
Causal Consistency of Structural Equation Models

Paul K. Rubenstein^{*12}, Sebastian Weichwald^{*13}, Stephan Bongers⁴, Joris M. Mooij⁴
Dominik Janzing¹, Moritz Grosse-Wentrup¹, Bernhard Schölkopf¹

^{*}Equal contribution

¹Empirical Inference, MPI for Intelligent Systems, ²Machine Learning Group, University of Cambridge,

³Max Planck ETH Center for Learning Systems, ⁴Informatics Institute, University of Amsterdam

Abstract

Complex systems can be modelled at various levels of detail. Ideally, causal models of the same system should be consistent with one another in the sense that they agree in their predictions of the effects of interventions. We formalise this notion of consistency in the case of Structural Equation Models (SEMs) by introducing *exact transformations* between SEMs. This provides a general language to consider, for instance, the different levels of description in the following three scenarios: (a) models with large numbers of variables versus models in which the ‘irrelevant’ or unobservable variables have been marginalised out; (b) micro-level models versus macro-level models in which the macro-variables are aggregate features of the micro-variables; (c) dynamical time series models versus models of their stationary behaviour. Our analysis stresses the importance of well specified interventions in the causal modelling process and sheds light on the interpretation of cyclic SEMs.

1 INTRODUCTION

Physical systems or processes in the real world are complex and can be understood at various levels of detail. For instance, a gas in a volume consists of a large number of molecules. But instead of modelling the motions of each particle individually (micro-level), we may choose to consider macroscopic properties of their motions such as temperature and pressure. Our decision to use such macroscopic properties is first necessitated by practical considerations. Indeed, for all but extremely simple cases, making a measurement of all the individual molecules is practically impossible and our resources insufficient for modelling the $\sim 10^{22}$ particles present per litre of ideal gas.

Furthermore, the decision for a macroscopic description level is also a pragmatic one: if we only wish to reason about temperature and pressure, a model of 10^{22} particles is ill-suited.

Statistical physics explains how higher-level concepts such as temperature and pressure arise as statistical properties of a system of a large number of particles, justifying the use of a macro-level model as a useful transformation of the micro-level model [1]. However, in many cases aggregate or indirect measurements of a complex system form the basis of a macroscopic description of the system, with little theory to explain whether this is justified or how the micro- and macro-descriptions stand in relation to each other.

Due to deliberate modelling choice or the limited ability to observe a system, differing levels of model descriptions are ubiquitous and occur, amongst possibly others, in the following three settings:

- (a) Models with large numbers of variables versus models in which the ‘irrelevant’ or unobservable variables have been marginalised out [3]; e. g. modelling blood cholesterol levels and risk of heart disease while ignoring other blood chemicals or external factors such as stress.
- (b) Micro-level models versus macro-level models in which the macro-variables are aggregate features of the micro-variables [4, 5, 9, 11, 17]; e. g. instead of modelling the brain as consisting of 100 billion neurons it can be modelled as averaged neuronal activity in distinct functional brain regions.
- (c) Dynamical time series models versus models of their stationary behaviour [6, 8, 11–13, 15]; e. g. modelling only the final ratios of reactants and products of a time evolving chemical reaction.

In the context of causal modelling, such differing model levels should be consistent with one another in the sense that they agree in their predictions of the effects of interventions. The particular causal models we focus on in this

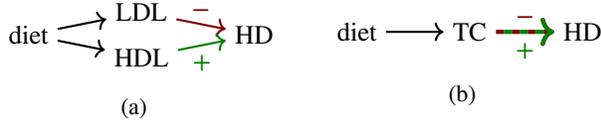


Figure 1: As illustrated by (a), the current consensus is that LDL (resp. HDL) has a negative (resp. positive) effect on heart disease (HD). Considering $TC = LDL + HDL$ to be a causal variable as in (b) leads to problems: two diets promoting raised LDL levels and raised HDL levels have the same effect on TC but opposite effects on heart disease. Hence different studies come to contradictory conclusions about the effect of TC on heart disease.

paper are Structural Equation Models (SEMs, Section 2, Section 3) [16, 19].

In Section 4, we introduce the notion of an exact transformation between two SEMs, providing us with a general framework to evaluate when two models can be thought of as causal descriptions of the same system. An important novel idea of this paper is to explicitly make use of a natural ordering on the set of interventions. On a high level, if an SEM can be viewed as an exact transformation of another SEM, we are provided with an explicit correspondence between the two models in such a way that causal reasoning on both levels is consistent. We discuss this notion of consistency in detail in Sections 4.4 and 4.5.

In Section 5 we apply this mathematical framework and prove the exactness of transformations belonging to each of the three categories listed above, with practical implications for the following questions in causal modelling: When can we model only a subsystem of a more complex system? When does a micro-level system admit a causal description in terms of macro-level features? How do cyclic SEMs arise? The fact that these distinct problems can all be considered using the language of transformations between SEMs demonstrates the generality of our approach. We close in Section 6 with a discussion.

1.1 A HISTORICAL MOTIVATION: CHOLESTEROL AND HEART DISEASE

In the following we give an example of the problems that can arise when there exists no consistent correspondence between two causal models, i. e. neither model can be viewed as an exact transformation of the other. This example falls into category (b) of the differing model levels listed above and was used by [18] to illustrate problems in the causal modelling process.

Historically, the level of total cholesterol in the blood (TC) was thought to be an important variable in determining risk of heart disease (HD). To investigate this, differ-

ent experiments were carried out in which patients were assigned to different diets in order to raise or lower TC. Conflicting evidence was found by different experiments: some found that higher TC had the effect of lowering HD, while others found the opposite (cf. Figure 1b) [20, 21].

From our point of view, this problem (seemingly conflicting studies) arose from trying to perform an ‘invalid’ transformation of the ‘true’ underlying model (cf. Figure 1a). According to the American Heart Association, the current scientific consensus is that the two types of blood cholesterol, low-density lipoprotein (LDL) and high-density lipoprotein (HDL), have a negative and positive effect on HD respectively. Assigning diets that raise LDL or HDL both raise TC but have different effects on HD. It is therefore not possible to transform the model in Figure 1a into the model in Figure 1b without leading to conflict: in order to reason about the causes of HD we need to consider the variables LDL and HDL separately.

2 STRUCTURAL EQUATION MODELS

SEMs are a widely used framework in causal modelling, with applications in neuroscience, economics and the social sciences [2, 16]. In this section we introduce them as an abstract mathematical object; in Section 3 we describe their use as a causal modelling tool. Readers already familiar with SEMs should note that our definition is more general and deviates from the standard definition of SEMs in the following ways: we do not require that all possible perfect interventions be modelled; we do not assume independence of exogenous variables;¹ and we do not require acyclicity.

Definition 1 (Structural Equation Model (SEM)). *Let \mathbb{I}_X be an index set. An SEM \mathcal{M}_X over variables $X = (X_i : i \in \mathbb{I}_X)$ taking value in \mathcal{X} is a triple $(S_X, \mathcal{I}_X, \mathbb{P}_E)$ where*

- S_X is a set of structural equations, i. e. it is a set of equations $X_i = f_i(X, E_i)$ for $i \in \mathbb{I}_X$;
- (\mathcal{I}_X, \leq_X) is a subset of all perfect interventions equipped with a natural partial ordering (see below), i. e. it is an index set where each index corresponds to a particular perfect intervention on some of the X variables;
- \mathbb{P}_E is a distribution over the exogenous variables $E = (E_i : i \in \mathbb{I}_X)$;
- with \mathbb{P}_E -probability one, under any intervention $i \in \mathcal{I}_X$ there is a unique solution $x \in \mathcal{X}$ to the intervened structural equations. This ensures that for any

¹Exogenous variables are also referred to as *noise variables* in the literature. Our relaxation of the assumption of independent exogenous variables means our models may be considered a type of semi-Markovian causal model.

intervention $i \in \mathcal{I}_X$, \mathcal{M}_X induces a well-defined distribution over \mathcal{X} .²

In an SEM, each X_i is a function of the X -variables and the exogenous variable E_i . In this mathematical model, a perfect intervention on a single variable $\text{do}(X_i = x_i)$ is realised by replacing the structural equation for variable X_i in S_X with $X_i = x_i$. Perfect interventions on multiple variables, e.g. $\text{do}(X_i = x_i, X_j = x_j)$, are similarly realised by replacing the structural equations for each variable individually. Elements of \mathcal{I}_X correspond to perfectly intervening on a subset of the X variables, setting them to some particular combination of values.

\mathcal{I}_X has a natural partial ordering in which, for interventions $i, j \in \mathcal{I}_X$, $i \leq_X j$ if and only if i intervenes on a subset of the variables that j intervenes on and sets them equal to the same values as j . For example, $\text{do}(X_i = x_i) \leq_X \text{do}(X_i = x_i, X_j = x_j)$.³ The observation that this structure is important is a contribution of this paper. We make crucial use of it in the next section.

The purpose of the following example is to illustrate how SEMs are written in our notation and to provide an example of a restricted set of interventions \mathcal{I}_X .

Example 2. Consider the following SEM defined over the variables $\{B_1, B_2, L\}$

$$\begin{aligned} S_X &= \{B_1 = E_1, B_2 = E_2, L = \text{OR}(B_1, B_2, E_3)\} \\ \mathcal{I}_X &= \{\emptyset, \text{do}(B_1 = 0), \text{do}(B_2 = 0), \\ &\quad \text{do}(B_1 = 0, B_2 = 0)\}, \\ \{E_1, E_2, E_3\} &\stackrel{iid}{\sim} \text{Bernoulli}(0.5) \end{aligned}$$

where by the element $\emptyset \in \mathcal{I}$ we denote the null-intervention corresponding to the unintervened SEM.

3 SEMS FOR CAUSAL MODELLING

In addition to being abstract mathematical objects, SEMs are used in causal modelling to describe distributions of variables and how they change under interventions [16]. The do-interventions as abstract manipulations of SEMs are understood as corresponding to actual (or potentially

²That is, with probability one over the exogenous variables E , for each draw $E = e$ there exists a unique value $x \in \mathcal{X}$ such that e and x satisfy the intervened structural equations. The distribution of E in conjunction with S_X then implies a distribution over \mathcal{X} for each intervention $i \in \mathcal{I}_X$ via these unique solutions. If the SEM is acyclic, this is always satisfied; we impose this condition because we also consider *cyclic* SEMs [3].

³Informally, this means that j can be performed after i without having to change or undo any of the changes to the structural equations made by i . Not all pairs of elements must be comparable: for instance, if $i = \text{do}(X_1 = x_1)$ and $j = \text{do}(X_2 = x_2)$, then neither $i \leq_X j$ nor $j \leq_X i$.

only hypothetical) physical implementations in the real world, i. e. the model is ‘rooted in reality’. For instance, if a binary variable B_1 in an SEM reflects whether a light bulb is emitting light, then $\text{do}(B_1 = 0)$ could be achieved by flipping the light switch or by removing the light bulb.

The SEM in Example 2 could be thought of as a simple causal model of two light bulbs B_1 and B_2 and the presence of light L in a room with a window. Suppose that we have no access to the light switch and there are no curtains in the room but that we can intervene by removing the light bulbs. We can model this restricted set of interventions by \mathcal{I}_X , i. e. the do-intervention on the SEM side $\text{do}(B_1 = 0)$ corresponds to removing the light bulb B_1 .

The partial ordering of \mathcal{I}_X corresponds to the ability to compose physical implementations of interventions. The fact that we can first remove light bulb B_1 ($\text{do}(B_1 = 0)$) and then afterwards remove light bulb B_2 (resulting in the combined intervention $\text{do}(B_1 = 0, B_2 = 0)$) is reflected in the partial ordering via the relation $\text{do}(B_1 = 0) \leq_X \text{do}(B_1 = 0, B_2 = 0)$.

4 TRANSFORMATIONS BETWEEN SEMS

We now work towards our definition of an exact transformation between SEMs. Our core idea is to analyse the correspondence between different levels of modelling by considering one model to be a transformation of the other. We discuss in Section 4.4 how causal reasoning in two SEMs relate when one SEM can be viewed as an exact transformation of the other and in Section 4.5 we illustrate what can go wrong when this is not the case.

4.1 DISTRIBUTIONS IMPLIED BY AN SEM

Usually, a statistical model implies a single joint distribution over all variables once its parameters are fixed. SEMs are different in that, once the parameters are fixed, an SEM implies a family of joint distributions over the random variables, one for each intervention. That is, for each intervention $i \in \mathcal{I}_X$, the SEM \mathcal{M}_X defines a distribution over \mathcal{X} which we denote by $\mathbb{P}_X^{\text{do}(i)}$. Throughout, we will denote the null-intervention corresponding to the unintervened setting by $\emptyset \in \mathcal{I}_X$. We can write the poset of all distributions implied by the SEM \mathcal{M}_X as

$$\mathcal{P}_X := \left(\left\{ \mathbb{P}_X^{\text{do}(i)} : i \in \mathcal{I}_X \right\}, \leq_X \right)$$

where \leq_X is the partial ordering inherited from \mathcal{I}_X , i. e. $\mathbb{P}_X^{\text{do}(i)} \leq_X \mathbb{P}_X^{\text{do}(j)} \iff i \leq_X j$.⁴

⁴More formally, one would need to define \mathcal{P}_X to be the poset of *tuples* $(i, \mathbb{P}_X^{\text{do}(i)})$ to avoid problems in the case that $\mathbb{P}_X^{\text{do}(i)} =$

Note that \mathcal{P}_X contains all of the information in \mathcal{M}_X about the different distributions implied by the SEM and, importantly, how they are related via the interventions.⁵

4.2 TRANSFORMATIONS OF RANDOM VARIABLES

Suppose we have a function $\tau : \mathcal{X} \rightarrow \mathcal{Y}$ which maps the variables of the SEM \mathcal{M}_X to another space \mathcal{Y} . Observe that since X is a random variable, $\tau(X)$ is also a random variable. For any distribution \mathbb{P}_X on \mathcal{X} we thus obtain the distribution of the variable $\tau(X)$ on \mathcal{Y} as $\mathbb{P}_{\tau(X)} = \tau(\mathbb{P}_X)$ via the push-forward measure.

In particular, for each intervention $i \in \mathcal{I}_X$ we can define the induced distribution $\mathbb{P}_{\tau(X)}^i = \tau(\mathbb{P}_X^{\text{do}(i)})$. We can write the poset of distributions on \mathcal{Y} that are induced by the original SEM \mathcal{M}_X and the transformation τ as

$$\mathcal{P}_{\tau(X)} := \left(\left\{ \mathbb{P}_{\tau(X)}^i : i \in \mathcal{I}_X \right\}, \leq_X \right)$$

where \leq_X is the partial ordering inherited from \mathcal{P}_X (and in turn from \mathcal{I}_X).

$\mathcal{P}_{\tau(X)}$ is just a structured collection of distributions over \mathcal{Y} , indexed by interventions \mathcal{I}_X on the \mathcal{X} -level; importantly, the indices are *not* interventions on the \mathcal{Y} -level.

4.3 EXACT TRANSFORMATIONS BETWEEN SEMS

Although $\mathcal{P}_{\tau(X)}$ is a poset of distributions over \mathcal{Y} , there does not necessarily exist an SEM \mathcal{M}_Y over \mathcal{Y} that implies it. For instance, if there is some intervention $i \in \mathcal{I}_X \setminus \{\emptyset\}$ such that none of the variables Y_i is constant under the distribution $\mathbb{P}_{\tau(X)}^i$, then $\mathbb{P}_{\tau(X)}^i$ could not possibly be expressed as arising from a do-intervention $j \in \mathcal{I}_Y \setminus \{\emptyset\}$ in any SEM over \mathcal{Y} .⁶

The case in which there *does* exist an SEM \mathcal{M}_Y that implies $\mathcal{P}_{\tau(X)}$ is special, motivating our main definition.

Definition 3 (Exact Transformations between SEMs). *Let \mathcal{M}_X and \mathcal{M}_Y be SEMs and $\tau : \mathcal{X} \rightarrow \mathcal{Y}$ be a function. We say \mathcal{M}_Y is an exact τ -transformation of \mathcal{M}_X if there exists a surjective order-preserving map $\omega : \mathcal{I}_X \rightarrow \mathcal{I}_Y$ such that*

$$\mathbb{P}_{\tau(X)}^i = \mathbb{P}_Y^{\text{do}(\omega(i))} \quad \forall i \in \mathcal{I}_X$$

⁵For example, the distribution over the variables X in the observational setting, \mathbb{P}_X^\emptyset , changes to $\mathbb{P}_X^{\text{do}(i)}$ if we implement the intervention $\text{do}(i)$, and the partial ordering contains all information about which interventions can be composed.

⁶This problem is elaborated upon in [7].

where $\mathbb{P}_{\tau(X)}^i$ is the distribution of the \mathcal{Y} -valued random variable $\tau(X)$ with $X \sim \mathbb{P}_X^{\text{do}(i)}$.

Order-preserving means that $i \leq_X j \implies \omega(i) \leq_Y \omega(j)$. It is important that the converse need not in general hold as this would imply that ω is injective,⁷ and hence also bijective. This would constrain the ways in which \mathcal{M}_Y can be ‘simpler’ than \mathcal{M}_X .⁸ That ω is surjective ensures that for any do-intervention $j \in \mathcal{I}_Y$ on \mathcal{M}_Y there is at least one corresponding intervention on the \mathcal{M}_X level, namely an element of $\omega^{-1}(\{j\}) \subseteq \mathcal{I}_X$. The following two results follow immediately from the definition (cf. proofs in Appendix A).

Lemma 4. *The identity mapping and permuting the labels of variables are both exact transformations.*

This is a good sanity check; it would be problematic if this were not the case and the labelling of our variables mattered. Similarly, compositions of exact transformations are also exact.

Lemma 5 (Transitivity of exact transformations). *If \mathcal{M}_Z is an exact τ_{ZY} -transformation of \mathcal{M}_Y and \mathcal{M}_Y is an exact τ_{YX} -transformation of \mathcal{M}_X , then \mathcal{M}_Z is an exact $(\tau_{ZY} \circ \tau_{YX})$ -transformation of \mathcal{M}_X .*

The following theorem is a consequence of the fact that ω is order-preserving. This is a mathematical formalisation of the sense in which an exact transformation preserves causal reasoning, which will be elaborated upon in the next subsection.

Theorem 6 (Causal consistency under exact transformations). *Suppose that \mathcal{M}_Y is an exact τ -transformation of \mathcal{M}_X and ω is a corresponding surjective order-preserving mapping between interventions. Let $i, j \in \mathcal{I}_X$ be interventions such that $i \leq_X j$. Then the following diagram commutes:*

$$\begin{array}{ccccc} & & \text{do}(i) & & \text{do}(j) & & \\ & & \longrightarrow & & \longrightarrow & & \\ \mathbb{P}_X & & \longrightarrow & \mathbb{P}_X^{\text{do}(i)} & \longrightarrow & \mathbb{P}_X^{\text{do}(j)} & \\ & \downarrow \tau & & \downarrow \tau & & \downarrow \tau & \\ & & \text{do}(\omega(i)) & & \text{do}(\omega(j)) & & \\ & & \longrightarrow & \mathbb{P}_Y^{\text{do}(\omega(i))} & \longrightarrow & \mathbb{P}_Y^{\text{do}(\omega(j))} & \end{array}$$

Proof. Let $i, j \in \mathcal{I}_X$ be interventions with $i \leq_X j$. The

⁷Since $\omega(i) = \omega(j) \iff (\omega(i) \leq_Y \omega(j)) \wedge (\omega(j) \leq_Y \omega(i))$, which, if the converse held, would imply that $(i \leq_X j) \wedge (j \leq_X i)$, which is equivalent to $i = j$.

⁸For instance, if it were necessary that ω be bijective, Theorems 9 and 11 would not hold.

commutativity of the left square of the diagram follows immediately from the definition of an exact transformation. It remains to be shown that the right square of the diagram commutes. By definition we have that $\tau\left(\mathbb{P}_X^{\text{do}(i)}\right) = \mathbb{P}_Y^{\text{do}(\omega(i))}$ and $\tau\left(\mathbb{P}_X^{\text{do}(j)}\right) = \mathbb{P}_Y^{\text{do}(\omega(j))}$. Thus, we only have to show that $\mathbb{P}_Y^{\text{do}(\omega(i))} \leq_Y \mathbb{P}_Y^{\text{do}(\omega(j))}$ as elements of \mathcal{P}_Y , i. e. that the arrow $\mathbb{P}_Y^{\text{do}(\omega(i))} \xrightarrow{\text{do}(\omega(j))} \mathbb{P}_Y^{\text{do}(\omega(j))}$ exists. This follows from the order-preservingness of ω . \square

4.4 CAUSAL INTERPRETATION OF EXACT TRANSFORMATIONS

The notion of an exact transformation between SEMs was motivated by the desire to analyse the correspondence between two causal models describing the same system at different levels of detail. The purpose of this section is to show that if one SEM can be viewed as an exact transformation of the other, then both can sensibly be thought of as causal models of the same system. In the following, we assume that \mathcal{M}_Y is an exact τ -transformation of \mathcal{M}_X with ω the corresponding map between interventions.

Surjectivity of ω ensures that any intervention in \mathcal{I}_Y can be viewed as an \mathcal{M}_Y -level representative of some intervention on the \mathcal{M}_X -level. Consequently, if do-interventions on the \mathcal{M}_X -level are in correspondence with physical implementations, then surjectivity of ω ensures that do-interventions on the \mathcal{M}_Y -level have at least one corresponding physical implementation, i. e. if \mathcal{M}_X is ‘rooted in reality’, then so is \mathcal{M}_Y .

Commutativity of the left hand part of the diagram ensures that the effects of interventions are consistently modelled by \mathcal{M}_X and \mathcal{M}_Y . Suppose we want to reason about the effects on the \mathcal{M}_Y -level caused by the intervention $j \in \mathcal{I}_Y$. For example, we may wish to reason about how the temperature and pressure of a volume of gaseous particles is affected by being heated. We could perform this reasoning by considering any corresponding \mathcal{M}_X -level intervention $i \in \omega^{-1}(\{j\})$ and considering the distribution this implies over \mathcal{Y} via τ . In our example, this would correspond to considering how heating the volume of gas could be modelled by changing the motions of all the gaseous particles and then computing the temperature and pressure of the volume of particles. Commutativity of the left hand part of the diagram implies that \mathcal{M}_X and \mathcal{M}_Y are consistent in the sense that \mathcal{M}_Y allows us to immediately reason about the effect of the intervention $j \in \mathcal{I}_Y$ while being equivalent to performing the steps above. That is, we can reason directly about temperature and pressure when heating a volume of gas without having to perform the intermediate steps that involve the microscopic description of the system.

Commutativity of the right hand side of the diagram en-

ures that once an intervention that fixes a subset of the variables has been performed, we can still consistently reason about the effects of further interventions on the remaining variables in \mathcal{M}_X and \mathcal{M}_Y . Furthermore, it ensures that compositionality of do-interventions on the \mathcal{M}_X -level carries over to the \mathcal{M}_Y -level, i. e. if the intervention j on the \mathcal{M}_X -level can be performed additionally to the intervention i in \mathcal{M}_X —that is, $i \leq_X j$ —, then the same is true of their representations in \mathcal{M}_Y .

If \mathcal{M}_X and \mathcal{M}_Y are models of the same system and it has been established that \mathcal{M}_Y is an exact τ -transformation of \mathcal{M}_X for some mapping τ , then the commutativity of the whole diagram in Theorem 6 ensures that they are causally consistent with one another in the sense described in the preceding paragraphs. If we wish to reason about the effects of interventions on the \mathcal{Y} -variables then it suffices to use the model \mathcal{M}_Y , rather than the (possibly more complex) model \mathcal{M}_X . In particular, this means that we can view the \mathcal{Y} -variables as causal entities, rather than only functions of underlying ‘truly’ causal entities. Only if this is the case, causal statements such as ‘raising temperature increases pressure’ or ‘LDL causes heart disease’ are meaningful.

4.5 WHAT CAN GO WRONG WHEN A TRANSFORMATION IS NOT EXACT?

In the previous section we argued that our definition of exact transformations between SEMs is a sensible formalisation of causal consistency. In this section we will try to give the reader an intuition for why weakening the conditions of our definition would be problematic. In particular we focus on the requirement that ω be order-preserving, which we view as one of the core ideas of our paper.

The requirement that ω be surjective is, as discussed above, required so that all interventions on the \mathcal{M}_Y -level have a corresponding intervention on the \mathcal{M}_X -level. If we were to only require that ω be surjective (but not order-preserving), the observational distribution of \mathcal{M}_X may be mapped to an interventional distribution of \mathcal{M}_Y , as illustrated by the following example (cf. Figure 2 for an illustration).

Example 7. Consider the SEM $\mathcal{M}_X = \{\mathcal{S}_X, \mathcal{I}_X, \mathbb{P}_E\}$ over $\mathcal{X} = \mathbb{R}^3$ where

$$\begin{aligned} \mathcal{S}_X &= \{X_1 = E_1, X_2 = E_2, X_3 = X_1 + X_2 + E_3\} \\ \mathcal{I}_X &= \{\emptyset, \text{do}(X_2 = 0), \text{do}(X_1 = 0, X_2 = 0)\}, \\ E_1 &\sim \mathbb{P}_{E_1}, E_2 = -E_1, E_3 \sim \mathbb{P}_{E_3} \end{aligned}$$

where \mathbb{P}_{E_1} and \mathbb{P}_{E_3} are arbitrary distributions. Let $\tau : \mathcal{X} \rightarrow \mathcal{Y} = \mathbb{R}^2$ be the mapping such that

$$\tau(x_1, x_2, x_3) = (y_1, y_2) = (x_1 + x_2, x_3)$$

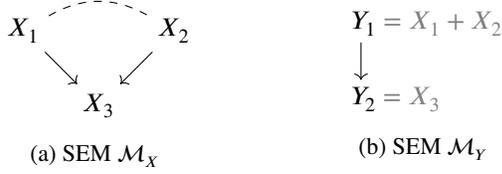


Figure 2: Graphical illustration of parent-child relationships for the examples in Section 4.5. The micro-level model \mathcal{M}_X depicted in (a) is to be transformed into the macro-level model \mathcal{M}_Y depicted in (b) which is a coarser descriptions as in it only considers the sum of X_1 and X_2 . In Section 4.5 we give examples of what can go wrong if the transformation is not exact.

Let $\mathcal{M}_Y = \{S_Y, \mathcal{I}_Y, \mathbb{P}_F\}$ be an SEM over \mathcal{Y} with

$$\begin{aligned} S_Y &= \{Y_1 = F_1, Y_2 = Y_1 + F_2\} \\ \mathcal{I}_Y &= \{\emptyset, \text{do}(Y_1 = 0)\}, \\ F_1 &\sim \mathbb{P}_{E_1}, F_2 \sim \mathbb{P}_{E_3} \end{aligned}$$

Let $\omega : \mathcal{I}_X \rightarrow \mathcal{I}_Y$ be defined by

$$\omega : \begin{cases} \emptyset & \mapsto \text{do}(Y_1 = 0) \\ \text{do}(X_2 = 0) & \mapsto \emptyset \\ \text{do}(X_1 = 0, X_2 = 0) & \mapsto \text{do}(Y_1 = 0) \end{cases}$$

Then it is true that $\mathbb{P}_{\tau(X)}^i = \mathbb{P}_Y^{\text{do}(\omega(i))}$ for all $i \in \mathcal{I}_X$, while ω is not order-preserving and $\omega(\emptyset) \neq \emptyset$.

If the SEMs in the above example were used to model the same system, it would be problematic that the observational setting of \mathcal{M}_X —a description of the system when not having physically performed any intervention—would correspond to an interventional setting in \mathcal{M}_Y , conversely suggesting that the system *had* been intervened upon.

To avoid the above conflict, we could demand in addition to surjectivity that ω map the null intervention of \mathcal{M}_X to the null intervention of \mathcal{M}_Y . This additional assumption would ensure commutativity of the left-hand part of the diagram in Theorem 6. However, as the following example shows, this would not ensure that the right-hand part of the diagram commutes for all pairs of interventions $i \leq_X j$, since in this case the arrow from $\mathbb{P}_Y^{\text{do}(\omega(i))}$ to $\mathbb{P}_Y^{\text{do}(\omega(j))}$ may not exist.⁹

Example 8. Let \mathcal{X}, \mathcal{Y} and τ be as in Example 7. Consider the SEM $\mathcal{M}_X = \{S_X, \mathcal{I}_X, \mathbb{P}_E\}$ where

$$\begin{aligned} S_X &= \{X_1 = E_1, X_2 = E_2, X_3 = X_1 + X_2 + E_3\} \\ \mathcal{I}_X &= \{\emptyset, \text{do}(X_2 = 0), \text{do}(X_1 = 0, X_2 = 0)\}, \\ E_1 &= 1, E_2 \sim \mathbb{P}_{E_2}, E_3 \sim \mathbb{P}_{E_3} \end{aligned}$$

⁹By definition of the poset \mathcal{P}_Y , this arrow exists if and only if $\omega(i) \leq_Y \omega(j)$.

where \mathbb{P}_{E_2} and \mathbb{P}_{E_3} are arbitrary distributions. Let $\mathcal{M}_Y = \{S_Y, \mathcal{I}_Y, \mathbb{P}_F\}$ be the SEM over \mathcal{Y} with

$$\begin{aligned} S_Y &= \{Y_1 = 1 + F_1, Y_2 = Y_1 + F_2\} \\ \mathcal{I}_Y &= \{\emptyset, \text{do}(Y_1 = 0), \text{do}(Y_1 = 1)\}, \\ F_1 &\sim \mathbb{P}_{E_2}, F_2 \sim \mathbb{P}_{E_3} \end{aligned}$$

Let $\omega : \mathcal{I}_X \rightarrow \mathcal{I}_Y$ be defined by

$$\omega : \begin{cases} \emptyset & \mapsto \emptyset \\ \text{do}(X_2 = 0) & \mapsto \text{do}(Y_1 = 1) \\ \text{do}(X_1 = 0, X_2 = 0) & \mapsto \text{do}(Y_1 = 0) \end{cases}$$

Then it is true that $\mathbb{P}_{\tau(X)}^i = \mathbb{P}_Y^{\text{do}(\omega(i))}$ for all $i \in \mathcal{I}_X$ and $\omega(\emptyset) = \emptyset$, although ω is not order-preserving.

If the above SEMs were used as models of the same system, they would not suffer from the problem illustrated in Example 7. Suppose now, however, that we have performed the intervention $\text{do}(X_2 = 0)$ in \mathcal{M}_X , corresponding to the intervention $\text{do}(Y_1 = 1)$ in \mathcal{M}_Y . If we wish to reason about the effect of the intervention $\text{do}(X_1 = 0, X_2 = 0)$ in \mathcal{M}_X , we run into a problem. \mathcal{M}_X suggests that $\text{do}(X_1 = 0, X_2 = 0)$ could be implemented by performing an additional action on top of $\text{do}(X_2 = 0)$. In contrast, \mathcal{M}_Y suggests that implementing the corresponding intervention $\text{do}(Y_1 = 0)$ would conflict with the already performed intervention $\text{do}(Y_1 = 1)$.

5 EXAMPLES OF EXACT TRANSFORMATIONS

In the introduction we motivated the problem considered in this paper by listing three settings in which differing model levels naturally occur. Having now introduced the notion of an exact transformation between SEMs, we provide in this section examples of exact transformations falling into each of these categories. The fact that a single framework can be used to draw an explicit correspondence between differing model levels in each of these settings demonstrates the generality of our framework.

Observe that in each of the following examples, the particular set of interventions considered is important. If we were to allow larger sets of interventions \mathcal{I}_X in the SEM \mathcal{M}_X , the transformations given would not be exact. This highlights the importance to the causal modelling process of carefully considering the set of interventions. All proofs are found in the Appendix.

5.1 MARGINALISATION OF VARIABLES

In the following two Theorems we consider two operations that can be performed on SEMs, namely marginalisation of childless or non-intervened variables, and prove

that these are exact transformations. That is, an SEM can be simplified into an SEM with fewer variables by either of these operations without losing any causal content concerning the remaining variables.

Thus if the SEM \mathcal{M}_Y can be obtained from another SEM \mathcal{M}_X by successively performing the operations in the following theorems, then \mathcal{M}_Y is an exact transformation of \mathcal{M}_X and hence the two models are causally consistent. This formally explains why we can sensibly consider causal models that focus on a subsystem \mathcal{M}_Y of a more complex system \mathcal{M}_X (cf. Figure 3). For a measurement-theoretic treatment of marginalisation in SEMs, see [3].

Theorem 9 (Marginalisation of childless variables). *Let $\mathcal{M}_X = (S_X, \mathcal{I}_X, \mathbb{P}_E)$ be an SEM and suppose that $\mathbb{I}_Z \subset \mathbb{I}_X$ is a set of indices of variables with no children, i. e. if $i \in \mathbb{I}_Z$ then X_i does not appear in the right-hand side of any structural equation in S_X . Let \mathcal{Y} be the set in which $Y = (X_i : i \in \mathbb{I}_X \setminus \mathbb{I}_Z)$ takes value. Then the transformation $\tau : \mathcal{X} \rightarrow \mathcal{Y}$ mapping*

$$\tau : (x_i : i \in \mathbb{I}_X) = x \mapsto y = (x_i : i \in \mathbb{I}_X \setminus \mathbb{I}_Z)$$

naturally gives rise to an SEM \mathcal{M}_Y that is an exact τ -transformation of \mathcal{M}_X , corresponding to marginalising out the childless variables X_i for $i \in \mathbb{I}_Z$.

Theorem 10 (Marginalisation of non-intervened variables). *Let $\mathcal{M}_X = (S_X, \mathcal{I}_X, \mathbb{P}_E)$ be an acyclic SEM and suppose that $\mathbb{I}_Z \subset \mathbb{I}_X$ is a set of indices of variables that are not intervened upon by any intervention $i \in \mathcal{I}_X$. Let \mathcal{Y} be the set in which $Y = (X_i : i \in \mathbb{I}_X \setminus \mathbb{I}_Z)$ takes value. Then the transformation $\tau : \mathcal{X} \rightarrow \mathcal{Y}$ mapping*

$$\tau : (x_i : i \in \mathbb{I}_X) = x \mapsto y = (x_i : i \in \mathbb{I}_X \setminus \mathbb{I}_Z)$$

naturally gives rise to an SEM \mathcal{M}_Y that is an exact τ -transformation of \mathcal{M}_X , corresponding to marginalising out the never-intervened-upon variables X_i for $i \in \mathbb{I}_Z$.

The assumption of acyclicity made in Theorem 10 can be relaxed to allow marginalisation of non-intervened variables in cyclic SEMs, at the expense of extra technical conditions (see Section 3 of [3]).

We remind the reader that our definition of an SEM does not require that the exogenous E -variables be independent. Theorem 10 would not hold if this restriction were made (which is usually the case in the literature); marginalising out a common parent node will in general result in its children having dependent exogenous variables.

5.2 MICRO- TO MACRO-LEVEL

Transformations from micro- to macro-levels may arise in situations in which the micro-level variables can be observed via a ‘coarse’ measurement device, represented by

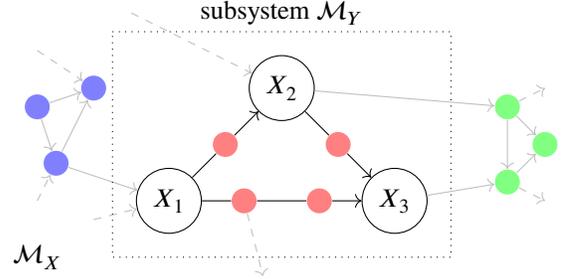


Figure 3: Suppose that there is a complex model \mathcal{M}_X but that we only wish to model the distribution over X_1, X_2, X_3 and how it changes under some interventions on X_1, X_2, X_3 . By Theorem 9, we can ignore downstream effects (●) after grouping them together as one multivariate variable and by Theorem 10 we can ignore intermediate steps of complex mechanisms (●) and treat upstream causes as noise fluctuations (●). That is, we can exactly transform the complex SEM \mathcal{M}_X into a simpler model \mathcal{M}_Y by marginalisation.

the function τ , e. g. we can use a thermometer to measure the temperature of a gas, but not the motions of the individual particles. They may also arise due to deliberate modelling choice when we wish to describe a system using higher level features, e. g. viewing the motor cortex as a single entity responsible for movements, rather than as a collection of individual neurons.

In such situations, our framework of exact transformations allows one to investigate whether such a macro-level model admits a causal interpretation. The following theorem provides an exact transformation between a micro-level model \mathcal{M}_X and a macro-level model \mathcal{M}_Y in which the variables are aggregate features of variables in \mathcal{M}_X obtained by averaging (cf. Figure 4).

Theorem 11 (Micro- to macro-level). *Let $\mathcal{M}_X = (S_X, \mathcal{I}_X, \mathbb{P}_{E,F})$ be a linear SEM over the variables $W = (W_i : 1 \leq i \leq n)$ and $Z = (Z_i : 1 \leq i \leq m)$ with*

$$\begin{aligned} S_X &= \{W_i = E_i : 1 \leq i \leq n\} \\ &\cup \left\{ Z_i = \sum_{j=1}^n A_{ij} W_j + F_i : 1 \leq i \leq m \right\} \\ \mathcal{I}_X &= \left\{ \emptyset, \text{do}(W = w), \text{do}(Z = z), \right. \\ &\quad \left. \text{do}(W = w, Z = z) : w \in \mathbb{R}^n, z \in \mathbb{R}^m \right\} \end{aligned}$$

and $(E, F) \sim \mathbb{P}$ where \mathbb{P} is any distribution over \mathbb{R}^{n+m} and A is a matrix.

Assume that there exists an $a \in \mathbb{R}$ such that each column of A sums to a . Consider the following transformation

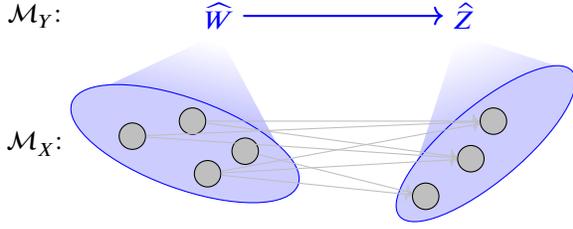


Figure 4: An illustration of the setting considered in Theorem 11. The micro-variables W_1, \dots, W_n and Z_1, \dots, Z_m in the SEM \mathcal{M}_X can be averaged to derive macro-variables \widehat{W} and \widehat{Z} in such a way that the resulting macro-level SEM \mathcal{M}_Y is an exact transformation of the micro-level SEM \mathcal{M}_X .

that averages the W and Z variables:

$$\tau : \mathcal{X} \rightarrow \mathcal{Y} = \mathbb{R}^2$$

$$\begin{pmatrix} W \\ Z \end{pmatrix} \mapsto \begin{pmatrix} \widehat{W} \\ \widehat{Z} \end{pmatrix} = \begin{pmatrix} \frac{1}{n} \sum_{i=1}^n W_i \\ \frac{1}{m} \sum_{j=1}^m Z_j \end{pmatrix}$$

Further, let $\mathcal{M}_Y = (\mathcal{S}_Y, \mathcal{I}_Y, \mathbb{P}_{\widehat{E}, \widehat{F}})$ over the variables $\{\widehat{W}, \widehat{Z}\}$ be an SEM with

$$\mathcal{S}_Y = \left\{ \widehat{W} = \widehat{E}, \widehat{Z} = \frac{a}{m} \widehat{W} + \widehat{F} \right\}$$

$$\mathcal{I}_Y = \left\{ \emptyset, \text{do}(\widehat{W} = \widehat{w}), \text{do}(\widehat{Z} = \widehat{z}), \right. \\ \left. \text{do}(\widehat{W} = \widehat{w}, \widehat{Z} = \widehat{z}) : \widehat{w} \in \mathbb{R}, \widehat{z} \in \mathbb{R} \right\}$$

$$\widehat{E} \sim \frac{1}{n} \sum_{i=1}^n E_i, \quad \widehat{F} \sim \frac{1}{m} \sum_{i=1}^m F_i$$

Then \mathcal{M}_Y is an exact τ -transformation of \mathcal{M}_X .

5.3 STATIONARY BEHAVIOUR OF DYNAMICAL PROCESSES

In this section we provide an example of an exact transformation between an SEM \mathcal{M}_X describing a time-evolving system and another SEM \mathcal{M}_Y describing the system after it has equilibrated. In this setting, τ could be thought of as representing our ability to only measure the time-evolving system at a single point in time, after the transient dynamics have taken place.

In particular, we consider a discrete-time linear dynamical system with identical noise and provide the explicit form of an SEM that models the distribution of the equilibria under each intervention (cf. Figure 5).¹⁰

¹⁰Note that the assumption that the transition dynamics be

Theorem 12 (Discrete-time linear dynamical process with identical noise). Let $\mathcal{M}_X = (\mathcal{S}_X, \mathcal{I}_X, \mathbb{P}_E)$ over the variables $\{X_t^i : t \in \mathbb{Z}, i \in \{1, \dots, n\}\}$ be a linear SEM with

$$\mathcal{S}_X = \left\{ X_{t+1}^i = \sum_{j=1}^n A_{ij} X_t^j + E_t^i : i \in \{1, \dots, n\}, t \in \mathbb{Z} \right\}$$

$$i. e. X_{t+1} = AX_t + E_t$$

$$\mathcal{I}_X = \left\{ \text{do}(X_t^j = x_j \forall t \in \mathbb{Z}, \forall j \in J) : x \in \mathbb{R}^{|J|}, J \subseteq \{1, \dots, n\} \right\}$$

$$E_t = E \forall t \in \mathbb{Z} \text{ where } E \sim \mathbb{P}$$

where \mathbb{P} is any distribution over \mathbb{R}^n and A is a matrix.

Assume that the linear mapping $v \mapsto Av$ is a contraction. Then the following transformation is well-defined under any intervention $i \in \mathcal{I}_X$:¹¹

$$\tau : \mathcal{X} \rightarrow \mathcal{Y}$$

$$(x_t)_{t \in \mathbb{Z}} \mapsto y = \lim_{t \rightarrow \infty} x_t$$

Let $\mathcal{M}_Y = (\mathcal{S}_Y, \mathcal{I}_Y, \mathbb{P}_F)$ be the (potentially cyclic) SEM over the variables $\{Y^i : i \in \{1, \dots, n\}\}$ with

$$\mathcal{S}_Y = \left\{ Y^i = \frac{\sum_{j \neq i} A_{ij} Y^j}{1 - A_{ii}} + \frac{F^i}{1 - A_{ii}} : i \in \{1, \dots, n\} \right\}$$

$$\mathcal{I}_Y = \left\{ \text{do}(Y^j = y_j \forall j \in J) : y \in \mathbb{R}^{|J|}, J \subseteq \{1, \dots, n\} \right\}$$

$$F \sim \mathbb{P}$$

Then \mathcal{M}_Y is an exact τ -transformation of \mathcal{M}_X .

The above theorem demonstrates how a linear additive SEM can arise as a result of making observations of a dynamical process. This supports one interpretation of SEMs as a description of a dynamical process that equilibrates quickly compared to its external environment.¹² The framework of exact transformations allows us to explain in a precise way the sense in which such equilibrium models can be used as causal descriptions of an underlying dynamical process.

linear can be relaxed to more general non-linear mappings. In this case, however, the structural equations of \mathcal{M}_Y can only be written in terms of implicit solutions to the structural equations of \mathcal{M}_X . For purposes of exposition, we stick here to the simpler case of linear dynamics.

¹¹In Appendix D.1 we show that A being a contraction mapping ensures that the sequence $(X_t)_{t \in \mathbb{Z}}$ defined by \mathcal{M}_X converges everywhere under any intervention $i \in \mathcal{I}_X$. That is, for any realisation $(x_t)_{t \in \mathbb{Z}}$ of this sequence, its limit $\lim_{t \rightarrow \infty} x_t$ as a sequence of elements of \mathbb{R}^n exists.

¹²This interpretation corresponds to the assumption that the noise in the dynamical model is constant through time, and is used by e.g. [10, 12, 14, 15] and [13] to meaningfully interpret cyclic SEMs.

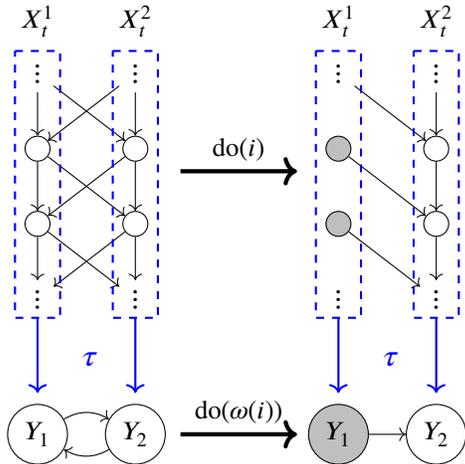


Figure 5: An illustration of the setting considered in Theorem 12. The discrete-time dynamical process is exactly transformed into a model describing its equilibria.

This result also sheds light on the interpretation of cyclic causal models. One interpretation of the structural equations of an acyclic SEM is that they represent a temporally ordered series of mechanisms by which data are generated. This is not possible in the case that the SEM exhibits cycles: there does not exist a partial ordering on the variables and hence one cannot think of each variable being generated temporally downstream of its parents. By showing that cyclic SEMs can arise as exact transformations of *acyclic* SEMs, we provide an interpretation of cyclic SEMs that does not suffer from the above problem.

6 DISCUSSION AND FUTURE WORK

It’s turtles all the way down! There is no such thing as a ‘correct’ model, but in this paper we introduced the notions of exact transformations between SEMs to evaluate when two SEMs can be viewed as causally consistent models of the same system. Illustrating how these notions can be used in order to relate differing model levels, we proved in Section 5 the exactness of transformations occurring in three different settings. These have implications for the following questions in causal modelling: When can we model only a subsystem of a more complex system? When does a micro-level system admit a causal description in terms of macro-level features? How do cyclic causal models arise?

Our work has implications for other problems in causal modelling. It suggests that ambiguous manipulations [18] may be thought of as arising due to the application of an inexact transformation to an SEM \mathcal{M}_X . This was illustrated in Section 1.1 in which LDL and HDL cholesterol were only measured via their sum TC, resulting in a model that

suffered from the problem of ambiguous manipulations (cf. Figure 1b) since it was not an exact transformation of the underlying model (cf. Figure 1a). This is related to the problem of causal variable definition as studied by [7].

A future line of enquiry would be to generalise the notion of an exact transformation in order to analyse the trade-off between model accuracy and model complexity for causal modelling using SEMs. For a transformation to be exact, we require that the posets $\mathcal{P}_{\tau(X)}$ and \mathcal{P}_Y be equal. One could imagine a ‘softening’ of this requirement such that the distributions in the posets are required to be only approximately equal. A slightly inaccurate model with a small number of variables may be preferable to an accurate but complex model.

We discussed the importance of an order-preserving ω to ensure a notion of causal consistency between two SEMs. It would be interesting to better understand the conditions under which different properties of consistency between causal models hold – for instance, counterfactual reasoning, which we have not discussed in this paper.

While we have introduced the notion of an exact transformation, we have not provided any criterion to choose from amongst the set of all possible exact transformations of an SEM. Foundational work in a similar direction to ours has been done by [4, 5], who consider a particular discrete setting. They provide algorithms to learn a transformation of a micro-level model to a macro-level model with desirable information-theoretic properties. We conjecture that our framework may lead to extensions of their work, e. g. to the continuous setting.

Finally, suppose that we have made observations of an underlying system \mathcal{M}_X via a measurement device τ , and that we want to fit an SEM \mathcal{M}_Y from a restricted model class to our data. By using our framework, asking whether or not \mathcal{M}_Y admits a causal interpretation consistent with \mathcal{M}_X reduces to asking whether the transformation is exact. More generally, by fixing any two of \mathcal{M}_X , τ and \mathcal{M}_Y , we can ask what properties must be fulfilled by the third in order for the two models to be causally consistent. We hope that this may lead to the practical use of SEMs being theoretically grounded.

Acknowledgements

We thank Tobias Mistele for valuable early feedback. Stephan Bongers was supported by NWO, the Netherlands Organization for Scientific Research (VIDI grant 639.072.410). This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement n° 639466).

References

- [1] R. Balian. *From microphysics to macrophysics*. Springer, 1992.
- [2] K. A. Bollen. *Structural equations with latent variables*. John Wiley & Sons, 2014.
- [3] S. Bongers, J. Peters, B. Schölkopf, and J. M. Mooij. Structural causal models: Cycles, marginalizations, exogenous reparametrizations and reductions. *arXiv preprint arXiv:1611.06221*, 2016.
- [4] K. Chalupka, P. Perona, and F. Eberhardt. Visual causal feature learning. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*, pages 181–190. AUAI Press, 2015.
- [5] K. Chalupka, P. Perona, and F. Eberhardt. Multi-level cause-effect systems. In *The 19th International Conference on Artificial Intelligence and Statistics*, 2016.
- [6] D. Dash and M. J. Druzdzel. Caveats for causal reasoning with equilibrium models. *Lecture notes in computer science*, pages 192–203, 2001.
- [7] F. Eberhardt. Green and grue causal variables. *Synthese*, 193(4):1029–1046, 2016.
- [8] F. M. Fisher. A correspondence principle for simultaneous equation models. *Econometrica: Journal of the Econometric Society*, pages 73–92, 1970.
- [9] E. P. Hoel, L. Albantakis, and G. Tononi. Quantifying causal emergence shows that macro can beat micro. *Proceedings of the National Academy of Sciences*, 110(49):19790–19795, 2013.
- [10] A. Hyttinen, F. Eberhardt, and P. O. Hoyer. Learning linear cyclic causal models with latent variables. *Journal of Machine Learning Research*, 13(Nov):3387–3439, 2012.
- [11] Y. Iwasaki and H. A. Simon. Causality and model abstraction. *Artificial Intelligence*, 67(1):143–194, 1994.
- [12] G. Lacerda, P. L. Spirtes, J. Ramsey, and P. O. Hoyer. Discovering cyclic causal models by independent components analysis. In *Proceedings of the Twenty-Fourth Conference Annual Conference on Uncertainty in Artificial Intelligence*, 2008.
- [13] J. M. Mooij and T. Heskes. Cyclic causal discovery from continuous equilibrium data. In *Proceedings of the Twenty-Ninth Conference Annual Conference on Uncertainty in Artificial Intelligence*, 2013.
- [14] J. M. Mooij, D. Janzing, T. Heskes, and B. Schölkopf. On causal discovery with cyclic additive noise models. In J. Shawe-Taylor, R. Zemel, P. Bartlett, F. Pereira, and K. Weinberger, editors, *Advances in Neural Information Processing Systems 24 (NIPS*2011)*, pages 639–647, 2011.
- [15] J. M. Mooij, D. Janzing, and B. Schölkopf. From Ordinary Differential Equations to Structural Causal Models: the deterministic case. In *Proceedings of the Twenty-Ninth Conference Annual Conference on Uncertainty in Artificial Intelligence*, pages 440–448, 2013.
- [16] J. Pearl. *Causality*. Cambridge University Press, 2009.
- [17] H. A. Simon and A. Ando. Aggregation of variables in dynamic systems. *Econometrica: journal of the Econometric Society*, pages 111–138, 1961.
- [18] P. Spirtes and R. Scheines. Causal inference of ambiguous manipulations. *Philosophy of Science*, 71(5):833–845, 2004.
- [19] P. Spirtes, C. N. Glymour, and R. Scheines. *Causation, prediction, and search*. MIT press, 2000.
- [20] D. Steinberg. *The Cholesterol Wars: The Skeptics vs the Preponderance of Evidence*. Academic Press, 2011.
- [21] A. S. Truswell. *Cholesterol and beyond: the research on diet and coronary heart disease 1900-2000*. Springer Science & Business Media, 2010.