Causality Across the Levels

British Academy Research Development Award 2009-11

Contributors: Jon Williamson (PI), Federica Russo (RA)

Context. Perhaps the chief concern of the health sciences is that of establishing, evaluating, and using causal claims: e.g., establishing what the causes of cancer are, to what extent a given substance is carcinogenic, establishing the extent to which the environment and social factors affect health, deciding whether chemotherapy is likely to be an effective treatment for a particular patient, etc. These examples suggest that causal inferences fall into two categories. (1) Generic, where data is analysed by means of various statistical models to infer causal relations from large data sets. (2) Single-case, where the question is to combine causal knowledge gathered from data with specific knowledge about a particular individual. Diagnosis, for instance, combines generic causal knowledge of diseases with an individual's symptoms and medical history to come up with a particularised diagnosis.

Currently, the relationship between causality at the generic level and causality at the single-case level is poorly understood. Crucial questions remain open for which this project seeks to provide answers.

Objectives. To develop an account of causality that answers conceptual and methodological questions about the levels of causation. Specifically: Does each level require a distinct notion of cause? If so, how are these notions related? How should one draw inferences about the individual from generic knowledge and vice versa? How can socio-economic variables affect health variables and vice-versa? What are 'aggregate' mechanisms in health contexts? What are mechanisms in which variables of different sorts (biological and social) operate.

Publications:

Phyllis McKay Illari, Federica Russo & Jon Williamson (eds): **Causality in the sciences**, Oxford University Press, [Amazon UK US], 2011.
Introduction:

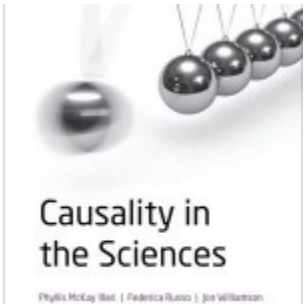

There is a need for integrated thinking about causality, probability and mechanisms in scientific methodology. Causality and probability are long-established central concepts in the sciences, with a corresponding philosophical literature examining their problems. On the other hand, the philosophical literature examining mechanisms is not long-established, and there is no clear idea of how mechanisms relate to causality and probability. But we need some idea if we are to understand causal inference in the sciences: a panoply of disciplines, ranging from epidemiology to biology, from econometrics to physics, routinely make use of probability, statistics, theory and mechanisms to infer causal relationships.

These disciplines have developed very different methods, where causality and probability often seem to have different understandings, and where the mechanisms involved often look very different. This variegated situation raises the question of whether the different sciences are really using different concepts, or whether progress in understanding the tools of causal inference in some sciences can lead to progress in other sciences. The book tackles these questions as well as others concerning the use of causality in the sciences.

Federica Russo & Jon Williamson: **EnviroGenomarkers: the interplay between mechanisms and difference making in establishing causal claims**, *Medicine Studies: International Journal for the History, Philosophy and Ethics of Medicine & Allied Sciences*, 3:249–262, 2012.

According to Russo and Williamson (2007, 2011a,b), in order to establish a causal claim of the form 'C is a cause of E', one needs evidence that there is an underlying mechanism between C and E as well as evidence that C makes a difference to E. This thesis has been used to argue that hierarchies of evidence, as championed by evidence-based movements, tend to give primacy to evidence of difference making over evidence of mechanism, and are flawed because the two sorts of evidence are required and they should be treated on a par.

An alternative approach gives primacy to evidence of mechanism over evidence of difference making. In this paper we argue that this alternative approach is equally flawed, again because both sorts of evidence need to be treated on a par. As an illustration of this parity we explain how scientists working in the EnviroGenomarkers' project constantly make use of the two evidential components in a dynamic and intertwined way. We argue that such an interplay is needed not only for causal assessment but also for policy purposes.

Federica Russo and Jon Williamson: **Epistemic causality and evidence-based medicine**, *History and*

Philosophy of the Life Sciences 33(4):563-582, 2011.

Causal claims in biomedical contexts are ubiquitous albeit that they are not always made explicit. This paper addresses the question of what causal claims mean in the context of disease. It is argued that in medical contexts causality ought to be interpreted according to the epistemic theory. The epistemic theory offers an alternative to traditional accounts that cash out causation either in terms of 'difference-making' relations or in terms of mechanisms. According to the epistemic approach, causal claims tell us about which inferences (e.g., diagnoses and prognoses) are appropriate, rather than about the presence of some physical causal relation analogous to distance or gravitational attraction. It is shown that the epistemic theory has important consequences for medical practice, in particular with regard to the evidencebased causal assessment.

Federica Russo: **Causal webs in epidemiology**, *Paradigmi* XXXIX (1), 67-98, 2011.

The notion of =causal web' emerged in the epidemiological literature in the early Sixties and had to wait until the Nineties for a thorough critical appraisal. Famously, Nancy Krieger argued that such a notion isn't helpful unless we specify what kind of spiders create the webs. This means, according to Krieger, (i) that the role of the spiders is to provide an explanation of the yarns of the web and (ii) that the sought spiders have to be biological and social. This paper contributes to the development of the notion of causal web, elaborating on the two following points: (i) to catch the spiders we need multi-fold evidence—specifically, mechanistic and difference-making—and (ii) for the eco-social to be explanatory, the web has to be mechanistic in a sense to be specified.

Barbara Osimani, Federica Russo and Jon Williamson: **Scientific evidence and the law: a Bayesian formalisation of the precautionary principle in pharmaceutical regulation**, *Journal of Philosophy, Science and Law* 11, 2011;

The paper considers the legal tools that have been developed in German pharmaceutical regulation as a result of the precautionary attitude inaugurated by the Contergan decision (1970). These tools are (i) the notion of "well-founded suspicion", which attenuates the requirements for safety intervention by relaxing the requirement of a proved causal connection between danger and source, and the introduction of (ii) the reversal of proof burden in liability norms. The paper focuses on the first and proposes seeing the precautionary principle as an instance of the requirement that one should maximise expected utility. In order to maximise expected utility certain probabilities are required and it is argued that objective Bayesianism offers the most plausible means to determine the optimal decision in cases where evidence supports diverging choices.

Federica Russo and Jon Williamson: **Generic versus single-case causality: the case of autopsy**, *European Journal for Philosophy of Science* 1(1): 47-69, 2011.

This paper addresses questions about how the levels of causality (generic and single case causality) are related. One question is epistemological: can relationships at one level be evidence for relationships at the other level? We present three kinds of answer to this question, categorised according to whether inference is top-down, bottom-up, or the levels are independent. A second question is metaphysical: can relationships at one level be reduced to relationships at the other level? We present three kinds of answer to this second question, categorised according to whether single-case relations are reduced to generic, generic relations are reduced to single-case, or the levels are independent.

We then explore causal inference in autopsy. This is an interesting case study, we argue, because it refutes all three epistemologies and all three metaphysics. We close by sketching an account of causality that survives autopsy—the epistemic theory.

Lorenzo Casini, Phyllis McKay Illari, Federica Russo and Jon Williamson**: recursive Bayesian networks, *Theoria* 26(1):5-33, 2011.**

The Recursive Bayesian Net (RBN) formalism was originally developed for modelling nested causal relationships. In this paper we argue that the formalism can also be applied to modelling the hierarchical structure of mechanisms. The resulting network contains quantitative information about probabilities, as well as qualitative information about mechanistic structure and causal relations. Since information about probabilities, mechanisms and causal relations is vital for prediction, explanation and control respectively, an RBN can be applied to all these tasks. We show in particular how a simple two-level RBN can be used to model a mechanism in cancer science. The higher level of our model contains variables at the clinical level, while the lower level maps the structure of the cell's mechanism for apoptosis.

Guillaume Wunsch, Federica Russo & Michel Mouchart: **Do we necessarily need longitudinal data to infer causal relations?**, *Bulletin of Sociological Methodology* 106 (1), 5-18, 2010;

It is generally admitted that causes precede their effects in time. This usually justifies the preference for longitudinal studies over cross-sectional ones, because the former allow the modelling of the dynamic process generating the outcome, while the latter cannot. Supporters of the longitudinal view make two interrelated claims: (i) causal inference requires following the same individuals over time, and (ii) no causal inference can be drawn from cross-sectional data. In this paper, we challenge this view and offer counterarguments to both claims. We also argue that the possibility of establishing causal relations does not so much depend upon whether we use longitudinal or cross-sectional data, but rather on whether or not the modelling strategy is structural.

George Darby and Jon Williamson: **Imaging Technology and the Philosophy of Causality**, *Philosophy and Technology* 24(2): 115-136, 2011.

Russo and Williamson (2007) put forward the thesis that, at least in the health sciences, to establish the claim that C is a cause of E one normally needs evidence of an underlying mechanism linking C and E as well as evidence that C makes a difference to E. This epistemological thesis poses a problem for most current analyses of causality which, in virtue of analysing causality in terms of just one of mechanisms or difference making, cannot account for the need for the other kind of evidence. Weber (2009) has suggested to the contrary that Giere's probabilistic analysis of causality survives this criticism. In this paper we respond to Weber's suggestion, arguing that Giere's account does not survive the criticism, and we look in detail at the case of medical imaging technology, which, we argue, supports the thesis of Russo and Williamson (2007).

Guillaume Wunsch, Federica Russo & Michel Mouchart: **Inferring Causality through Counterfactuals in**

Observational Studies. Some epistemological issues, *Bulletin of Sociological Methodology* in press;

This paper contributes to the debate on the virtues and vices of counterfactuals as a basis for causal inference. The goal is to put the counterfactual approach in an epistemological perspective. We discuss a number of issues, ranging from its non-observable basis to the parallelisms drawn between the counterfactual approach in statistics and in philosophy. We argue that the question is not to oppose or to endorse the counterfactual approach as a matter of principle, but to decide what modelling framework is best to adopt depending on the research context.

Federica Russo: **On empirical generalisations**, in D. Dieks, W.J. Gonzalez, S. Hartmann, M. Stoeltzner, M. Weber (eds), Probabilities, Laws, and Structures, Springer, 2012;

Federica Russo: **Public health policy, evidence, and causation: lessons from the studies on obesity**, *Medicine, Health Care and Philosophy*, 2011;

The paper addresses the question of how different types of evidence ought to inform public health policy. By analysing case studies on obesity, the paper draws lessons about the different roles that different types of evidence play in setting up public health policies. More specifically, it is argued that evidence of difference-making supports considerations about 'what works for whom in what circumstances', and that evidence of mechanisms provides information about the 'causal pathways' to intervene upon.

Jon Williamson: **Why Frequentists and Bayesians Need Each Other**, *Erkenntnis* 78:293-318, 2013..

The orthodox view in statistics has it that frequentism and Bayesianism are diametrically opposed – two totally incompatible takes on the problem of statistical inference. This paper argues to the contrary that the two approaches are complementary and need to mesh if probabilistic reasoning is to be carried out correctly.

Jon Williamson: **Mechanistic theories of causality**, *Philosophy Compass* 6(6): 421-447, 2011. Part 1: ; Part II: ; Teaching and learning guide: ; Local combined copy:

Part I of this paper introduces a range of mechanistic theories of causality, including process theories and the complex-systems theories, and some of the problems they face. Part II argues that while there is a decisive case against a purely mechanistic analysis, a viable theory of causality must incorporate mecha-

nisms as an ingredient, and describes one way of providing an analysis of causality which reaps the rewards of the mechanistic approach without succumbing to its pitfalls.

Federica Russo: **Explaining causal modelling. Or, what a causal model ought to explain**, in M. D'Agostino, G. Giorello, F. Laudisa, T. Pievani and C. Sinigaglia (eds), *New Essays in Logic and Philosophy of Science*, SILF Series, Volume I, College Publications, London, pp. 347-361, 2011;

One of the goals of the social sciences is to explain social phenomena. In this explanatory enterprise, causal relations, established by means of the so-called 'causal models', play a major role. This paper advances the view that causal models, by modelling causal mechanisms, (ought to) provide an explanation of social phenomena and should be seen as a model of explanation.

Guillaume Wunsch, Federica Russo & Michel Mouchart: **Inferring causal relations by modelling structures**, *Statistica* in press;

This paper contributes to the debate on the virtues and vices of counterfactuals as a basis for causal inference. The goal is to put the counterfactual approach in an epistemological perspective. We discuss a number of issues, ranging from its non-observable basis to the parallelisms drawn between the counterfactual approach in statistics and in philosophy. We argue that the question is not to oppose or to endorse the counterfactual approach as a matter of principle, but to decide what modelling framework is best to adopt depending on the research context.

Federica Russo: **Correlational data, causal hypotheses, and validity**, *Journal for General Philosophy of Science* 152:167-179, 2011;

The main focus of this paper is the question as to what it is for an individual to think of her environment in terms of a concept of causation, or causal concepts, in contrast to some more primitive ways in which an individual might pick out or register what are in fact causal phenomena. I show how versions of this question arise in the context of two strands of work on causation, represented by Elizabeth Anscombe and Christopher Hitchcock, respectively. I then describe a central type of reasoning that, I suggest, a subject has to be able to engage in, if we are to credit her with causal concepts. I also point out that this type of reasoning turns on the idea of a physical connection between cause and effect, as articulated in recent singularist approaches of causation.

Acknowledgements:

We are very grateful to the British Academy for financial support, and to Lorenzo Casini, David Corfield, George Darby, Phyllis Illari, Jan Lemeire and Bert Leuridan who helped make our regular causality seminars so productive.