
Monash University
Email:
Kevin.Korb@infotech.monash.edu.au

David Lagnado Department of Psychology,
University College London
Email: d.lagnado@ucl.ac.uk

Jan Lemeire
Vrije Universiteit Brussel
Email:jlemeire@info.vub.ac.be

Andrea L'Episcopo
?
Email: andrea.lepiscopo@libero.it

Bert Leuridan
Ghent University
Email: Bert.Leuridan@ugent.be

Francoise Longy
IHPST - Paris 1
Email: longy@idf.ext.jussieu.fr

Sam Maes
Vrije Universiteit Brussel
Email: sammaes@vub.ac.be

Phyllis Mckay
University of Bristol
Email: Phyllis.McKay@bristol.ac.uk

Alessio Moneta
Laboratory of Economics and Management, Pisa
Email: amoneta@sssup.it

Robert Northcott University of Missouri-St Louis
Email: northcottr@umsl.edu

Amit Pundik
Oxford University
Email: amitpu@yahoo.com

Julian Reiss
University of Madrid
Email: julian.nyc@gmail.com

Federica Russo
Philosophy, University of Kent
Email: f.russo@kent.ac.uk

Nuala Sheehan
University of Leicester
Email: nas11@leicester.ac.uk
Duncan Smith
University of Manchester
Email:
duncan.g.smith@manchester.ac.uk
Mauricio Suárez
University of Madrid
Email: msuarez@filos.ucm.es
Aviezer Tucker Queens University,
Belfast
Email: avitucker@yahoo.com
Erik Weber
Ghent University
Email: Erik.Weber@UGent.be
Friedel Weinert
University of Bradford
Email: F.Weinert@Bradford.ac.uk
Jon Williamson
Philosophy, University of Kent
Email: j.williamson@kent.ac.uk