

THE RATIONALE OF VARIATION

IN METHODOLOGICAL AND EVIDENTIAL PLURALISM

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

To appear in *Philosophica*

Draft of 14th February 2007

Abstract. Causal analysis in the social sciences takes advantage of a variety of methods and of a multi-fold source of information and evidence. This pluralistic methodology and source of information raises the question of whether we should accordingly have a pluralistic metaphysics and epistemology. This paper focuses on epistemology and argues that a pluralistic methodology and evidence don't entail a pluralistic epistemology. It will be shown that causal models employ a single rationale of testing, based on the notion of variation. Further, I shall argue that this monistic epistemology is also involved in alternative philosophical theories of causation.

1. Introduction

Different social sciences study society from different angles and perspectives. Sociology studies the structure and development of human society, demography studies variations in populations due to mortality, fertility and migration behaviours, economics studies the management of goods and services, epidemiology studies the distribution of disease in human populations and the factors determining that distribution, etc. In spite of these differences, the social sciences share a common objective: to understand, predict and intervene on society. In these three moments of the scientific demarche, *knowledge of causes* is in many cases a necessary ingredient in order to provide an explanation of social phenomena as well as of individual behaviours.

This causalist perspective, however, is not always explicit. Causal vocabulary is sometimes replaced with more vague terms such as factor, determinant, risk, but not *cause* and *effect*. Also, it is said that in spite of the heavy formalism of modern models, the social sciences cannot establish *causal* relations but only make *associational* claims. However, if the social sciences merely *described* phenomena, it would be useless to design policies or prescribe treatments that rely on the results of research. Adopting an explicit causalist stance is motivated by two distinct but related objectives: cognitive and action-oriented. We pursue a *cognitive goal* in detecting causes and thus in gaining general causal knowledge of the causal mechanisms that govern the development of society, and such general causal knowledge is meant to inform and guide social policy, that is we also pursue an *action-oriented goal*.

In the social sciences, causal analysis takes advantage of a variety of methods and of a multi-fold source of information and evidence. In this paper I raise the question of whether such methodological and evidential pluralism also entails epistemological pluralism. In a

nutshell, I shall give a negative answer and argue in favour of a monistic rationale of causality based on the notion of *variation*.

The paper is organised as follows. I first give an overview of methodological and evidential pluralism by presenting different causal models and the variety of types of evidence and of information used in causal analysis, and then spell out the question of whether this form of pluralism entails epistemological pluralism. Afterwards, I present the rationale of variation and support it with methodological arguments; I also offer a taxonomy of variations and discuss some possible objections. Finally, I show that this rationale is consistent with or even adopted in alternative philosophical accounts of causation.

2. Methodological and evidential pluralism

The first developments of quantitative causal analysis in the social sciences are due to the pioneering works of A. Quetelet (1869) and E. Durkheim (1895 and 1897) in demography and sociology respectively. Significant improvements are due to H. Blalock (1964) and Duncan (1975), and since then causal analysis has shown noteworthy progress in the formal methods of analysis. In the following, I shall just give some examples of different methods through which contemporary causal analysis is carried out.

Structural equation models


Arguably, structural equation models (SEM) are the most widespread methodology. Originators of SEM were mainly geneticists, such as S. Wright (1921, 1934), and economists, such as T. Haavelmo (1943, 1944) and T.C. Koopmans (1950). SEM consist of a set of equations, which can be used to determine a causal graph. SEM are designed in order to combine qualitative causal information with statistical data to provide quantitative assessment of cause-effect relationships among variables of interest. Other classes of models, e.g. covariance structure or hierarchical models, rely on SEM, so it is worth spending some time on their structure, assumptions and hypothetico-deductive methodology.¹

To illustrate, Figure 1 represents a canonical econometric model relating price and demand through two equations²:

FIG. 1 Econometric model relating price and demand

¹ Nowadays, in the SEM-literature there isn't an unanimous consensus as to whether structural equations can be given a causal interpretation. However, as Judea Pearl argues at length (Pearl 2000, ch.5), the original interpretation of SEM was eminently causal and it is a trend of contemporary researchers to require extra ingredients for the causal interpretation.

² I borrow this example from Pearl (2000 : 27-28).



$$Q = \beta_1 P + \delta_1 I + \varepsilon_1$$

$$P = \beta_2 Q + \delta_2 W + \varepsilon_2$$

Q is the quantity of household demand for a certain product, P is the unit price of the same product, I is the household income, W is the wage rate for producing the product, ε_1 and ε_2 are the error terms, β and δ are the parameters. The first equation states that demand depends on – or is causally determined by – the unit price of the product and the household income, while the second states that the unit price of the product depends on – or is causally determined by – the demand and the wage rate. In the graph, each variable is represented by a node and relationships between variables are represented by arrows.

The basic idea underlying SEM is that in a system of equations we can test whether variables are interrelated through a set of relationships, by examining the variances and covariances of variables. Sewall Wright, as early as 1934, has taught us to write the covariance of any pair of variables in terms of path coefficients. The path coefficient quantifies the (direct) causal effect of a putative cause X on the putative effect Y ; given the numerical value of the path coefficient β , the structural equation $Y = \beta X + \varepsilon$ claims that a unit increase in X would result in β units increase in Y .

SEM rely upon a number of assumptions, some of which have mere statistical importance (for instance, normality) whereas others have a fundamental bearing on causality (for instance, the non-correlation of errors terms, covariate sufficiency or no confounding). Among causal assumptions the crucial one is the so-called invariance condition or structural stability. This condition states that parameters have to be stable across a large number of interventions or environmental changes. Stability of parameters is usually taken as *the* condition ensuring the causal interpretation of structural equations.

In SEM, the process of model building involves a continuous interaction between background knowledge and a sequence of statistical procedures for elaborating and testing hypotheses. This is the bulk of the hypothetico-deductive methodology (H-D). H-D methodology is a procedure that accounts for data obtained through observations and/or experimentation and that confirms or disconfirms a given causal structure by confrontation with empirical evidence. Empirical testing is performed through two stages:

- (i) prior theorizing of out-of-sample information, including in particular the selection of variables deemed to be of interest, the formulation of a causal hypothesis, *etc.*;
- (ii) iteratively:
 - a. building the statistical model;

b. testing the adequacy between the model and the data to accept the empirical validity or non-validity of the causal hypothesis.

Causal modelling requires accurate knowledge of the causal context: previous studies, well confirmed scientific theories or background knowledge are essential. The causal hypothesis states a hypothesized causal structure to be put forward for empirical testing. Thus, causality is a matter of confirmation, or borrowing the statistical vocabulary, a matter of accepting or rejecting a given hypothetical causal structure. This strategy is hypothetico-deductive because the causal claim is not inferred from the data, as in inductive methods³, but confirmed or disconfirmed in the given causal context and relative to the structural model. Elsewhere (Russo 2005, Russo et al. 2006) I argued in more detail that several elements participate in justifying the causal interpretation of these models, for instance the causal context in which they are built, the specific set of assumptions having causal compass and their peculiar H-D methodology.

Covariance structure models

Covariance structure models (CSM) attempt to explain the relationships among a set of observed variables in terms of a generally smaller number of unobserved or latent variables.⁴ Formally, CSM consist of an analysis of the covariances of the observed variables in two conceptually distinct steps. A *measurement* model links observed variables to unobserved variables and a *structural* model links unobserved variables. In turn, the measurement component of CSM consists of a confirmatory factor model (CFM) explaining the covariations in a set of observed variables in terms of a smaller number of common factors. The idea behind CFM is that, although some variables of theoretical interest cannot be observed directly, information about them can be obtained indirectly from their effects on observed variables.

Thus, CSM consist in the simultaneous specification of the factor model and of the structural model. The task is to explain the interrelationships among the observed variables as indicated by the covariances among them, in terms of the relationships among the unobserved variables used in the structural equations. It is worth noting that what specifies causal relations in CSM is the structural equation model. On the other hand, the measurement model allows us to estimate latent variables from observed variables.

Multilevel models

Recently, multilevel or hierarchical models⁵ are used in a variety of disciplines, ranging from education to demography. Multilevel analysis is a methodology for the analysis of data with complex patterns of variability, the underlying assumption being that data shows a hierarchy that cannot be neglected in the analysis.

The object of a discipline does not straight specify the level of aggregation at which analyses have to be carried out. For instance, economics is interested in the production, distribution and consumption of wealth, however, there is no a priori specification of whether analyses have to concern individuals, markets, firms, or nations. Thus, multilevel analysis recognises the existence of a multiplicity of levels and tries – within the framework of a single

³ An example of an inductive method, that tries to infer causal structures from data, is TETRAD, the algorithm developed by Spirtes, Glymour and Scheines (1993).

⁴ For a very clear introduction see Long (1983).

⁵ Very good introductions to multilevel modelling are those of Goldstein (2003), Snijders and Bosker (2004), and Courgeau (2003) in which epistemological and methodological problems are also discussed in detail.

model – to specify the relations holding among individuals and/or among different levels of aggregation. In other words, this approach recognises that the grouping of individuals introduces an influence of the group on its members, and, conversely, that members have an influence on the group's behaviour. Failure to recognise this twofold source of influence and variability may lead to two types of fallacy: the atomistic fallacy and the ecological fallacy.⁶

Contingency tables

Causal analysis is also performed by means of contingency tables or cross tabs. When variables involved are non-metric, categorical data analysis (CDA) is more often employed instead. CDA has a very long history. It began in the early 1900s, when K. Pearson and U. Yule were debating on measures of associations, and two decades later CDA took advantage of significant contributions by R. Fisher. The first lucid exposition of the use of contingency tables in sociology is due to Boudon and Lazarfeld (1966) and in recent years, clear presentations and further improvements of CDA are available in the works of Hellevik (1984) or Agresti (1996).

A categorical variable is one for which the measurement scale consists of a set of categories. Categorical scales are very often used in the social sciences to measure attitudes and opinions on several issues. Categorical data consist of frequency counts of observations occurring in the response categories. Consider the simplest case, where only two variables X and Y are involved. X has i levels, and Y has j levels, according to the number of categories that X and Y involve. The ij possible combinations of outcomes are then displayed in a rectangular table having i rows and j columns. The cells of the table in fact represent the ij possible outcomes and contain frequency counts of outcomes. Tables thus construed are called contingency tables or crosstabulations. Such ordinary percentage tables can be analysed by adopting an explicit causal framework.

Such a causal framework is offered, for instance, by Hellevik (1984), where an explicit causal terminology is adopted. In the bivariate case, the independent variable Y represents the cause and the dependent variable X represents the effect. Different levels of X are then compared with regard to the proportion having a specific value on Y . The difference in proportion will then measure the degree of association of the two variables and, in this framework, it will be interpreted as the causal effect of Y on X . In contingency tables differences in proportions play the analogue of regression coefficients in SEM, giving highly similar results. Thus, to some extent, the causal framework for contingency tables rests on the same features as SEM, namely on background knowledge, choice of variables, issues of confounding and control, etc.

Data and evidence

By means of these different methodologies, social scientists try to make sense of observations and to infer causal relations between variables of interest with reasonable confidence. Observations, however, first have to be collected. Data comes from a variety of

⁶ The motivation for developing hierarchical models lies in the so-called ecological fallacy, firstly recognised by Robinson (1950). The ecological fallacy consists of inferring individual behaviours from aggregate measures. Robinson pointed out, for instance, that correlations between two characteristics measured on a binary basis among individuals (e.g. being black and illiterate in the US), or by proportions in regions (e.g. proportions of black and illiterate people in the population) were generally not identical and could even carry opposite signs. Conversely, the atomistic fallacy arises when, analysing individual behaviours, the context in which such behaviours occur is neglected.

different sources: surveys, census, experiments, interviews, etc. Analogously, evidence of causal relations can come from different sources: previous studies, background knowledge, knowledge of mechanisms or of probabilistic relations, etc.

Previous studies often make it plausible to investigate a given causal relation in a different population or at a different time. Background knowledge gives a (causal) context to causal models, for instance by providing the socio-political context of a population or socio-demographic differences across different populations, etc. A different type of background knowledge is constituted by the iterated application of some methods for the analysis of a given relationship.

Evidence for causal relations can be of two different sorts. We infer causal relationships from probabilistic evidence: causes have to be statistically relevant for their effects, but we also require repetition of similar studies and coherence in their results, namely covariations among variables of interest have to show some stability. Yet, probabilistic evidence is not sufficient as correlations may be spurious, as is well known. To infer causal relationships we also have to exhibit a plausible mechanism.

It seems then clear that scientific practice in the social sciences takes advantage of a pluralistic methodology, evidence and source of information. Consequently, the following question arises: does this methodological and evidential pluralism entail ontological and epistemological pluralism? In other words, does the fact that in practice social scientists use different models, different sources of evidence and of information entail that there are different *concepts* or different *rationales* of causality? In the remainder of the paper I shall focus on epistemological pluralism and argue that a single rationale of causality, based on the notion of variation, is used in the social sciences.

3. Epistemological monism: the rationale of variation

The rationale of variation

Epistemology studies the origin, nature and limits of human knowledge. Epistemology of causality then wonders how we *know* about causal relations, what epistemic access we have to causal relations, under what conditions we can interpret correlations causally, whether specific conditions such as invariance under intervention, structural stability or the Markov condition in Bayes nets guarantee causality. More specifically, I address the question: what *rationale* of causality governs causal models in the social sciences? Is there a unique rationale or different ones depending on the model at hand?

A rationale is the principle or notion underlying some opinion, action, phenomenon, reasoning, model, or the like. A rationale of causality in causal modelling is then the principle or the notion that guides causal reasoning (in causal modelling). It is worth emphasising that a rationale is not a definition of what causality is. A definition is a description of a thing by its properties; thus, a definition of causality states what causality in fact is. Whilst providing a definition of causality is a job for metaphysics, the development of a rationale is a matter of epistemology. In this paper, I am concerned with the latter problem but not with the former.

The received view, an heritage of Hume, sees this rationale in the notion of regularity and this standpoint still pervades contemporary philosophy of science. As is well known, Hume believed that causality lies in the constant conjunction of causes and effects. In the *Treatise* Hume says that, in spite of the impossibility of providing rational foundations for the existence of objects, space, or causal relations, to believe in their existence is a “built in” habit of human nature. In particular, belief in causal relations is granted by experience. For Hume,

simple impressions always precede simple ideas in our mind, and by introspective experience we also know that simple impressions are always associated with simple ideas. Simple ideas are then combined in order to form complex ideas. This is possible thanks to imagination, which is a normative principle that allows us to order complex ideas according to (i) resemblance, (ii) contiguity in space and time, and (iii) causality. Of the three, causation is the only principle that takes us beyond the evidence of our memory and senses. It establishes a link or connection between past and present experiences with events that we predict or explain, so that all reasoning concerning matters of fact seems to be founded on the relation of cause and effect.

The causal connection is thus part of a principle of association that operates in our mind. Regular successions of impressions are followed by regular successions of simple ideas, and then imagination orders and conceptualizes successions of simple ideas into complex ideas, thus giving birth to causal relations. The famed problem is that regular successions so established by experience clearly lack the necessity we would require for *causal* successions (otherwise successions would be merely casual). Hume's solution is that if causal relations cannot be established *a priori*, then they must be grounded in our experience, in particular, in our psychological habit of witnessing effects that regularly follow causes in time and space.

My suggestion is, instead, to depict the rationale of causality as the *measure of variation or change*. The study of change is the study of factors which produce change. Thus, measuring variations conveys the following idea: to test causal models means to measure suitable changes or variations. Causal models apply a H-D methodology: causal hypotheses are first formulated, and then put forward for empirical testing. In turn, empirical tests are designed to assess the presence of a variation, and to assess whether this variation satisfies certain conditions. Therefore, the point I want to make is twofold: (i) causal modelling is not governed by a rationale of regularity but by a *rationale of variation*, and (ii) there is only one rationale. In other words, methodological and evidential pluralism do not imply epistemological pluralism.⁷ The monistic epistemology I offer turns around a rationale of causality which is based on the notion of variation.

Methodological arguments

Let us consider structural equation models (SEM) first. Recall that the basic idea of SEM is that in a system of equations we can test whether variables are interrelated through a set of linear relationships, by examining the variances and covariances of variables, and that, given the numerical value of the path coefficient β , the structural equation claims that a unit increase in X would result in β units increase of Y . This means that β quantifies the *variation* on Y accompanied by the *variation* on X . The equality sign in structural equations does not state an algebraic equivalence; jointly with the associated graph, the structural equation is meant to describe the causal relationship implied by the data generating process. The path coefficient β , in turn, is meant to quantify the (direct) causal effect of X on Y . β quantifies the *variation* on Y produced by the *variation* of X , hence the structural equation does not merely

⁷ This paper is mainly concerned with the scientific literature on causal modelling. In Russo (2005) I also show that a number of authors in the philosophical literature employ or presuppose the notion of variation in their accounts. For instance, in Woodward's account (Woodward 2003) causal generalisations are "*change-relating*" or "*variation-relating*" and variation is a necessary condition for interventions. In Hausman's account (Hausman 1998), modal invariance presupposes that intervening on the putative cause will produce a *variation* on the putative effect, this variation being possible to compute, or at least possible to estimate.

describe a regular concomitant occurrence of Y and X , but how the dependent variable *varies* depending on the variation of the independent one.

It is commonly agreed that, in structural equations, variations in the independent variables, i.e. the X s, explain the variation in the dependent variable, i.e. Y . Witness, for instance, Haavelmo (1944 : 23):

In other words, we hope that, for each variable, y , to be “explained,” there is a relatively small number of explaining factors the *variations* of which are practically decisive in determining the *variations* of y . (My emphasis)

So, one might suggest that the explanatory power of a causal model is given by the inverse of the unexplained variance in the dependent variable. Namely, because the dependent variable is determined to a certain extent by variations in the independent variables, the more we can account for those variations, the higher the explanatory power of the model. What is not determined by variations in independent variables, depends on the errors. In fact, the squared path coefficient r^2 measures the portion of the variance in the dependent variable the independent variable is responsible for, or, differently put, the squared path coefficient represents the proportion of common variation in both the dependent and in the independent variable. Again, reasoning about the squared path coefficient involves the notion of variation rather than regularity.

CSM are governed by a rationale of variation too. By way of reminder, CSM have two models: a measurement model and a structural model: the covariance matrix in the measurement model is explained by the (causal) relations as indicated in the structural model. On the one hand, to analyse covariances means exactly to measure (joint) *variations*. The covariance matrix is the matrix of the covariances between elements of a vector and represents the natural generalization to higher dimensions of the concept variance of a scalar-valued random variable. The variance measures how much a single variable varies around the mean, and the covariance measures how much two variables vary together. On the other, the structural component of CSM is not meant to explain regularities, but joint *variations* and, as we just saw, is regimented by the variation rationale.

Analogously, as they are based on structural equations, multilevel models too rely on the rationale of variation. Furthermore, variations are measured not only at one level of aggregation, but also across different levels. For instance, a multilevel model can assess how educational achievement varies among students in the same class and across classes in a school.

Contingency tables also rely on the rationale of variation. The question is, in fact, the extent to which each of the independent variables contributes to the *variation* between the categories in the dependent variable. In the scheme of reasoning behind contingency tables it is not hard to recognize the *statistical-relevance model* (S-R) of explanation advanced by W. Salmon (1971, 1984). And in fact, not surprisingly, the rationale of variation is clearly involved in the S-R model.

According to Salmon, to explain a fact, i.e. to identify its causes, one has to identify the correct cell in the reference class the fact to be explained belongs to. Consider Salmon’s example (1984 : 37). To understand why a particular individual – say, Albert – committed a delinquency– say, stealing a car, we first consider the broadest reference class Albert belongs to (American teenager); then, this class is partitioned into subclasses based on the number of (all and only) relevant factors. Sociological theories suggest taking gender, religious background, marital status of parents, type of residential community, socioeconomic status and education, etc. into account. We will thus obtain a large number of cells, each of which will be assigned a probability of the degree of delinquent behaviour. This fact, i.e. why Albert

committed a delinquency, will be explained once the narrowest class Albert belongs to is identified, e.g. male and parents divorced and living in a suburban area and low education ...

As the name suggests, statistical *relevance* relations are used in the S-R model in order to isolate *relevant* causal factors. Let me explain how it works – the rationale of variation will then become apparent. Let A denote American teenagers, and B_i various degrees of juvenile delinquency. What we are interested in is not just $P(B_i|A)$, that is, the probability of committing a delinquency in the population of American teenagers, but in a more specific probability, say $P(B_i | A \cap C_j \cap D_k \cap E_n)$, where C_j , D_k and E_n are all *relevant factors*, for instance gender, religious background, marital status of parents, etc. The crucial point is that if conditioning on a further factor, say F_m , does not *change* the previous conditional probability, then F_m is not a relevant factor and hence should not be considered in the explanation. So all factors entering the S-R model are statistically relevant, i.e. *responsible for variations*, in the probability of the fact to be explained.

Varieties of variations

So far I argued that a monistic epistemology, based on the notion of variation, regiments different types of causal models. However, the rationale of variation as just described is still very general. In practice, social scientists may look for different types of variations depending on the case at hand. A taxonomy of variations can be sketched according to the following criteria:⁸

1. variation across time;
2. variation across individuals;
3. variation across characteristics;
4. counterfactual and control group variations;
5. observational vs. interventional variations.

We might be interested in whether the same characteristic, say unemployment rate, varies across time (taxon 1) – e.g. in two successive surveys, or across individuals (taxon 2) – e.g. individuals in the survey may show radically different employment histories, or across characteristics (taxon 3) – e.g. unemployment rate may be different according to different levels of education. In observational studies we can model counterfactual variations (taxon 4), for instance the individual probability of finding a job given certain characteristics, or, in experimental studies, we can check whether variations hold between the test and control group. Finally (taxon5), variations can be merely observed – when we deal with observational data, or can be the result of interventions – if we can manipulate and operate directly on data.

Objections, or further evidence

The rationale of variation is not exempt from possible objections. The first doubt might be that *variation* means something stronger, namely *dependence*. If so, Pearl has argued instead that *independence* is more basic:

“[...] conditional independence is not a ‘restrictive assumption’ made for mathematical convenience; nor it is an occasional grace of nature for which we must passively wait. Rather, it is a mental construct that we actively create, a psychological necessity that our culture labours to satisfy.” (1988 : 385).

In other words, *independence* is an essential feature for causality. Nonetheless, a few pages later, Pearl seems to hold quite a different view, when he draws some conclusions about

⁸ These taxa are detailed and fully exemplified in Russo (2005).

causal poly trees. He addresses the old question: causation or covariation? According to him, the threshold is in the notion of *control*: causal directionality between X and Y can only be tested through the introduction of a third variable Z . This is because by introducing Z we test whether:

“by activating Z we can create *variations* in Y and none in X , or alternatively, if *variations* in Z are accompanied by *variations* in X while Y remains unaltered. [...] the construct of causality is merely a tentative, expedient device for encoding complex structures of dependencies in the closed world of a predefined set of variables. It serves to highlight useful independencies at a given level of abstraction, but *causal relationships undergo change* upon the introduction of new variable.” (1988 : 397, my emphasis)

In the ultimate analysis, Pearl did use the rationale of variation, and this rationale seemed to precede the notion of independence, contrary to what he himself stated, i.e. that independence is the basic notion for causal learning.⁹

Another obvious objection to the rationale of variation is that this rationale is nothing but a reformulation of Humean regularist accounts. This is only partly true. Let me deal with the non-true part first. The crucial step in Hume’s argument is significantly different from the rationale I propose. My claim is that we look for variations, not for regularities. Once variations are detected, a condition of invariance or structural stability (among others) is imposed on them. What does structural stability give us? Not logical or nomic necessity, nor mere constant conjunction. Structural stability is a condition imposed on joint variations, in order to ensure that the model correctly specifies the data generating process and to ensure that the model does not confuse accidental variations with causal ones. Although the invariance condition is the most important one, other conditions – e.g. no confounding – grant the causal interpretation of statistical models. The true side of the objection is that in observational studies attention is mainly directed to variations that happen to be regular, at least regular enough not to be accidental.

Let me underline why so much importance is bestowed on the notion of variation. Variation is conceptually a *precondition*. If causality is not set in the notion of variation, then it will be lodged in the invariance condition, which is conceptually misleading. The same holds for regularity. In both cases there is a further question to ask: invariance *of what?*, regularity *of what?* The answer is, in both cases, *of a variation*. In particular, invariance – the queen of the causal conditions – only makes sense *within* a causal model, whereas variation is exactly what motivates testing invariance. In other words, without variation, invariance is devoid of meaning. This is why variation conceptually precedes invariance.¹⁰

To provide a rationale of causality means, to put it otherwise, to give the bottom-line concept – namely, variation – and the constraints to put on this variation – namely, invariance and regularity. Neither invariance nor regularity are apt to accomplish that task. But variation is. Hume inferred causation from regularity, whereas my claim is that we infer causation from variation because variation conceptually and empirically comes before regularity. Of course, both notions – regularity and variation – don’t guarantee a straight causal interpretation, but

⁹ It is worth noting that Pearl changed his mind about causality between 1988 and 2000 : “Ten years ago, when I began writing *Probabilistic Reasoning in Intelligent Systems* (1988), I was working within the empiricist tradition. In this tradition, probabilistic relations constitute the foundations of human knowledge, whereas causality simply provides useful ways of abbreviating and organizing intricate patterns of probabilistic relationships. Today, my view is quite different. I now take causal relationships to be the fundamental building blocks both of physical reality and of human understanding of that reality, and I regard probabilistic relationships as but the surface phenomena of the causal machinery that underlies and propels our understanding of the world.” (2000 : xiii-xiv).

¹⁰ For a thorough discussion of variation as a *precondition*, see Russo (2005).

the rationale of variation puts us on the right track because it makes causality an empirical issue rather than a psychological fact or a mere reduction to statistical conditions to be satisfied.

The Humean paradigm of regularity still dominates contemporary philosophy of science. Regularity views of causation are clearly an heritage of the Humean account. Defenders of regularist accounts claim, roughly, that to assert a causal relation between two events x and y means to assert the existence of a regular succession such that every time an event of type X occurs, then an event of type Y will invariably follow. Humeans like J.S. Mill or J. Mackie have advanced more sophisticated versions of the regularist view (e.g. the I.N.U.S. condition) and tried to characterise the kind of regularity that can underpin causal relations by tying causation to laws of nature. Even Lewis' counterfactual approach uses regularities as means to capture the conditions under which counterfactual assertions are true. Also, the probability raising requirement often advocated in probabilistic accounts, has been backed by a *ceteris paribus* condition – i.e. a condition of homogeneity – that makes things regular enough to let the cause raise the probability of the effect.

The variation rationale profoundly breaks down this received view. The variation rationale argues, in the first place, that this emphasis on regularity is not well-founded, for regularities themselves require a prior notion, which is variation. Secondly, the difficulty or impossibility of establishing causal laws in the social sciences is usually taken as a structural weakness or even as an intrinsic impossibility for social science to reach the kingdom of “hard sciences”. The variation rationale is the first step for a radical change in the dominant paradigm: if, in the ultimate analysis, causal modelling aims at measuring variations rather than establishing regularities, this might be due to the fact that the regularist rationale is not, after all, well founded as empiricists claim since Hume. This calls for a change of paradigm in causal modelling, rather than giving up our endeavour to establish causal claims or questioning the rigorous scientific character of the social sciences.

4. The rationale of variation in philosophical accounts

A number of accounts have been proposed in the last decades. The probabilistic, mechanist and counterfactual approach take causal relations to be objective, in the sense that causality is defined independently of the agent, the first relying on statistical relevance, the second on the notion of physical process and interaction, and the third on counterfactual logic. Agency theories, instead, define causality in terms of an agent's ability to operate on causal relations; manipulability theories try to get rid of anthropomorphism and to regain objectivity by developing a notion of intervention that fits causal modelling. A different attempt to give causality an objective character is epistemic causality, where objectivity is understood as “non arbitrary” rather than “mind-independent”. The goal of this section is to disclose how the rationale of variation is consistent with or (more or less explicitly) adopted in those accounts.

Variation in probabilistic theories

Probabilistic theories (PT) of causality have been developed in slightly different manners by different philosophers in the last decades.¹¹ In spite of the significant differences in these accounts, a core of agreement can be found in the pioneering works of Good and Suppes that, roughly speaking, turn around the probability raising requirement: *ceteris paribus*, causes

¹¹ See for instance Good (1961-62), Suppes (1970), Cartwright (1979) and (1989), Eells (1991).

make their effects more probable. Prima facie, C is a cause of E if, and only if, (i) C occurs before E and (ii) C is positively, statistically relevant to E , that is $P(E_{t'} | C_t) > P(E_{t'})$ ($t < t'$).

Thus PT focus on the difference between the conditional probability of the effect given the cause $P(E|C)$ and the marginal probability of the effect $P(E)$. To compare the conditional and marginal probability means to analyse a statistical relevance relation. The underlying idea is that if C is a cause of E , then C is also statistically relevant for E . To evaluate a statistical relevance relation exactly means to measure a *variation*, in particular, a variation in the conditional probability of E given C with respect to the marginal probability of E . That is to say, the change hereby produced by C in the effect E will be detected because the conditional and the marginal probability *differ*.

Variation in mechanist accounts

The mechanist approach¹² takes physical processes and interactions between them to be the fundamental concepts for causation. Simply put, the Salmon-Dowe theory is based on three definitions: (i) causal interaction, (ii) causal process, and (iii) causal transmission.

First, a *causal interaction* is an intersection of world-lines which involve exchange of a conserved quantity. In this definition, the meaning of *exchange* is worth stressing: at least one outgoing process manifests a *change* in the value of the conserved quantity and the exchange is governed by the conservation law. Second, a *causal process* is a world-line of an object that transmits a nonzero amount of an invariant quantity at each moment of its history (each space-time point of its trajectory). Last, the principle of mark transmission is formulated as follows: a process transmits an invariant (or conserved) quantity from A to B ($A \neq B$) if it possesses this quantity at A and at B and at every stage of the process between A and B without any interactions in the half-open interval $(A, B]$ that involves an exchange of the particular invariant (or conserved) quantity.

The notion of variation plays a fundamental role in the definition of causal interaction. In fact, an exchange of invariant or conserved quantities between processes actually produces a *modification* or *variation* in them and this is what makes the interaction causal. Of course, the way in which the rationale of variation is here employed differs from the quantitative one depicted in causal modelling, but the qualitative claim still holds: the bottom-line concept of causality is in the concept of variation, not in regularity, stability or invariance.

Variation in counterfactuals

D. Lewis (1973) is the main proponent of the counterfactual theory of causation. Causal relations are analysed in terms of subjunctive conditionals, also called counterfactuals: “ A caused B ” is interpreted as “ B would not have occurred if it were not for A ”. Counterfactuals are subjunctive conditionals where the antecedent is known or supposed to be false and are regimented by a possible-world semantics.

Possible-world semantics rest on the assumption of the existence of a plurality of worlds, among which there is also our actual world.¹³ Worlds are compared with each other on the basis of their similarity or closeness. The relation of comparative over-all similarity among possible worlds is taken as primitive and we say that one world is closer to actuality than another if the first resembles our actual world more than the second does. The truth of a counterfactual is then ascertained by an “inspection” of what happens in other possible

¹² Salmon (1984) and (1994), Dowe (1992).

¹³ This position is also known as modal realism.

worlds. Given any two propositions A and B , the counterfactual $A \square \rightarrow B$ reads: “if A were true, then B would also be true”. The counterfactual operator $\square \rightarrow$ is defined by the following rule of truth:

The counterfactual $A \square \rightarrow B$ is true (at a world w) if, and only if:

- (i) there no possible A -worlds¹⁴, or
- (ii) some A -world where B holds is closer to w than is any A -world where B does not hold.

The second case is the interesting one, for in the former the counterfactual is just vacuously true. Causality comes in because by asking whether a counterfactual is true, we wonder whether the occurrence of A is the *cause* of the occurrence of B . So, the counterfactual, if true, states that if the cause had not occurred, the effect would not have occurred either.

How the rationale of variation is involved in the counterfactual approach will become apparent once the motivation for its development is spelled out. Lewis wants to go beyond standard regularity theories as they failed to notice the second definition of cause Hume gave in the *Enquiry* (Section VII):

“We may define a cause an object followed by another, and where all the objects, similar to the first, are followed by objects similar to the second. Or, in other words, where, if the first object had not been, the second had never existed.”

If the cause had not been, the effect had never existed. For Lewis, this is not just a restatement of the regularist view, but a clear input and encouragement to take up the counterfactual path. In Lewis’ words (1986 : 160-161):

“We think of a cause as something that makes a difference, and the difference it makes must be a difference from what would have happened without it. Had it been absent, its effects – some of them, at least, and usually all – would have been absent as well.”

Causes are supposed to make a difference, i.e. are responsible for *variations*. The bottom-line concept is, again, not in regularity, uniformity, or invariance, but is in difference, change, and variation.

Agreed, it would be misleading to maintain the counterfactual approach *relies* on the rationale of variation; yet, the variation idea is definitively consistent with it, as Lewis’ words clearly show. The rationale of variation here involved is not quantitative, as is the case in causal modelling. Instead, a *qualitative* notion of variation is here at stake. Surely Lewis’ account is of little help in testing causality over large data sets, but counterfactuals do grasp, at least, our intuitions about how the causal relation works: *ceteris paribus*, if the cause had not occurred, the effect had never existed either. In other words, we expect the cause to be responsible for the *change* leading to the effect.

Variation in agency-manipulability theories

Agency theories¹⁵ analyse causal relations in terms of the ability of agents to achieve goals by manipulating causes. In a nutshell, C is said a cause of E if bringing about C would be an effective way to bring about E . To explain what counts as an effective strategy, Menzies and Price invoke the means-end relation, which is characterised in terms of agent probabilities. Agent probabilities are conditional probabilities assessed from an agent perspective as follows: the agent probability of E conditional on C is the probability that should enter in the calculations of a rational agent, whose abilities consist in the capacity to realise or prevent C , and whose goals entirely concern E . So a strategy to bring about E is effective if a rational decision theory prescribes it as a way of bringing about E . That is to say, agents probabilities

¹⁴ In Lewis’ theory “A-world” means “the world in which A is true”.

¹⁵ Price (1991) and (1992), Menzies and Price (1992).

are defined in terms of their role in rational decision-making and this is why they embody a basis for a formal analysis of the means-end relation.

This account is consistent with the central idea of the probabilistic theory of causality. In fact, C constitutes a means for achieving E only in the case that the agent probability $P_C(E)$ is greater than $P_{not-C}(E)$, where P_C denotes the agent probability that E would hold, were one to realise C . Because the agency theory turns out to be consistent with probabilistic theories of causality, *a fortiori* it is consistent with the variation rationale.

D. Hausman and J. Woodward¹⁶ propose a manipulationist account of causation. Their purpose is to overcome the objection of anthropomorphism raised against the agency theory, by developing a notion of intervention which is not agency-dependent. In their manipulationist or interventionist account, causal relations have essentially two features: (i) they are potentially exploitable for purposes of manipulation and control, and (ii) they are invariant under intervention. Everything turns around the specification of the notions of intervention and invariance.

Briefly put, an intervention on X with respect to Y changes the value of X in such a way that, if any change occurs in Y , it occurs only as a result of the change in the value of X and not from other sources. On the other hand, the notion of invariance is closely related to the notion of intervention and takes advantage of the notion of generalization. A generalization G is invariant if it would continue to hold under some intervention that changes the value of X in such a way that the value of Y would change. “Continue to hold” means that G correctly describes how the value of Y would change under this intervention. For Hausman and Woodward, equations in SEM are correct descriptions of the causal relationship between X and Y if, and only if, were one to intervene in the right way to change the value of X , then Y should change in the way indicated by the equation.

Invariance, or structural stability, is then a necessary ingredient for avoiding equations describing contingent or spurious relations. Nonetheless, it is worth asking – and here is where the rationale of variation emerges, *what* exactly remains invariant. The answer is: invariance of a detected *variation*. As we have seen earlier, structural equations describe how variations in X accompany variations in Y and structural parameters quantify the causal effect of X on Y . Eventually, in characterising causality as invariance under intervention, manipulability theories rely on the variation rationale in an essential manner.

Variation in epistemic causality

According to Williamson’s epistemic theory of causality¹⁷, causality is an objective mental construct. Causality is mental because it is a feature of an agent’s mental state, as opposed to physical causality which is a feature of the world “out there”, and it is objective because if two agents differ as to causal structures, then at least one of them must be wrong. Causal relations belong to an agent’s representation of the world, more precisely, epistemic causality deals with *causal beliefs*. It is convenient to represent the world in terms of causes and effects because such causal representations, if correct, enable accurate predictions, diagnosis, decisions and interventions.

This metaphysical stance about causality is also accompanied by an account of the epistemology of causality: causal relations are discovered by an hybrid of the inductive and hypothetico-deductive method. Because the variation rationale belongs to epistemology, the

¹⁶ Hausman (1998), Woodward (2003), Hausman and Woodward (1999), (2004).

¹⁷ Williamson (2005).

question is whether or not the rationale is compatible with Williamson's learning strategies for epistemic causality, which involves four stages:

- (i) hypothesise;
- (ii) predict;
- (iii) test;
- (iv) update.

The *first* stage – the inductive one– requires a procedure for obtaining a causal graph from data and standard artificial intelligence techniques allow us to induce a minimal causal graph that satisfies the Causal Markov Condition. In the *second* stage, predictions are drawn from the induced graph and those predictions will be tested in the *third* stage. By renewed information or by performing experiments, predictions will be confirmed or disconfirmed. Finally, the *fourth* step represents a radical change in the hypothetico-deductive method: in case predictions fail, we do not start from the very beginning at step one, but we update the causal graph according to new evidence and information gathered.

The variation rationale permeates Williamson's learning strategy in the same sense as it permeates causal modelling. Causal models are tested by measuring suitable variations among variables, and this is exactly what happens in the test stage. Witness Williamson (2005 : 152):

If, for instance, the hypothesised model predicts that *C* causes *E*, and an experiment is performed which shows that intervening to *change* the value of *C* does not *change* the distribution of *E*, controlling for *E*'s other direct causes, then this evidence alone may be enough to warrant removing the arrow from *C* to *E* in the causal model.” (p. 149)

Again, the rationale for testing causal relations is based on variation and not on regularity or invariance.

6. Conclusion

The social sciences perform causal analyses by means of a variety of methods and rely on several sources of information and of evidence. This methodological and evidential pluralism raises the question of whether ontological and epistemological pluralism ought to be accordingly adopted.

This paper focused on the epistemological side and argued in favour of a monistic epistemology based on a rationale of variation. The rationale of variation is shown to be the basic notion employed in causal reasoning in different causal models, e.g. structural equation models, covariance structure models, multilevel analysis and contingency tables. I have argued that the variation rationale goes against the regularity view and constitutes the bottom-line concept of causality because it is a *precondition* both for regularity and invariance. Finally, I also showed how this rationale is consistent with or even adopted in various philosophical accounts, from probabilistic theories to mechanist approaches, from agency-manipulability theories to epistemic causality.

We have seen that the inheritance of the received view is a rationale of regularity. The rationale of *variation* profoundly breaks down with this conception. The philosophical gain in adopting the rationale of variation is fourfold. *First*, causality is not merely lodged in a psychological habit of observing regular successions of events. Agreed, we do experience such regular sequences but, I argue, it is not because of regularity that we interpret them causally. Instead, this is because certain variational relations hold. *Second*, causality is not reduced to statistics either. Correlation, as is well known, does not prove causation. Further, to claim that variation is a *precondition* for regularity and invariance has the undoubted advantage of not confusing the rationale of causality with the conditions that allow to interpret

variations causally. *Third*, along the same line, the adoption of the rationale of variation avoids confusing (i) what causality *is* (metaphysics) with the notion employed in testing (epistemology) and (ii) with the conditions – e.g. invariance – to impose on the variation to interpret it causally (methodology). *Last*, the rationale of variation is a first step in redeeming the social sciences as *sciences*. In fact, as individuals and societies are too mutable, the social sciences cannot establish universal and necessary regularities as physics does. For this reason their scientific status has been often questioned. This indictment is ill-founded because the social sciences aim at establishing causal variations rather than regularities.

Acknowledgments

I wish to thank Daniel Courgeau, Rober Franck, Jon Williamson, and an anonymous referee for their helpful comments and suggestions.

References

- Agresti A. (1996), *An Introduction to Categorical Data Analysis*, Wiley, New York.
- Blalock H. M. (1964), *Causal Inference in Nonexperimental Research*, University of North Carolina Press, Chapel Hill.
- Boudon R. and Lazarfeld P. (1966) (eds), *L'analyse empirique de la causalité*, Mouton&Co, Paris.
- Cartwright N. (1979), "Causal Laws and Effective Strategies", *Noûs*, 13, pp. 419-37.
- Cartwright N. (1989), *Nature's Capacities and their Measurement*, Clarendon Press, Oxford.
- Courgeau D. (2003) (ed), *Methodology and Epistemology of Multilevel Analysis. Approaches from Different Social Sciences*, Kluwer, Dordrecht.
- Dowe P. (1992), "Wesley Salmon's Process Theory of Causality and the Conserved Quantity Theory", *Philosophy of Science*, 59, pp. 195-216.
- Duncan O. D. (1975), *Introduction to structural equation models*, Academic Press, New York.
- Durkheim E. (1895 - 1912), *Les règles de la méthode sociologique*, Libraire Félix Arcan, Paris, 6th edition.
- Durkheim E. (1897 - 1960), *Le suicide*, Presses Universitaires de France, Paris.
- Eells E. (1991), *Probabilistic Causality*, Cambridge University Press, Cambridge.
- Goldstein H. (2003), *Multilevel statistical analysis*, Arnold, London, 3rd Edition.
- Good I. J. (1961), "A Causal Calculus I", *British Journal for the Philosophy of Science*, 11, pp. 305-18.
- Good I. J. (1962), "A Causal Calculus II", *British Journal for the Philosophy of Science*, 12, pp. 43-51.
- Haavelmo T. (1943), "The statistical implications of a system of simultaneous equations", *Econometrica*, 11, pp. 1-12.
- Haavelmo T. (1944), "The Probability Approach in Econometrics", *Econometrica*, 12, iii-vi + 1-115.
- Hausman D. (1998), *Causal Asymmetries*, Cambridge University Press.
- Hausman D. and Woodward J. (1999), "Independence, Invariance, and the Causal Markov Condition", *British Journal for the Philosophy of Science*, 50, pp. 521-583.
- Hausman D. and Woodward J. (2004), "Modularity and the Causal Markov Condition: a Restatement", *British Journal for the Philosophy of Science*, 55, pp. 147-161.
- Hellevik O. (1984), *Introduction to Causal Analysis*, Allen & Unwin, London.
- Hume D. (1739-1740), *A treatise of human nature*, edited by David Fate Norton and Mary J. Norton, Oxford University Press, Oxford/New York, 2000.

- Hume D. (1748), *An enquiry concerning human understanding*, edited by Tom L. Beauchamp, Oxford University Press, Oxford/New York, 1999.
- Koopmans T. C. (1950), “When is an equation system complete for statistical purposes?”, in T.C. Koopmans (ed), *Statistical inference in dynamic models* (Cowles Commission Monograph n.10), Wiley, New York.
- Lewis D. (1973), “Causation”, reprinted with postscripts in Lewis D. (1986), *Philosophical Papers II*, Oxford University Press, Oxford, pp. 159-213.
- Long J. S. (1983), *Covariance Structure Models*, Sage, Beverly Hills.
- Mackie J. L. (1980), *The cement of the universe: A study on causation*, Clarendon Press, Oxford.
- Menzies P. and Price H. (1993), “Causation as a Secondary Quality”, *British Journal for the Philosophy of Science*, 44, pp. 187-203.
- Mill J. S. (1843), *A System of Logic, Ratiocinative and Inductive: Being a Connected View of the Principles of Evidence and the Methods of Scientific Investigation*, Harper & Brothers, New York, 8th edition.
- Pearl J. (1988), *Probabilistic Reasoning in Intelligent Systems*, Morgan Kaufman, San Mateo (Cal.).
- Pearl J. (2000), *Causality. Models, Reasoning, and Inference*, Cambridge University Press.
- Price H. (1991), “Agency and Probabilistic Causality”, *British Journal for the Philosophy of Science*, 42, pp. 157-176.
- Price H. (1992), “Agency and Causal Asymmetry”, *Mind*, 101, pp. 501-520.
- Quetelet A. (1869), *Physique sociale. Ou Essai sur le developpement des facultés de l’homme*, Muquardt, Bruxelles.
- Robinson W. (1950), “Ecological correlations and the behaviour of individuals”, *American Sociological Review*, 15, 351-357.
- Russo F. (2005), Measuring variations. An epistemological account of causality and causal modelling, PhD Thesis, Université catholique de Louvain, Belgium.
- Russo F., Mouchart M., Ghins M., Wunsch G. (2006), “Statistical Modelling and Causality”, Discussion Paper 0601, Institut de Statistique, Université catholique de Louvain, Belgium.
- Salmon W. C. et al. (1971), *Statistical explanation and statistical relevance*, Pittsburgh University Press, Pittsburgh.
- Salmon W. C. (1984), *Scientific explanation and the causal structure of the world*, Princeton University Press.
- Spirtes P., Glymour C. and Scheines R. (1993), *Causation, Prediction, and Search*, New York, N.Y.: Springer-Verlag. 2nd Edition, MIT Press (2001).
- Suppes P. (1970), *A Probabilistic Theory of Causality*, North Holland Publishing Company, Amsterdam.
- Snijders T. and Bosker R. (2004), *Multilevel analysis. An introduction to basic and advanced multilevel modelling*, Sage Publications, London, Fourth edition.

- Williamson J. (2005), *Bayesian Nets and Causality. Philosophical and Computational Foundations*, Clarendon Press, Oxford.
- Wright S. (1921), "Correlation and causation", *Journal of Agricultural Research*, 20, pp. 557-585.
- Wright S. (1934), "The Method of Path Coefficient", *Annals of Mathematical Statistics*, 5(3), pp. 161-215.
- Woodward J. (2003), *Making Things Happen: A Theory of Causal Explanation*, Oxford University Press, New York.