

INSTITUT DE STATISTIQUE
BIOSTATISTIQUE ET
SCIENCES ACTUARIELLES
(ISBA)

UNIVERSITÉ CATHOLIQUE DE LOUVAIN



DISCUSSION
PAPER

2011/07

**INFERRING CAUSAL RELATIONS
BY MODELLING STRUCTURES**

MOUCHART, M., RUSSO, F. and G. WUNSCH

Inferring Causal Relations by Modelling Structures ^{*}

MICHEL MOUCHART ^a, FEDERICA RUSSO ^b, GUILLAUME WUNSCH^c

^a *ISBA, Université catholique de Louvain, Belgium*

^b *Philosophy, University of Kent at Canterbury, UK.*

^c *Demography, Université catholique de Louvain, Belgium.*

February 17, 2011

Abstract

This paper provides an overview of structural modelling in its close relation to explanation and causation. It stems from previous works by the authors and stresses the role and importance of the notions of invariance, recursive decomposition, exogeneity and background knowledge. It closes with some considerations about the importance of the structural approach for practicing scientists.

Keywords: Background knowledge, Causality, Exogeneity, Invariance, Latent Variables, Structural Modelling.

Corresponding Author: Michel Mouchart, Institute of Statistics, Biometrics and Actuarial Sciences (ISBA), 20 Voie du Roman Pays, B-1348 Louvain-la-Neuve (Belgium).

e-mail: Michel.Mouchart@uclouvain.be

Contents

1	Introduction and background	3
2	Modelling structures	5
2.1	Hypothetico-deductivism	5

^{*}Michel Mouchart gratefully acknowledges financial support from IAP research network grant *nr.P6/03* of the Belgian government (Belgian Science Policy). Federica Russo wishes to thank the British Academy for funding her research. The authors gratefully acknowledge ongoing stimulating discussions with Robert Franck.

2.2	The conditional model	7
2.3	Exogeneity and recursive decomposition	8
2.4	Structural models, exogeneity and causality	9
2.5	Partial observability and latent variables	11
3	Explaining through structures	12
3.1	Meanings of ‘structural’	12
3.2	The mechanistic interpretation of the recursive decomposition	14
4	On the counterfactual approach to causal analysis	16
4.1	Counterfactual questions as an aid to inferring causal relations	16
4.2	Counterfactuals as a basis for causality	17
5	Discussion and conclusion	19
	References	21

1 Introduction and background

A common objective across the sciences is to look for causes. Knowledge of causes is helpful in order to explain phenomena, to make predictions and to control for bias and confounding. The social sciences are no exception in this respect. This is not to say that we are *always* able to find and to use causes, but this is a reason why causal knowledge is so central in science.

One difficulty the social sciences face is the large variability of its units of observation—that is, at bottom, humans, but going up in the level of aggregation we may also find families, firms, nations—and the large size of data we can collect on them. Statistics proved so helpful in the practice of quantitative social science to the point of becoming a necessary tool. Yet, statistics alone cannot find causes; social scientists instead need expertise in modelling procedures and wise use of background knowledge. A very brief history of causal analysis will make the previous point clear.

Causal analysis has a long tradition, starting with the pioneering and seminal works of the demographer Adolphe Quetelet and of the sociologist Emile Durkheim. Major improvements have been done by, to name a few, Sewall Wright and Otis Dudley Duncan, up to the most recent advancement by econometricians such as Jim Heckman or Kevin Hoover and computer scientists such as Judea Pearl or Peter Spirtes, Clark Glymour and Richard Scheines. What is peculiar to this development is that, whilst the early methodologists were overtly and explicitly causalist, the most recent generations have shown some skepticism as to whether we can infer causation from statistics.

There is another way the development of causal analysis can be read. Whilst early methodologists used basic statistics only, as time passed researchers aimed to *model structures* in order to analyse complex networks of causal relations. This led to the so-called structural modelling approach, which is endorsed, with some differences to be discussed later in the paper, by a number of scholars nowadays.

The last reading of this short history of causal analysis concerns the way researchers dealt with *modelling*. In fact, despite the great advancement in complexity and sophistication of the statistical models, what went lost is an overarching view of modelling *qua* scientific practice. That is to say, what went lost is a global view on modelling that does not confine to the *statistical* model and does not narrowly focus on the (statistical) properties of the distributions and on (statistical) tests. As we shall discuss later in detail, modelling includes, within a hypothetico-deductive approach, making background and field knowledge explicit and taking into account the type of data to be analysed, and on *this* grounds to put in motion the whole statistical machinery.

This last point is a very sketchy formulation of the so-called *hypothetico-deductive* (H-D) methodology. According to the H-D methodology, causal relations cannot be inferred directly from data nor with the sole aid of an algorithmic procedure. Nevertheless, there is also a florid tradition that uses statistics in an inductive way, *e.g.* data mining or exploratory data analysis. This is, for instance, the approach of Spirtes *et al.* (2000). Inductivist approaches claim that causal relations

can be drawn from data without the burden of extra-statistical and causal assumptions made in their hypothetico-deductive counterparts. Unfortunately, it goes far beyond the goal of this paper to discuss the success of inductive causal models. Consequently the scope will be limited to causal models that employ a hypothetico-deductive methodology.

Thus, following the seminal works of *e.g.* Wright, Haavelmo, Blalock, Pearl and others, we present the main features of structural modelling. There are in the literature competing ‘structural accounts’, and we will discuss in section 3.1 in what respect these accounts are similar or different and what are the characterising features of our own approach. In essence, a model is deemed structural if it uncovers a structure underlying the data generating process. In section 2.1 we present hypothetico-deductivism, a general methodology according to which causal analysis can be schematically represented stepwise thus: 1. formulate the causal hypothesis out of background knowledge and preliminary analyses of data; 2. build the statistical model; 3. test the model; 4. check congruence between the results and available background knowledge. Sections 2.2–2.4 argue that a structural approach systematically blends two ingredients. First, the model must be congruent with background knowledge: modelling the data generating process must be operated in the light of the current information on the relevant field. Second, the model must show stability in a wide sense: both the structure of the model and the parameters have to be stable or invariant with respect to a large class of interventions or of modifications of the environment. Section 2.4 introduces causality in terms of exogeneity within structural models. Section 2.5 discusses some difficulties and threats to structural modelling in the presence of latent variables.

It is crucial to note that this concept of structural modelling is wider than the framework of structural equation models, also known as covariance structure models or LISREL type models, widely used in psychology or in sociology, and of simultaneous equations models, widely used in econometrics.

A first consequence of this approach is that the notion of causality becomes relative to the model itself, rather than to the data, as is the case, for instance, in the Granger-type concept of causality. Also, this means that we do not aim at making metaphysical claims about causal relations, but rather at saying when we have enough reasons—specifically, reasons about background knowledge and about structural stability—to believe that we modelled a causal mechanism able to explain a given social phenomenon. A second consequence of this model-based concept of causality, involving both background knowledge and stability, is that the model does not simply derive from theory as is often the case in the econometric tradition. Therefore structural modelling is much more than a sophisticated statistical tool that translates (economic) theory into mathematical equations. Good structural modelling ought to be accompanied by a broad and sensible account of what a statistical model is and represents, of what statistical inference is.

An important aspect of structural modelling is its explanatory import. In section 3 we discuss

the feature of structural modelling that is supposed to do the explanatory job. This is the recursive decomposition, appropriately interpreted as a mechanism.

Another important debate in the broad area of causal analysis concerns the relations and competitions between the structural and the counterfactual approach, which we have thoroughly done elsewhere. Space is limited and we cannot engage in a thorough discussion of the counterfactual approach. Nevertheless, we shall try to set the tracks that we think should guide this debate in section 4. We close the paper, in section 5, with some conclusions and discussions.

In this paper we attempt to present our views building on our previous works on structural modelling, explanation, causality, and the relations they stand with each other. The attempt here is to blend viewpoints coming from statistics, social sciences and philosophy in order to provide a comprehensive approach to structural modelling, along with its practical aspects.

2 Modelling structures

In this section, we develop the formal framework of structural modelling. We start by presenting the hypothetico-deductive methodology and three notions : (i) conditional model, (ii) exogeneity, (iii) recursive decomposition. We then introduce the structural model as the model enabling a causal interpretation of exogenous variables and a mechanistic interpretation of the recursive decomposition. We close the section with some remarks on partial observability and latent variables.

2.1 Hypothetico-deductivism

As anticipated in the Introduction, there are, broadly speaking, two approaches in causal analysis. On the one hand, inductivist approaches, put it very simply, try to infer causal relations from data, with very minimal prior assumptions and virtually no explicit background knowledge supporting the specific algorithm used (such as principal components analysis). On the other hand, hypothetico-deductive approaches heavily rely on the prior specification of assumptions and of background knowledge in order to infer causal relations. In this section we discuss this hypothetico-deductive methodology, which we endorse, in detail.

Hypothetico-deductivism (H-D) is a general methodology that prescribes to formulate hypotheses and to derive consequences in order to test whether the hypotheses obtain or not. Famously, the philosopher of science Karl Popper (1959) was the first to theorise the H-D methodology motivated by the need of providing a methodology alternative to inductive ones (that is, Baconian). Popper's H-D methodology was characterised by (i) a strict meaning of 'deducing the consequences from the hypotheses', and (ii) a complete rejection of the hypothesis in case tests disconfirm it. This implied, in Popper's view, to start every time from scratch. However, the form of hypothetico-deductivism more recently endorsed by philosophers of sciences (see, *e.g.* Cartwright (2007, ch. 2) and Russo

(2009, ch. 3.2)) is much less strict concerning deduction and does not imply starting from complete scratch any time a hypothesis is disconfirmed. Yet, it does retain from the Popperian account the primary role of the hypothesis-formulation stage. We shall return to these points later.

The H-D methodology can be presented as a stepwise procedure for model building and model testing: 1. formulate the causal hypothesis out of background knowledge and preliminary analyses of data; 2. build the statistical model; 3. test the model; 4. check congruence between the results and available background knowledge.

In the first step (hypothesis formulation) background knowledge—from knowledge concerning the phenomenon at stake to preliminary analyses of data—looms large. But background knowledge is also very important for building the model, that is, as we shall discuss in detail in section 2.3 and 2.4, choosing the particular statistical model and the recursive decomposition. Model testing, performed in step 3, concerns various aspects: estimation, goodness of fit, exogeneity, stability. Yet the results of those tests alone do not allow yet inferring a causal structure as the results need to be checked, again, against background knowledge. This is done in step 4. This last stage is very important because, even if positive results are not obtained—that is we are not able to successfully infer the presence of a causal relation—it is not all lost. In fact, this all may feed research later on. Differently put, we also learn from failure.

As just said, hypothetico-deductivism in causal modelling does not involve making deductive inferences strictly speaking. What is instead at stake in H-D methodology is a weaker inferential step of ‘drawing consequences’ from the hypothesis. This means that, after we formulate the causal hypothesis taking into account available background knowledge and meaningful co-variations between the putative cause and effect, we do not require data to be *implied* by the hypothesis but just that data conform to it. ‘Conform’ means that the available indicators in the data set adequately represent the conceptual variables appearing in the causal hypothesis.

It is in this sense that the confirmation of the causal hypothesis is not done by ‘deductive inferences’ strictly interpreted but does involve a ‘deductive procedure’ loosely speaking. To be more precise, it is a *hypothetico*-deductive procedure in the sense that it goes the opposite direction of inductive methodologies: not from data to causation, but from causal *hypotheses* down to data, so to speak.

The important role of background knowledge is worth emphasising. The notion of background knowledge belongs to most quoted and least explicated concepts in causal analysis. It is so broad that it is hard to discriminate between knowledge and information that does or does not count as *background knowledge*. Let’s try to be more specific. Background knowledge may include various aspects: (i) general knowledge about socio-demo-political contexts, (ii) knowledge of physical-biological-physiological mechanisms, (iii) institutional knowledge (such as of the procedure of a central bank), (iv) evidence supporting similar mechanisms in different populations, (v) use of sim-

ilar or different methodologies or data to study the phenomenon of interest. It is important to carefully take these points into account because studies in social science typically consider different populations. It is a proper and explicit use of background knowledge that justifies the choice of variables and of the model, and that allows a sound interpretation of results. Any empirical study based on sound structural modelling is in turn a contribution to background knowledge, that is toward the process of gathering together knowledge and information coming from different sources.

The H-D methodology hereby presented is general and is not restricted to quantitative causal approaches but also includes qualitative methodologies. In the following, we shall restrict our attention to quantitative approaches, and more particularly to the so-called structural modelling approach in statistics.

2.2 The conditional model

Let us start with *an unconditional parameterized* statistical model $\mathbf{M}_{\mathbf{X}}$ given in the following form:

$$\mathbf{M}_{\mathbf{X}} = \{p_X(x | \omega) : \omega \in \Omega\} \quad (1)$$

where for each $\omega \in \Omega$, $p_X(x | \omega)$ is a (sampling) probability density on an underlying sample space corresponding to a (well-defined) random variable X and Ω is the parameter space, aimed at describing the set of sampling distributions considered to be of interest. The basic idea is that the data can be analyzed as if they were a realization of one of those distributions. Thus, a statistical model is based on a stochastic representation of the world. Its randomness delineates the frontier or the internal limitation of the statistical explanation, since the random component represents what is not explained by the model.

A conditional model is constructed through embedding this concept into the usual concept of an unconditional statistical model (1). For expository purposes, we only consider the case where a random vector X of observations is decomposed into $X' = (Y', Z')$ (where $'$ denotes transposition) and where the model is conditional on Z .

The basic idea of a conditional model is the following: starting from a global model $\mathbf{M}_{\mathbf{X}}$ as given in (1), each sampling density $p_X(x | \omega)$ is first decomposed through a marginal-conditional product:

$$p_X(x | \omega) = p_Z(z | \theta_Z) p_{Y|Z}(y | z, \theta_{Y|Z}) \quad \omega = (\theta_Z, \theta_{Y|Z}) \quad (2)$$

where $p_Z(z | \theta_Z)$ is the marginal density of Z , parametrized by θ_Z , and $p_{Y|Z}(y | z, \theta_{Y|Z})$ is the conditional density of $(Y | Z)$, parametrized by $\theta_{Y|Z}$. Next, one makes specific assumptions on the conditional component leaving virtually unspecified the marginal component. Thus a conditional model may be represented as follows :

$$\mathbf{M}_{\mathbf{Y}}^Z = \{p_X(x | \omega) = p_Z(z | \theta_Z) p_{Y|Z}(y | z, \theta_{Y|Z}) \quad \omega = (\theta_Z, \theta_{Y|Z}) \in \Omega = \Theta_Z \times \Theta_{Y|Z} \} \quad (3)$$

where Θ_Z parametrizes a typically large family of sampling probabilities on Z only and for each $\theta_{Y|Z} \in \Theta_{Y|Z}$, $p_{Y|Z}(y | z, \theta_{Y|Z})$ represents a conditional density of $(Y | Z)$.

A conditional model, as in (3), endows the global model (1) with two properties. Firstly, the parameters characterizing the marginal (θ_Z) and the conditional ($\theta_{Y|Z}$) components are independent. Here, ‘independence’ means ‘variation-free’ in a sampling theory framework, *i.e.* $\omega = (\theta_Z, \theta_{Y|Z}) \in \Omega = \Theta_Z \times \Theta_{Y|Z}$, or independent in the (prior) probability in a Bayesian framework, *i.e.* $\theta_Z \perp \theta_{Y|Z}$. Secondly, a conditional model leaves almost unspecified the marginal component, *i.e.* the set Θ_Z represents a ‘very large’ set of possible distributions for Z .

2.3 Exogeneity and recursive decomposition

Formally, the condition of exogeneity can be stated as follows: the parameter of interest should only depend on the parameters identified by the conditional model and the parameters identified by the marginal process should be “independent” of the parameters identified by the conditional process. It should be stressed that the independence among parameters has no bearing on a (sampling) independence among the corresponding variables.

Exogeneity is a condition of separation of inference. The (partial) explanation of the statistician is cast in the framework of a statistical model, in terms of parameters that characterise the distribution of interest. Originally, the concept of exogeneity appears with regression models. A first, and naive, approach was to consider an exogenous variable as a non-random variable, the endogenous variable being the only random one. That this approach was unsatisfactory became clear when considering complex models where the same variable could be exogenous in one equation and endogenous in another one. A first progress came through a proper recognition of the nature of a conditional model; for a more formal presentation, see *e.g.* Mouchart and Oulhaj (2000) and Oulhaj and Mouchart (2003).

The concept of exogeneity has a long history in econometrics. The works of the Cowles Commission in the late Forties and the early Fifties have been path-breaking and are still influential nowadays; in particular, Koopmans (1950) puts emphasis on exogeneity in dynamic models. Barndorff-Nielsen (1978) is significant in the development of conditions for separation of inference. Florens and Mouchart (1980, 1985) and Florens, Mouchart and Rolin (1980) bridge Koopmans (1950) and Barndorff-Nielsen (1978) works and provide a coherent account of exogeneity integrating the separation of inference in dynamic and non-dynamic models. Engle, Hendry and Richard (1983) present a list of different concepts from the econometric literature and display their connections with exogeneity through the introduction of supplementary conditions.

What are the consequences of a failure of exogeneity? There may be a loss of efficiency in the inference if the failure comes from a restriction (equality or inequality), or a lack of independence in a Bayesian framework, between the parameters of the marginal model and those of the conditional

model. There may also be an impossibility of finding a suitable, *e.g.* unbiased or consistent, estimator if the parameter of interest is not a function of $\theta_{Y|Z}$ only. A typical example, well known in the field of simultaneous equations in econometrics, is that the parameter of interest in a structural equation may not be a function of the parameters identified by the model conditional on the explanatory variables corresponding to that specific equation.

Let us now consider a decomposition of X into p components: $X = (X_1, X_2, \dots, X_p)$. Suppose that the components of X have been ordered in such a way that in the complete marginal-conditional decomposition:

$$\begin{aligned} p_X(x | \omega) &= p_{X_p|X_1, X_2, \dots, X_{p-1}}(x_p | x_1, x_2, \dots, x_{p-1}, \theta_{p|1, \dots, p-1}) \\ &\quad \cdot p_{X_{p-1}|X_1, X_2, \dots, X_{p-2}}(x_{p-1} | x_1, x_2, \dots, x_{p-2}, \theta_{p-1|1, \dots, p-2}) \cdots \\ &\quad \cdot p_{X_j|X_1, X_2, \dots, X_{j-1}}(x_j | x_1, x_2, \dots, x_{j-1}, \theta_{j|1, \dots, j-1}) \cdots p_{X_1}(x_1 | \theta_1) \end{aligned} \quad (4)$$

each component of the right hand side have mutually independent parameters, *i.e.* in a sampling theory framework:

$$\omega = (\theta_{p|1, \dots, p-1}, \theta_{p-1|1, \dots, p-2}, \dots, \theta_1) \in \Theta_{p|1, \dots, p-1} \times \Theta_{p-1|1, \dots, p-2} \cdots \times \Theta_1 \quad (5)$$

Under the condition (5), the decomposition (4) is called a *recursive decomposition* and the conditioning variables of each term are the exogenous variables of their corresponding component.

Once the number of components p increases, we shall see in section 2.4 that background knowledge, possibly substantiated by statistical tests, typically provides a simplification of the factors in the form of conditional independence properties. More specifically, it is often the case that the distribution of $(X_j | X_1, \dots, X_{j-1})$ is known not to depend on some of the conditioning variables. Thus there is a subset $\mathcal{I}_j \subset \{X_1, \dots, X_{j-1}\}$, of variables actually relevant for the conditional process generating $X_j | X_1, \dots, X_{j-1}$ as defined by the property

$$X_j \perp\!\!\!\perp X_1, \dots, X_{j-1} | \mathcal{I}_j \quad (6)$$

implying that the factor $p_{X_j|X_1, X_2, \dots, X_{j-1}}$ in (4) is actually simplified into $p_{X_j|\mathcal{I}_j}$ and \mathcal{I}_j may be called the *relevant information of the j -th factor*. Once \mathcal{I}_j has been specified for each factor, (4) is condensed into

$$p_{X_1, X_2, \dots, X_p} = \prod_{1 \leq j \leq p} p_{X_j|\mathcal{I}_j} \quad (7)$$

This form is accordingly called a *condensed recursive decomposition*.

2.4 Structural models, exogeneity and causality

Structural models are a class of models enabling the interpretation of exogeneous variables as causes of the phenomenon to be explained. In this subsection, we show how the concepts of structural model, exogeneity and causality are connected.

A *structural model* conveys the idea of a representation of the world that is stable, or invariant, under a large class of interventions or of modifications of the environment. Structural models are also called ‘causal models’. Here, the concept of causality is internal to a model which is itself stable, in the sense of structurally stable. Thus a structural model aims at capturing an underlying structure; modelling this underlying structure requires taking into account the contextual knowledge of the field of application.

The invariance, or stability, requirement is however not a sufficient condition for making a structural model palatable. A structural model should also help to understand the data generating process; more precisely its characteristics, or parameters, should be interpretable. This is typically achieved by decomposing a model representing the generation of a set (vector) of variables into an ordered sequence of subprocesses representing sub-mechanisms that are congruent with field knowledge. This decomposition corresponds to a systematic *marginal-conditional decomposition*; the ordered feature corresponds to a recursive decomposition. As developed in section 2.3, the conditioning variables of each conditional component of the decomposition are *exogenous* variables for the corresponding sub-process.

For this purpose, we look for a recursive decomposition such that each factor of the right-hand side of (4) is structurally valid. As argued in Mouchart, Russo and Wunsch (2009), *causes* may then be viewed as exogenous variables in the condensed recursive decomposition, alternatively as the relevant information of a structurally valid conditional distribution. Readers familiar with the literature on graph models may recognise that a directed acyclic graph (DAG) is a graphic representation of a condensed recursive decomposition and that the causal structure is depicted by the set of ancestors.

Why interpreting *exogenous* variables as *causal* factors? The main reason is that structural modelling is also meant to explain a phenomenon—an issue that we just briefly mentioned in the Introduction and that we will develop in detail in section 3. In order to do that, we have to model ‘structures’, that is mechanisms where *causal* factors play a role. This ties a knot between structural modelling, explanation, and causation.

Thus the variation-free condition (5) does not only allow us to separate the inferences on $\theta_{j|1,\dots,j-1}$ and on $\theta_{1,\dots,j-1}$, but it also allows us to distinguish the process generating the causes, characterised by $\theta_{1,\dots,j-1}$, and the process generating the effects, characterised by $\theta_{j|1,\dots,j-1}$. Separating causes from effects mirrors the asymmetry of causation.

The goal of structural modelling is to characterise clearly identified and interpretable mechanisms. The choice of the marginal-conditional decomposition is therefore not arbitrary; we need background knowledge and invariance to make a selection among the various possible decompositions. In other words, the marginal-conditional decomposition *alone* does not provide a (causal) explanation of a given phenomenon, but the whole modelling procedure does. This is indeed mirrored in the step-wise H-D methodology presented in section 2.1 which seeks to decompose a vector of variables

into structurally valid components. Next section shows that when latent variables are present in a structural model, causal attribution becomes substantially more complex.

2.5 Partial observability and latent variables

Consider a three-variate completely recursive system, represented in Figure 1, for data in the form $X = (Y, Z, U)$:

$$p_X(x | \theta) = p_{Y|Z,U}(y | z, u, \theta_{Y|Z,U}) p_{Z|U}(z | u, \theta_{Z|U}) p_U(u | \theta_U) \quad (8)$$

where each of the three components of the right hand side may be considered as structural models with mutually independent parameters,

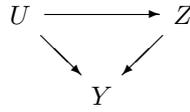


Figure 1: *3-component completely recursive system*

This diagram suggests that U causes Z and (U, Z) cause Y . Here, U is a confounding variable for the effect of Z on Y ; for more detail see *e.g.* Wunsch(2007).

Now suppose that U is not observable. It might be tempting to collapse the diagram in Figure 1 into that of Figure 2. Formally, Figure 2 may be obtained by integrating the latent variable U out of (8):

$$p_{Y|Z}(y | z, \theta_{Y|Z}) = \frac{\int p_{Y|Z,U}(y | z, u, \theta_{Y|Z,U}) p_{Z|U}(z | u, \theta_{Z|U}) p_U(u | \theta_U) du}{\int \int p_{Y|Z,U}(y | z, u, \theta_{Y|Z,U}) p_{Z|U}(z | u, \theta_{Z|U}) p_U(u | \theta_U) du dy} \quad (9)$$

$$p_Z(z | \theta_Z) = \int p_{Z|U}(z | u, \theta_{Z|U}) p_U(u | \theta_U) du \quad (10)$$



Figure 2: *2-component system*

Therefore:

$$\theta_{Y|Z} = f_1(\theta_{Y|Z,U}, \theta_{Z|U}, \theta_U) \quad \theta_Z = f_2(\theta_{Z|U}, \theta_U) \quad (11)$$

Two remarks are in order:

1. In general, Z is not exogenous anymore because (11) shows that the parameter $\theta_{Y|Z}$ and θ_Z are, in general, not independent; indeed some components of $\theta_{Z|U}$ and of θ_U may be common to $\theta_{Y|Z}$ and θ_Z . Therefore, Figure 2 is an inadequate simplification of Figure 1 (see however next remark);

2. the non-observability of U typically implies a loss of identification: the functions f_1 and f_2 are *not* one-to-one; thus Z might still be exogenous because potentially common parameters in $\theta_{Y|Z}$ and θ_Z might not be identified.

3 Explaining through structures

In the Introduction, we mentioned that causal analysis is important in order to explain, predict, and intervene. We also expressed skepticism about methodologies that pretend to infer causal relations on the basis of sole data under analysis. In this section we therefore address two issues of a more theoretical concern. One is the specific meaning of ‘structural’ adopted in our account. We will argue that calling the approach ‘structural’ goes well beyond employing structural *equations* and has instead to do with the goal of modelling *structures*. The second theoretical issue is the explanatory import of structural models. We will argue that structural models explain insofar as they model structures, that is causal mechanisms.

3.1 Meanings of ‘structural’

As the name suggests, structural modelling has to do with structures. What we take to be peculiar to structural modelling is that the whole modelling procedure aims to uncovering (or modelling) structures, *i.e.* mechanisms. The literature on ‘structural modelling’ abounds. One conception (and perhaps the most widespread) takes ‘structural modelling’ and ‘structural equation modelling’ as synonyms. In our view, however, those are not coextensive terms.

‘Structural *equation* modelling’ is a particular type of statistical model used in quantitative social science, especially in econometrics. ‘Structural modelling’, instead, does not denote a particular (statistical) model (*e.g.* structural equation models, covariance models, multilevel models, *etc.*) but refers to a general methodological account of model-building and model-testing. In this sense, we take structural modelling to be a *general* methodological framework for causal analysis. This distinction between a particular statistical model and a general methodological framework is often not clear. Some examples of how scholars both in statistics and in philosophy of science characterise structural modelling clarifies the point.

A notable example is Pearl (2000). In his seminal book on causation, he deals with structural models sometimes meaning structural modelling and sometimes more specifically structural equation modelling. In a more recent contribution, Pearl (2011) develops a ‘structural theory of causation’, *i.e.* in his terms, a ‘general theory of causation’. A possible interpretation of what he means is a ‘general methodological framework’ for causal analysis. In that paper, he develops a formalism that is general enough to subsume, as special cases, particular models such as structural equation models, potential outcome models, and graphical models. The leading idea of his formalism is

that we have to evaluate whether the probability distributions over a set of variables would differ *if* external conditions were to change. Such information is provided by causal assumptions made in the model. Those assumptions allow us to identify relationships that remain invariant when external conditions change. We do not think that Pearl's latest approach is in opposition with our approach, at least as far as the *formal* tools of causal analysis are concerned. However, what distinguishes Pearl's from our approach is an explicit explanation of what makes a structural model *structural*. This we attempt to do later in section 3.2, where we advance an explicit mechanistic interpretation of structural modelling. There is another difference between Pearl's and our approach. Pearl usually distinguishes three types of queries: (i) about the effects of potential interventions, (ii) about counterfactuals, and (iii) about direct and indirect effects. According to the methodology of structural modelling detailed throughout section 2, some structure needs to be identified beforehand in order to answer any of these queries. Such structure, or mechanism, is identified by the recursive decomposition, as explained in section 2.4.

Among philosophers of science, Woodward (2003, ch.7) provides a lengthy discussion of 'structural models'. To begin with, Woodward focuses on structural *equation* models; in particular, he explains the underlying regression techniques and provides a particular definition of invariance condition. However, in this way his account neglects the whole variety of statistical models used in quantitative causal analysis and that could rightly be called 'structural' under our understanding but not under his understanding. Thus the main difference is that whilst our discussion is quite general in scope, Woodward's discussion narrows down to one specific type of statistical model. Other differences between our approaches exist, for instance about the use of counterfactuals in causal analysis or about the role of manipulation and intervention. A thorough discussion falls beyond the scope of the present paper. However, counterfactuals will be covered later in section 4.2. As for the role of manipulation, the reader is referred to Russo (2011a and b).

In econometrics, another example is Hoover (2011) who considers that the goal of quantitative causal analysis is to represent causal relationships by invariant parametrizations of a system of equations. It is worth noting that this position does not automatically imply endorsing a counterfactual approach. In fact, quite to the contrary, Hoover criticises Woodward because, if causal analysis is reduced to counterfactual analysis, it is eventually impoverished. Hoover is also concerned with approaches that bestow much importance to counterfactual manipulability. As it will be clear from section 4, we do share Hoover's criticism of manipulation. So, there is a lot of positive overlap between Hoover's and our approach. Nevertheless, when Hoover presents the structural account as based on Herbert Simon's causal ordering (see, *e.g.* Simon (1953 and 1954) and Fennell (2011)) and on the condition of exogeneity, the question that is left unanswered is, again, what makes a structural account *structural*.

Going back to philosophy of science again, Nancy Cartwright has also extensively written on

quantitative causal analysis. Her worries and concerns are however different from ours. For instance Cartwright (2007) discusses at length econometric techniques and theoretical models used for the purposes of causal analysis in economics. Thus she restricts on purpose the discussion to one social science area, whereas our arguments apply also to disciplines outside economics. But, interestingly enough, the conclusion of one of her arguments indirectly supports ours. Cartwright is concerned that controlled experiments are considered the ‘ideal’ test for causal hypotheses. She challenges the claim that sometimes we do not need to run the experiment, as *Nature* does it for us. Cartwright is particularly worried because we are seldom in the conditions to appeal to such an argument. In fact, even if in principle ‘ideal’ tests done by Nature are conceivable, the reality of things is that causal hypotheses *are* difficult to test. Consequently, we need reasons that are outside the regime over which the test is conducted in order to draw causal conclusions. These reasons are, at bottom, the background knowledge that we have been invoking so far.

Thus, we believe, none of the available approaches explains how we get a mechanism out of a statistical model, which is exactly what the modelling procedure hereby presented instead is meant to do. In the next sub-section, we focus on what we take to be the main characterising aspect of structural models: that is, structural models model mechanisms, and those carry explanatory power.

It is interesting to point out that some discussion on the relations between structural (equation) modelling and mechanisms exists in the philosophical and methodological literature. Yet, the connection has been drawn differently from the one done here. Consider classical economists such as Adam Smith, David Ricardo, Thomas Malthus and John Stuart Mill or the Chicago School of Economics. They developed economic theory so that the theory *dictates* the mechanism, but it does not *model* it. Here is another example. We mentioned earlier Kevin Hoover’s approach, according to which the causal structure represented by a set of structural equations is a “network of counterfactual relations that maps out the underlying mechanisms through which one thing is used to control or manipulate another” (Hoover 2001, p.24). This, as it will clear in the discussion later, is at variance with our own understanding of mechanisms.

3.2 The mechanistic interpretation of the recursive decomposition

As explained earlier in section 2.4, the basic idea behind the recursive decomposition is to factorise the initial joint probability distribution into a sequence of products of marginal and conditional components, where the variables to condition upon play the explicit role of the causes. In this subsection we address the question of the explanatory import of structural models. We argue, in a nutshell, that the recursive decomposition interpreted in mechanistic terms does the explanatory job; for a detailed discussion see Russo(2011c) and Mouchart and Russo(2011).

The argument is that the whole recursive decomposition *can* be interpreted as characterising a global mechanism, whereas each conditional distribution within the recursive decomposition *can*

be interpreted as characterising a sub-mechanism within the global one. If we can identify sub-mechanisms within a global one, this means that we are able to decompose the global mechanism and thus disentangle the action (or function) of each component. Notice, however, that this does not mean that *all* recursive decompositions are mechanisms, but that *some* are interpretable in mechanistic terms; a thorough discussion of this issue is however too broad to be discussed here.

The point we make is that explanatory power is provided by the specification of the decomposition and its interpretation in mechanistic terms. Why? Because the decomposition specifies, as much as possible, the functioning of a phenomenon; the articulation of the mechanism in this sense does the explaining. Some remarks about the meaning of ‘mechanism’ is now in order.

To begin with, in the philosophical literature, a very vivid debate on the concept of mechanisms is happening. Many questions concern what a mechanism *is* and consequently what definition captures its essential features. Notwithstanding the importance of these discussions for those who have an interest in the metaphysics of science, our claim that structural models carry explanatory power insofar as the recursive decomposition is mechanistically interpreted does not depend on the specific definition of mechanism.

What is most important for our claim is *why* mechanisms carry explanatory power (rather than what is *the right* definition of mechanism). Thus, in saying that structural modelling looks for structures, we are in line with the characterisation of Machamer, Darden and Craver (2000), in that the structure being modelled assembles ‘things’ (entities, in MDC vocabulary; variables, in the jargon of statistical models) that interact with and influence each other in a specified way (the activities, in MDC vocabulary; the statistical relevance relations, in the jargon of statistical models). We are also in line with the characterisation of Bechtel and Abrahamsen (2005), who put a lot more emphasis on the structure or the organisation of the mechanism. Our understanding is however not equivalent to the one of Woodward (2003), who conceives of mechanisms as chains of invariance relations.

In our approach we do not make any ontological commitment as to the (degree of) physical existence of mechanisms. In other words, the choice of a particular ontological account of mechanisms is perpendicular to the epistemological issue of the explanatory power of mechanisms. The understanding of mechanisms in structural modelling is rather epistemic—it is more concerned with how we can (causally) make sense of the phenomenon to be explained. Structural modelling achieves this goal by offering a story about a mechanism. Such epistemic understanding is akin to ‘mechanism schemata’, as discussed in Machamer (2004), Machamer, Darden and Craver (2000), and Darden and Craver (2002).

Needless to say, this is not to suggest that structural models provide immutable and eternal causal explanations of social phenomena. Explanation is intrinsically relative and partial, that is relative to the specific conceptual framework and dependent on available empirical and theoretical

information. This means that nothing prevents future explanations to discard previous ones. Also, such causal explanations involve a stopping rule in order to avoid an otherwise ad infinitum chain of ‘explaining the explanatory’.

4 On the counterfactual approach to causal analysis

In the present paper, we defend a structural modelling approach to causality. Other frameworks have however been proposed in the literature. In particular an approach based on counterfactuals has also gained recognition. In this section, we first point out that counterfactuals may be of help for causal reasoning. Next we question counterfactuals as a basis for causality.

4.1 Counterfactual questions as an aid to inferring causal relations

In observational studies especially, one is never sure that an observed association between a putative cause and an effect reflects a causal relation between the two, as the association might also be due to the presence of unknown latent confounding variables which have not been controlled for. This important issue has been tackled in particular by epidemiologists several decades ago; they have developed a series of criteria which have been recommended for drawing causal inferences. The criteria have been systemised by Bradford Hill (1966); more recent versions can be found *e.g.* in Beaglehole, Bonita and Kjellström (1994) or in Rothman and Greenland (1998). Briefly, these criteria refer to the strength of the association, to the consistency of the observation of the association in different populations and settings, to the temporality of the relation (a cause must precede its effect in time), to the dose-response gradient, to the plausibility and coherence of the cause-and-effect interpretation with background knowledge, and to the reversibility of the association: in the absence of the cause one should not observe the effect. This last criterion can be stated in counterfactual terms: if the cause had not been, the effect would not have occurred.

Except most probably for the criterion of temporality (see Wunsch, Russo and Mouchart, 2010), all the other criteria suffer from exceptions and reservations (see Rothman and Greenland, *op. cit.*). For example, an implausible explanation according to prior beliefs and current knowledge might actually be correct and possibly lead to progress in science. Lack of consistency might be due to the fact that the cause produces its effect only under particular circumstances. As to the counterfactual criterion, several counterarguments will be outlined in the following sub-section. To take but one issue, suppose one is interested in the effect of ‘Education’ (putative cause) on ‘Self-rated health’ (outcome). Following the counterfactual criterion, one would ask what would happen to self-rated health for counterfactual values of education. In order to assess the causal effect of, *e.g.* higher education on self-rated health, we have to ask what would happen to the individuals not benefiting from a higher education. This counterfactual is however highly ambiguous. Not having a higher

education may mean different things: secondary schooling or less, just secondary schooling, primary schooling or less, just primary schooling, technical education, or no formal schooling at all. Which one is the counterfactual to evaluate? Each counterfactual would correspond to a different model, leading to a different measure of the effect.

Finally, all these criteria can be of help in inferring causality but they are not foolproof. In actual social situations, taking into account the complexity of the network of causes and effects, some of these criteria will be satisfied and others very often not. For example, the effect might appear even in the absence of the putative cause due to the presence of other causes, in a multiple-cause multiple-effect framework. It remains for the scientist to check all the evidence available and then to decide upon the existence or not of a possible causal relation. No set of criteria—and in particular no single criterion—will automatically lead to this result.

4.2 Counterfactuals as a basis for causality

We have seen that the counterfactual argument is but one of several criteria to be of help in inferring causal relations. As with the other criteria, it is neither necessary nor sufficient for this purpose. Nevertheless, notwithstanding its shortcomings, the counterfactual argument has been taken as a basis for causality in what is now known as Rubin’s causal model (Holland, 1986), which has become influential in various spheres of social research.

Rubin (1974) formalised the basic idea behind the counterfactual model as follows. Consider comparing two ‘treatments’, E and C , in the case of a headache. Let E represent taking two aspirins and C drinking just a glass of water. The potential outcomes Y relating to these two treatments may then be written as two random variables, namely $Y(E)$ and $Y(C)$. The causal effect of treatment E versus treatment C on Y for a particular subject j (say, Jones) observed at time $t+k$ is then defined as $Y_j(E) - Y_j(C)$, *i.e.* the differential headache response to taking the aspirins or just drinking a glass of water at time t . If we consider n subjects instead of only one subject, we have one causal effect $Y_j(E) - Y_j(C)$ per subject j . The average causal effect for this group of n persons can then be written $\sum_j [Y_j(E) - Y_j(C)]/n$, the sum extending from $j = 1$ to n . Rubin’s solution is often called the potential outcome (or response) approach, the two potential outcomes being in this simple case $Y_j(E)$ and $Y_j(C)$ for each j . The causal effect may differ from one individual to the other; thus a ‘typical’ causal effect (Rubin’s own term) is obtained as above by taking the average (or any other summary measure) of the individual causal effects.

In the actual world, one never observes at the same time for the same individual both $Y(E)$ and $Y(C)$. People are assigned or experience either E or C but not both at the same time, *i.e.* not the fact and its counterfact. Rubin has nevertheless shown that randomization and matching are two approaches measuring the causal effect in experimental and observational studies, though randomization cannot often be used in the social sciences and perfect matching is hardly possible

in practice. In many actual situations in observational research, the assignment of units to the case and control groups is often prone to selection bias. Thus the likelihood of treatment on the one hand and of the outcome on the other hand are not independent. In this case, one must control as best as possible for the assignment factors which have an impact on the outcome. Clearly the actual challenge for the researcher is to ensure that that ‘all relevant factors’ have been controlled for.

The counterfactual basis of Rubin’s potential outcome framework raises some important epistemological issues, which we have discussed in another paper to which the reader is referred (Russo, Wunsch and Mouchart, 2010). The first two issues are quite often discussed in the literature. One concerns the soundness of the counterfactual approach given that the counter-fact is not observed, thus resulting in a lack of sound empirical basis. The other concerns the alternative between a counterfactual model measuring effects of causes and other models concerned instead with the causes of effects, the counterfactual model being ill-suited for searching for the causes of an effect. The next two issues concern the concepts that back up the experimental method and which are subsumed by the counterfactual model: manipulation and randomization. Though useful of course, neither is necessarily required for inferring causal relations. Our arguments hinge on the idea that manipulation (contra the view of counterfactualist theorists) is not the methodological (nor the epistemological) basis of causal analysis.

This can be seen analysing different readings of an equation of the type $Y = \beta X + \epsilon$. It is important to notice that such an equation can be given a *variational* and a *manipulationist* reading. Let us explain further. At bottom, this equation is read thus: variations in the putative cause X are accompanied by variations in the putative effect Y . How much Y varies in response to the variation in X is quantified by the parameter β . The manipulationist reading is then derived from this basic variational reading as follows. In an experimental setting, manipulations on X make X varying, such that Y varies accordingly. In a controlled experiment, therefore, co-variations in X and Y are due to manipulations, unlike in observational studies. However, for causal inference it is not necessary that X has been manipulated. It could well be, as is typically the case in observational studies in social science, that statements about *co-variations* are based on calculated statistical correlations between the variables. In such cases causal inference has to be supported by further considerations about structural stability, exogeneity, possible confounders, the chosen recursive decomposition and background knowledge. Differently put, we are in a position of making causal inference even in the absence of manipulation, provided that the model is deemed *structural*, to the best of our knowledge.

The next issue deals with complex mechanisms. Even in seemingly simple situations one has to face an issue of multiple causes and multiple effects, involving more than one mechanism at a time. In practice, it is usually not sufficient, to come back to our example, to compare Jones taking the aspirin to Jones not taking the aspirin. One must control for the factors possibly confounding the relationship between aspirin and headache. The two Jones should be matched on all the relevant

covariates which could lead to confounding. However if there are many covariates including latent ones, as is most often the case in social sciences, it will often be impossible to match on all these covariates. Finally, our last argument makes a critical assessment of the simplistic analogies and parallelisms that have oft been made between the counterfactual model developed in statistics and the counterfactual analysis of causation developed by philosophers. The former, although based on individual-level data, is generic, whilst the latter is single-case, that is it concerns a particular causal relation taking place in a given time and place.

Following the arguments developed in Russo, Wunsch and Mouchart (op. cit.), our conclusion is that though counterfactual questions can help the researcher in drawing causal connections, counterfactuality per se cannot serve as the basis for inferring causal relations.

5 Discussion and conclusion

Structural modelling, in the way it has been presented in this paper, aims to provide an approach to causal explanation of social phenomena. We made the effort to go beyond structural modelling intended as a mere statistical machinery and we attempted to provide a view on *modelling*, that is a scientific procedure that goes from data collection and hypothesis formulation to statistical testing and interpretation of results. Of course the implementation of scientific procedures, no matter how well and clearly defined, is easier said than done. Structural modelling may indeed be hard to put in place and is not immune from criticisms.

Many scientists remain sceptical about the practical usefulness of structural modelling, even if they recognise that “understanding and identifying causal mechanisms is, perhaps, the primary driving force of science” (Holland, 2001, p. 224). For Holland, for instance, the danger lies in the fact that almost ‘anything’ can be considered as a cause “because we are just talking rather than doing”, *i.e.* setting up ‘treatments’ or ‘interventions’ (Holland 2001 p. 225). Actually, a causal mechanism does not appear from nowhere. Nor does it necessarily result from adding more and more variables to the predictive set (Sobel, 2000). As we have argued in this paper, a structural model should be based on the best available knowledge one has of the field; all postulated relations should be accounted for. In particular, to avoid loss of exogeneity the model should incorporate those variables deemed to be responsible for possible assignment bias.

The structural modelling framework also has its problems, of course. First of all, known confounders can be incorporated into the model only upon the condition that indicators of these confounders are available in the data set. In many situations, especially when one uses secondary data (*i.e.* data collected by others), no information is available for some of the variables in the model. Confounding bias may not be avoidable then, though in some cases omitted variable bias can be controlled for by fixed effects regression or by instrumental variables regression (Stock and Watson,

2003). Unknown latent confounders may however still bias the results.

A second issue relates to the temporal ordering of variables. In many cases, the observation window is too wide to observe the exact ordering of the events in time. Depending on the length of the window, it is not always possible to state if the putative cause occurred before or after the effect. For example, if migration and occupational change are recorded on a yearly basis in the data set, one does not know on an annual basis whether migration has occurred before or after occupational change, if both have happened during the year. To illustrate further, let us take the example of commercial contracts. Consider the typical situation in the tour-operator market. Suppose that the tour-operator, acting as a price setter, prints in January a catalogue for the coming season. If the price is not altered within the year, the quantity observed on an annual basis may be assumed to have been contracted by a price-taker demand side, and the price construed as an exogenous factor of the demand. But if the price is modified, say around Easter time, annual data will not allow disentangling the demand from the supply, even under a price-taker demand, as long as different quantities have been contracted under different prices and the price changes have been operated under the pressure of the demand. Recursivity of the decomposition is jeopardized in these examples, as causes and effects appear to be simultaneous.

Another issue concerns invariance or stability. It will often be difficult, in the social field, to repeat the research on comparable populations, in order to check if the results remain invariant under changes of context, even if the reference population is well-defined. For example, it has been shown (Laplante and Flick, 2010) that cohabitation and health patterns are drastically different between Quebec and Ontario, two neighbouring provinces in Canada with rather similar populations. Very often, results differ from one study to another, and it will usually hardly be possible to determine if this is due to an inadequate model, insufficient data, or differential contexts.

Moreover, structural modelling requires reliable prior information on the putative recursive decomposition of the data matrix. A major drawback is that in many cases one only has a scant knowledge of the underlying mechanism. Some of the sub-mechanisms thus remain black boxes and, in this situation, several alternative models may fit the data equally well. It is then difficult to discriminate between alternatives that equally account for the same data set. In this case, a good descriptive analysis or exploratory data analysis might be more useful than poor structural modelling.

Yet another danger is that because background knowledge plays a central role, it might be given too much weight. But background knowledge is not meant to be infallible nor immutable. Instead, background knowledge has to be the bridge between established knowledge and establishing knowledge. Established scientific knowledge is (and ought to be) used to formulate the causal hypothesis and to evaluate the plausibility of results on theoretical grounds. But causal analysis also participates in establishing new knowledge by proposing new conceptual frameworks and testing

them against new data or new interpretations of existing data. This reflects the idea that science is far from being monolithic, discovering immutable and eternal truths. If the model fits the data, if the relations are sufficiently invariant and congruent with background knowledge, then we can say, to the best of our knowledge, that we modelled a causal mechanism that explains a given social phenomenon. But what if one of these conditions fails? A negative result may trigger further research by improving the modelling strategies, or by collecting new data, thus leading to new discoveries that, perhaps, discard background knowledge.

Actually, these various criticisms should be evaluated from three perspectives. Firstly, poor knowledge of the underlying structure, or the lack of consensus about it, should remind us that causality attribution is only relative to a well-specified structural model: if the modelling procedure is weak, causal attribution is weak too. Secondly, an analysis confined to statistical associations (such as correlations *etc.*) does not allow causal attribution. Statistical associations and regularities are surely an essential ingredient for building a structural model but, we insist, causality can only be relative to a structural model. Thirdly, a major role of structural modelling, in particular of the recursive decomposition, is precisely to tentatively provide a structure in a world where apparently ‘everything depends upon everything’, and consequently ‘almost everything is a cause’, which would lead to a situation where no action is possible or where effects of treatments are not identifiable.

What is the contribution of a structural modelling approach to causal attribution? Firstly, a structural approach can deal, within a same framework, both with the effects of causes and with the causes of effects. Secondly, the specification of the recursive decomposition allows distinguishing and ordering the various sub-mechanisms within a global one. From this, exogenous variables can eventually be given a causal interpretation.

And what are the implications of a structural modelling approach for the practicing scientist? In such an approach, the quality of the results crucially depends upon the quality of the process of model building and model testing. Thus, from the very start, the scientist should make the causal hypotheses explicit until s/he obtains a recursive decomposition in agreement with all the available background knowledge. Then the scientist should carefully check for structural stability, or invariance, of each component of the recursive decomposition and try to condense the latter as much as possible. Finally, the scientist should take into account the possible non-observability of some of the relevant variables and evaluate its impact, in the resulting model, on exogeneity and on the identifiability of the parameters of interest.

References

- BARNDORFF-NIELSEN O. (1978), *Information and Exponential Families in Statistical Theory*, New York: John Wiley.

- BRADFORD HILL A. (1966), *Principles of Medical Statistics*, The Lancet Limited, London.
- BEAGLEHOLE R., BONITA R., KJELLSTRÖM T. (1994), *Éléments d'Épidémiologie*, World Health Organization, Geneva.
- BECHTEL, W., ABRAHAMSEN, A. (2005), Explanation: a mechanist alternative, *Studies in the History and Philosophy of the Biological and Biomedical Sciences*, **36**, 421-441.
- CARTWRIGHT, N. (2007), *Hunting causes and using them: approaches in philosophy and economics*. Cambridge University Press, Cambridge.
- DARDEN, L., CRAVER, C. (2002), Strategies in the interfield discovery of the mechanism of protein synthesis, *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, **33**(1), 1-28.
- ENGLE R.F., HENDRY D.F., RICHARD J.-F. (1983), Exogeneity, *Econometrica*, **51**(2), 277-304.
- FENNELL, D.J. (2011), The structural theory of causation. In P. McKay Illari, F. Russo, J. Williamson (eds), *Causality in the Sciences*, Oxford University Press.
- FLORENS J.-P., MOUCHART M. (1980), Initial and sequential reduction of Bayesian experiments, CORE Discussion Paper 8015, Université catholique de Louvain, Louvain-la-Neuve.
- FLORENS J.-P., MOUCHART M. (1985), Conditioning in Dynamic models, *Journal of Times Series Analysis*, **53** (1), 15-35.
- FLORENS J.-P., MOUCHART M., ROLIN J.-M. (1980), Réduction dans les expériences bayésiennes séquentielles, *Cahiers du Centre d'Études de Recherche Opérationnelle*, **22**(3-4), 353-362.
- FLORENS J.-P., MOUCHART M., ROLIN J.-M. (1990), *Elements of Bayesian Statistics*, New York: Marcel Dekker.
- HOLLAND P.W. (1986), Statistics and causal inference, *Journal of the American Statistical Association*, **81**(396), 945-960.
- HOLLAND P.W. (2001), The false linking of race and causality: lessons from standardized testing, *Race & Society*, **4**, 219-233.
- HOOVER, K. (2001), *Causality in macroeconomics*, Cambridge: Cambridge University Press.
- HOOVER, K. (2011), Counterfactuals and Causal Structure. In P. McKay Illari, F. Russo, J. Williamson (eds), *Causality in the Sciences*, Oxford University Press.

- KOOPMANS T.C. (1950), When is an equation system complete for statistical purposes?, in Koopmans T.C. (ed.), *Statistical Inference in Dynamic Economic Models*, New-York: John Wiley, 393-490.
- LAPLANTE B., FLICK C. (2010), Le mariage, l'union de fait et la santé des conjoints au Québec et en Ontario, *Cahiers Québécois de Démographie*, **39**(1), 1-26.
- MACHAMER, P. (2004), Activities and Causation: The Metaphysics and Epistemology of Mechanisms, *International Studies in the Philosophy of Science*, **18**(1), 27-39.
- MACHAMER, P., DARDEN, L., CRAVER, C. F. (2000), Thinking about Mechanisms, *Philosophy of Science*, **67**, 1-25.
- MOUCHART M., OULHAJ A. (2000), On Identification in Conditional Models, *Discussion Paper DP0015*, Institut de statistique, Université catholique de Louvain, Louvain-la-Neuve.
- OULHAJ A., MOUCHART M. (2003), The Role of the Exogenous Randomness in the Identification of Conditional Models, *Metron*, **LXI**(2), 267-283.
- MOUCHART, M., RUSSO, F. (2011), Causal explanation: mechanisms and recursive decomposition. In P. McKay Illari, F. Russo, J. Williamson (eds), *Causality in the sciences*, Oxford University Press.
- MOUCHART, M., RUSSO, F., WUNSCH, G. (2009), Structural modelling, exogeneity and causality. In H. Engelhardt, H-P Kohler, A. Prskawetz (Ed.), *Causal analysis in population studies: concepts, methods, applications*, Dordrecht: Springer, 59-82.
- PEARL J. (2000), *Causality. Models, Reasoning, and Inference*, Cambridge University Press, Cambridge.
- PEARL J. (2011), The structural theory of causation. In P. McKay Illari, F. Russo, J. Williamson (eds), *Causality in the Sciences*. Oxford University Press.
- POPPER, K. (1959), *The logic of scientific discovery*, London: Hutchinson.
- ROTHMAN K.J., GREENLAND S. (1998), *Modern Epidemiology*, 2nd edition, Lippincott-Raven, Philadelphia.
- RUBIN D. (1974), Estimating causal effects of treatments in randomized and non randomized studies, *Journal of Educational Psychology*, **66**(5), 688-701.
- RUSSO F.(2009), *Causality and Causal Modelling in the Social Sciences: Measuring Variations*, Methodos Series Vol.5, Springer, New York.

- RUSSO, F. (2011a), Correlational data, causal hypotheses, and validity, *Journal for General Philosophy of Science*, In press.
- RUSSO, F. (2011b), On empirical generalisations, submitted.
- RUSSO, F. (2011c), Explaining causal modelling. Or, what a causal model ought to explain, in D'Agostino, M., Giorello G., Laudisa F., Pievani T., Sinigaglia, C. (eds), *New Essays in Logic and Philosophy of Science*, College Publications, London.
- RUSSO F., WUNSCH G., MOUCHART M. (2010), Inferring Causality through Counterfactuals in Observational Studies, *Discussion Paper* 1029, Institute of Statistics, Biostatistics and Actuarial Sciences (ISBA), Université catholique de Louvain, Louvain-la-Neuve.
- SIMON, H. (1953), Causal ordering and identifiability. In W. C. Hood, T. C. Koopmans (Ed.), *Studies in econometric method* (pp. 49-74). New York: Wiley.
- SIMON, H. (1954), Spurious correlation: a causal interpretation, *Journal of the American Statistical Association*, 49(267), 467-479.
- SOBEL M.E. (2000), Causal inference in the social sciences, *Journal of the American Statistical Association*, **95**(450), 647-651.
- SPIRITES P., GLYMOUR C., SCHEINES R. (2000), *Causation, Prediction and Search*, second edition, Cambridge (Mass.): The MIT Press.
- STOCK J.H., WATSON M.W. (2003), *Introduction to Econometrics*, Addison-Wesley, Boston.
- WOODWARD, J. (2003), *Making things happen: a theory of causal explanation*, Oxford University Press, Oxford.
- WUNSCH, G. (2007), Confounding and control, *Demographic Research*, **16**:15-35.
- WUNSCH G., RUSSO F., MOUCHART M. (2010), Do we necessarily need longitudinal data to infer causal relations?, *Bulletin de Méthodologie Sociologique*, **106**(1), 1-14.