

Research Article

S. Bongers*, T. Blom, and J.M. Mooij

Causal Modeling of Dynamical Systems

Abstract: Dynamical systems are widely used in science and engineering to model systems consisting of several interacting components. Often, they can be given a *causal* interpretation in the sense that they not only model the evolution of the states of the system's components over time, but also describe how their evolution is affected by external interventions on the system that perturb the dynamics. We introduce the formal framework of structural dynamical causal models (SDCMs) that explicates the causal semantics of the system's components as part of the model. SDCMs represent a dynamical system as a collection of stochastic processes and specify the basic causal mechanisms that govern the dynamics of each component as a structured system of random differential equations of arbitrary order. SDCMs extend the versatile causal modeling framework of structural causal models (SCMs), also known as structural equation models (SEMs), by explicitly allowing for time-dependence. An SDCM can be thought of as the stochastic-process version of an SCM, where the static random variables of the SCM are replaced by dynamic stochastic processes and their derivatives. We provide the foundations for a theory of SDCMs, by (i) formally defining SDCMs, their solutions, stochastic interventions, and a graphical representation; (ii) studying existence and uniqueness of the solutions for given initial conditions; (iii) discussing under which conditions SDCMs equilibrate to SCMs as time tends to infinity; (iv) relating the properties of the SDCM to those of the equilibrium SCM. This correspondence enables one to leverage the wealth of statistical tools and discovery methods available for SCMs when studying the causal semantics of a large class of stochastic dynamical systems. The theory is illustrated with several well-known examples from different scientific domains.

Keywords: causal models, structural causal models, dynamical systems, random differential equations, interventions, causal graph, cycles, equilibrium

MSC: 37H05, 34F05, 62A09, 68T30

1 Introduction

Continuous dynamical systems consisting of differential equations are widely used in science and engineering to model the time-dependent behavior of certain phenomena. A classical example is the modeling of the trajectory of a die that is thrown, by means of Newton's equations of motion. Initial conditions or parameters of the dynamics may be stochastic, which can be modeled mathematically by making use of random differential equations (RDEs). These provide a natural extension of ordinary differential equations (ODEs) to the stochastic setting [10, 55, 64, 65]. For example, the initial position of the die is often not known, and varies from throw to throw, which leads to a probability distribution over the possible trajectories of the die (and eventually, to an uncertain outcome of the throw).

Many dynamical systems can be considered to consist of several interacting subsystems or components, for example, mass-spring systems in physics, predator-prey systems in biology, and mass-action law kinetics in chemistry. These dynamical systems are often implicitly given a *causal* interpretation in the sense that they are not only supposed to model the evolution of the state of the system over time, but also describe how the evolution of the system's components is affected by external interventions on the system that perturb the dynamics. For example, when applying an external force to a particle, the change in the force

*Corresponding Author: S. Bongers: Pattern Recognition Laboratory, Delft University of Technology; E-mail: s.r.bongers@tudelft.nl

T. Blom: Informatics Institute, University of Amsterdam; E-mail: t.blom2@uva.nl

J.M. Mooij: Korteweg-de Vries Institute, University of Amsterdam; E-mail: j.m.mooij@uva.nl

term in Newton's second law of motion results in a changed acceleration, and hence a changed position, of the particle. Another example is that hunting wolves may lead to an increase in the population of sheep. The ensuing causal semantics of the system is usually only treated in an implicit and intuitive fashion, rather than that it is formally specified by (or derivable from) the mathematical model. Indeed, a system of (random) differential equations simply expresses symmetric relations between the components, without any preferred order or asymmetry. On the other hand, causal relations may be asymmetric, as they distinguish cause from effect. Thus, while dynamical systems may describe how the state of a system consisting of several components evolves over time, by themselves they do not express the inherent "causal structure" of the system's components.

An apparently rather different modeling framework that allows to represent the causal semantics of a system composed of components is provided by structural causal models (SCMs), also known as (non-parametric) structural equation models (SEMs) [8, 9, 57, 67]. First introduced in genetics by Wright [71], they became popular over the years in econometrics [34], the social sciences [20, 32], and more recently in AI [56]. SCMs express causal relationships between variables corresponding to "autonomous" subsystems or components in the form of deterministic, functional relationships, and stochasticity is introduced through the assumption that certain variables are exogenous (latent) random variables. Their predictive power stems from the assumption that the equations of these models are organized in a structural way: each equation represents a distinct autonomous causal mechanism, where distinctness of the mechanisms means that they can be changed independently of one another by targeted interventions—at least in principle. While SCMs explicate the causal semantics of a system composed of different components in this specific way, they have no built-in notion of time. A commonly used workaround for this limitation is to introduce multiple "copies" of the variables, corresponding to observations at different points (or intervals) in time. This workaround only applies to discrete time, and SCMs cannot be used to model causal semantics of continuous-time systems without somehow discretizing time.

In this work, we propose the modeling framework of *structural dynamical causal models* (SDCMs), which on the one hand explicates the causal relationships between components of continuous dynamical systems, and on the other hand extends structural causal models to explicitly allow for time-dependence. SDCMs represent a dynamical system as a collection of stochastic processes (each one referring to a causally "autonomous" component) subject to a "structured" dynamics, which specifies the causal mechanisms that govern the dynamics of the components by means of random differential equations of arbitrary order. An SDCM can be thought of as the stochastic-process version of an SCM, where the static (time-independent) random variables of the SCM are replaced by dynamic (time-dependent) stochastic processes and their derivatives. Our contributions can be considered as the first steps towards a theory of SDCMs. More specifically, we:

- (i) formally define SDCMs, their solutions, stochastic interventions, and a graphical representation;
- (ii) study existence and uniqueness of the solutions for given initial conditions;
- (iii) discuss under which conditions SDCMs equilibrate to SCMs as time tends to infinity;
- (iv) relate the properties of the SDCM to those of the equilibrium SCM.

The correspondence between SCMs and equilibrated SDCMs enables one to leverage the wealth of statistical tools and discovery methods available for SCMs when studying the causal semantics of a large class of stochastic dynamical systems. We illustrate the theory with several well-known examples from different scientific domains.

Related work

Over the years, several efforts have been made to develop a notion of causality for stochastic processes, both in discrete and continuous time.

For discrete time, Granger causality [23, 24, 33, 69], simultaneous equation models [25, 45], vector autoregressive (VAR) models [48, 63] and dynamic Bayesian networks [14, 31] have been studied extensively.

More recently, there has been some work on learning difference-based causal models [68] and structural equation models [39]. In principle, all these models fit directly into the framework of SCMs by labeling the random variables with time.

For continuous time, there has been substantial work in the graphical modeling community [1, 16–19] based on the concept of local independence, which was introduced by Schweder [62]. However, none of these approaches explicitly takes into account that dynamical models are often based on differential equations. In parallel, several attempts have been made to arrive at causal interpretations of processes described by ordinary and stochastic differential equations. Many of these approaches start from the assumption of a first-order system of ODEs written in canonical form, and implicitly (or explicitly) attribute a causal interpretation to this [5, 38, 52, 59]. The notion of causality in ODEs has also been studied using Simon’s causal ordering algorithm [38]. Relations between a certain class of causally interpreted ODEs and deterministic SCMs at equilibrium have been established under the strong assumption that all the solutions of the ODE converge to a single static equilibrium state [52], independent of the initial condition. This assumption can be relaxed to allow for asymptotic dynamics [61] such as periodic oscillations, but this still requires the assumption that the asymptotic dynamics does not depend on the initial condition. Another way to relax the assumption of [52] is taken in the framework of causal constraints models [6], which can model static equilibrium states as long as the dynamical system has a unique static equilibrium state corresponding to each initial condition, for every intervention. These models can give a more complete causal description of these static equilibrium states than SCMs can [6], but this comes at the cost that they appear to be too “flexible” in general. Finally, several approaches in terms of stochastic differential equations, which are differential equations with an additive white noise term, have been developed over the years [12, 28, 37, 50, 58]. The stochastic differential equations have the advantage that they can deal with “instantaneous” stochasticity in the dynamics, but solving them usually requires a considerable mathematical effort using Itô calculus.

Compared with existing work, the framework of structural dynamical causal models that we propose here has the novel combination of features that it extends the semantics of continuous dynamical systems by formally encoding the causal structure into the model, it allows for stochasticity due to uncertainty over initial conditions or parameters of the dynamics without relying on strong stability assumptions, and it does not force one to consider time derivatives of processes as being “causally independent” of the processes themselves (that is, time derivatives of processes are considered to describe the same subsystem or component as the process itself). Our framework reconciles the traditional intuitive treatment of causality in the context of deterministic dynamical systems as practiced in many exact sciences with the treatment of causality of stochastic systems that is nowadays very popular in AI, statistics and other scientific disciplines. An attractive feature is that it naturally accommodates many causally interpreted continuous dynamical systems that appear “in the wild”.

Contributions

In this paper, we introduce the framework of *structural dynamical causal models* (SDCMs),¹ which allows to model the causal semantics of stochastic processes for a large class of continuous dynamical systems by means of a “structured” system of random differential equations of arbitrary order (including zeroth-order). One can consider SCMs as special cases of SDCMs that only contain zeroth-order equations. The proposed modeling framework enables modeling of stochasticity, time-dependence and causality in a natural way. We study the existence and uniqueness of solutions of SDCMs, define an idealized notion of stochastic

¹ Not to be confused with the *dynamic causal models* of [30] or the *dynamic structural causal models* of [61]. The dynamic causal models of [30] have been developed to infer the causal relations between the activities of different brain regions, where each neuronal state is modeled by a first order differential equation. These much more restricted models could in principle be represented by SDCMs. The dynamic structural causal models of [61] have been developed to model the asymptotic behavior of an ordinary differential equation under non-constant interventions and assume that the asymptotic behavior does not depend on the initial condition.

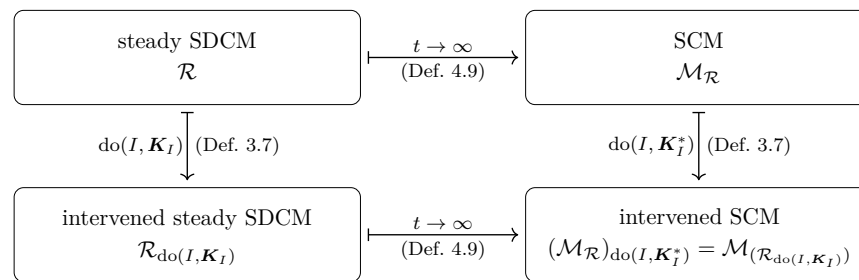


Figure 1. This diagram shows that, under certain convergence assumptions, equilibration (left-to-right in the diagram) commutes with intervention (top-to-bottom in the diagram). The precise statement is made explicit in Theorem 4.18.

interventions, and propose a convenient graphical representation of the model structure. We define a notion of equilibration of an SDCM to an SCM, which corresponds with letting a system converge towards equilibrium as time tends to infinity, and relate the properties of the SDCM to those of the equilibrium SCM. In the next paragraphs, we describe our contributions in more detail.

Intuitively, an SDCM can be thought of as an SCM where the notion of time is added to the structural equations by replacing the random variables of the SCM by stochastic processes and their (higher-order) derivatives. In the presence of these derivative processes, these equations, which we coin *dynamic structural equations*, can be read as random differential equations. The dynamic structural equations have the property that they are organized in a structural way, similar to how the structural equations of an SCM are organized by associating a distinct causal mechanism to each observed variable. This distinguishes SDCMs from other “non-causal” (random) dynamical systems, and allows to define idealized stochastic interventions on these models, similarly to how this is usually done for SCMs. We also derive a graphical representation of SDCMs with a causal interpretation in analogy with that of SCMs.

The framework of SDCMs on the one hand allows one to specify the causal semantics of a system of RDEs, and on the other hand it enables temporal extensions for SCMs. In particular, we show when and how we can equilibrate an SDCM to an SCM, such that the static solutions of the SCM contain the equilibrium states of the SDCM. Our equilibration operation, inspired by the one of Mooij et al. [52], has the key property that it preserves the structure of the endogenous processes. Intuitively, the idea is that in the limit as time tends to infinity, the dynamic structural equations converge to those equations for which the higher-order derivatives of the processes have been set to zero, yielding the structural equations of an SCM. This allows us to use SCMs to model the equilibrium states of dynamical systems, including cases that were previously considered to fall outside their scope, such as the price, supply and demand model in econometrics. In addition, we show that this equilibration operation commutes with intervention (as in Figure 1), and naturally maps the graph of the SDCM to the graph of the SCM. This provides a different perspective on what Dash [15] calls the “violation of the equilibration-manipulation commutability property”. Our formalism allows us to generalize the main result of Mooij et al. [52], which states that certain causally interpreted systems of ODEs can be equilibrated to SCMs, in several directions: (i) we replace the deterministic setting with a more general stochastic setting, that is, we can deal with randomness in the initial conditions and in the parameters, (ii) we allow the order of the equations of the dynamical model to be arbitrary, including zeroth-order, rather than restricting to first-order differential equations only, and (iii) we drop the strong assumption that the dynamical model needs to have a single static equilibrium that is independent of the initial condition.

By no longer restricting to first-order dynamical systems, we arrive at a more natural causal interpretation of systems of higher-order RDEs, like the coupled harmonic oscillator. Thereby, we circumvent questions like “does position cause velocity, or does velocity cause position, or both?”. However, allowing for zeroth-order dynamic structural equations leads to additional technical challenges that are absent when solving first-order RDEs. Indeed, the initial conditions of the solutions may be constrained by the zeroth-order dynamic structural equations, and possibly even by additional “hidden” constraints. We provide sufficient conditions

under which the existence and uniqueness of a solution of an SDCM with a given initial condition can be guaranteed. We also provide stronger conditions under which this still holds after certain interventions.

Even if the existence and uniqueness of a solution of an SDCM can be guaranteed, not all solutions of an SDCM equilibrate, in general. For example, a coupled harmonic oscillator may oscillate indefinitely in the absence of friction. Moreover, the solutions that equilibrate may not always equilibrate to the same equilibrium state. For example, a freely moving particle subject to friction may end up anywhere, depending on its initial position and velocity. In other words, equilibrium states may depend on the initial condition. This is compatible with the recently proposed framework of cyclic SCMs of Bongers et al. [9], which allows for the absence of (or, the presence of multiple) solutions of the structural equations. The intricate connection between the dependence of the equilibrium states of an SDCM on the initial conditions and the solvability properties of the equilibrated SCM sheds new light on the counterintuitive “nonancestral” causal effects in certain “pathological” cyclic SCMs with self-cycles that were first observed by Neal [54].

The scope of this paper is limited to establishing the framework of SDCMs and its bridge to SCMs at equilibrium. The importance of this bridge is that, although SDCMs can be used for modeling causal relationships between stochastic processes, inferring such causal models from data may pose certain difficulties. One significant practical drawback of using SDCMs for modeling systems with an unknown dynamics is that obtaining time series data with sufficiently high temporal resolution can be costly, impractical or even impossible.² The results of this work enable one to study the causal semantics of the equilibrium states of a large class of random dynamical models in terms of SCMs. In particular, this allows to infer properties of these dynamical models by employing the statistical tools and discovery methods available for static SCMs on equilibrium data.

Outline

The paper is organized as follows: In Section 2, we provide the necessary concepts of stochastic processes and random differential equations. In Section 3, we introduce the class of structural dynamical causal models, define SCMs as special cases of SDCMs, define interventions, define the graph of an SDCM, discuss initial conditions, and provide existence and uniqueness results of solutions. In Section 4, we define the equilibration operation on steady SDCMs, define the graph of the equilibrated SDCM, describe the commutation of the intervention and the equilibration operation, study the inverse problem of finding steady SDCMs with non-trivial dynamics for which all the solutions equilibrate to solutions of the SCM, and discuss subtleties in the causal interpretation of the graph of the equilibrated SDCM. We conclude with a discussion and some open problems in Section 5. Proofs are provided in Appendix A.

2 Preliminaries

2.1 Stochastic processes

In this section, we introduce the basic definitions and terminology for stochastic processes [see also 10, 55]. A *stochastic process* is an \mathbb{R}^n -valued function $\mathbf{X} : T \times \Omega \rightarrow \mathbb{R}^n$, where T is some index set, such that \mathbf{X}_t (which denotes $\mathbf{X}(t, \cdot)$, also sometimes denoted as $\mathbf{X}(t)$) is for each $t \in T$ a random variable³ on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. A random variable $\mathbf{X} : \Omega \rightarrow \mathbb{R}^n$ can itself be seen as a stochastic process that is constant in time, that is, as the process $\mathbf{X} : T \times \Omega \rightarrow \mathbb{R}^n$ defined by $\mathbf{X}_t(\omega) := \mathbf{X}(\omega)$. We always assume

² For example, modern measurement techniques in biology, like RNA sequencing and mass cytometry, enable simultaneous measurements of multiple variables at once in single cells, but at the cost of destroying the cells during the measurement process. This means that it is impossible to obtain time-series measurements for individual cells, although one can take a “snapshot” of the internal states of many single cells at the same point in time.

³ Assuming the Borel σ -algebra $\mathfrak{B}(\mathbb{R}^n)$ on \mathbb{R}^n , that is, the smallest σ -algebra on \mathbb{R}^n that contains all open n -balls.

that there exists some background probability space $(\Omega, \mathcal{F}, \mathbb{P})$ on which all random variables and processes are defined. Furthermore, we only consider processes where $T = [t_0, t_1]$ or $T = [t_0, \infty)$ for $t_0 < t_1$ with $t_0, t_1 \in \mathbb{R}$, and the points of T are thought of as representing *time*. For each $\omega \in \Omega$ we have an \mathbb{R}^n -valued function $T \rightarrow \mathbb{R}^n$ mapping t to $\mathbf{X}_t(\omega)$, which is called a *sample path*, or just a *path*, of \mathbf{X} . We call two stochastic processes \mathbf{X} and \mathbf{Y} *a.s. equal* to each other, denoted by $\mathbf{X} = \mathbf{Y}$ a.s., if \mathbb{P} -almost surely all sample paths are equal, that is, if there exists a \mathbb{P} -null set⁴ $N \subseteq \Omega$ such that for all $\omega \in \Omega \setminus N$ and for all $t \in T$ we have $\mathbf{X}_t(\omega) = \mathbf{Y}_t(\omega)$. We consider stochastic processes, and random variables in particular, only up to a.s. equality.

A family $(\mathbf{X}_i)_{i \in \mathcal{I}}$ of stochastic processes for some finite index set \mathcal{I} is called *independent* if for all $k \in \mathbb{N}$ and all k -tuples (t_1, \dots, t_k) of distinct elements of T the family

$$(\tilde{\mathbf{X}}_i)_{i \in \mathcal{I}}$$

of random variables $\tilde{\mathbf{X}}_i := ((\mathbf{X}_i)_{t_1}, \dots, (\mathbf{X}_i)_{t_k})$ is independent.

We call a stochastic process \mathbf{X} *continuous*, if its paths are continuous almost surely, that is, for \mathbb{P} -almost every $\omega \in \Omega$ and for all $t \in T$ we have

$$\lim_{s \rightarrow t} \mathbf{X}_s(\omega) = \mathbf{X}_t(\omega).$$

We call a stochastic process \mathbf{X} *differentiable*, if its paths are differentiable almost surely, that is, for \mathbb{P} -almost every $\omega \in \Omega$ and for all $t \in T$ the derivative

$$\mathbf{X}'_t(\omega) := \frac{d\mathbf{X}_t}{dt}(\omega) := \lim_{h \rightarrow 0} \frac{\mathbf{X}_{t+h}(\omega) - \mathbf{X}_t(\omega)}{h}$$

exists. The mapping $\mathbf{X}' : T \times \Omega \rightarrow \mathbb{R}^n$ defines a stochastic process and is called the *derivative* of \mathbf{X} . Similarly, one can define, if it exists, the n^{th} -order derivative of \mathbf{X} as the derivative of the $(n-1)^{\text{th}}$ -order derivative of \mathbf{X} , which we also write as $\mathbf{X}^{(n)}$, where the zeroth-order derivative of \mathbf{X} is $\mathbf{X}^{(0)} := \mathbf{X}$.

We call a stochastic process \mathbf{X} *continuously differentiable* or a C^1 -stochastic process, if its derivative \mathbf{X}' exists and is continuous. Similarly, we call \mathbf{X} a C^n -stochastic process, if its derivatives $\mathbf{X}', \mathbf{X}'', \dots, \mathbf{X}^{(n)}$ exist and are continuous. In particular, \mathbf{X} is a C^0 -stochastic process if it is continuous.

2.2 Random differential equations

In this section, we give a brief overview of some key aspects of random differential equations [for more details, see 10, 55]. Random differential equations (RDEs) are similar to ordinary differential equations (ODEs), but can deal with randomness in the initial conditions and in the parameters. Due to their close connection to ODEs they can be analyzed by use of methods that are analogous to those in the theory of ODEs [10]. Their formalism is conceptually easier than the formalism of the white-noise driven stochastic differential equations (SDEs), while still being applicable to those systems via the generalized Doss-Sussmann correspondence [see 40, 55]. They have been used for many years in a wide range of applications [see, for example, 10, 36, 46, 55, 64, 65].

A stochastic process $\mathbf{X} : T \times \Omega \rightarrow \mathbb{R}^d$ is a *solution* of a (first-order) random differential equation

$$\mathbf{X}' = \mathbf{f}(\mathbf{X}, \mathbf{E}), \quad (1)$$

where $\mathbf{f} : \mathbb{R}^d \times \mathbb{R}^e \rightarrow \mathbb{R}^d$ is a measurable function and $\mathbf{E} : T \times \Omega \rightarrow \mathbb{R}^e$ a stochastic process, if for \mathbb{P} -almost every $\omega \in \Omega$ the (first-order) ordinary differential equation⁵

$$\mathbf{X}'_t(\omega) = \mathbf{f}(\mathbf{X}_t(\omega), \mathbf{E}_t(\omega))$$

⁴ Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. A set $N \subseteq \Omega$ is called a \mathbb{P} -null set if there exists a measurable set $\tilde{N} \in \mathcal{F}$ with $N \subseteq \tilde{N}$ and $\mathbb{P}(\tilde{N}) = 0$.

⁵ These ordinary differential equations are also called *explicit ordinary differential equations* [2]. Similarly, the random differential equations (1) are also called *explicit random differential equations*.



Figure 2. Two masses coupled by a spring, freely drifting in space (left, see Example 2.1) and with one of the masses attached to a fixed point (right, see Example 2.2).

holds for all $t \in T$. An *initial condition* of the RDE (1) is a tuple $(t_0, \mathbf{X}_{[0]})$ that specifies those solutions \mathbf{X} of the RDE (1) that satisfy for \mathbb{P} -almost every $\omega \in \Omega$

$$\mathbf{X}_{t_0}(\omega) = \mathbf{X}_{[0]}(\omega)$$

at the initial time t_0 . Since every n^{th} -order ODE can be rewritten as a system of first-order ODEs, the general form of the random differential equation (1) can be used to express analogously all the n^{th} -order *random differential equations*.⁶

The inclusion of randomness in the equations can be classified into two basic types. The first type consists of *randomness in the initial conditions*, that is, the initial conditions are not a.s. equal to a constant deterministic process. The second type consists of *randomness in the parameters*, that is, the process \mathbf{E} is not a.s. equal to a deterministic stochastic process. Of course, a combination of both types can hold. In particular, an RDE together with an initial condition reduces to an *initial value problem* for ODEs if it has no randomness in both the initial conditions and the parameters.

If the stochastic process \mathbf{E} is continuous, sufficient conditions that guarantee the existence and uniqueness of solutions for any initial condition can be found in Bunke [10] and Kloeden and Platen [43]. These results are similar to the uniqueness and existence theorems for ODEs [11].

Example 2.1 (Two masses coupled by a spring). *Consider a one-dimensional system of two point masses m_1 and m_2 with positions X_1 and X_2 respectively that are coupled by an ideal spring with spring constant $\kappa_1 > 0$ and equilibrium length $L_1 > 0$ under influence of friction with friction coefficients $b_1, b_2 \geq 0$ respectively (see Figure 2 (left)). The equations of motion of this system, whose derivation can be found in physics textbooks, are given by the second-order random differential equations*

$$\begin{cases} X_1'' = \frac{\kappa_1}{m_1}(X_2 - X_1 - L_1) - \frac{b_1}{m_1}X_1' \\ X_2'' = \frac{\kappa_1}{m_2}(X_1 - X_2 + L_1) - \frac{b_2}{m_2}X_2'. \end{cases}$$

Randomness may enter the system via the initial condition $(t_0, (X_1(t_0), X_1'(t_0), X_2(t_0), X_2'(t_0)))$ or via the parameters. For example, instead of assuming that the length L_1 has a fixed value, we can assume that it is an exogenous random variable distributed according to some distribution. The system of equations then forms an RDE.

In this paper, we propose a modeling class that allows to model the causal semantics of stochastic processes with RDEs in an unambiguous way. The following example illustrates that modeling interventions on RDEs, and thereby grounding their causal semantics, is not a completely trivial matter.

Example 2.2 (Two masses coupled by a spring, continued). *Consider again the RDE that describes the two masses coupled by an ideal spring from Example 2.1. These equations denote a symmetric relation, that is, for both equations X_1 can be expressed in terms of X_2 , and vice versa. The causal relations between the processes X_1 and X_2 are not inherently implied by the form of the equations. For example, what happens to*

⁶ Furthermore, explicit time-dependence of \mathbf{f} can be incorporated by adding a dummy variable with t with dynamics $t' = 1$ and initial condition $t_{[0]} = 0$.

X_2 if we fix the mass m_1 to a fixed wall, say, at $X_1 = 0$ (see Figure 2 (right))? The corresponding RDE for X_1 and X_2 is then given by

$$\begin{cases} X_1 = 0 \\ X_2'' = \frac{\kappa_1}{m_2}(X_1 - X_2 + L_1) - \frac{b_2}{m_2}X_2'. \end{cases}$$

In both cases we implicitly assumed that each mass has its own equation of motion, that is, the first and second equation determine the motion of the mass m_1 and m_2 , in terms of the processes X_1 and X_2 , respectively. Therefore, the intervention of fixing the mass m_1 to the wall is accomplished by changing only the equation for m_1 to the equation $X_1 = 0$. If instead we had changed the other equation to $X_1 = 0$, then, as one can easily verify, X_2 would always be fixed, which does not correspond to the expected physical behavior. This additional “structure” of knowing which RDE determines the dynamics of which process is not “intrinsically” defined by RDE. Moreover, RDEs usually do not include zeroth-order equations (also referred to as “algebraic equations”), such as $X_1 = 0$. Allowing for RDEs of arbitrary order, including zeroth, allows to model a wide range of interventions on these models. For example, instead of fixing the mass m_1 to the fixed wall at $X_1 = 0$ we could fix it to a wall that is driven by some external force, such as $X_1 = A \sin(2\pi ft)$ for some $A, f > 0$.

3 Structural dynamical causal models

In this section, we introduce the class of structural dynamical causal models (SDCMs) that allows to formally specify causal semantics for any RDE of arbitrary order (including zeroth-order). We organize the differential equations of the RDEs in a structural way, similar to how this is done for structural causal models, such that each differential equation expresses the causal mechanism that governs the dynamics of a single stochastic process (corresponding to a single component of the system). This allows us to model stochastic idealized interventions targeting certain components in dynamical models, similarly to how this is done for SCMs.

We start in Section 3.1 with introducing the notation and terminology that will be used throughout the paper. In Section 3.2, we formally define SDCMs and their solutions. In Section 3.3, we formalize the causal semantics of SDCMs in terms of stochastic “perfect” interventions. In Section 3.4, we introduce and discuss a graphical representation for SDCMs. In Section 3.5, we discuss the initial conditions and how these relate to the existence of solutions. We finish in Section 3.6 with stating results about the existence and uniqueness of solutions of an SDCM.

3.1 Notation and terminology

Let $\mathcal{I} = \{1, \dots, d\}$ be a finite index set and $\mathcal{X} = \prod_{i \in \mathcal{I}} \mathcal{X}_i$ the product of the domains of the components of a system, where domain $\mathcal{X}_i = \mathbb{R}^{d_i}$ encodes the range of possible values that the i^{th} component can take. The stochastic process $\mathbf{X} = (X_1, \dots, X_d) : T \times \Omega \rightarrow \mathcal{X}$ has component processes $X_i : T \times \Omega \rightarrow \mathcal{X}_i$.

Let $i \in \mathcal{I}$ and $n_i \in \mathbb{N}_0$. If for the i^{th} component X_i the n_i^{th} -order derivative exists, then the *complete n_i^{th} -order derivative* of X_i , defined as the stochastic process $\bar{X}_i^{(n_i)} := (X_i, X_i', X_i'', \dots, X_i^{(n_i)}) : T \times \Omega \rightarrow \mathcal{X}_i^{n_i+1}$, is the tuple of all the derivatives of X_i up to and including order n_i . We adopt a similar notation for the values in $\mathcal{X}_i^{n_i+1}$, that is, $\bar{x}_i^{(n_i)} \in \mathcal{X}_i^{n_i+1}$. Each component $X_i^{(k_i)}$ of $\bar{X}_i^{(n_i)}$, or similarly $x_i^{(k_i)}$ of $\bar{x}_i^{(n_i)}$, corresponds to an index $i^{(k_i)}$, which gives the index set $\bar{i}^{(n_i)} := \{i^{(k_i)} : 0 \leq k_i \leq n_i\}$ for $\bar{X}_i^{(n_i)}$, where the index $i^{(0)}$ is also written as i .

Let $\mathbf{n} = (n_1, \dots, n_d) \in \mathbb{N}_0^{\mathcal{I}}$ be a tuple. If the n_i^{th} -order derivative of X_i exists for every $i \in \mathcal{I}$, then the *\mathbf{n}^{th} -order derivative* of \mathbf{X} is defined as the stochastic process $\mathbf{X}^{(\mathbf{n})} := (X_1^{(n_1)}, \dots, X_d^{(n_d)}) : T \times \Omega \rightarrow \mathcal{X}$ and the *complete \mathbf{n}^{th} -order derivative* of \mathbf{X} is defined as the stochastic process $\bar{\mathbf{X}}^{(\mathbf{n})} := (\bar{X}_1^{(n_1)}, \bar{X}_2^{(n_2)}, \dots, \bar{X}_d^{(n_d)}) :$

$T \times \Omega \rightarrow \mathcal{X}^{n+1}$, where $\mathcal{X}^{n+1} := \prod_{i=1}^d \mathcal{X}_i^{n_i+1}$. We adopt a similar notation for the values in \mathcal{X}^{n+1} , that is, $\bar{x}^{(n)} \in \mathcal{X}^{n+1}$. Similarly, each component $X_i^{(k_i)}$ of $\bar{X}^{(n)}$ corresponds to an index $i^{(k_i)}$ which gives the index set $\bar{\mathcal{I}}^{(n)} := \bigcup_{i \in \mathcal{I}} \bar{i}^{(n_i)}$ for $\bar{X}^{(n)}$.

For a subset $I := \{i_1, \dots, i_k\} \subseteq \mathcal{I}$ we will use the notation $\mathbf{n}_I := (n_{i_1}, \dots, n_{i_k})$ and write $\mathcal{X}_I = \prod_{i \in I} \mathcal{X}_i$ and $\mathcal{X}_I^{n_I+1} = \prod_{i \in I} \mathcal{X}_i^{n_i+1}$. For the I^{th} components of the process \mathbf{X} and the complete \mathbf{n}^{th} -order derivative $\bar{X}^{(n)}$, we write $\mathbf{X}_I := (X_{i_1}, \dots, X_{i_k})$ and $\bar{\mathbf{X}}_I^{(n_I)} := (\bar{X}_{i_1}^{(n_{i_1})}, \dots, \bar{X}_{i_k}^{(n_{i_k})})$ respectively. Similarly, for the values in \mathcal{X}_I and $\mathcal{X}_I^{n_I+1}$, we write $\mathbf{x}_I := (x_{i_1}, \dots, x_{i_k}) \in \mathcal{X}_I$ and $\bar{\mathbf{x}}_I^{(n_I)} := (\bar{x}_{i_1}^{(n_{i_1})}, \dots, \bar{x}_{i_k}^{(n_{i_k})}) \in \mathcal{X}_I^{n_I+1}$ respectively.

In this notation, a stochastic process X is a C^n -stochastic process, if its complete \mathbf{n}^{th} -order derivative $\bar{X}^{(n)}$ exists and is continuous. Similarly, we call a stochastic process \mathbf{X} a C^n -stochastic process, if its complete \mathbf{n}^{th} -order derivative $\bar{\mathbf{X}}^{(n)}$ exists and is continuous.

3.2 Structural dynamical causal models and their solutions

Informally, we think of an SDCM as an SCM where we replace the random variables of the SCM by stochastic processes and their derivatives, and where each structural equation of the SCM becomes a random differential equation of arbitrary order. This generalizes the class of SCMs to the continuous time domain and enables a causal semantics for a broad range of random dynamical models. In this paper, we closely follow the terminology of Bongers et al. [9] for SCMs and extend it to SDCMs.

Definition 3.1 (Structural dynamical causal model). *A structural dynamical causal model (SDCM) is a tuple⁷*

$$\mathcal{R} := \langle \mathcal{I}, \mathcal{J}, \mathcal{X}, \mathcal{E}, \mathbf{n}, \mathbf{f}, \mathbf{E} \rangle$$

where

- \mathcal{I} is a finite index set for endogenous processes,
- \mathcal{J} is a disjoint finite index set for exogenous processes,
- $\mathcal{X} = \prod_{i \in \mathcal{I}} \mathcal{X}_i$ is the product of the domains of the endogenous processes, where each domain $\mathcal{X}_i = \mathbb{R}^{d_i}$,
- $\mathcal{E} = \prod_{j \in \mathcal{J}} \mathcal{E}_j$ is the product of the domains of the exogenous processes, where each domain $\mathcal{E}_j = \mathbb{R}^{e_j}$,
- $\mathbf{n} = (n_i)_{i \in \mathcal{I}} \in \mathbb{N}_0^{\mathcal{I}}$ is the order tuple,
- $\mathbf{f} : \mathcal{X}^{n+1} \times \mathcal{E} \rightarrow \mathcal{X}$ is a measurable function that specifies the dynamic causal mechanism,
- $\mathbf{E} : T \times \Omega \rightarrow \mathcal{E}$ is an exogenous stochastic process with independent components, that is, $(E_j)_{j \in \mathcal{J}}$ is independent.

The solutions of a structural dynamical causal model in terms of stochastic processes are defined by the associated dynamic structural equations.

Definition 3.2 (Solution of an SDCM). *A stochastic process $\mathbf{X} : T \times \Omega \rightarrow \mathcal{X}$ is a solution of the dynamic structural equations (dynamic SEs) associated to SDCM \mathcal{R} ,*

$$\mathbf{X} = \mathbf{f}(\bar{\mathbf{X}}^{(n)}, \mathbf{E}),$$

if \mathbf{X} is a C^n -stochastic process, and for \mathbb{P} -almost every $\omega \in \Omega$ the ordinary differential equations⁸

$$\mathbf{X}_t(\omega) = \mathbf{f}(\bar{\mathbf{X}}_t^{(n)}(\omega), \mathbf{E}_t(\omega))$$

hold for all $t \in T$.

⁷ We often use boldface for variables that have multiple components, that is, which take values in a Cartesian product.

⁸ These equations are called *implicit ordinary differential equations* if the Jacobian matrix $\frac{\partial \mathbf{f}(\bar{\mathbf{x}}^{(n)}, \mathbf{e})}{\partial \bar{\mathbf{x}}^{(n)}}$ is nonsingular for all its argument values in an appropriate domain, otherwise they are called *differential-algebraic equations* [2].

The value n_i of the order tuple \mathbf{n} denotes the highest-order derivative of X_i that may occur in the dynamic structural equations. Note that taking higher n_i 's will in general reduce the set of possible solutions, due to additional imposed smoothness constraints on the solutions. In contrast to the common way of writing RDEs (see equation (1)), the (higher-order) derivatives of the endogenous processes of an SDCM always appear on the right-hand side of the dynamic SEs.⁹ This notation explicitly allows us to model zeroth-order dynamic structural equations, that is, equations that contain no derivatives of order one or higher, in other words, random algebraic equations.

In particular, if *all* dynamic structural equations are of zeroth order and the exogenous stochastic processes in the model are constant in time (that is, random variables), then the structural dynamical causal model reduces to a structural causal model [see 9]. In contrast to [9], we define an SCM here in terms of an exogenous random variable instead of an exogenous distribution.

Definition 3.3 (Structural causal model). *A structural causal model (SCM) is a tuple*

$$\mathcal{M} := \langle \mathcal{I}, \mathcal{J}, \mathbf{X}, \mathbf{E}, \mathbf{f}, \mathbf{E} \rangle,$$

such that $\langle \mathcal{I}, \mathcal{J}, \mathbf{X}, \mathbf{E}, \mathbf{0}, \mathbf{f}, \mathbf{E} \rangle$ is an SDCM with \mathbf{E} a random variable.

That is, we can identify SCMs with certain special cases of SDCMs. Similarly, we can identify the solutions of an SCM with the (constant) solutions of the corresponding SDCM. The following definition is equivalent to Definition 3.2 when the latter is applied to an SCM.

Definition 3.4 (Solution of an SCM). *A random variable $\mathbf{X} : \Omega \rightarrow \mathcal{X}$ is a solution of the structural equations associated to SCM \mathcal{M} ,*

$$\mathbf{X} = \mathbf{f}(\mathbf{X}, \mathbf{E}),$$

if for \mathbb{P} -almost every $\omega \in \Omega$

$$\mathbf{X}(\omega) = \mathbf{f}(\mathbf{X}(\omega), \mathbf{E}(\omega))$$

holds.

Similar to the structural equations of an SCM [56, 70], the dynamic structural equations of an SDCM model the underlying causal mechanisms in a structural way, that is, each dynamic structural equation expresses a specific endogenous process (on the left-hand side) in terms of a dynamic causal mechanism depending on certain processes and their derivatives (on the right-hand side). It is this additional structure, which allows us to explicitly model the causal semantics, that distinguishes structural dynamical causal models from dynamical models such as ODEs and RDEs.¹⁰ Allowing for zeroth and higher-order derivatives of X_i in the dynamic structural equations gives rise to a broad range of random dynamical models that can be described by an SDCM, ranging from ODEs (including first-order ODEs as in [52]), RDEs (as in Section 2.2) and more general random dynamical systems such as partially equilibrated systems (as in [38]).

Example 3.5 (Damped coupled harmonic oscillator). *Consider a one-dimensional system of d point masses $m_i > 0$ ($i = 1, \dots, d$) with positions $X_i \in \mathbb{R}$, which are coupled by ideal springs, with spring constants $\kappa_i > 0$ and equilibrium lengths $L_i > 0$ ($i = 1, \dots, d-1$), under influence of friction with friction coefficients $b_i \geq 0$ ($i = 1, \dots, d$) (see Figure 3 left). This system can be modeled by the SDCM¹¹ $\mathcal{R} = \langle \{1, \dots, d\}, \{1, \dots, d-1\}, \mathbb{R}^d, \mathbb{R}^{d-1}, \mathbf{n}, \mathbf{f}, \mathbf{E} \rangle$ with order tuple $\mathbf{n} := (2, \dots, 2)$, where the exogenous*

⁹ For every RDE of the form $\mathbf{X}' = \mathbf{f}(\mathbf{X}, \mathbf{E})$ with \mathbf{f} and \mathbf{E} continuous, there exists an SDCM with the same solutions: \mathbf{X} is a solution of the RDE if and only if it is a solution of the SDCM \mathcal{R} with the dynamic SE $\mathbf{X} = \mathbf{X} - \mathbf{X}' + \mathbf{f}(\mathbf{X}, \mathbf{E})$, as long as $\mathbf{n} = 1$ (since all solutions of the RDE must be continuously differentiable).

¹⁰ The importance of assigning a differential equation to an endogenous variable was already observed in [52].

¹¹ We abuse notation here; more formally, we should use an index set for \mathcal{J} that is disjoint from \mathcal{I} , for example, $\{\tilde{1}, \dots, \tilde{d}-1\}$.




Figure 3. Damped coupled harmonic oscillator model \mathcal{R} of Example 3.5 (left) and the intervened model $\mathcal{R}_{\text{do}(\{1,5\},(0,L))}$ of Example 3.8 (right), both for $d = 5$.

process $\mathbf{E} = \mathbf{L} := (L_1, \dots, L_{d-1})$ is constant in time (that is, a random variable), and the causal mechanism is specified by the dynamic structural equations

$$\begin{cases} X_1 = f_1(\bar{\mathbf{X}}^{(n)}, \mathbf{L}) := X_2 - L_1 - \frac{b_1}{\kappa_1} X_1' - \frac{m_1}{\kappa_1} X_1'' \\ X_i = f_i(\bar{\mathbf{X}}^{(n)}, \mathbf{L}) := \frac{\kappa_i}{\kappa_i + \kappa_{i-1}} (X_{i+1} - L_i) + \frac{\kappa_{i-1}}{\kappa_i + \kappa_{i-1}} (X_{i-1} + L_{i-1}) \\ \quad - \frac{b_i}{\kappa_i + \kappa_{i-1}} X_i' - \frac{m_i}{\kappa_i + \kappa_{i-1}} X_i'' \quad (i = 2, \dots, d-1) \\ X_d = f_d(\bar{\mathbf{X}}^{(n)}, \mathbf{L}) := X_{d-1} + L_{d-1} - \frac{b_d}{\kappa_{d-1}} X_d' - \frac{m_d}{\kappa_{d-1}} X_d'' . \end{cases}$$

The motion of the masses, in terms of their positions X_i , velocities X_i' and accelerations X_i'' , is described by a separate equation of motion for each mass. For the case $d = 2$, this SDCM \mathcal{R} has the same solutions as those described by the RDE in Example 2.1.

The following example motivates why we only consider processes as solutions of SDCMs in case they satisfy the smoothness conditions.

Example 3.6 (Sufficient smoothness of the solutions). Let $\mathcal{R} = \langle \{1\}, \emptyset, \mathcal{X}, \mathcal{E}, n, f, E \rangle$ be the SDCM with $\mathcal{X} = \mathbb{R}$, \mathcal{E} the singleton $\{*\}$, $n = 0$, the dynamic causal mechanism $f : \mathcal{X} \times \mathcal{E} \rightarrow \mathcal{X}$ given by $f(x, e) = x - x^2 + 1$, and E the trivial exogenous process. The zeroth-order dynamic structural equation associated to \mathcal{R} reads

$$X = X - X^2 + 1.$$

This dynamic structural equation does not depend on any exogenous process. The set of endogenous processes $X : T \times \Omega \rightarrow \mathbb{R}$ that satisfy the dynamic structural equation consists of all stochastic processes in $\{-1, 1\}^{T \times \Omega}$. Most of the stochastic processes in $\{-1, 1\}^{T \times \Omega}$ are not continuous. The solutions of \mathcal{R} are exactly those processes in $\{-1, 1\}^{T \times \Omega}$ that are C^0 -stochastic processes. These are the processes that are constant in time, that is, the random variables of $\{-1, 1\}^\Omega$. In particular, the solutions of the SDCM \mathcal{R} correspond exactly to the solutions of the SCM described by the above structural equation.

3.3 Interventions

Interventions on a structural dynamical causal model can be modeled in different ways. We consider here a stochastic version of *perfect* interventions¹² on the endogenous processes [22] that are analogous to stochastic perfect interventions in structural causal models [21, 56]. A stochastic perfect intervention on some endogenous process forces the intervened process to be equal to a given independent exogenous process. More generally, we model a stochastic perfect intervention on a subset $I := \{i_1, \dots, i_k\} \subseteq \mathcal{I}$ of the endogenous processes by forcing those processes \mathbf{X}_I to be equal to the intervened processes \mathbf{K}_I , by changing the model such that the corresponding dynamical structural equations become $\mathbf{X}_I = \mathbf{K}_I$. The

¹² These are also referred to as *ideal*, *hard*, *structural*, *surgical*, *atomic* [22] or *independent* [44] interventions.

process \mathbf{K}_I is treated as an independent exogenous process, such that all its components K_i are mutually independent and independent from all the other exogenous processes that were already present in the model in the absence of the intervention. The dynamic causal mechanisms of the other endogenous processes $\mathcal{I} \setminus I$ are untouched and their dynamics are still specified by the same dynamic structural equations associated to those processes in the absence of the intervention, that is¹³

$$\mathbf{X}_{\setminus I} = \mathbf{f}_{\setminus I}(\bar{\mathbf{X}}^{(n)}, \mathbf{E}).$$

This yields the following formal definition of an intervened structural dynamical causal model.

Definition 3.7 (Stochastic perfect intervention on an SDCM). *Let $\mathcal{R} = \langle \mathcal{I}, \mathcal{J}, \mathcal{X}, \mathcal{E}, \mathbf{n}, \mathbf{f}, \mathbf{E} \rangle$ be an SDCM, $I \subseteq \mathcal{I}$ a subset, and $\mathbf{K}_I : T \times \Omega \rightarrow \mathcal{X}_I$ a stochastic process such that $((K_i)_{i \in I}, (E_j)_{j \in \mathcal{J}})$ is independent. The stochastic perfect intervention $\text{do}(I, \mathbf{K}_I)$ maps \mathcal{R} to the SDCM¹⁴*

$$\mathcal{R}_{\text{do}(I, \mathbf{K}_I)} := \langle \mathcal{I}, I \cup \mathcal{J}, \mathcal{X}, \mathcal{X}_I \times \mathcal{E}, \mathbf{n}, \tilde{\mathbf{f}}, (\mathbf{K}_I, \mathbf{E}) \rangle,$$

where the intervened causal mechanism $\tilde{\mathbf{f}} : \mathcal{X}^{n+1} \times (\mathcal{X}_I \times \mathcal{E}) \rightarrow \mathcal{X}$ is given by

$$\tilde{f}_i(\bar{\mathbf{x}}^{(n)}, (e_I, e_{\mathcal{J}})) = \begin{cases} f_i(\bar{\mathbf{x}}^{(n)}, e_{\mathcal{J}}) & i \in \mathcal{I} \setminus I \\ e_i & i \in I. \end{cases} \quad (2)$$

We call a stochastic perfect intervention $\text{do}(I, \mathbf{K}_I)$ a perfect intervention if \mathbf{K}_I is a deterministic stochastic process (that is, if it does not depend on ω).

This definition explicitly exposes a hitherto implicit but crucial modeling assumption: exogenous processes are not caused by endogenous processes. Indeed, no stochastic perfect intervention on any subset of the endogenous processes will lead to a change in any of the exogenous processes.

Example 3.8. Consider the damped coupled harmonic oscillator represented by the SDCM \mathcal{R} of Example 3.5. Performing the perfect interventions on the masses m_1 and m_d by fixing m_1 and m_d to the walls at $X_1 = 0$ and $X_d = L > 0$, respectively, (see Figure 3 (right)) yields the model $\mathcal{R}_{\text{do}(\{1, d\}, (0, L))}$ with the dynamic structural equations

$$\begin{cases} X_1 = 0 \\ X_i = \frac{\kappa_i}{\kappa_i + \kappa_{i-1}}(X_{i+1} - L_i) + \frac{\kappa_{i-1}}{\kappa_i + \kappa_{i-1}}(X_{i-1} + L_{i-1}) \\ \quad - \frac{b_i}{\kappa_i + \kappa_{i-1}}X'_i - \frac{m_i}{\kappa_i + \kappa_{i-1}}X''_i \quad (i = 2, \dots, d-1) \\ X_d = L. \end{cases}$$

It is clear from the definition that performing stochastic perfect interventions on disjoint subsets of the endogenous processes commutes. In case of overlap, the dynamic structural equations of the overlapping intervention targets are determined by the most recent intervention applied to them.

As a special case, Definition 3.7 reduces to the usual notion of (stochastic) perfect intervention on SCMs [see 9].

Definition 3.9 (Stochastic perfect intervention on an SCM). *Let $\mathcal{M} = \langle \mathcal{I}, \mathcal{J}, \mathcal{X}, \mathcal{E}, \mathbf{f}, \mathbf{E} \rangle$ be an SCM, $I \subseteq \mathcal{I}$ a subset, and $\mathbf{K}_I : \Omega \rightarrow \mathcal{X}_I$ a random variable such that $((K_i)_{i \in I}, (E_j)_{j \in \mathcal{J}})$ is independent. The stochastic perfect intervention $\text{do}(I, \mathbf{K}_I)$ maps \mathcal{M} to the SCM*

$$\mathcal{M}_{\text{do}(I, \mathbf{K}_I)} := \langle \mathcal{I}, I \cup \mathcal{J}, \mathcal{X}, \mathcal{X}_I \times \mathcal{E}, \tilde{\mathbf{f}}, (\mathbf{K}_I, \mathbf{E}) \rangle,$$

¹³ For $I \subseteq \mathcal{I}$ we adopt the notation $\setminus I$ for $\mathcal{I} \setminus I$.

¹⁴ We abuse notation here; more formally, we should make a disjoint copy $\tilde{I} := \{\tilde{i} : i \in I\}$ and use $\tilde{I} \cup \mathcal{J}$ as the new exogenous index set instead of $I \cup \mathcal{J}$, to keep the endogenous indices \mathcal{I} and the exogenous indices $\tilde{I} \cup \mathcal{J}$ disjoint.

where \tilde{f} is defined by equation (2).

This provides SDCMs with a causal semantics that is analogous to that of SCMs. The following example illustrates how this resolves the ambiguity of the causal interpretation of the RDE of Example 2.2.

Example 3.10 (Ambiguous causal interpretation of RDEs). *Consider the SDCM \mathcal{R} of Example 3.5 for $d = 2$, with dynamic structural equations given by*

$$\begin{cases} X_1 = X_2 - L_1 - \frac{b_1}{\kappa_1} X_1' - \frac{m_1}{\kappa_1} X_1'' \\ X_2 = X_1 + L_1 - \frac{b_2}{\kappa_1} X_2' - \frac{m_2}{\kappa_1} X_2'' \end{cases}$$

The solutions of \mathcal{R} correspond exactly to the solutions of the RDE that describes the two masses attached to a spring in Example 2.1. Fixing the mass m_1 to the left wall at $X_1 = 0$ (see Figure 2 (right)) by performing the stochastic perfect intervention¹⁵ $\text{do}(1, K_1)$ with $K_1 = 0$ on \mathcal{R} gives the intervened model $\mathcal{R}_{\text{do}(1,0)}$ with dynamic structural equations given by

$$\begin{cases} X_1 = 0 \\ X_2 = X_1 + L_1 - \frac{b_2}{\kappa_1} X_2' - \frac{m_2}{\kappa_1} X_2'' \end{cases}$$

The intervened model $\mathcal{R}_{\text{do}(1,0)}$ has exactly the same solutions as the RDE in Example 2.2.

Consider now the SDCM $\tilde{\mathcal{R}}$ that is the same as \mathcal{R} except for its dynamic causal mechanism \tilde{f} , for which the associated dynamic structural equations are given by

$$\begin{cases} X_1 = X_2 - L_1 + \frac{b_2}{\kappa_1} X_2' + \frac{m_2}{\kappa_1} X_2'' \\ X_2 = X_1 + L_1 + \frac{b_1}{\kappa_1} X_1' + \frac{m_1}{\kappa_1} X_1'' \end{cases}$$

Both models \mathcal{R} and $\tilde{\mathcal{R}}$ have the same solutions as those described by the RDEs in Example 2.1. However, the intervened models $\mathcal{R}_{\text{do}(1,0)}$ and $\tilde{\mathcal{R}}_{\text{do}(1,0)}$ have different solutions. Only the model $\mathcal{R}_{\text{do}(1,0)}$ describes the expected physical behavior (see also Example 2.2).

Stochastic perfect interventions are only defined for the endogenous processes, but not for their higher-order derivatives. The higher-order derivative processes in an SDCM are always obtained by differentiation of the underlying endogenous processes and hence it suffices to define the stochastic perfect interventions only for those underlying endogenous processes. Allowing for stochastic perfect intervention on both the endogenous processes and some of their higher-order derivatives will generally lead to nonsensible causal behavior, as is illustrated in the following example.

Example 3.11 (Modeling higher-order derivatives as separate endogenous processes). *Suppose we model the velocities X_i' of the positions X_i of the masses between the walls in the damped coupled harmonic oscillator of Example 3.8 explicitly as separate endogenous processes $V_{i'}$. We could attempt to model this with an SDCM $\tilde{\mathcal{R}}$ for which the dynamic structural equations are given by $X_1 = 0$, $X_d = L$ and*

$$\begin{cases} X_i = \frac{\kappa_i}{\kappa_i + \kappa_{i-1}} (X_{i+1} - L_i) + \frac{\kappa_{i-1}}{\kappa_i + \kappa_{i-1}} (X_{i-1} + L_{i-1}) - \frac{b_i}{\kappa_i + \kappa_{i-1}} V_{i'} - \frac{m_i}{\kappa_i + \kappa_{i-1}} V_{i'}' \\ V_{i'} = X_i' \end{cases}$$

for $i = 2, \dots, d-1$. Performing a stochastic perfect intervention on both the position X_i and the velocity $V_{i'}$ of one of the masses between the walls ($i \in \{2, \dots, d-1\}$) can lead to unphysical behavior. For example, the perfect intervention $\text{do}(\{2, 2'\}, (0, 1))$ gives an intervened SDCM with a solution that is physically impossible if we keep interpreting X_i as the position and $V_{i'}$ the velocity of the i^{th} mass.

¹⁵ For convenience, we write $\text{do}(i, K_i)$ for a stochastic perfect intervention $\text{do}(I, \mathbf{K}_I)$ whenever $I = \{i\}$ for some $i \in \mathcal{I}$.

This observation constitutes strong motivation for considering the higher-order derivatives $X_i^{(k_i)}$ (up to and including order n_i) to be aspects of the endogenous process X_i rather than as “causally independent” processes. Thereby, we circumvent modeling velocity as the (instantaneous) cause of position [as in 38], or the other way around. The resulting modeling framework appears more natural than that of [52], which is explicitly limited to first-order dynamics and cannot accommodate systems like the damped harmonic oscillator as easily as SDCMs can, as it has to impose restrictions on the possible interventions to deal with this problem.

The higher-order derivatives $X_i^{(k_i)}$ do not always exist for a process X_i . For example, if we force the mass m_1 to follow a Brownian motion¹⁶ K_1 in the spring model \mathcal{R} of Example 3.10, then the intervened model $\mathcal{R}_{\text{do}(1, K_1)}$ does not yield a solution (because X_1'' needs to exist and be continuous, which is not the case for $X_1 = K_1$). In practice, we therefore only consider stochastic perfect interventions $\text{do}(I, \mathbf{K}_I)$ for which \mathbf{K}_I is a C^{n_I} -stochastic process.

3.4 Graph

We will now define a graphical representation of the structural properties of SDCMs that is inspired by the graphical representation of SCMs [9, 56]. Where the graph of an SCM describes the functional relationships between the random variables encoded by the structural equations, the graph of an SDCM expresses the functional dependencies between the stochastic processes encoded by the dynamic structural equations.

Typically, for $i \in \mathcal{I}$, the component f_i of the dynamic causal mechanism \mathbf{f} only depends on a subset of the (derivatives of the) endogenous and exogenous processes that we call the *functional parents* of i .

Definition 3.12 (Functional and integrated parents). *Let $\mathcal{R} = \langle \mathcal{I}, \mathcal{J}, \mathcal{X}, \mathcal{E}, \mathbf{n}, \mathbf{f}, \mathbf{E} \rangle$ be an SDCM. For $k \in \bar{\mathcal{I}}^{(n)} \cup \mathcal{J}$ and $i \in \bar{\mathcal{I}}^{(n)}$, we call*

1. k a functional parent of i if and only if $i \in \mathcal{I}$ and there does not exist a measurable function¹⁷ $\tilde{f}_i : (\mathcal{X}^{n+1})_{\setminus k} \times \mathcal{E}_{\setminus k} \rightarrow \mathcal{X}_i$ such that for all $\mathbf{e} \in \mathcal{E}$ and for all $\bar{\mathbf{x}}^{(n)} \in \mathcal{X}^{n+1}$

$$x_i = f_i(\bar{\mathbf{x}}^{(n)}, \mathbf{e}) \iff x_i = \tilde{f}_i((\bar{\mathbf{x}}^{(n)})_{\setminus k}, \mathbf{e}_{\setminus k});$$

2. k an integrated parent of i if and only if there exists an $\ell \in \mathcal{I}$ such that $k = \ell^{(m_\ell-1)}$ and $i = \ell^{(m_\ell)}$ for some $0 < m_\ell \leq n_\ell$.

Exogenous processes have no functional and integrated parents by definition. The integrated parents denote the differential relationships that are satisfied by the endogenous processes. That is, for every $\ell \in \mathcal{I}$ and $0 < m_\ell \leq n_\ell$ we have that $\ell^{(m_\ell-1)}$ is an integrated parent of $\ell^{(m_\ell)}$, which represents the differential relationship

$$X_\ell^{(m_\ell)} = \frac{d}{dt} X_\ell^{(m_\ell-1)}.$$

These differential relationships are absent for SCMs, because the endogenous variables are considered static. In contrast to [38], we express the differential relationships between the endogenous processes by the derivative operator, instead of the integration operator. In general, the integration operator of [38] is not uniquely defined, since for a particular process there may exist several integrated processes differing by a (possibly random) integration constant. The derivative of a process, however, is always a.s. uniquely defined, if it exists. Hence, for a solution \mathbf{X} of an SDCM we can always derive the higher-order derivatives of X_i

¹⁶ A stochastic process B on $T = [0, \infty)$ is called a *Brownian motion* if: (i) $B_0 = 0$; (ii) B has independent, stationary increments; (iii) $B_t \sim \mathcal{N}(0, t)$ for all $t > 0$; (iv) B is continuous. In particular, B is not differentiable [see, for example, Theorem 21.17 in 42].

¹⁷ For $\mathcal{X}^{n+1} = \prod_{i(k_i) \in \bar{\mathcal{I}}^{(n)}} \mathcal{X}_i$, some subset $I \subseteq \bar{\mathcal{I}}^{(n)}$ and $k \in \bar{\mathcal{I}}^{(n)}$, we denote $(\mathcal{X}^{n+1})_{\setminus I} = \prod_{i(k_i) \in \bar{\mathcal{I}}^{(n)} \setminus I} \mathcal{X}_i$ and $(\mathcal{X}^{n+1})_{\setminus k} = \prod_{i(k_i) \in \bar{\mathcal{I}}^{(n)} \setminus \{k\}} \mathcal{X}_i$, and similarly for their elements.

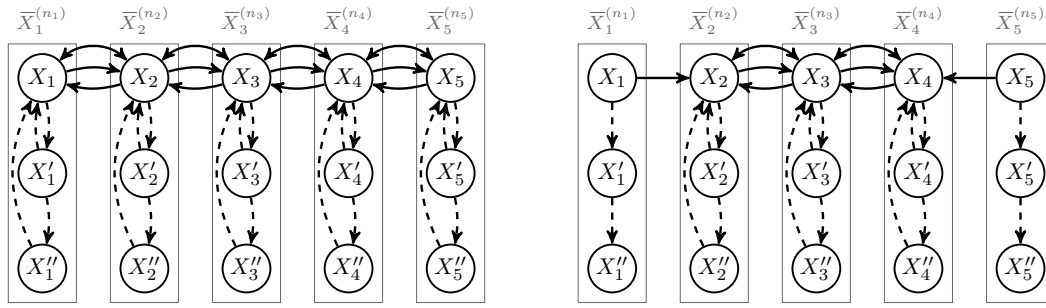


Figure 4. Graphs of the damped coupled harmonic oscillator model \mathcal{R} of Example 3.5 (left) and the intervened model $\mathcal{R}_{\text{do}(\{1,5\},(0,L))}$ of Example 3.8 (right), both for $d = 5$.

up to order n_i by repeatedly applying the derivative operator. In this way, we can consider the complete n_i^{th} -order derivative $\bar{X}_i^{(n_i)}$ to encode aspects of the same endogenous process X_i .

The different parental relations can be expressed in a clustered mixed graph,¹⁸ where each cluster represents a complete n_i^{th} -order derivative.

Definition 3.13 (Graph). Let $\mathcal{R} = \langle \mathcal{I}, \mathcal{J}, \mathcal{X}, \mathcal{E}, \mathbf{n}, \mathbf{f}, \mathbf{E} \rangle$ be an SDCM with order tuple \mathbf{n} . The graph $\mathcal{G}(\mathcal{R})$ of \mathcal{R} is the clustered mixed graph with nodes $\bar{\mathcal{I}}^{(\mathbf{n})}$ partitioned into clusters $\bar{i}^{(n_i)} = \{i^{(k_i)} : 0 \leq k_i \leq n_i\}$ for $i \in \mathcal{I}$, directed edges $i \rightarrow j$ if and only if i is functional parent of j in a different cluster, dashed directed edges $i \rightarrow j$ if and only if i is a functional or integrated parent of j in the same cluster, and bidirected edges $i \leftrightarrow j$ if and only if there exists a $k \in \mathcal{J}$ that is a functional parent of both i and j .

A cluster $\bar{i}^{(n_i)}$ of the graph of an SDCM refers to the complete n_i^{th} -order derivative $\bar{X}_i^{(n_i)}$ and is represented by a box around the nodes of the cluster. This is illustrated¹⁹ in Figure 4 (left) for the damped coupled harmonic oscillator model of Example 3.5. Between the nodes of different clusters there are only functional parental relations. Within a cluster, the higher-order derivatives $i^{(k_i)}$ for $k_i > 0$ of the endogenous processes $i \in \mathcal{I}$ have no functional parents, but have only integrated parents. However, any node $i^{(k_i)}$ with $k_i > 0$ may be a functional parent of another node $j \in \mathcal{I}$; see, for example, the graph of the SDCM $\tilde{\mathcal{R}}$ in Example 3.10.²⁰

In particular, this definition of the graph of an SDCM reduces to the usual notion of the graph of an SCM if we ignore the clusters. That is, the graph $\mathcal{G}(\mathcal{M})$ of an SCM $\mathcal{M} = \langle \mathcal{I}, \mathcal{J}, \mathcal{X}, \mathcal{E}, \mathbf{f}, \mathbf{E} \rangle$ is a mixed graph with nodes \mathcal{I} , directed edges $i \rightarrow j$ if and only if i is a functional parent of j with $i \neq j$, dashed directed edge $i \rightarrow i$ if and only if i is a functional parent of itself, and bidirected edges $i \leftrightarrow j$ if and only if there exists a $k \in \mathcal{J}$ that is a functional parent of both i and j , where we apply Definition 3.12 of a functional parent to \mathcal{M} (note that by definition, an SCM has no integrated parents).

On the graph of an SDCM, the operation of a stochastic perfect intervention acts in a simple way.

Proposition 3.14 (Graph of the intervened SDCM). Let \mathcal{R} be an SDCM and $\text{do}(I, \mathbf{K}_I)$ a stochastic perfect intervention for $I \subseteq \mathcal{I}$ a subset and \mathbf{K}_I an independent stochastic process. The graph $\mathcal{G}(\mathcal{R}_{\text{do}(I, \mathbf{K}_I)})$ of

¹⁸ A *mixed graph* is a pair $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where \mathcal{V} is a set of nodes and \mathcal{E} is a set of edges between the nodes of different types, in our case, $\rightarrow, \leftarrow, \leftrightarrow, \rightarrow, \leftarrow$. A *clustered mixed graph* is a triple $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathcal{P})$ where $(\mathcal{V}, \mathcal{E})$ is a mixed graph and \mathcal{P} is a partition of the set \mathcal{V} . Each element of \mathcal{P} is called a *cluster* of the clustered mixed graph.

¹⁹ For visualizing the graphs we stick to the common convention of using stochastic processes and random variables with the index as a subscript, instead of using the indices themselves (even when no solutions are defined).

²⁰ A more realistic example could be Faraday's law of induction. In terms of individual point charges: a moving point charge generates a magnetic field, which exerts a force on some other point charge that is proportional to the velocity of the moving point charge.




Figure 5. Graphs of the price, supply and demand model \mathcal{R} (left) of Example 3.15 and the corresponding equilibrated model $\mathcal{M}_{\mathcal{R}}$ (right) of Example 4.17.

the intervened SDCM $\mathcal{R}_{\text{do}(I, \mathbf{K}_I)}$ is the graph $\mathcal{G}(\mathcal{R})$, but without the edges that have an arrowhead pointing towards a node in the intervention target set I .

The graph of the damped coupled harmonic oscillator model of Example 3.8, where we performed the perfect intervention of fixing the endpoint masses to the walls, is illustrated in Figure 4 (right). Performing a stochastic perfect intervention on an endogenous process removes all the (bi-)directed edges that point towards the intervened process, including the dashed directed edges within the cluster. The dashed directed edges within the cluster that correspond to the integrated parents, that is, those pointing to a higher-order derivative, indicate that the higher-order derivatives of the intervened endogenous process need to exist for any solution of the model. Hence, we view a stochastic perfect intervention on an endogenous process as an intervention on the whole cluster of the intervened process. We say that there is a *directed edge from cluster I to cluster J* , if there exists a directed edge from some $i \in I$ to some $j \in J$. Since a stochastic perfect intervention can be seen as an intervention on the entire associated cluster, the directed edges between the clusters express the direct causal relationships between the clusters. We call a dashed directed edge $i \rightarrow i$ in the graph of an SDCM (that is, where i is a functional parent of itself) a *self-cycle at i* . An example of a model where a self-cycle arises is the well-known market equilibrium model from economics, which has been thoroughly discussed in the literature [see, for example, 60].

Example 3.15 (Price, supply and demand). Let X_P denote the price, X_S denote the supply and X_D the demand of a quantity of a product. The following dynamic structural equations specify an SDCM \mathcal{R} that describes how the demanded and supplied quantities are determined by the price, and how price adjustments occur in the market:

$$\begin{cases} X_P = X_P + \lambda(X_D - X_S) - X'_P \\ X_S = \beta_S X_P + E_S \\ X_D = \beta_D X_P + E_D, \end{cases}$$

where $\mathbf{n} := (n_P, n_S, n_D) = (1, 0, 0)$. Here, E_S and E_D are the exogenous influences on the supply and demand respectively, $\beta_S > 0$ is the reciprocal of the slope of the supply curve, $\beta_D < 0$ is the reciprocal of the slope of the demand curve, and $\lambda > 0$ models how fast the price adjusts to market conditions. The graph of this model is depicted in Figure 5 (left) and contains a self-cycle at P .

Until now, we encountered already several instances of *linear* SDCMs (for example, in Example 3.5 and 3.15).

Definition 3.16 (Linear SDCM). We call an SDCM \mathcal{R} linear, if the dynamic causal mechanism $\mathbf{f} : \mathcal{X}^{n+1} \times \mathcal{E} \rightarrow \mathcal{X}$ is of the form

$$\mathbf{f}(\bar{\mathbf{x}}^{(\mathbf{n})}, \mathbf{e}) := B\bar{\mathbf{x}}^{(\mathbf{n})} + \Gamma\mathbf{e},$$

where $B \in \mathbb{R}^{\mathcal{I} \times \bar{\mathcal{I}}^{(\mathbf{n})}}$ and $\Gamma \in \mathbb{R}^{\mathcal{I} \times \mathcal{J}}$ are matrices.

For a linear SDCM \mathcal{R} , a nonzero coefficient B_{ij} for $i, j \in \bar{\mathcal{I}}^{(\mathbf{n})}$ such that $i \neq j$ corresponds to a directed edge $j \rightarrow i$ in the graph $\mathcal{G}(\mathcal{R})$ if i lies in a different cluster than j , and a dashed directed edge $j \rightarrow i$ if i

lies in the same cluster as j . A coefficient $B_{ii} = 1$ for $i \in \mathcal{I}$ corresponds to a self-cycle $i \rightarrow i$. There is a bidirected edge $i \leftrightarrow j$ in the graph $\mathcal{G}(\mathcal{R})$ for $i, j \in \mathcal{I}$ with $i \neq j$ if and only if there exists a $k \in \mathcal{J}$ for which $\Gamma_{ik} \neq 0$ and $\Gamma_{jk} \neq 0$.

3.5 Initial conditions

In contrast to RDEs, SDCMs allow for both zeroth and higher-order differential equations. For this reason, the dynamic SEs of SDCMs admit problems that can be quite different from those of RDEs. For example, the order of the initial conditions for SDCMs does not directly relate to the order of the SDCM.

Definition 3.17 (Initial condition). *Let \mathcal{R} be an SDCM, $I \subseteq \mathcal{I}$ a subset of the endogenous variables, $\mathbf{m}_I = (m_i)_{i \in I} \in \mathbb{N}_0^I$ an order tuple, $t_0 \in T$ and $\bar{\mathbf{X}}_{I,[0]}^{(\mathbf{m}_I)}$ a random variable taking values in $\mathcal{X}_I^{\mathbf{m}_I+1}$. We say that a solution \mathbf{X} of \mathcal{R} has initial condition $(t_0, \bar{\mathbf{X}}_{I,[0]}^{(\mathbf{m}_I)})$ if $\bar{\mathbf{X}}_I^{(\mathbf{m}_I)}(t_0)$ exists and satisfies*

$$\bar{\mathbf{X}}_I^{(\mathbf{m}_I)}(t_0) = \bar{\mathbf{X}}_{I,[0]}^{(\mathbf{m}_I)}$$

almost surely. Here, \mathbf{m}_I is called the order of the initial condition; for $I = \mathcal{I}$ we also refer to the initial condition as a full initial condition, and for $I \subsetneq \mathcal{I}$ as a partial initial condition. A solution \mathbf{X} of \mathcal{R} with initial condition $(t_0, \bar{\mathbf{X}}_{I,[0]}^{(\mathbf{m}_I)})$ is called almost surely unique if for every solution \mathbf{Y} of \mathcal{R} with initial condition $(t_0, \bar{\mathbf{X}}_{I,[0]}^{(\mathbf{m}_I)})$ we have $\mathbf{X} = \mathbf{Y}$ a.s..

For an SDCM for which the dynamic SEs can be rewritten into the form of a system of n_i^{th} -order RDEs (with all $n_i \geq 1$), the full initial conditions of order $\mathbf{n} - 1$ of the SDCM correspond exactly with the usually considered initial conditions of this system of RDEs. For example, the solutions of the damped coupled harmonic oscillator of Example 3.5 can be a.s. uniquely determined by the full initial conditions of order $\mathbf{n} - 1$ (see also Corollary 3.26). In general, however, the solutions of an SDCM may not be a.s. uniquely determined by the full initial conditions of order $\mathbf{n} - 1$, as the following example illustrates.

Example 3.18 (The order of the SDCM and of the initial conditions). *Let $\mathcal{R} = \langle \{1\}, \emptyset, \mathcal{X}, \mathcal{E}, n, f, E \rangle$ be the SDCM with $\mathcal{X} = \mathbb{R}$, $\mathcal{E} = \{*\}$, $n = 1$, the dynamic causal mechanism $f : \mathcal{X}^2 \times \mathcal{E} \rightarrow \mathcal{X}$ given by $f(\bar{x}^{(1)}, e) = x - x^2 + (x')^2$, and E the trivial exogenous process. The dynamic structural equation associated to \mathcal{R} reads*

$$X = X - X^2 + (X')^2.$$

This dynamic SE cannot be written as a (first-order) RDE of the form (1), since it cannot be a.s. uniquely solved for X' . “Solving for” X' leads to two RDEs that are of the form (1), namely

$$X' = X \quad \text{or} \quad X' = -X.$$

The solutions of these RDEs are given by $X_t = X_{[0]}e^t$ and $X_t = X_{[0]}e^{-t}$ respectively, where $(0, X_{[0]})$ denotes the initial condition for both RDEs. These processes are also solutions of the SDCM, and one can show that all (continuously differentiable) solutions of \mathcal{R} are of this form. Note that, in principle, we could well have taken the order n arbitrarily high without restricting the set of solutions, because the solutions are C^∞ -stochastic processes. If we consider the solutions of \mathcal{R} with an initial condition $(0, \bar{X}_{[0]}^{(0)})$ of order 0, then there are always two solutions with this initial condition that are not a.s. equal to each other, unless $X_{[0]}^{(0)} = 0$. For the initial condition $(0, \bar{X}_{[0]}^{(1)})$ of order 1, we can specify the solution X a.s. uniquely, if it exists. Take for example $\bar{X}_{[0]}^{(1)} = (X_{[0]}, X_{[0]})$, then the solution X with this initial condition is a.s. uniquely given by $X_t = X_{[0]}e^t$. However, an arbitrary initial condition $(0, \bar{X}_{[0]}^{(m)})$ of order m greater or equal to 1 may well be inconsistent with the dynamic structural equations. For example, the initial condition $\bar{X}_{[0]}^{(1)} = (X_{[0]}, 2X_{[0]})$ will not have a solution for $X_{[0]} \neq 0$, since the initial condition $\bar{X}_{[0]}^{(1)} := (X_{[0]}^{(0)}, X_{[0]}^{(1)})$ does not satisfy $(X_{[0]}^{(0)})^2 = (X_{[0]}^{(1)})^2$.

This example illustrates that an arbitrary imposed initial condition may well be inconsistent with the dynamic structural equations.

Definition 3.19 (Consistent initial condition). *Let \mathcal{R} be an SDCM and $\mathbf{m} = (m_i)_{i \in I} \in \mathbb{N}_0^I$ an order tuple for $I \subseteq \mathcal{I}$. We call an initial condition $(t_0, \bar{\mathbf{X}}_{I,[0]}^{(\mathbf{m})})$ for \mathcal{R} consistent if there exists a solution of \mathcal{R} with this initial condition.*

In other words, for an initial condition there only exists a solution if and only if the initial condition is consistent. In particular, zeroth-order dynamic structural equations may constrain the initial conditions (of any order) for which a solution exists.

Example 3.20 (Zeroth-order dynamic structural equation constraint). *Consider the price, supply and demand model \mathcal{R} of Example 3.15 that has order tuple $\mathbf{n} = (n_P, n_S, n_D) = (1, 0, 0)$. The zeroth-order dynamic structural equations of \mathcal{R} are those associated with the supply X_S and the demand X_D processes. Since the solutions of \mathcal{R} satisfy these zeroth-order dynamic structural equations almost surely at every point in time, the consistent full initial conditions $(t_0, \bar{\mathbf{X}}_{[0]}^{(\mathbf{m})})$ also need to satisfy the zeroth-order dynamic structural equations almost surely, that is, $X_{[0],S} = \beta_S X_{[0],P} + (E_S)_{t_0}$ and $X_{[0],D} = \beta_D X_{[0],P} + (E_D)_{t_0}$ almost surely.*

By definition, the consistent full initial conditions always need to satisfy the zeroth-order dynamic structural equations of the SDCM. Initial conditions of an order greater than or equal to the order of the SDCM need to satisfy the corresponding dynamic structural equations of the SDCM, as we already saw in Example 3.18. Additionally, in general, SDCMs that have higher-order dynamic structural equations may contain *hidden* constraints²¹ as the following example illustrates.

Example 3.21 (Hidden constraint). *Consider the SDCM $\mathcal{R} = \langle \{1, 2\}, \{3\}, \mathbb{R}^2, \mathbb{R}, \mathbf{n}, \mathbf{f}, E \rangle$ with $\mathbf{n} = (0, 1)$, the dynamic causal mechanism $\mathbf{f} : \mathcal{X}^{n+1} \times \mathcal{E} \rightarrow \mathcal{X}$ given by $f_1(\bar{\mathbf{x}}^{(\mathbf{n})}, e) := x_2'$ and $f_2(\bar{\mathbf{x}}^{(\mathbf{n})}, e) := e$, and $E := E_3$ some exogenous process. The dynamic structural equations associated to \mathcal{R} read*

$$\begin{cases} X_1 = X_2' \\ X_2 = E. \end{cases}$$

This model cannot be written as an RDE,²² since the Jacobian matrix

$$\frac{\partial \mathbf{f}(\bar{\mathbf{x}}^{(\mathbf{n})}, e)}{\partial \mathbf{x}^{(\mathbf{n})}} := \begin{bmatrix} \frac{\partial f_1}{\partial x_1} & \frac{\partial f_1}{\partial x_2'} \\ \frac{\partial f_2}{\partial x_1} & \frac{\partial f_2}{\partial x_2'} \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$$

is singular everywhere. In order to solve the dynamic SEs we can differentiate the second equation with respect to time to get

$$X_1 = X_2' = E'.$$

This SDCM only has solutions if the derivative E' exists. If it exists, then the solutions are given by $X_1 = E'$ and $X_2 = E$. Thus, the solutions satisfy not only the obvious constraint $X_2 = E$, but also need to satisfy the “hidden” constraint $X_1 = E'$. That a solution of the model depends on a derivative of the exogenous variable E cannot happen in a system of RDEs. These constraints imply that every consistent full initial condition $(t_0, \bar{\mathbf{X}}_{[0]}^{(\mathbf{m})})$ of \mathcal{R} needs to satisfy $X_{[0],1} = E'_{t_0}$ and $X_{[0],2} = E_{t_0}$ almost surely.

After performing a stochastic perfect intervention $\text{do}(I, \mathbf{K}_I)$ on an SDCM \mathcal{R} , all consistent full initial conditions $(t_0, \bar{\mathbf{X}}_{[0]}^{(\mathbf{m})})$ must satisfy $\bar{\mathbf{X}}_{[0],I}^{(\mathbf{m}_I)} = \bar{\mathbf{K}}_I^{(\mathbf{m}_I)}$ almost surely. For example, the consistent initial

²¹ We refer the reader to the literature on differential-algebraic equations for more details on this, for example, [2].

²² Observe that a higher-order RDE is of the form $\mathbf{X}^{(\mathbf{n})} = \mathbf{g}(\bar{\mathbf{X}}^{(\mathbf{n}-1)}, \mathbf{E})$ for some measurable function $\mathbf{g} : \mathcal{X}^{\mathbf{n}} \times \mathcal{E} \rightarrow \mathcal{X}$ and stochastic process $\mathbf{E} : T \times \Omega \rightarrow \mathcal{E}$.

conditions $(t_0, (\bar{X}_{[0],0}^{(1)}, \dots, \bar{X}_{[0],d}^{(1)}))$ for the SDCM \mathcal{R} in Example 3.8 need to satisfy $\bar{X}_{[0],0}^{(1)} = (0, 0)$ and $\bar{X}_{[0],d}^{(1)} = (L, 0)$ after the perfect intervention $\text{do}(\{1, d\}, (0, L))$ on the model.

In summary, Examples 3.18, 3.20 and 3.21 show that the initial (random) value problems associated to dynamic SEs of an SDCM behave differently compared to those of RDEs, as not every initial condition is consistent, and the solutions may involve (higher-order) derivatives of the exogenous process \mathbf{E} .

3.6 Existence and uniqueness of the solutions

For RDEs, there exist sufficient conditions for the existence and uniqueness of the solutions with an initial condition, which are similar to the existence and uniqueness theorems for initial value problems for ODEs [10, 11, 43]. No similar theorem is known in such generality for dynamic SEs, although there are some weaker results of this type for differential-algebraic equations [2]. In this subsection, we provide sufficient conditions for the existence and uniqueness of solutions with a specified initial condition.

We start with an assumption on the form of the dynamic SEs. The class of SDCMs that satisfy this assumption are those SDCMs for which the dynamic SEs corresponding to a subset $I \subseteq \mathcal{I}$ of endogenous variables can be written as an RDE, while the remaining dynamic SEs for the complement $\mathcal{I} \setminus I$ can be solved uniquely for their corresponding endogenous processes. Additionally, smoothness conditions are imposed on exogenous processes and on dynamical causal mechanisms to ensure the required smoothness of the solution.²³

Assumption 1-I. For the SDCM \mathcal{R} and subset $I \subseteq \mathcal{I}$ with order tuple $\mathbf{n}_I \geq 1$, the exogenous process \mathbf{E} is continuous and there exist continuous functions $\mathbf{g}_I : \mathcal{X}_I^{\mathbf{n}_I} \times \mathcal{X}_{\setminus I} \times \mathcal{E} \rightarrow \mathcal{X}_I$ and $\mathbf{g}_{\setminus I} : \mathcal{X}_I^{\mathbf{n}_I} \times \mathcal{E} \rightarrow \mathcal{X}_{\setminus I}$ such that for all $\mathbf{e} \in \mathcal{E}$ and for all $\bar{\mathbf{x}}^{(\mathbf{n})} \in \mathcal{X}^{\mathbf{n}+1}$

$$\mathbf{x}_I^{(\mathbf{n}_I)} = \mathbf{g}_I(\bar{\mathbf{x}}_I^{(\mathbf{n}_I-1)}, \mathbf{x}_{\setminus I}, \mathbf{e}) \iff \mathbf{x}_I = \mathbf{f}_I(\bar{\mathbf{x}}^{(\mathbf{n})}, \mathbf{e})$$

and

$$\mathbf{x}_{\setminus I} = \mathbf{g}_{\setminus I}(\bar{\mathbf{x}}_I^{(\mathbf{n}_I-1)}, \mathbf{e}) \iff \mathbf{x}_{\setminus I} = \mathbf{f}_{\setminus I}(\bar{\mathbf{x}}^{(\mathbf{n})}, \mathbf{e}).$$

Furthermore, for each $j \in \mathcal{I} \setminus I$, either $n_j = 0$, or g_j only depends on \mathbf{e} (that is, $g_j(\bar{\mathbf{x}}_I^{(\mathbf{n}_I-1)}, \mathbf{e}) = \tilde{g}_j(\mathbf{e})$ for $\tilde{g}_j : \mathcal{E} \rightarrow \mathcal{X}_j$) and $\tilde{g}_j(\mathbf{E})$ is a C^{n_j} -stochastic process.

In particular, under Assumption 1-I the dynamic structural equations of \mathcal{R} are equivalent to an RDE. For an SDCM that satisfies Assumption 1-I with $I \subsetneq \mathcal{I}$, we can eliminate the processes $\mathbf{X}_{\setminus I}$ by substitution, giving an RDE for the endogenous processes I of the form

$$\mathbf{X}_I^{(\mathbf{n}_I)} = \mathbf{g}_I(\bar{\mathbf{X}}_I^{(\mathbf{n}_I-1)}, \mathbf{g}_{\setminus I}(\bar{\mathbf{X}}_I^{(\mathbf{n}_I-1)}, \mathbf{E}), \mathbf{E}). \quad (3)$$

Every solution of the original SDCM satisfies this RDE, and every solution of this RDE induces a solution of the SDCM, if it is sufficiently smooth.

Example 3.22. Consider the price, supply and demand model of Example 3.15. This model satisfies Assumption 1-I for $I = \{P\}$. Substituting the zeroth-order dynamic structural equations into the first-order equation of X_P yields the RDE

$$X'_P = \lambda(\beta_D - \beta_S)X_P + \lambda(E_D - E_S). \quad (4)$$

²³ The required smoothness of the solutions implies that we need to make assumptions about the smoothness of the exogenous processes and the dynamical causal mechanisms in the model. The assumption we made here is still rather crude in the sense that it suffices, but it is not at all necessary; if desired, one can arrive at weaker conditions by carefully tracing through the graph how the required smoothness of the solution can be guaranteed by demanding certain smoothness of each exogenous process and each dynamical causal mechanism individually.

Since SDCMs that satisfy Assumption 1- I have the property that they determine an RDE on the subset I , we can apply the existence and uniqueness results of RDEs on this subsystem, which leads to the following result.

Theorem 3.23. *Let \mathcal{R} be an SDCM that satisfies Assumption 1- I for a subset $I \subseteq \mathcal{I}$. If for almost all $\omega \in \Omega$ there exists a continuous function $K_\omega : T \rightarrow \mathbb{R}$ such that for all $i \in I$, $t \in T$ and for all $\bar{\mathbf{x}}_I^{(\mathbf{n}_I-1)}, \bar{\mathbf{y}}_I^{(\mathbf{n}_I-1)} \in \mathcal{X}_I^{\mathbf{n}_I}$ the condition*

$$\|g_i(\bar{\mathbf{x}}_I^{(\mathbf{n}_I-1)}, g_{\setminus I}(\bar{\mathbf{x}}_I^{(\mathbf{n}_I-1)}, \mathbf{E}), \mathbf{E})(t, \omega) - g_i(\bar{\mathbf{y}}_I^{(\mathbf{n}_I-1)}, g_{\setminus I}(\bar{\mathbf{y}}_I^{(\mathbf{n}_I-1)}, \mathbf{E}), \mathbf{E})(t, \omega)\| \leq K_\omega(t) \|x_i^{(0)} - y_i^{(0)}\|$$

is satisfied, where $\|\cdot\|$ denotes the Euclidean norm, then for any partial initial condition $(t_0, \bar{\mathbf{X}}_{I,[0]}^{(\mathbf{n}_I-1)}) \in T \times \mathcal{X}_I^{\mathbf{n}_I}$ there exists an a.s. unique solution \mathbf{X} of \mathcal{R} with initial condition

$$(\bar{\mathbf{X}}_I^{(\mathbf{n}_I-1)}(t_0), \mathbf{X}_{\setminus I}(t_0)) = (\bar{\mathbf{X}}_{I,[0]}^{(\mathbf{n}_I-1)}, g_{\setminus I}(\bar{\mathbf{X}}_{I,[0]}^{(\mathbf{n}_I-1)}, \mathbf{E}(t_0))).$$

This theorem guarantees the existence and uniqueness of solutions for a large class of SDCMs. In particular, it provides a sufficient condition for an initial condition to be consistent (see Definition 3.19).

In general, Assumption 1- I for an SDCM is not preserved under a stochastic perfect intervention. Consider for example the SDCM $\tilde{\mathcal{R}}$ in Example 3.10 which satisfies Assumption 1- \mathcal{I} . Performing the intervention $\text{do}(1, 0)$ on this model yields a model that does not satisfy Assumption 1- I for any $I \subseteq \mathcal{I}$. Under the following stronger assumption the SDCM will satisfy Assumption 1- I for some $I \subseteq \mathcal{I}$ after every (sufficiently smooth) stochastic perfect intervention.

Assumption 2- I . *For the SDCM \mathcal{R} and subset $I \subseteq \mathcal{I}$ with order tuple $\mathbf{n}_I \geq 1$, the exogenous process \mathbf{E} is continuous and there exist continuous functions $g_i : \mathcal{X}_i^{\mathbf{n}_i} \times \mathcal{X}_{\setminus i} \times \mathcal{E} \rightarrow \mathcal{X}_i$ for all $i \in I$ and $g_j : \mathcal{X}_I \times \mathcal{E} \rightarrow \mathcal{X}_j$ for all $j \in \mathcal{I} \setminus I$ such that for all $i \in I$, all $j \in \mathcal{I} \setminus I$, all $\mathbf{e} \in \mathcal{E}$ and all $\bar{\mathbf{x}}^{(\mathbf{n})} \in \mathcal{X}^{\mathbf{n}+1}$,*

$$x_i^{(\mathbf{n}_i)} = g_i(\bar{\mathbf{x}}_i^{(\mathbf{n}_i-1)}, \mathbf{x}_{\setminus i}, \mathbf{e}) \iff x_i = f_i(\bar{\mathbf{x}}^{(\mathbf{n})}, \mathbf{e})$$

and

$$x_j = g_j(\mathbf{x}_I, \mathbf{e}) \iff x_j = f_j(\bar{\mathbf{x}}^{(\mathbf{n})}, \mathbf{e}).$$

Furthermore, for each $j \in \mathcal{I} \setminus I$, either $n_j = 0$, or g_j only depends on \mathbf{e} (that is, $g_j(\mathbf{x}_I, \mathbf{e}) = \tilde{g}_j(\mathbf{e})$ for $\tilde{g}_j : \mathcal{E} \rightarrow \mathcal{X}_j$) and $\tilde{g}_j(\mathbf{E})$ is a C^{n_j} -stochastic process.

In particular, Assumption 2- I implies Assumption 1- I .

Proposition 3.24 (Assumption 2- I under stochastic perfect intervention). *Let \mathcal{R} be an SDCM that satisfies Assumption 2- I for a subset $I \subseteq \mathcal{I}$. Then, for a stochastic perfect intervention $\text{do}(J, \mathbf{K}_J)$ for $J \subseteq \mathcal{I}$ with \mathbf{K}_J a $C^{\mathbf{n}_J}$ -stochastic process, the intervened SDCM $\mathcal{R}_{\text{do}(J, \mathbf{K}_J)}$ satisfies Assumption 2- $(I \setminus J)$.*

This proposition shows the usefulness of Assumption 2- I , in that it gives a guarantee that after any sufficiently smooth stochastic perfect intervention, Assumption 1- \tilde{I} is satisfied for some $\tilde{I} \subseteq \mathcal{I}$, and hence Theorem 3.23 can be applied.

Linear SDCMs

Observe that a linear SDCM that satisfies Assumption 1- I is of the following form.

Proposition 3.25. *Let \mathcal{R} be a linear SDCM, $I \subseteq \mathcal{I}$ a subset, and let $J := \mathcal{I} \setminus I$. Then \mathcal{R} satisfies Assumption 1- I iff the dynamic causal mechanism \mathbf{f} of \mathcal{R} is of the form*

$$\begin{cases} \mathbf{f}_I(\bar{\mathbf{x}}^{(\mathbf{n})}, \mathbf{e}) := B_{II^{(\mathbf{n}_I)}} \mathbf{x}_I^{(\mathbf{n}_I)} + B_{I\tilde{I}^{(\mathbf{n}_I-1)}} \bar{\mathbf{x}}_I^{(\mathbf{n}_I-1)} + B_{IJ} \mathbf{x}_J + \Gamma_{IJ} \mathbf{e} \\ \mathbf{f}_J(\bar{\mathbf{x}}^{(\mathbf{n})}, \mathbf{e}) := B_{J\tilde{I}^{(\mathbf{n}_I-1)}} \bar{\mathbf{x}}_I^{(\mathbf{n}_I-1)} + B_{JJ} \mathbf{x}_J + \mathbf{x}_J + \Gamma_{J\mathcal{I}} \mathbf{e}, \end{cases}$$

where $B_{II^{(n_I)}}$ and B_{JJ} are invertible matrices, \mathbf{E} is continuous, and furthermore, for each $j \in J$, either $n_j = 0$, or $(B_{JJ}^{-1})_{jJ} B_{J\bar{I}^{(n_I-1)}} = \mathbf{0}$ and $(B_{JJ}^{-1})_{jJ} \Gamma_{J\mathcal{J}} \mathbf{E}$ is a C^{n_j} -stochastic process.

In particular, for linear SDCMs, Theorem 3.23 gives the following useful corollary.

Corollary 3.26. *Let \mathcal{R} be a linear SDCM that satisfies Assumption 1-I for a subset $I \subseteq \mathcal{I}$. Then for any $(t_0, \bar{\mathbf{X}}_{I,[0]}^{(n_I-1)}) \in T \times \mathcal{X}_I^{n_I}$ there exists an a.s. unique solution \mathbf{X} of \mathcal{R} with $\bar{\mathbf{X}}_I^{(n_I-1)}(t_0) = \bar{\mathbf{X}}_{I,[0]}^{(n_I-1)}$.*

Examples of linear SDCMs that satisfy Assumption 1-I for some subset I are the SDCMs \mathcal{R} of Example 3.5 and $\mathcal{R}_{\text{do}(\{1,d\},(0,L))}$ of Example 3.8, which satisfy Assumption 1- \mathcal{I} and 1- $(\mathcal{I} \setminus \{1,d\})$, respectively. Hence, they both have an a.s. unique solution for each respective partial initial condition.

In particular, for linear SDCMs that satisfy Assumption 2-I we have the following corollary.

Corollary 3.27. *Let \mathcal{R} be a linear SDCM that satisfies Assumption 2-I for a subset $I \subseteq \mathcal{I}$, and let $\text{do}(J, \mathbf{K}_J)$ be a stochastic perfect intervention for $J \subseteq \mathcal{I}$ with \mathbf{K}_J a C^{n_J} -stochastic process. Then for any partial initial condition $(t_0, \bar{\mathbf{X}}_{I \setminus J, [0]}^{(n_I \setminus J - 1)}) \in T \times \mathcal{X}_{I \setminus J}^{n_I \setminus J}$ there exists an a.s. unique solution \mathbf{X} of $\mathcal{R}_{\text{do}(J, \mathbf{K}_J)}$ with that partial initial condition.*

Examples of linear SDCMs that satisfy Assumption 2-I for some subset I are the damped coupled harmonic oscillator of Example 3.5 and the price, supply and demand model of Example 3.15. Hence, the existence of solutions is guaranteed for both models after any (sufficiently smooth) stochastic perfect intervention, and the solutions are a.s. uniquely determined by the respective partial initial conditions.

Nonlinear SDCMs

An example of an SDCM that is not linear but satisfies Assumption 2-I is the bathtub model discussed in [38]. The existence and uniqueness conditions apply to this particular model.

Example 3.28 (Bathtub model). *Water enters a bathtub from the faucet at a certain rate X_{Q_i} and exits the bathtub via the drain at a rate X_{Q_o} . The drain has a diameter of X_K , the depth of the water is X_D and the pressure at the base of the drain is X_P . Iwasaki and Simon [38] propose to model this as a dynamical system with (random) differential equations given by*

$$\begin{cases} X_K = k_0 \\ X_{Q_i} = q_0 \\ X'_P = \alpha_2(\alpha_4 X_D - X_P) \\ X'_{Q_o} = \alpha_3(\alpha_1 X_K X_P - X_{Q_o}) \\ X'_D = \alpha_0(X_{Q_i} - X_{Q_o}), \end{cases} \quad (5)$$

where $k_0, q_0 \in \mathbb{R}_{>0}$ and $\boldsymbol{\alpha} = (\alpha_0, \alpha_1, \dots, \alpha_4) \in \mathbb{R}_{>0}^5$ are some constants. We consider the dynamic causal mechanism

$$\begin{cases} f_K(\bar{\mathbf{x}}^{(n)}, e) := e_K \\ f_{Q_i}(\bar{\mathbf{x}}^{(n)}, e) := e_{Q_i} \\ f_P(\bar{\mathbf{x}}^{(n)}, e) := \alpha_4 x_D - \alpha_2^{-1} x_P \\ f_{Q_o}(\bar{\mathbf{x}}^{(n)}, e) := \alpha_1 x_K x_P - \alpha_3^{-1} x_{Q_o} \\ f_D(\bar{\mathbf{x}}^{(n)}, e) := x_D + \alpha_0(x_{Q_i} - x_{Q_o}) - x_{D'} \end{cases}$$




Figure 6. Graphs of the bathtub model: original model \mathcal{R} of Example 3.28 (top left), the equilibrated model $\mathcal{M}_{\mathcal{R}}$ (top right), the intervened model $\mathcal{R}_{\text{do}(D, K_D)}$ (bottom left), and the intervened and equilibrated model $\mathcal{M}_{\mathcal{R}_{\text{do}(D, K_D)}}$ (bottom right) of Example 4.19.

with order tuple $\mathbf{n} := (n_K, n_{Q_i}, n_P, n_{Q_o}, n_D) = (0, 0, 1, 1, 1)$ and the exogenous processes are given by $E_K(t, \omega) := k_0$, $E_{Q_i}(t, \omega) := q_0$. The dynamic structural equations of this SDCM, denoted by \mathcal{R} , read

$$\begin{cases} X_K = E_K \\ X_{Q_i} = E_{Q_i} \\ X_P = \alpha_4 X_D - \alpha_2^{-1} X'_P \\ X_{Q_o} = \alpha_1 X_K X_P - \alpha_3^{-1} X'_{Q_o} \\ X_D = X_D + \alpha_0 (X_{Q_i} - X_{Q_o}) - X'_D, \end{cases}$$

and have the same solutions as the system of equations (5) (see also Footnote 9). The corresponding SDCM graph is depicted in Figure 6 (top left). This SDCM of the bathtub model satisfies Assumption 2- $(\{P, Q_o, D\})$, and hence, after any sufficiently smooth stochastic perfect intervention $\text{do}(J, \mathbf{K}_J)$ with $J \subseteq \{K, Q_i, P, Q_o, D\}$, the intervened bathtub model $\mathcal{R}_{\text{do}(J, \mathbf{K}_J)}$ satisfies Assumption 2- $(\{P, Q_o, D\} \setminus J)$. Since the induced RDE of the intervened model $\mathcal{R}_{\text{do}(J, \mathbf{K}_J)}$ on the endogenous processes $\{P, Q_o, D\} \setminus J$ is linear in these endogenous processes, it follows from Theorem 3.23 that (for sufficiently smooth exogenous process \mathbf{K}_J) $\mathcal{R}_{\text{do}(J, \mathbf{K}_J)}$ has an a.s. unique solution for any partial initial condition $(t_0, \bar{\mathbf{X}}_{\{P, Q_o, D\} \setminus J, [0]}^{(\mathbf{n}_{\{P, Q_o, D\} \setminus J})})$.

4 Equilibration of SDCMs

In this section, we will take $T = [t_0, \infty)$ and study the equilibrium states of SDCMs and, in particular, of steady SDCMs, which are SDCMs for which the dynamic structural equations and exogenous processes become explicitly time-independent asymptotically as $t \rightarrow \infty$. We introduce an equilibration operation on a

steady SDCM, which equilibrates the model to an SCM such that all the equilibrium states of the SDCM are described by the solutions of the SCM. Intuitively, this equilibration operation separately equilibrates each dynamic causal mechanism, which corresponds mathematically to transforming each dynamic structural equation into a structural equation of the SCM. We show that this equilibration operation commutes with perfect stochastic interventions, without requiring the strong global stability assumption of [52], which assumes that all the solutions equilibrate to the same static equilibrium state. This allows to study the causal semantics of the equilibrium states of steady SDCMs within the framework of SCMs.

We start in Section 4.1 with the definition of equilibrating solutions and their corresponding equilibrium states. In Section 4.2, we define the class of steady SDCMs which have several convenient convergence properties. In Section 4.3, we show how we can equilibrate a steady SDCM to an SCM. In Section 4.4, we show how the equilibration acts on the graph of an SDCM. In Section 4.5, we show that the equilibration operation commutes with intervention. We discuss in Section 4.6 the inverse problem of finding steady SDCMs for which all the solutions equilibrate to solutions of the SCM independently of the initial condition. We provide sufficient conditions under which one can construct a first-order steady SDCM such that its equilibration coincides with a given linear SCM. This establishes a class of linear SCMs that model the causal equilibrium semantics of certain linear dynamical systems. In Section 4.7, we discuss some subtleties in the causal interpretation of the graph of the equilibrated SDCM.

4.1 Equilibrating solutions and equilibrium states

In this subsection, we define the equilibrating solutions of an SDCM as those solutions for which all the higher-order derivatives that are considered in the model converge to zero a.s.. For a stochastic process \mathbf{X} we say that it *converges almost surely* to a random variable \mathbf{X}^* , if the limit $\lim_{t \rightarrow \infty} \mathbf{X}_t$ exists almost surely²⁴ and is a.s. equal to \mathbf{X}^* . In this case, we call \mathbf{X} *almost surely convergent*.

Definition 4.1 (Equilibrating solution, equilibrium state). *Let \mathbf{X} be a solution of an SDCM \mathcal{R} . We call \mathbf{X} an equilibrating solution, if $\bar{\mathbf{X}}^{(n)}$ is a.s. convergent. In particular, an equilibrating solution \mathbf{X} converges almost surely to a random variable \mathbf{X}^* , and we say that \mathbf{X} equilibrates to \mathbf{X}^* and call \mathbf{X}^* an equilibrium state of \mathcal{R} .*

An example of an SDCM with equilibrium states is the price, supply and demand model of Example 3.15, where the equilibrium states correspond to “market equilibrium”, as illustrated in the following example.

Example 4.2 (Market equilibrium). *Consider the price, supply and demand model of Example 3.15 with E_S and E_D constant exogenous processes. Market equilibrium for this model is reached if*

$$X_D^* - X_S^* = 0,$$

that is, if the demanded and supplied quantities become equal asymptotically. The solutions that satisfy this condition are equilibrating solutions for which

$$X_P^* = 0, \quad X_P^* = \frac{E_D - E_S}{\beta_S - \beta_D}, \quad X_S^* = X_D^* = \frac{\beta_S E_D - \beta_D E_S}{\beta_S - \beta_D}.$$

In fact, for every solution \mathbf{X} that equilibrates, the higher-order derivatives of \mathbf{X} must converge to zero almost surely.

Proposition 4.3. *Let \mathbf{X} be a solution of an SDCM \mathcal{R} . If \mathbf{X} equilibrates, then $\lim_{t \rightarrow \infty} \bar{X}_i^{(n_i)} = (X_i^*, 0, \dots, 0)$ a.s. for all $i \in \mathcal{I}$, where X_i^* is the i^{th} component of the corresponding equilibrium state \mathbf{X}^* .*

²⁴ In that case, it defines a random variable, because $\lim_{t \rightarrow \infty} \mathbf{X}_t = \lim_{t \in \mathbb{N}} \mathbf{X}_t$ a.s., and the latter is a random variable.

In particular, for linear SDCMs we can show that all the solutions of the SDCM equilibrate under certain conditions.

Proposition 4.4. *Let \mathcal{R} be a linear SDCM that satisfies Assumption 1-I for a subset $I \subseteq \mathcal{I}$ with an order tuple $\mathbf{n}_I = 1$ and an exogenous process \mathbf{E} that is constant in time.²⁵ By Proposition 3.25, the dynamical causal mechanism \mathbf{f} is of the form*

$$\begin{cases} \mathbf{f}_I(\bar{\mathbf{x}}^{(\mathbf{n})}, \mathbf{e}) := B_{II'}\mathbf{x}'_I + B_{II}\mathbf{x}_I + B_{IJ}\mathbf{x}_J + \Gamma_{IJ}\mathbf{e} \\ \mathbf{f}_J(\bar{\mathbf{x}}^{(\mathbf{n})}, \mathbf{e}) := B_{JI}\mathbf{x}_I + B_{JJ}\mathbf{x}_J + \mathbf{x}_J + \Gamma_{JJ}\mathbf{e}, \end{cases}$$

where $J := \mathcal{I} \setminus I$ and $B_{II'}$ and B_{JJ} are invertible matrices. If the matrix $B_{II'}^{-1}(B_{IJ}B_{JJ}^{-1}B_{JI} - B_{II} + \mathbb{I}_I)$, where \mathbb{I}_I denotes the identity matrix, is Hurwitz (that is, every eigenvalue has a strictly negative real part), then every solution \mathbf{X} of \mathcal{R} equilibrates to the same equilibrium state, irrespective of the initial condition.

This proposition allows us to derive a condition for which the price, supply and demand model always reaches market equilibrium.

Example 4.5 (Market equilibrium, continued). *Applying Proposition 4.4 to the price, supply and demand model of Example 3.15 shows that $B_{II'}^{-1}(B_{IJ}B_{JJ}^{-1}B_{JI} - B_{II} + \mathbb{I}_I) = \lambda(\beta_D - \beta_S)$. This matrix is Hurwitz if and only if $\lambda(\beta_D - \beta_S) < 0$. Thus, since $\lambda > 0$, the price X_P , supply X_S and demand X_D equilibrate for constant exogenous processes E_D and E_S if $\beta_S > \beta_D$.*

4.2 Steady SDCMs

In this subsection, we define the class of steady SDCMs which have the convenient property that their dynamics become explicitly time-independent asymptotically for $t \rightarrow \infty$.

Definition 4.6 (Steady SDCM). *We call an SDCM \mathcal{R} steady, if it has a dynamic causal mechanism \mathbf{f} that is continuous and an exogenous process \mathbf{E} that is a.s. convergent.*

The continuity of the dynamic causal mechanism and the convergence assumption on the exogenous process assure us that the equilibrium states satisfy asymptotic dynamic structural equations.

Lemma 4.7. *Let \mathcal{R} be a steady SDCM and let \mathbf{E}^* be the random variable to which the exogenous process \mathbf{E} converges a.s.. If \mathbf{X} is an equilibrating solution of a steady SDCM \mathcal{R} , then the random variable $\bar{\mathbf{X}}^{(\mathbf{n})*}$ to which the complete \mathbf{n}^{th} -order derivative $\bar{\mathbf{X}}^{(\mathbf{n})}$ converges satisfies*

$$\mathbf{X}^* = \mathbf{f}(\bar{\mathbf{X}}^{(\mathbf{n})*}, \mathbf{E}^*) \quad \text{a.s..}$$

In general, not all solutions of a steady SDCM have to be equilibrating solutions, as one sees for example in Example 3.18.

The class of steady SDCMs is not closed under stochastic perfect interventions, since performing a stochastic perfect intervention that is not a.s. convergent yields an SDCM that is not steady. However, the class of steady SDCMs is closed under the following class of interventions.

²⁵ In general, we can let \mathbf{E} be an exogenous process that depends on time as long as both \mathbf{E}_t and $\exp(At) \int_0^t \exp(-As) C \mathbf{E}_s ds$ converge almost surely for $t \rightarrow \infty$, where $A := B_{II'}^{-1}(B_{IJ}B_{JJ}^{-1}B_{JI} - B_{II} + \mathbb{I}_I)$ and $C := B_{II'}^{-1}(B_{IJ}B_{JJ}^{-1}\Gamma_{JJ} - \Gamma_{IJ})$. In that case, the order tuple may matter, and it must be checked whether the solutions are sufficiently smooth.

Definition 4.8 (Steady stochastic perfect intervention). *We call a stochastic perfect intervention $\text{do}(I, \mathbf{K}_I)$ a steady stochastic perfect intervention if the process \mathbf{K}_I converges a.s. to a random variable \mathbf{K}_I^* . We call it a steady perfect intervention if in addition $\mathbf{K}_I^* \in \mathcal{X}_I$ (that is, it does not depend on ω).*

4.3 Equilibration of a steady SDCM

In this subsection, we show how we can equilibrate a steady SDCM to an SCM, such that the equilibrium states of the SDCM are described by the SCM. In the previous subsections, we saw that for an equilibrating solution of a steady SDCM, all the higher-order derivatives converge to zero, and the corresponding equilibrium state satisfies the asymptotic dynamic structural equations. Hence, we can construct an SCM from a steady SDCM such that every equilibrium state of the steady SDCM is a solution of this SCM.

Definition 4.9 (Equilibration of an SDCM). *Let $\mathcal{R} = \langle \mathcal{I}, \mathcal{J}, \mathcal{X}, \mathcal{E}, \mathbf{n}, \mathbf{f}, \mathbf{E} \rangle$ be a steady SDCM and let \mathbf{E}^* be a random variable such that \mathbf{E} converges a.s. to it. We call the SCM $\mathcal{M}_{\mathcal{R}} := \langle \mathcal{I}, \mathcal{J}, \mathcal{X}, \mathcal{E}, \mathbf{f}^*, \mathbf{E}^* \rangle$ with the equilibrated dynamic causal mechanism $\mathbf{f}^* : \mathcal{X} \times \mathcal{E} \rightarrow \mathcal{X}$ given by*

$$\mathbf{f}^*(\mathbf{x}, \mathbf{e}) := \mathbf{f}(\bar{\mathbf{t}}(\mathbf{x}), \mathbf{e}),$$

an equilibration of \mathcal{R} , where the mapping $\bar{\mathbf{t}} : \mathcal{X} \rightarrow \mathcal{X}^{n+1}$ defined by

$$\bar{t}_i^{(k_i)}(\mathbf{x}) = \begin{cases} x_i & \text{if } k_i = 0 \\ 0 & \text{otherwise,} \end{cases}$$

is the embedding that sets all the higher-order derivatives of the endogenous processes to 0.

In other words, the equilibration of an SDCM sets all the higher-order derivative entries in its dynamic causal mechanism to zero and replaces its exogenous process by its limiting random variable. In particular, linearity is preserved under equilibration, that is, a steady linear SDCM equilibrates to a linear SCM.

The equilibration of an SDCM is well defined due to the following result, which shows that the independence property for the family of exogenous processes $(\mathbf{E}_j)_{j \in \mathcal{J}}$ is preserved in the limit when time tends to infinity.

Proposition 4.10. *Let $(\mathbf{E}_j)_{j \in \mathcal{J}}$ be a family of stochastic processes, where \mathcal{J} is some finite index set, such that \mathbf{E}_j converges almost surely to the random variable \mathbf{E}_j^* , for every $j \in \mathcal{J}$. Then, if $(\mathbf{E}_j)_{j \in \mathcal{J}}$ is independent, so is the family of random variables $(\mathbf{E}_j^*)_{j \in \mathcal{J}}$.*

This equilibration of an SDCM to an SCM leads to the main insight that SCMs are capable of modeling all the equilibrium states of steady SDCMs.

Theorem 4.11. *If \mathbf{X} is an equilibrating solution of a steady SDCM \mathcal{R} , then its limit \mathbf{X}^* is a solution of the corresponding equilibration $\mathcal{M}_{\mathcal{R}}$.*

Intuitively, the equilibration of a steady SDCM to an SCM can be seen as the approximation of the dynamic structural equations by the structural equations of the SCM, which becomes exact at equilibrium. This is illustrated in the following example.

Example 4.12 (Equilibrated damped coupled harmonic oscillator). *Consider the intervened damped coupled harmonic oscillator of Example 3.8 for which the dynamic structural equations are specified by*

$$\begin{cases} X_1 = 0 \\ X_i = \frac{\kappa_i}{\kappa_i + \kappa_{i-1}}(X_{i+1} - L_i) + \frac{\kappa_{i-1}}{\kappa_i + \kappa_{i-1}}(X_{i-1} + L_{i-1}) \\ \quad - \frac{b_i}{\kappa_i + \kappa_{i-1}}X'_i - \frac{m_i}{\kappa_i + \kappa_{i-1}}X''_i \quad (i = 2, \dots, d-1) \\ X_d = L, \end{cases}$$

and where the exogenous processes \mathbf{L} are random variables. In the limit, as time tends to infinity, the equilibrating solutions of the SDCM converge to the equilibrium states of the equilibrated SDCM, which can be obtained by setting the higher-order derivatives to zero. This yields the equations

$$\begin{cases} X_1^* = 0 \\ X_i^* = \frac{\kappa_i(X_{i+1}^* - L_i) + \kappa_{i-1}(X_{i-1}^* + L_{i-1})}{\kappa_i + \kappa_{i-1}} \quad (i = 2, \dots, d-1) \\ X_d^* = L, \end{cases}$$

which describe the equilibrium states for the positions of the masses. Not all solutions necessarily equilibrate to an equilibrium, which happens for example in the case when there is no friction, that is, $b_i = 0$ for all $i \in \{2, \dots, d-1\}$. In this case, if any mass m_i starts at an off-equilibrium position (that is, if $X'_i(t_0) \neq 0$ or $X_i(t_0) \neq X_i^*$ for some $i \in \{2, \dots, d-1\}$), the solution will not equilibrate, but will keep on oscillating forever.

In case there is friction and the exogenous processes \mathbf{L} are fixed to constant values, the equilibrated damped coupled harmonic oscillator exactly coincides with the deterministic SCM derived in [52]. In Section 4.5 we will show that the equilibration operation, as defined in Definition 4.9, also preserves the causal semantics. The next example illustrates that our equilibration operation can also be applied to models that cannot be treated with the theory of [52].

Example 4.13 (Equilibrated price, supply and demand model). *Setting the higher-order derivatives of the price, supply and demand model \mathcal{R} of Example 3.15 to zero yields the structural equations:*

$$\begin{cases} X_P^* = X_P^* + \lambda(X_D^* - X_S^*) \\ X_S^* = \beta_S X_P^* + E_S^* \\ X_D^* = \beta_D X_P^* + E_D^*. \end{cases}$$

The equations describe the market equilibrium states. In Figure 7, we simulate the solutions of the SDCM \mathcal{R} for random constant exogenous influences E_S and E_D and random consistent initial conditions. The dispersion of X_P , X_S and X_D at large t illustrates that the equilibrium state is not unique and depends on the initial condition. Hence, this example cannot be treated with the theory of [52].

Richardson and Robins [60] argue that the price, supply and demand model cannot be modeled at equilibrium as an SCM without self-cycles. We conclude that it can be modeled by an SCM that contains self-cycles, with the corresponding graph depicted in Figure 5 (right).

A consequence of Theorem 4.11 is that if the SCM $\mathcal{M}_{\mathcal{R}}$ has no solutions, then the SDCM \mathcal{R} has no equilibrating solutions. However, the converse does not hold in general, as the following example illustrates.

Example 4.14. *Let $\mathcal{R} = \langle \{1, 2\}, \{3\}, \mathcal{X}, \mathcal{E}, \mathbf{n}, \mathbf{f}, E \rangle$ be the steady SDCM with $\mathcal{X} = \mathbb{R}^2$, $\mathcal{E} = \mathbb{R}$, $\mathbf{n} = (0, 1)$, the dynamic causal mechanism \mathbf{f} given by $f_1(\bar{\mathbf{x}}^{(\mathbf{n})}, e) = x_2$, and $f_2(\bar{\mathbf{x}}^{(\mathbf{n})}, e) = e$, and the exogenous process E given by $E_t = \sin(t^3)/t$. The dynamic structural equations associated to \mathcal{R} are given by*

$$X_1 = X_2', \quad X_2 = E.$$




Figure 7. Simulation of solutions of the SDCM of the price, supply and demand model of Example 4.13 under different steady perfect interventions.

This model can be equilibrated to the model $\mathcal{M}_{\mathcal{R}}$ with structural equations

$$X_1^* = 0, \quad X_2^* = E^*,$$

and exogenous variable $E^* = 0$. Although the SCM $\mathcal{M}_{\mathcal{R}}$ clearly has a solution, the SDCM \mathcal{R} has no equilibrium states, since $X_1 = X_2' = E'$ is not a.s. convergent.

The following result shows that if the exogenous process is constant in time, this cannot happen.

Proposition 4.15. *Let \mathcal{R} be a steady SDCM such that the exogenous process \mathbf{E} is a random variable (i.e., \mathbf{E} is constant in time). If the SDCM \mathcal{R} has no equilibrating solution, then its equilibration $\mathcal{M}_{\mathcal{R}}$ has no solutions.*

4.4 Graph of the equilibrated SDCM

In this subsection, we show how the equilibration operation acts on the graph of the SDCM.

Proposition 4.16 (Graph of the equilibrated SDCM is a subgraph²⁶ of the original mixed graph). *Let \mathcal{R} be a steady SDCM. The graph $\mathcal{G}(\mathcal{M}_{\mathcal{R}})$ of the equilibrated SDCM $\mathcal{M}_{\mathcal{R}}$ is the mixed graph obtained from the graph $\mathcal{G}(\mathcal{R})$ of \mathcal{R} by removing the partition into clusters and removing the nodes $i^{(k_i)}$ for $i \in I$ and $k_i > 0$ together with their adjacent edges.*

The following example illustrates this for the equilibrated price, supply and demand model.

Example 4.17 (Price, supply and demand, continued). *Consider the price, supply and demand model \mathcal{R} of Example 3.15 for a very large λ , that is, for which the price adjusts very quickly to changes in supply and demand. This system can be approximated by the equilibrated price, supply and demand model $\mathcal{M}_{\mathcal{R}}$. The graph of this equilibrated model $\mathcal{M}_{\mathcal{R}}$ is a subgraph of the graph of the original model \mathcal{R} , as can be seen in Figure 5.*

4.5 Equilibration commutes with intervention

Theorem 4.11, our first main result, states that the equilibrium states of a steady SDCM are solutions of the SCM to which the SDCM equilibrates. In the previous subsection, we showed, moreover, that the functional relationships between the endogenous processes that are encoded in the dynamic structural equations are

²⁶ Let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a mixed graph. A mixed graph $\tilde{\mathcal{G}} = (\tilde{\mathcal{V}}, \tilde{\mathcal{E}})$ is a *subgraph* of \mathcal{G} if $\tilde{\mathcal{V}} \subseteq \mathcal{V}$ and $\tilde{\mathcal{E}} \subseteq \mathcal{E}$.

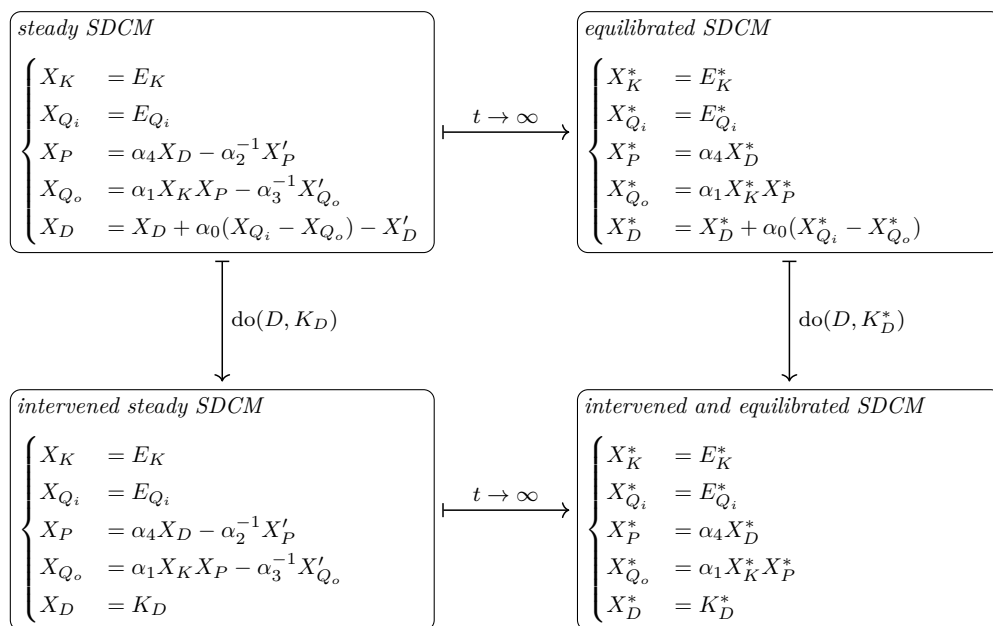
preserved under equilibration. This leads to our second main result: the equilibration operation preserves the causal semantics of the equilibrium states, as is illustrated in Figure 1 in Section 1.

Theorem 4.18. *Let \mathcal{R} be a steady SDCM and let $\text{do}(I, \mathbf{K}_I)$ be a steady stochastic perfect intervention for some subset $I \subseteq \mathcal{I}$ and stochastic process \mathbf{K}_I that converges a.s. to a random variable \mathbf{K}_I^* . Then the steady stochastic perfect intervention commutes with equilibration, that is*

$$(\mathcal{M}_{\mathcal{R}})_{\text{do}(I, \mathbf{K}_I^*)} = \mathcal{M}_{(\mathcal{R}_{\text{do}(I, \mathbf{K}_I)})}.$$

This result allows us to perform causal reasoning on the equilibrium states of the SDCM by considering only the equilibrated model, as is illustrated in the following example.

Example 4.19 (Bathtub model, continued). *In Example 3.28 we defined the SDCM for the bathtub model. The equilibrium states of this model can be described by the structural equations of the equilibrated model, as depicted in the top row of the following diagram.*



After equilibration, one can perform causal reasoning on the level of the equilibrated SDCM, without needing to resort to the original SDCM description. Indeed, we see in the above diagram that it doesn't matter whether we first perform the steady stochastic perfect intervention $\text{do}(D, K_D)$, and then let the system equilibrate, or the other way around. The graphs of the SDCM, the equilibrated SDCM and their corresponding intervened models are depicted in Figure 6. Choosing different a.s. convergent processes for K_D yields different solution processes X_P and X_{Q_o} of the intervened SDCM, but the solution processes for X_{Q_i} and X_K stay unchanged. Similarly, the perfect intervention $\text{do}(D, K_D^*)$ on the equilibrated SDCM yields different solutions X_P^* and $X_{Q_o}^*$ of the intervened SCM depending on the value of K_D^* , but does not change the solutions $X_{Q_i}^*$ and X_K^* . This behavior is also reflected in the graphs depicted in the bottom row of Figure 6.

Intuitively, one would indeed expect the chosen intervention value for the depth to have an effect on pressure and outflow (but not on inflow or drain size) at equilibrium. For example, one could (approximately) implement such a perfect intervention by adding a water level control device that constantly monitors the level and that can pump water in and out of the bathtub via a hose, regulating the depth at K_D at all times by using an optimal control feedback loop, independently of the exogenous processes E_K and E_{Q_i} . Indeed, the depth directly determines the pressure X_P exerted by the water in the bathtub at the drain, and the outflow rate X_{Q_o} is a direct consequence of that. Once the other processes in the system have equilibrated, the processes X_P and X_{Q_o} will also equilibrate to random variables that depend on K_D^* . The inflow $X_{Q_i}^*$ of

water through the faucet no longer needs to be equal to the outflow $X_{Q_o}^*$ through the drain at equilibrium because water is also constantly added or removed via the hose by the water level control device in order to maintain the (eventually) constant depth K_D^* .²⁷

This sheds some new light on the violation of the equilibration-manipulation commutability property (the “EMC-property”) of Dash [15], who shows the—at first sight contradictory—result that equilibration does not always commute with intervention. The paradox is resolved by noting that Dash [15] defines a different notion of “equilibration”, inspired by Iwasaki and Simon [38], for which commutativity with perfect intervention indeed does not always hold. One can readily verify that the “equilibration” operation of Dash [15] does not preserve the functional relationships between the endogenous processes that are encoded in the equations under the equilibration. Recently, Blom and Mooij [5] showed that the “equilibration” operation of Dash [15] maps an SDCM \mathcal{R} to a Markov ordering graph that encodes the conditional independencies in solutions of $\mathcal{M}_{\mathcal{R}}$ instead of the functional relationships. In contrast, our equilibration operation, defined in Definition 4.9, preserves the functional relationships between the endogenous processes, since each dynamic structural equation equilibrates to a structural equation associated to the same endogenous process/variable. This is also reflected in Proposition 4.16 where we showed that the graph of the equilibrated SDCM is a subgraph of the mixed graph of the SDCM.

Theorem 4.11 and 4.18 together imply that our equilibration operation preserves the equilibrium states of a steady SDCM while also preserving the causal semantics. In particular, we do not require that all solutions of the steady SDCM have to equilibrate. As a consequence, the equilibrium states of the model may depend on the (consistent) initial conditions. This is in contrast to the work of Mooij et al. [52], who assume that the equilibrium state of the dynamical system is unique and independent of the initial condition. This is a strong assumption that limits the applicability of the theory, since this does not allow for any stochasticity at equilibrium. Indeed, many random dynamical systems have multiple equilibrium states that depend on the chosen initial condition, as is illustrated in the following example.

Example 4.20 (Bathtub model, continued). *Consider again the bathtub model \mathcal{R} of Example 3.28. Figure 8 (top left) illustrates some numerical solutions of the dynamic SEs, with $\alpha = (1, 1, 1, 1, 4/5)$, $E_K = 1/2$, $E_{Q_i} = 1$ and for randomly drawn consistent initial conditions $(0, \mathbf{X}_{[0]})$ of order $\mathbf{0}$. We see that the solutions equilibrate to the a.s. unique equilibrium state $(X_K^*, X_{Q_i}^*, X_P^*, X_{Q_o}^*, X_D^*) = (1/2, 1, 2, 5/2, 1)$ corresponding to the solution of the equilibrated SDCM $\mathcal{M}_{\mathcal{R}}$. If we now perform the perfect intervention $\text{do}(Q_o, K_{Q_o})$ on the system \mathcal{R} , where we force the water outflow X_{Q_o} to be equal to the water inflow X_{Q_i} at all time, that is, $K_{Q_o} = E_{Q_i}$, then this does not give an a.s. unique equilibrium state, but the equilibrium state that is obtained depends on the initial condition, as can be seen in Figure 8 (top center). Indeed, the depth X_D^* at equilibrium will equal the initial depth $X_{[0],D}$ at $t_0 = 0$, if the inflow X_{Q_i} equals the outflow X_{Q_o} . This example cannot be treated with the theory of [52], which assumes that the equilibrium state is unique and does not depend on the initial condition. However, if instead we perform the perfect intervention $\text{do}(Q_o, K_{Q_o})$ on \mathcal{R} where $K_{Q_o} < E_{Q_i}$, then the depth X_D will not reach equilibrium, but will increase indefinitely, since the rate of water flowing into the bathtub is larger than the outflow rate. This is illustrated in Figure 8 (top right). This is also reflected in the equilibrated SDCM $\mathcal{M}_{\mathcal{R}}$, which does not have any solution after the corresponding perfect intervention $\text{do}(Q_o, K_{Q_o})$.*

Similar behavior is observed for the equilibrium states of the price, supply and demand model \mathcal{R} of Example 4.13. For example, the model \mathcal{R} will reach market equilibrium if one holds the price fixed at all times by the perfect intervention $\text{do}(P, K_P)$, but will not reach equilibrium if the supply and demand are fixed at all times by the perfect intervention $\text{do}(\{S, D\}, (K_S, K_D))$ for which $K_S < K_D$ (see Figure 7 center and right respectively). In all the cases depicted in Figure 7 we see a dependence of the equilibrium states on the initial condition.

²⁷ At equilibrium, the total inflow of water through the faucet and the hose has to be equal to the total outflow through the drain and the hose.

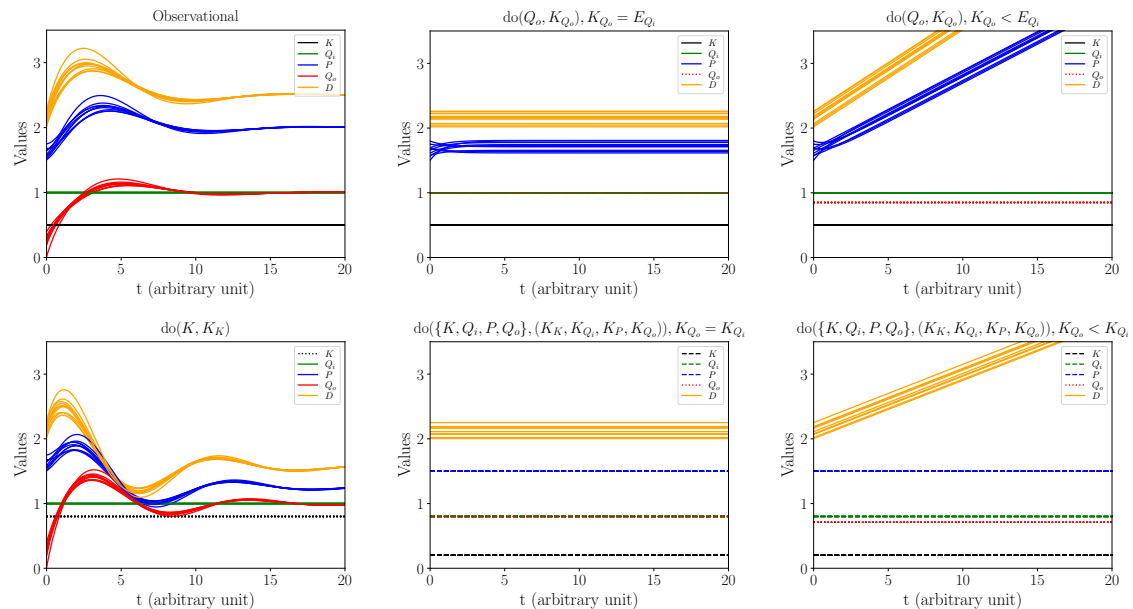


Figure 8. Simulation of solutions of the SDCM of the bathtub model of Example 4.20 and 4.25 under different steady perfect interventions.

In summary, the equilibration of a steady SDCM to an SCM generalizes the work of [52] in three directions: (i) the deterministic setting is replaced with a more general stochastic setting, (ii) the dynamic structural equations can be of arbitrary order (including zeroth-order), rather than only first-order, which prevents complications with the causal interpretation (see, for example, Example 3.11), and (iii) the equilibrium state is allowed to depend on initial conditions. Together, this substantially extends the applicability of the theory.

4.6 Realizing a given SCM as a stable SDCM

Although each steady SDCM equilibrates to an SCM, not all solutions of the SDCM need to equilibrate to solutions of the corresponding SCM (see, for example, Example 4.12). In this subsection, we address the inverse problem of finding steady SDCMs with non-trivial dynamics for which all solutions equilibrate to solutions of a specified SCM. This can be thought of as realizing the given SCM as a “stable” SDCM. In Proposition 4.4 we provided certain conditions under which all the solutions of a linear SDCM equilibrate. Based on this result and some results in the linear systems theory literature, we show that for a certain class of SCMs one can construct a first-order SDCM such that *all* its solutions equilibrate to the solutions of the SCM. Moreover, we show that under certain stronger conditions, the SDCM can be chosen such that its solutions still equilibrate to the solutions of the intervened SCM after any constant stochastic perfect intervention. Hence, the constructed SDCM realizes the causal semantics of the SCM at equilibrium.

First, we observe that one cannot uniquely recover an SDCM from its equilibration in general.

Example 4.21. Consider the linear SDCM \mathcal{R} with dynamic SE given by

$$\mathbf{X} = B\mathbf{X} - \mathbf{X}' + \Gamma\mathbf{E},$$

where the matrix $A := \mathbb{I} - B$ is invertible and the exogenous process \mathbf{E} is a random variable. Consider another SDCM $\tilde{\mathcal{R}}$ which differs only in its dynamic causal mechanism, and has the dynamic SE

$$\mathbf{X} = B\mathbf{X} - \Lambda\mathbf{X}' + \Gamma\mathbf{E}, \quad (6)$$

where Λ is some invertible diagonal matrix. The equilibrated SDCMs $\mathcal{M}_{\mathcal{R}}$ and $\mathcal{M}_{\tilde{\mathcal{R}}}$ coincide, and have structural equations of the form

$$\mathbf{X}^* = B\mathbf{X}^* + \Gamma\mathbf{E}^*. \quad (7)$$

Hence, the equilibrium states \mathbf{X}^* of the SDCMs \mathcal{R} and $\tilde{\mathcal{R}}$ are indistinguishable, since both have to satisfy $\mathbf{X}^* = A^{-1}\Gamma\mathbf{E}^*$ a.s.. Furthermore, if the matrix $\mathbb{I}_I - B_{II}$ is invertible for some subset $I \subseteq \mathcal{I}$, then also the intervened equilibrium states of $\mathcal{R}_{\text{do}(I, \mathbf{K}_I)}$ and $\tilde{\mathcal{R}}_{\text{do}(I, \mathbf{K}_I)}$ are indistinguishable for any sufficiently smooth steady stochastic intervention $\text{do}(I, \mathbf{K}_I)$.

Although the equilibrated SDCM in Example 4.21 describes the possible equilibrium state of both SDCMs, it is not necessarily guaranteed that the solutions of both SDCMs equilibrate. One might hope that for any given linear SCM of the form (7), one can always find an invertible diagonal matrix Λ such that one can construct a steady SDCM of the form (6) for which all solutions of the SDCM equilibrate to the (a.s. unique) solution of the SCM (see Proposition 4.4). Such a “stabilization matrix” Λ does not always exist. A sufficient condition for its existence was given in [26, 27], leading to the following result.

Corollary 4.22. *Let \mathcal{M} be a linear SCM with structural equations*

$$\mathbf{X} = B\mathbf{X} + \Gamma\mathbf{E},$$

where $\mathcal{I} = \{1, \dots, d\}$, $\mathcal{J} = \{1, \dots, e\}$, $B \in \mathbb{R}^{d \times d}$, $\Gamma \in \mathbb{R}^{d \times e}$, and with \mathbf{E} a random variable. Write $A := \mathbb{I} - B$. For a diagonal matrix $\Lambda \in \mathbb{R}^{d \times d}$, consider the linear SDCM $\mathcal{R}_{\mathcal{M}, \Lambda}$ with dynamic SE of the form

$$\mathbf{X} = B\mathbf{X} - \Lambda\mathbf{X}' + \Gamma\mathbf{E}.$$

If there exists a sequence of matrices M_d, M_{d-1}, \dots, M_1 with $M_d = A$ such that for $k = 2, \dots, d$ each M_{k-1} is a principal $(k-1) \times (k-1)$ submatrix of M_k , with $\det M_k \neq 0$ for all $k = 1, \dots, d$, then there exists a diagonal stabilization matrix $\Lambda \in \mathbb{R}^{d \times d}$ such that the linear SDCM $\mathcal{R}_{\mathcal{M}, \Lambda}$ has the properties that (i) its equilibrated SDCM is $\mathcal{M}_{\mathcal{R}_{\mathcal{M}, \Lambda}} = \mathcal{M}$, and (ii) all its solutions equilibrate to an a.s. unique equilibrium state that satisfies the structural equations of the SCM \mathcal{M} , independent of the initial condition.

While this sufficient condition guarantees the existence of a stabilization matrix Λ such that the *observational* equilibrium distribution of the SCM is recovered as the distribution of the equilibrium state of the SDCM, it does not guarantee that after a stochastic perfect intervention on the SDCM, all solutions will equilibrate to an (a.s. unique) equilibrium solution of the corresponding intervened SCM. Indeed, a certain Λ that stabilizes the dynamics in the absence of the intervention may no longer stabilize the dynamics after the intervention has been carried out. Can we, under some conditions, find a single Λ that will stabilize the dynamics after *any* stochastic perfect intervention? The answer is affirmative, as was shown by Locatelli and Schiavoni [47] who provide a necessary and sufficient condition for the existence of an invertible diagonal stabilization matrix Λ that simultaneously stabilizes all subsystems.²⁸ This leads to the following result on how one can “realize” a given linear SCM as a stable linear first-order SDCM.

Corollary 4.23. *Let \mathcal{M} be a linear SCM with structural equations*

$$\mathbf{X} = B\mathbf{X} + \Gamma\mathbf{E},$$

where $\mathcal{I} = \{1, \dots, d\}$, $\mathcal{J} = \{1, \dots, e\}$, $B \in \mathbb{R}^{d \times d}$, $\Gamma \in \mathbb{R}^{d \times e}$, and with \mathbf{E} a random variable. Write $A := \mathbb{I} - B$. For a diagonal matrix $\Lambda \in \mathbb{R}^{d \times d}$, consider the linear SDCM $\mathcal{R}_{\mathcal{M}, \Lambda}$ with dynamic SE of the

²⁸ Locatelli and Schiavoni [47] consider an extension of the stabilization problem studied by Fisher and Fuller [27]. Whereas Fisher and Fuller [27] consider the problem of finding a diagonal matrix $\Lambda \in \mathbb{R}^{d \times d}$ for a matrix $A \in \mathbb{R}^{d \times d}$ such that the matrix ΛA is Hurwitz (for which they provide a sufficient condition), Locatelli and Schiavoni [47] consider the case where *all* the principal submatrices of ΛA should be Hurwitz, and provide a condition that is both sufficient and necessary, as well as a construction of such a stabilization matrix Λ .

form

$$\mathbf{X} = B\mathbf{X} - \Lambda\mathbf{X}' + \Gamma\mathbf{E}.$$

If

$$\det(A_{II}) \det(\text{diag}(A_{II})) > 0 \quad \forall I \subseteq \mathcal{I}, \quad (8)$$

then there exists an invertible diagonal stabilization matrix $\Lambda \in \mathbb{R}^{d \times d}$ such that the linear SDCM $\mathcal{R}_{\mathcal{M}, \Lambda}$ has the properties that (i) its equilibrated SDCM is $\mathcal{M}_{\mathcal{R}_{\mathcal{M}, \Lambda}} = \mathcal{M}$, and (ii) under every stochastic perfect intervention $\text{do}(J, \mathbf{K}_J)$ with the exogenous process \mathbf{K}_J constant in time, all solutions of $(\mathcal{R}_{\mathcal{M}, \Lambda})_{\text{do}(J, \mathbf{K}_J)}$ equilibrate to an a.s. unique equilibrium state that is the unique solution of the SCM $\mathcal{M}_{\text{do}(J, \mathbf{K}_J^*)}$, independent of the initial condition.

Condition (8) implies that the matrices $\mathbb{I}_I - B_{II}$ are invertible for every subset $I \subseteq \mathcal{I}$. Such linear SCMs are special cases of the class of *simple* SCMs (see [9]). Simple SCMs have the convenient property that their solutions are a.s. unique after any stochastic perfect intervention. We conclude that for the subclass of simple linear SCMs that satisfy condition (8), we can construct a linear first-order SDCM whose causal semantics at equilibrium “realizes” that described by the SCM. We speculate that this result can be extended to higher-order and nonlinear systems, but we will not pursue these questions here.

Example 4.24. We show that the equilibrated SDCM of Example 4.12 (see also Example 3.5), modeling the equilibrium states of a damped coupled harmonic oscillator, satisfies condition (8). Indeed, taking $\mathcal{I} = \{1, \dots, d\}$, the matrix B of this linear SCM is tridiagonal, given as

$$B = \begin{bmatrix} 0 & \frac{\kappa_1}{\kappa_0 + \kappa_1} & & & \\ \frac{\kappa_1}{\kappa_1 + \kappa_2} & 0 & \frac{\kappa_2}{\kappa_1 + \kappa_2} & & \\ & \frac{\kappa_2}{\kappa_2 + \kappa_3} & 0 & \ddots & \\ & & \ddots & \ddots & \frac{\kappa_{d-1}}{\kappa_{d-2} + \kappa_{d-1}} \\ & & & \frac{\kappa_{d-1}}{\kappa_{d-1} + \kappa_d} & 0 \end{bmatrix},$$

where $\kappa_0 = \kappa_d = 0$. Hence $A = \mathbb{I} - B = DC$ with diagonal

$$D = \begin{bmatrix} \frac{1}{\kappa_0 + \kappa_1} & & & & \\ & \frac{1}{\kappa_1 + \kappa_2} & & & \\ & & \frac{1}{\kappa_2 + \kappa_3} & & \\ & & & \ddots & \\ & & & & \frac{1}{\kappa_{d-1} + \kappa_d} \end{bmatrix}$$

and tridiagonal

$$C = \begin{bmatrix} \kappa_0 + \kappa_1 & -\kappa_1 & & & \\ -\kappa_1 & \kappa_1 + \kappa_2 & -\kappa_2 & & \\ & -\kappa_2 & \kappa_2 + \kappa_3 & \ddots & \\ & & \ddots & \ddots & -\kappa_{d-1} \\ & & & -\kappa_{d-1} & \kappa_{d-1} + \kappa_d \end{bmatrix}.$$

The determinants of D and C can be expressed in closed form as

$$\det D = \prod_{i=1}^d \frac{1}{\kappa_{i-1} + \kappa_i}, \quad \det C = \sum_{i=0}^d \prod_{\substack{j=0 \\ j \neq i}}^d \kappa_j.$$

Hence, since $\kappa_i > 0$ for $i = 1, \dots, d-1$, $\det A = (\det C)(\det D) > 0$. Also, we clearly have $\det \text{diag}(A) > 0$. Hence, condition (8) holds for $I = \mathcal{I}$. A similar calculation (and exploiting the block structure of the principal submatrices) shows that condition (8) holds for all $I \subseteq \mathcal{I}$.

Remarkably, we can thus apply Corollary 4.23 to the damped harmonic oscillator SCM to obtain a realization of this causal equilibrium model as a *first-order* linear SDCM (remember that the original SDCM is a second-order linear SDCM).

4.7 Causal interpretation of the graph of the equilibrated SDCM

While the graph of an acyclic SCM has a straightforward causal interpretation, this need not be the case for general SCMs with cycles [9].²⁹ While an acyclic SCM induces a unique “observational” distribution, cyclic SCMs may induce none, one or several different observational distributions [35]. Similarly, after performing a perfect intervention on some of the variables, a cyclic SCM may induce none, one or several different corresponding interventional distributions. In general, one has to be careful in how to causally interpret the graph of an SCM if cycles are present; in particular, this caveat holds for SCMs that are obtained as the equilibration of a SDCM. First, not all directed edges and directed paths in the graph can easily be identified from differences in interventional distributions in case cycles are present [9]. Second, if cycles are present, “nonancestral” effects may exist [9, 54], that is, an intervention on a variable may change the distribution of some of its nondescendants in the graph. In this subsection, we show how these subtleties and counterintuitive nonancestral effects in cyclic equilibrated SDCMs can be explained in terms of properties of the underlying SDCM.

In general, the presence or absence of a directed edge or path in the graph of an SCM \mathcal{M} cannot always be identified from the observational and interventional distributions. In the cyclic setting, the following sufficient condition can be used to identify such directed edges or paths between nodes i and j (see Proposition 7.1 in [9] for the exact formulation).

- A *direct causal effect* of i on j can be identified, that is, there exists a $i \rightarrow j \in \mathcal{G}(\mathcal{M})$, if (i) the structural equation of j can be solved a.s. uniquely for X_j in terms of the other variables that appear in the equation, and (ii) there exist values $\mathbf{K}_I \in \mathcal{X}_I$ and $K_i \neq \tilde{K}_i \in \mathcal{X}_i$, where $I = \mathcal{I} \setminus \{i, j\}$, and a measurable set $\mathcal{B}_j \subseteq \mathcal{X}_j$ such that the following probabilities are uniquely defined and do not coincide:

$$\mathbb{P}_{(\mathcal{M}_{\text{do}(I, \mathbf{K}_I)})_{\text{do}(i, K_i)}}(X_j \in \mathcal{B}_j) \neq \mathbb{P}_{(\mathcal{M}_{\text{do}(I, \mathbf{K}_I)})_{\text{do}(i, \tilde{K}_i)}}(X_j \in \mathcal{B}_j);$$

- An *indirect causal effect* of i on j can be identified, that is, there exists a directed path $i \rightarrow \dots \rightarrow j$ in $\mathcal{G}(\mathcal{M})$, if (i) the structural equations of the ancestors of j in $\mathcal{G}(\mathcal{M})_{\setminus i}$ (that is, the graph $\mathcal{G}(\mathcal{M})$ where we removed the node i and its adjacent edges) can be solved a.s. uniquely for their associated variables in terms of the other variables that appear in these equations, and (ii) there exist values $K_i \neq \tilde{K}_i \in \mathcal{X}_i$ and a measurable set $\mathcal{B}_j \subseteq \mathcal{X}_j$ such that the following probabilities are uniquely defined and do not coincide:

$$\mathbb{P}_{\mathcal{M}_{\text{do}(i, K_i)}}(X_j \in \mathcal{B}_j) \neq \mathbb{P}_{\mathcal{M}_{\text{do}(i, \tilde{K}_i)}}(X_j \in \mathcal{B}_j).$$

In the following example, we illustrate how we can interpret the directed edges and paths of the equilibrated bathtub model that cannot be identified by this sufficient condition from an SDCM perspective.

Example 4.25 (Bathtub model, continued). *Consider again the bathtub model \mathcal{R} of Example 3.28. We simulated some numerical solutions, with parameters as given in Example 4.20, shown in Figure 8 (top left). In Table 1 (bottom left) one can read off all the indirect causal effects that can be identified by comparing different interventional distributions from the equilibrated model $\mathcal{M}_{\mathcal{R}}$ with the help of Proposition 7.1 in [9]. The indirect causal effects of P and Q_o cannot be identified by comparing interventional distributions, since the intervened equilibrated models $(\mathcal{M}_{\mathcal{R}})_{\text{do}(P, K_P)}$ and $(\mathcal{M}_{\mathcal{R}})_{\text{do}(Q_o, K_{Q_o})}$ do not have a solution (except*

²⁹ The straightforward causal interpretation of acyclic SCMs actually extends to a much more general class of possibly cyclic SCMs, referred to by Bongers et al. [9] as *simple* SCMs.

		directed path to					
from		K	Q_i	P	Q_o	D	
	K	-	×	✓	✓	✓	
	Q_i	×	-	✓	✓	✓	
	P	×	×	-	✓	✓	
	Q_o	×	×	✓	-	✓	
	D	×	×	✓	✓	✓	

		directed edge to					
from		K	Q_i	P	Q_o	D	
	K	-	×	×	✓	×	
	Q_i	×	-	×	×	✓	
	P	×	×	-	✓	×	
	Q_o	×	×	×	-	✓	
	D	×	×	✓	×	✓	

		indirect effect					
cause		K	Q_i	P	Q_o	D	
	K	-	?	✓	?	✓	
	Q_i	?	-	✓	✓	✓	
	P	??	??	-	??	??	
	Q_o	??	??	??	-	??	
	D	?	?	✓	✓	-	

		direct effect					
cause		K	Q_i	P	Q_o	D	
	K	-	?	?	✓	??	
	Q_i	?	-	?	?	??	
	P	?	?	-	✓	??	
	Q_o	?	?	?	-	??	
	D	?	?	✓	?	-	

according to $\mathcal{G}(\mathcal{M}_{\mathcal{R}})$ (see Figure 6 top right)

Identifiable from the observational and interventional distributions with Prop. 7.1 in [9]

Table 1. The directed paths/edges (top tables) of the equilibrated bathtub model $\mathcal{M}_{\mathcal{R}}$ and the (in)direct causal effects that can be identified by Proposition 7.1 in [9] (bottom tables) are denoted by a “✓”. Those that cannot be identified are denoted by the question marks “?” and “?”. A single question mark “?” denotes that condition (i) is satisfied, but not condition (ii), while a double question mark “??” denotes that condition (i) is not satisfied.

for one special choice of K_P respectively K_{Q_o}), and hence condition (i) is not satisfied. This was already illustrated for the perfect intervention $\text{do}(Q_o, K_{Q_o})$ in Figure 8 (top center/right) of Example 4.20.

The direct causal effects that can be identified from $\mathcal{M}_{\mathcal{R}}$ are given in Table 1 (bottom right). The direct causes of D cannot be identified due to the self-cycle at D , which means that condition (i) is not satisfied, that is, the structural equation of D cannot be a.s. uniquely solved for D in terms of the other variables. Indeed, the depth D will not equilibrate, but will increase indefinitely, if the rate of water into the bathtub is larger than the outflow rate, that is, $K_{Q_o} < K_{Q_i}$ (see Figure 8 bottom right). On the other hand, it will reach an equilibrium state only if the rate of water into and out of the bathtub are equal, that is, $K_{Q_o} = K_{Q_i}$. In this case, the depth D will remain constant over all times, as illustrated in Figure 8 (bottom center).

The directed path from K to Q_o in the graph of the equilibrated model $\mathcal{M}_{\mathcal{R}}$ cannot be straightforwardly identified as an indirect causal effect at equilibrium, because the equilibrium distribution of Q_o does not change due to perfect interventions on K (this corresponds to the single question mark in the Table 1, bottom left), as explicit calculations reveal. However, at some finite time point one does observe changes in the distribution of Q_o when performing perfect interventions on K (Figure 8 (left)). Together, this implies that this system is capable of perfect adaptation [5]. Interestingly, the direct edge $K \rightarrow Q_o$ in the graph of the equilibrated model $\mathcal{M}_{\mathcal{R}}$ can be identified by changes in the equilibrium distribution of Q_o under perfect interventions on K, Q_i, P, D (which then also implies that there is a directed path from K to Q_o in the graph of the equilibrated model).

In particular, this example illustrates that one can run into several problems when one attempts to identify directed edges and paths of the graph of the SCM from the differences in equilibrium distributions under interventions on the SDCM:

- if the intervened SCM has no solutions, then the descendants of the intervention targets cannot be easily identified;
- if the graph of the SCM has a self-cycle at some variable, then the parents of that variable cannot be easily identified;
- if the equilibrium distribution of some descendants of the intervention target variable remain insensitive to the intervention (for example, when the dynamical system exhibits perfect adaptation [5]), these descendants cannot be easily identified.




Figure 9. Graphs of the SDCM \mathcal{R} (left) and the corresponding equilibrated model $\mathcal{M}_{\mathcal{R}}$ (right) of Example 4.26.

In Example 4.25, the identified indirect causal relationships are a subset of the ancestral relationships. This can be seen from observing that each “✓” in Table 1 (bottom left) has a corresponding “✓” in Table 1 (top left). In other words, performing a perfect intervening on a variable can only change the distribution of its descendants in the graph. In general, however, it can happen that an intervention on a nonancestor of a variable can change the distribution of that variable [9, 54]. This counterintuitive behavior of “nonancestral” effects in an equilibrated SDCM can be explained by the dependence of the equilibrium states on the initial conditions in combination with the fact that not each initial condition corresponds to an equilibrating solution. The following example illustrates this.

Example 4.26 (Selection bias leading to nonancestral effects in an equilibrated SDCM). *Consider the SDCM \mathcal{R} with dynamic structural equations given by*

$$\begin{cases} X_1 = X_1 - X'_1 + 2X_2 - X_3 \\ X_2 = X_2 - X'_2 \\ X_3 = E, \end{cases}$$

with order tuple $\mathbf{n} = (1, 1, 0)$ and E some constant in \mathbb{R} . Denote $I = \{1, 2\}$ and note that \mathcal{R} satisfies Assumption 1-I. The equilibrated model $\mathcal{M}_{\mathcal{R}}$ is given by

$$\begin{cases} X_1^* = X_1^* + 2X_2^* - X_3^* \\ X_2^* = X_2^* \\ X_3^* = E. \end{cases}$$

The graphs of \mathcal{R} and $\mathcal{M}_{\mathcal{R}}$ are depicted in Figure 9. First observe that the induced equilibrium distribution of X_2^ differs for two constant perfect interventions $\text{do}(3, K_3)$ and $\text{do}(3, \tilde{K}_3)$ with $K_3 \neq \tilde{K}_3$, since the equilibrium state has to satisfy $X_2^* = X_3^*/2$ a.s.. However, there is no directed path from the variable X_3 to the variable X_2 in the graph of the SCM $\mathcal{M}_{\mathcal{R}}$. This counterintuitive behavior can be explained by taking the initial conditions of the solutions of the SDCM into account, as we shall now explain.*

In Figure 10, we plot the solutions of the SDCM \mathcal{R} for different partial initial conditions $(t_0, \mathbf{X}_{I, [0]}^i)$ at $t_0 = 0$ (for $i = a, b, \dots, g$) under two steady perfect interventions, namely $\text{do}(3, K_3 = 1.0)$ and $\text{do}(3, \tilde{K}_3 = 0.6)$. For illustration purposes, we consider here only non-random initial conditions, because we can then identify the initial conditions with “individual” solutions, as depicted in Figure 10 (note that Corollary 3.26 applies). Observe that the set of partial initial conditions that correspond to equilibrating solutions differs for the two interventions. For the intervened model $\mathcal{R}_{\text{do}(3, K_3)}$, the only solution that equilibrates is the one with initial condition $(t_0, \mathbf{X}_{I, [0]}^a)$ (denoted by the dark solid lines Figure 10 (top left)), whereas for the intervened model $\mathcal{R}_{\text{do}(3, \tilde{K}_3)}$ the only solution that equilibrates is the one with initial condition $(t_0, \mathbf{X}_{I, [0]}^b)$ (denoted by the dark dotted lines in Figure 10 (top right)). This explains the counterintuitive behavior of nonancestral effects in the equilibrium SCM: The chosen value for X_3 affects which solutions will equilibrate, and thereby affects the equilibrium distribution of X_2 .

Note that at any finite point in time, these “nonancestral” effects do not occur; indeed, Figure 10 shows that the distribution of X_1 differs for the two interventions at finite time, while that of X_2 remains unaffected.

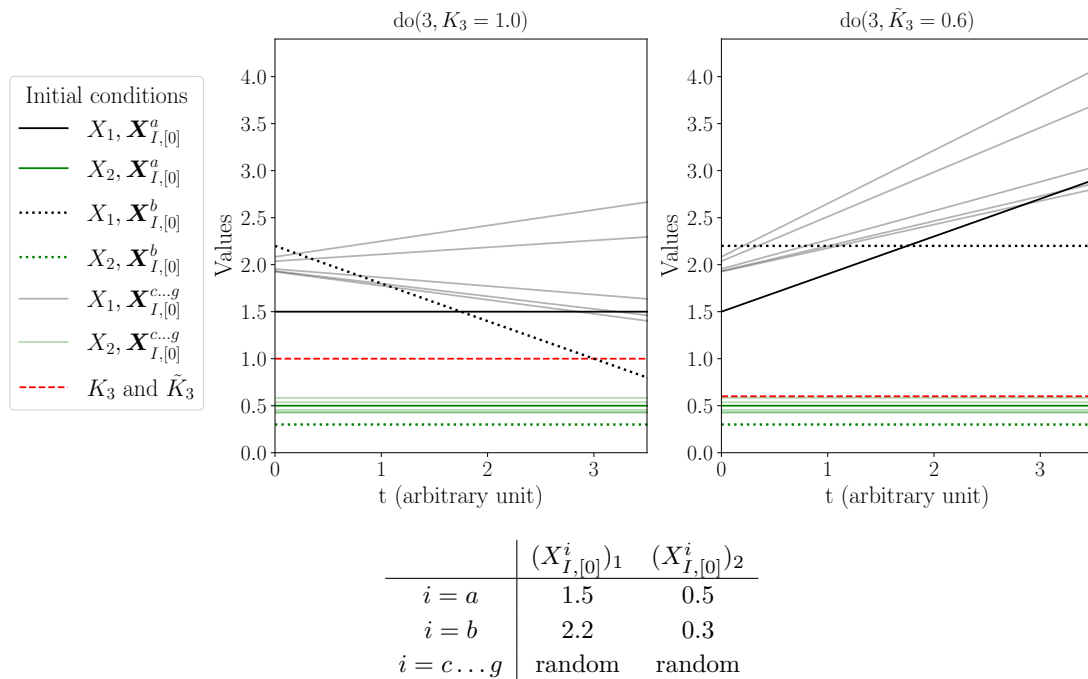


Figure 10. Simulation of solutions of the SDCM of Example 4.26 under different steady perfect interventions on X_3 (top left and right). The simulations in the top left and right plots are performed under the same set of initial conditions, summarized in the bottom table, but under different interventions.

This example shows that the nonancestral effects in an equilibrated SDCM can be explained by the dependence of the equilibrium states on the initial conditions, in combination with the fact that not each initial condition corresponds to an equilibrating solution. Another way to think about this is as selection bias due to the assumption that the system has reached equilibrium. An intervention targeting a certain variable may change the set of equilibrating initial conditions of the system, and it can even change initial conditions for non-ancestors of the intervention target. By only considering these equilibrating initial conditions, this may appear as a causal effect of a variable on some of its non-ancestors at equilibrium. When seen from this perspective, these “causal effects” can be considered to be spurious as they do not appear on an “individual level”, that is, for individual trajectories (at finite time t), but only appear on a “population level” when selecting on some later event (namely, the system being at equilibrium). One can indeed think of this as selection bias due to equilibration.

5 Discussion

Dynamical models consisting of (ordinary or random) differential equations are widely applied in science and engineering to model the dynamics of systems that are composed of several components. These differential equations by themselves do not have a clearcut causal interpretation. Although they may implicitly explain a particular phenomenon in terms of its causes, the causal semantics of the constituent components are generally not explicitly defined without additional assumptions.

In this work, we introduced structural dynamical causal models that formally encode causal semantics of stochastic processes by means of a structured set of random differential equations. SDCMs can be seen as stochastic-process versions of structural causal models, where the random variables are replaced by stochastic processes and their derivatives. By viewing the (higher-order) derivatives $X_i^{(k_i)}$ to be aspects

of the process X_i we arrive at a natural causal interpretation, where it is not necessary to even consider questions like “does position cause velocity, or does velocity cause position, or both?”.

For steady SDCMs (for which the explicit time-dependence of the dynamics vanishes as $t \rightarrow \infty$) we introduced an equilibration operation that equilibrates the dynamic causal mechanism of each component separately. This led to the important result that intervention and equilibration commute, thus connecting the causal semantics at equilibrium with the causal semantics of the dynamics. It generalizes the analogous result of Mooij et al. [52] in three directions: (i) we replaced the deterministic setting with a more general stochastic setting, which allows us to address both cycles and confounders, (ii) we allowed the order of the dynamic structural equations to be arbitrary, including zeroth-order, rather than only allowing for first-order differential equations, and (iii) we have dropped the strong assumption that the dynamical model needs to have a single globally attractive equilibrium state. This allows us to study the causal semantics of the equilibrium states of a plethora of dynamical systems subject to time-varying random disturbances encountered in science and engineering within the framework of structural causal models.

Our commutation result may appear to be at odds with the possible “violation of the equilibration-manipulation commutability property” pointed out by Dash [15]. Under our notion of equilibration—contrary to that of Dash—each dynamic structural equation of the SDCM becomes a structural equation of the SCM. This one-to-one correspondence between the equations leads to the preservation of the causal semantics under equilibration. We can reinterpret the phenomenon that Dash observed as the fact that the equilibrium distributions of certain dynamical systems (for example ones that exhibit perfect adaptation) are not faithful to the graph of the equilibrated SCM [5], in the sense that they can contain conditional independences not explained by this graph. For dynamical systems exhibiting perfect adaptation, these faithfulness violations are due to the structure of the dynamics, rather than “accidental” parameter cancellations. This has serious repercussions for attempts at inferring the causal structure from (conditional independences in) equilibrium data [5]. Thus, in a different way we arrive at the same conclusion as Dash obtained.

In comparison with the causal constraints models of Blom et al. [6], our modeling framework is more “agnostic” as we decided not to incorporate the initial conditions into the model.³⁰ This allowed us to causally model *all* the equilibrium states of a steady SDCM with a single SCM. However, that single SCM may not provide a *complete* description of the causal semantics at equilibrium [6]. This is indeed a modeling tradeoff: the simpler structure of SCMs compared to that of causal constraints models can come at the cost of a less complete description of the equilibrium behavior of certain dynamical systems. On the other hand, the connection between the structure of the SCM and that of the underlying SDCM is straightforward, whereas it is not well understood at present how one can easily derive a concise yet complete representation of an equilibrated SDCM (and a corresponding initial condition) as a causal constraints model.

However, allowing for multiple (or no) solutions also comes at a cost: the causal interpretation of the SCM is more subtle than that of acyclic (or more generally, simple) SCMs, and in particular, does not straightforwardly relate to properties of its graph. We illustrated for the bathtub model how one can causally interpret the directed edges and paths of the graph of the SCM that models the equilibrium states of the underlying SDCM. We saw that one may run into several problems when attempting to identify aspects of the SCM graph from comparing differences in equilibrium distributions after intervening on some of the variables:

- if the intervened SCM has no solution (which may happen if the intervened SDCM does not converge to a finite equilibrium state, but instead diverges to infinity, or reaches a periodic limit cycle, for example), descendants of the intervention targets cannot be easily identified;
- if the SCM graph has a self-cycle at some variable (which may happen if the causal mechanism for that variable does not equilibrate for certain values of its parents), then the parents of that variable cannot be easily identified;

³⁰ This is analogous to the difference between an ODE and an initial-value problem.

- if the equilibrium distributions of some descendants of the intervention target variable remain insensitive to the intervention (which may happen in dynamical systems exhibiting perfect adaptation), these descendants cannot be easily identified.

Even worse, the equilibrium SCM may entail distribution changes under interventions that appear to be of a causal nature, while no corresponding causal relations are present in the dynamics (and therefore, no corresponding ancestral relations are present in the SCM graph), as we pointed out in Example 4.26. These counterintuitive “nonancestral causal effects” can be understood as arising from the implicit selection bias due to conditioning on the system having reached an equilibrium state. Indeed, the solutions of the equilibrium SCM correspond to those solutions of the SDCM that have equilibrated, while the non-equilibrating solutions of the SDCM are ignored. In other words, the SCM provides the “population-level” causal semantics of the population of equilibrating SDCM solutions (at $t = \infty$), which can deviate from the “individual-level” causal semantics of (possibly non-equilibrating) SDCM solutions (at finite t). The phenomenon that population-level causality may differ from individual-level causality due to post-intervention selection bias is well-known in other contexts. For example, a car mechanic who only observes cars that don’t start may conclude that replacing the battery causes start engines to fail. While this appears as a genuine causal effect on the population level, it would be foolish to conclude that this causal effect also pertains to individual cars. Intuitively, one might prefer to interpret such phenomena as not representing “truly causal” relations. On the other hand, if one is only interested in the effects of interventions on a population level, there seems to be no harm in considering these distribution changes as causal. Thus, as long as one is explicit whether one refers to population-level or individual-level causality, both notions of causality can meaningfully co-exist. The important take-away, from our point of view, is that focussing on equilibrated systems may lead to selection bias.

As a side note, Example 4.26 also shows that SCMs may not fully capture such population-level causal relations graphically. We note that the recently proposed framework of Blom et al. [7] is better suited in general to read off such population-level causal effects graphically from the structure of the equilibrium equations, under certain “local” solvability assumptions on these equations (rather than having to study global solutions of intervened equilibrium equations, as we did here).

Apart from these subtleties regarding their causal semantics, SCMs with cycles bring about several other challenges in general. For example, they generally do not have a Markov property, and the class of cyclic SCMs is not closed under marginalization. The subclass consisting of simple SCMs [9] allows for cycles, but simple SCMs share many of the convenient properties of acyclic SCMs. Hence, these convenient properties are directly applicable to the equilibrium states of those steady SDCMs that equilibrate to a simple SCM. This enables one to study the equilibrium states of those SDCMs by statistical tools and discovery methods available for simple SCMs. For example, one can apply adjustment criteria and Pearl’s do-calculus [29]. Several causal discovery algorithms, originally designed for acyclic SCMs, like Local Causal Discovery (LCD) [13], Y-structures [49], and the Fast Causal Inference (FCI) algorithm [51, 66, 72], are directly applicable to simple SCMs as well [53]. Furthermore, the Joint Causal Inference (JCI) framework can be applied to combine data from different contexts (for example, observational and interventional) for causal discovery and inference purposes [53].

Given that steady SDCMs for which all solutions equilibrate give rise to SCMs at equilibrium, the inverse problem becomes interesting as well: given an SCM, can we find an SDCM (with non-trivial dynamics) that equilibrates to this SCM and for which all solutions equilibrate? This question was answered affirmatively for a certain class of linear simple SCMs with additional constraints on the parameters by leveraging existing results from linear systems theory. We speculate that this result can be further generalized to allow for non-linearity. Perhaps surprisingly, this result allows to start from a second-order SDCM modeling a system of damped coupled harmonic oscillators, equilibrate it to obtain an SCM, and from that then construct a *first-order* SDCM with the same equilibrium SCM that describes all equilibrium states under any constant stochastic perfect intervention. This shows that the order of the dynamic structural equations is not necessarily constrained by the equilibrium SCM. Thus, the properties of the system at equilibrium may contain not enough information to identify the order of the dynamical equations.

We hope that the framework of SDCMs provides a natural starting point for modeling the causal mechanisms that underly the dynamics of various systems, which could, in principle, be inferred from observations and experiments [see, for example, 3, 46, 59]. We believe that most of this work can also easily be adapted to discrete time by replacing the differential equations by difference equations. Future work might consist of (i) investigating the notion of local independence in SDCMs, (ii) studying how SDCM graphs can be interpreted causally, in particular if self-cycles or zeroth order equations are present, (iii) developing structure and parameter learning algorithms for SDCMs, and (iv) investigating possible extensions to stochastic dynamics by means of stochastic differential equations.

Acknowledgements

We are grateful to Fabian Wirth for pointing us to the work of Fisher and Fuller [27]. Stephan Bongers was supported by NWO, the Netherlands Organization for Scientific Research (VIDI grant 639.072.410). Tineke Blom and Joris Mooij were supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement 639466).

Appendix

A Proofs

Proof of Theorem 3.23. For every $i \in I$ we can write the random differential equations

$$X_i^{(n_i)} = g_i(\bar{X}_I^{(n_I-1)}, \mathbf{X}_{\setminus I}, \mathbf{E})$$

as a system of first-order random differential equations

$$\frac{d}{dt} \bar{X}_i^{(n_i-1)} = \tilde{g}_i(\bar{X}_i^{(n_i-1)}, \bar{X}_{I \setminus i}^{(n_{I \setminus i}-1)}, \mathbf{X}_{\setminus I}, \mathbf{E}),$$

where $\tilde{g}_i : \mathcal{X}_i^{n_i} \times \mathcal{X}_{I \setminus i}^{n_{I \setminus i}} \times \mathcal{X}_{\setminus I} \times \mathcal{E} \rightarrow \mathcal{X}_i^{n_i}$ is the mapping defined by

$$\tilde{g}_i(\bar{x}_i^{(n_i-1)}, \bar{x}_{I \setminus i}^{(n_{I \setminus i}-1)}, \mathbf{x}_{\setminus I}, \mathbf{e}) := (x_i^{(1)}, \dots, x_i^{(n_i-1)}, g_i(\bar{x}_I^{(n_I-1)}, \mathbf{x}_{\setminus I}, \mathbf{e})).$$

Note that $\bar{X}_i^{(n_i-1)} = (X_i, X_i^{(1)}, \dots, X_i^{(n_i-1)})$ and $\frac{d}{dt} \bar{X}_i^{(n_i-1)} = (X_i^{(1)}, \dots, X_i^{(n_i-1)}, X_i^{(n_i)})$.

Substituting the functions $g_{\setminus I}$ yields the following first-order RDE:

$$\frac{d}{dt} \bar{X}_i^{(n_i-1)} = \tilde{g}_i(\bar{X}_i^{(n_i-1)}, \bar{X}_{I \setminus i}^{(n_{I \setminus i}-1)}, g_{\setminus I}(\bar{X}_I^{(n_I-1)}, \bar{X}_{I \setminus i}^{(n_{I \setminus i}-1)}, \mathbf{E}), \mathbf{E}).$$

Let $\tilde{h}_i(\bar{x}_I^{(n_I-1)}, \mathbf{e}) := \tilde{g}_i(\bar{x}_I^{(n_I-1)}, g_{\setminus I}(\bar{x}_I^{(n_I-1)}, \mathbf{e}), \mathbf{e})$. Continuity of g_i and $g_{\setminus I}$ and continuity of the exogenous process \mathbf{E} implies that for almost all $\omega \in \Omega$ the function $(t, \bar{x}_I^{(n_I-1)}) \mapsto \tilde{h}_i(\bar{x}_I^{(n_I-1)}, \mathbf{E}(t, \omega))$ is continuous on $T \times \mathcal{X}_I^{n_I}$. Moreover, for each $\bar{x}_I^{(n_I-1)} \in \mathcal{X}_I^{n_I}$ the function $\mathbf{e} \mapsto \tilde{h}_i(\bar{x}_I^{(n_I-1)}, \mathbf{e})$ is continuous in \mathbf{e} and in particular measurable in \mathbf{e} . Hence, for all $(t, \bar{x}_I^{(n_I-1)}) \in T \times \mathcal{X}_I^{n_I}$ the function $\omega \mapsto \tilde{h}_i(\bar{x}_I^{(n_I-1)}, \mathbf{E}(t, \omega))$ is \mathcal{F} -measurable. Hence, we can write the dynamic SEs as:

$$\frac{d}{dt} \bar{X}_I^{(n_I-1)} = \tilde{h}_I(\bar{X}_I^{(n_I-1)}, \mathbf{E}). \quad (9)$$

Under the assumed condition, the following inequality holds for \mathbb{P} -almost all ω , for all $\bar{x}_I^{(n_I-1)}, \bar{y}_I^{(n_I-1)} \in \mathcal{X}_I^{n_I}$ and all $t \in T$:

$$\begin{aligned} & \sum_{i \in I} \left\| \tilde{h}_i(\bar{x}_I^{(n_I-1)}, \mathbf{E}(t, \omega)) - \tilde{h}_i(\bar{y}_I^{(n_I-1)}, \mathbf{E}(t, \omega)) \right\|^2 \\ &= \sum_{i \in I} \left\| \tilde{g}_i(\bar{x}_I^{(n_I-1)}, g_{\setminus I}(\bar{x}_I^{(n_I-1)}, \mathbf{E}), \mathbf{E})(t, \omega) - \tilde{g}_i(\bar{y}_I^{(n_I-1)}, g_{\setminus I}(\bar{y}_I^{(n_I-1)}, \mathbf{E}), \mathbf{E})(t, \omega) \right\|^2 \\ &= \sum_{i \in I} \left[\|x_i^{(1)} - y_i^{(1)}\|^2 + \dots + \|x_i^{(n_i-1)} - y_i^{(n_i-1)}\|^2 + \right. \\ & \quad \left. \|g_i(\bar{x}_I^{(n_I-1)}, g_{\setminus I}(\bar{x}_I^{(n_I-1)}, \mathbf{E}), \mathbf{E})(t, \omega) - g_i(\bar{y}_I^{(n_I-1)}, g_{\setminus I}(\bar{y}_I^{(n_I-1)}, \mathbf{E}), \mathbf{E})(t, \omega)\|^2 \right] \\ &\leq \sum_{i \in I} [K_\omega(t)^2 \|x_i - y_i\|^2 + \|x_i^{(1)} - y_i^{(1)}\|^2 + \dots + \|x_i^{(n_i-1)} - y_i^{(n_i-1)}\|^2] \\ &\leq (1 + K_\omega(t)^2) \|\bar{x}_I^{(n_I-1)} - \bar{y}_I^{(n_I-1)}\|^2. \end{aligned}$$

Hence the conditions of Theorem 1.2 in Bunke [10] (or Theorem 3.2 in Neckel and Rupp [55]) are satisfied, which proves that there exists an a.s. unique solution \mathbf{X}_I of the system (9) of first-order RDEs for any partial initial condition $(t_0, \bar{X}_{I, [0]})$. Note that $\mathbf{X}_I(t)$ is a C^{n_I} -stochastic process. Extend this to a global solution \mathbf{X} by setting $\mathbf{X}_{\setminus I} = g_{\setminus I}(\bar{X}_I^{(n_I-1)}, \mathbf{E})$. The result satisfies the smoothness requirement; indeed, from the assumptions it follows for each $j \in \mathcal{I} \setminus I$ that X_j is a C^{n_j} -stochastic process. \square

Proof of Proposition 3.24. Let $g_i : \mathcal{X}_i^{n_i} \times \mathcal{X}_{\setminus i} \times \mathcal{E} \rightarrow \mathcal{X}_i$ and $g_j : \mathcal{X}_I \times \mathcal{E} \rightarrow \mathcal{X}_j$ for $i \in I$ and $j \in \mathcal{I} \setminus I$ be continuous mappings that make \mathcal{R} satisfy Assumption 2-I. Consider the stochastic perfect intervention $\text{do}(J, \mathbf{K}_J)$ with \mathbf{K}_J a C^{n_J} -stochastic process. Then, the mappings $h_i : \mathcal{X}_i^{n_i} \times \mathcal{X}_{\setminus i} \times (\mathcal{X}_J \times \mathcal{E}_{\mathcal{J}}) \rightarrow \mathcal{X}_i$ for $i \in I \setminus J$ defined by

$$h_i(\bar{x}_i^{(n_i-1)}, \mathbf{x}_{\setminus i}, (\tilde{\mathbf{e}}_J, \mathbf{e}_{\mathcal{J}})) := g_i(\bar{x}_i^{(n_i-1)}, \mathbf{x}_{\setminus i}, \mathbf{e}_{\mathcal{J}})$$

and the mappings $h_j : \mathcal{X}_{I \setminus J} \times (\mathcal{X}_J \times \mathcal{E}) \rightarrow \mathcal{X}_j$ for $j \in \mathcal{I} \setminus (I \setminus J)$ defined by

$$h_j(\mathbf{x}_{I \setminus J}, (\tilde{\mathbf{e}}_J, \mathbf{e}_{\mathcal{J}})) := \begin{cases} g_j((\mathbf{x}_{I \setminus J}, \tilde{\mathbf{e}}_J), \mathbf{e}_{\mathcal{J}}) & \text{if } j \notin J \\ \tilde{\mathbf{e}}_j & \text{if } j \in J \end{cases}$$

make $\mathcal{R}_{\text{do}(J, \mathbf{K}_J)}$ satisfy Assumption 2-($I \setminus J$).

□

Proof of Proposition 3.25. If the causal mechanism \mathbf{f} is defined as in the proposition, then the mappings $g_I : \mathcal{X}_I^{n_I} \times \mathcal{X}_J \times \mathcal{E} \rightarrow \mathcal{X}_I$ and $g_J : \mathcal{X}_I^{n_I} \times \mathcal{E} \rightarrow \mathcal{X}_J$ are given by

$$\begin{aligned} g_I(\bar{\mathbf{x}}_I^{(n_I-1)}, \mathbf{x}_J, \mathbf{e}) &= -B_{II}^{-1}(B_{II}^{(n_I-1)} \bar{\mathbf{x}}_I^{(n_I-1)} - \mathbf{x}_I + B_{IJ} \mathbf{x}_J + \Gamma_{IJ} \mathbf{e}) \\ g_J(\bar{\mathbf{x}}_I^{(n_I-1)}, \mathbf{e}) &= -B_{JJ}^{-1}(B_{JI}^{(n_I-1)} \bar{\mathbf{x}}_I^{(n_I-1)} + \Gamma_{JJ} \mathbf{e}). \end{aligned}$$

The converse is shown by taking for $B_{II}^{(n_I)}$ and B_{JJ} the identity matrices.

□

Proof of Corollary 3.26. For a linear SDCM \mathcal{R} that satisfies Assumption 1-I there always exists a $K_\omega \in \mathbb{R}$, which is independent of ω , such that the condition of Theorem 3.23 holds (see also Bunke [10] or Neckel and Rupp [55]).

□

Proof of Corollary 3.27. This follows directly from Corollary 3.26 and Proposition 3.24.

□

Proof of Proposition 4.3. We show that if \mathbf{X} is an equilibrating solution and $i \in \mathcal{I}$, then $\bar{X}_i^{(n_i)*} = (X_i^*, 0, \dots, 0)$ almost surely. For all $0 \leq k_i \leq n_i$ we have for almost all $\omega \in \Omega$

$$\lim_{t \rightarrow \infty} X_i^{(k_i)}(t, \omega) = X_i^{(k_i)*}(\omega).$$

Let $0 \leq m_i < n_i$. Let $\omega \in \Omega$ such that $\bar{X}_i^{(n_i)*}(t, \omega)$ converges. If $X_i^{(m_i+1)*}(\omega) > 0$, then there exists a $\bar{t} \in T$ such that $X_i^{(m_i+1)}(t, \omega) > \frac{1}{2} X_i^{(m_i+1)*}(\omega)$ for $t > \bar{t}$. From the mean value theorem, it follows that there exists a $c \in (\bar{t}, t)$ such that

$$X_i^{(m_i)}(t, \omega) - X_i^{(m_i)}(\bar{t}, \omega) = X_i^{(m_i+1)}(c, \omega)(t - \bar{t}) > \frac{1}{2} X_i^{(m_i+1)*}(\omega)(t - \bar{t})$$

and hence $X_i^{(m_i)}(t, \omega)$ cannot converge to $X_i^{(m_i)*}(\omega)$. We get a similar contradiction under the assumption $X_i^{(m_i+1)*}(\omega) < 0$, and hence $X_i^{(m_i+1)*}(\omega) = 0$. We conclude that $\bar{X}_i^{(n_i)*} = (X_i^*, 0, \dots, 0)$ almost surely. □

Proof of Proposition 4.4. We can rewrite the dynamic structural equations of \mathcal{R} as

$$\begin{cases} \mathbf{X}'_I = -B_{II'}^{-1}(B_{II} - \mathbb{I}_I) \mathbf{X}_I - B_{II'}^{-1} B_{IJ} \mathbf{X}_J - B_{II'}^{-1} \Gamma_{IJ} \mathbf{E} \\ \mathbf{X}'_J = -B_{JJ}^{-1} B_{JI} \mathbf{X}_I - B_{JJ}^{-1} \Gamma_{JJ} \mathbf{E}. \end{cases}$$

Eliminating \mathbf{X}_J from the right-hand side by substitution yields the RDE

$$\mathbf{X}'_I = A \mathbf{X}_I + C \mathbf{E},$$

where $A := B_{II'}^{-1}(B_{IJ} B_{JJ}^{-1} B_{JI} - B_{II} + \mathbb{I}_I)$ and $C := B_{II'}^{-1}(B_{IJ} B_{JJ}^{-1} \Gamma_{JJ} - \Gamma_{IJ})$. The matrix A is a Hurwitz matrix by assumption and thus invertible (note $\det(A) \neq 0$). The solutions of the ODE $\mathbf{x}' = A \mathbf{x} + C \mathbf{e}$, where the vector \mathbf{e} does not depend on time, are of the form $\mathbf{x} = \exp(At) \mathbf{x}_0 - A^{-1} C \mathbf{e}$, where \mathbf{x}_0 is some vector. For any matrix A there exists a nonsingular matrix P (possibly complex) that transforms A into its

Jordan normal form, that is, $P^{-1}AP = \Lambda$ is a block diagonal matrix where each block Λ_i is a Jordan block associated with the eigenvalue λ_i of A , and is a square matrix of order m_i of the form

$$\Lambda_i = \begin{bmatrix} \lambda_i & 1 & 0 & \cdots & \cdots & 0 \\ 0 & \lambda_i & 1 & 0 & \cdots & 0 \\ \vdots & & \ddots & \ddots & & \vdots \\ \vdots & & & \ddots & \ddots & 0 \\ \vdots & & & & \ddots & 1 \\ 0 & \cdots & \cdots & \cdots & 0 & \lambda_i \end{bmatrix}.$$

Therefore,

$$\begin{aligned} (\mathbf{X}_I)_t &= \exp(At)\mathbf{X}_{I,[0]} - A^{-1}C\mathbf{E}_t \\ &= \exp(P\Lambda P^{-1}t)\mathbf{X}_{I,[0]} - A^{-1}C\mathbf{E}_t \\ &= \sum_{i=1}^n \sum_{j=1}^{m_i} t^{j-1} \exp(\lambda_i t) R_{ij} \mathbf{X}_{I,[0]} - A^{-1}C\mathbf{E}_t \end{aligned}$$

with $\mathbf{X}_{I,[0]}$ some random variable, n the total number of block diagonal matrices, and the R_{ij} 's certain block matrices that depend on P and Λ [41]. Since A is a Hurwitz matrix by assumption and \mathbf{E} is constant in time, we conclude that for all solutions \mathbf{X} of \mathcal{R} ,

$$\lim_{t \rightarrow \infty} (\mathbf{X}_I)_t = -A^{-1}C\mathbf{E}$$

and

$$\lim_{t \rightarrow \infty} (\mathbf{X}_J)_t = (B_{JJ}^{-1}B_{JI}A^{-1}C - B_{JJ}^{-1}\Gamma_{JJ})\mathbf{E}$$

almost surely.

At last, we consider replacing the condition that the exogenous process \mathbf{E} is constant in time by the assumption that \mathbf{E} may depend on time and that both \mathbf{E}_t and $\exp(At) \int_{t_0}^t \exp(-As)C\mathbf{E}_s ds$ converge almost surely. Observe that the general solutions of $\mathbf{x}' = A\mathbf{x} + C\mathbf{e}$, where we allow \mathbf{e} to be a time-dependent vector, are of the form $\mathbf{x} = \exp(At)\mathbf{x}_0 + \exp(At) \int_{t_0}^t \exp(-As)C\mathbf{E}_s ds$. Then, replacing the term $-A^{-1}C\mathbf{E}_t$ in the equation above for $(\mathbf{X}_I)_t$ by $\exp(At) \int_{t_0}^t \exp(-As)C\mathbf{E}_s ds$ implies also that $(\mathbf{X}_I)_t$ converges a.s., from which the result follows, provided the smoothness requirements are met. \square

Proof of Lemma 4.7. Let \mathbf{X} be an equilibrating solution and let \mathbf{E} converge a.s. to the random variable \mathbf{E}^* . Then

$$\mathbf{X}^* = \lim_{t \rightarrow \infty} \mathbf{X}_t = \lim_{t \rightarrow \infty} \mathbf{f}(\bar{\mathbf{X}}_t^{(n)}, \mathbf{E}_t) = \mathbf{f}\left(\lim_{t \rightarrow \infty} \bar{\mathbf{X}}_t^{(n)}, \lim_{t \rightarrow \infty} \mathbf{E}_t\right) = \mathbf{f}(\bar{\mathbf{X}}^{(n)*}, \mathbf{E}^*)$$

almost surely, where in the third equality we used the continuity of \mathbf{f} . \square

Proof of Proposition 4.10. Consider the finite index set $\mathcal{J} = \{1, \dots, e\}$ for some $e \in \mathbb{N}$. The independence of $(\mathbf{E}_j)_{j \in \mathcal{J}}$ implies that, in particular, for every $t \in T$ the family of random variables $\tilde{\mathbf{E}} := ((\mathbf{E}_j)_t)_{j \in \mathcal{J}}$ is independent, that is, we have $\mathbb{P}^{\tilde{\mathbf{E}}_t} = \prod_{j \in \mathcal{J}} \mathbb{P}^{(\mathbf{E}_j)_t}$, where $\tilde{\mathbf{E}}_t := ((\mathbf{E}_1)_t, \dots, (\mathbf{E}_e)_t)$.

Because $\lim_{t \rightarrow \infty} \tilde{\mathbf{E}}_t = \lim_{n \in \mathbb{N}} \tilde{\mathbf{E}}_n$ a.s., we have $\lim_{n \rightarrow \infty} \tilde{\mathbf{E}}_n = \tilde{\mathbf{E}}^*$ a.s., where $\tilde{\mathbf{E}}^* := (\mathbf{E}_1^*, \dots, \mathbf{E}_e^*)$. This implies that $\tilde{\mathbf{E}}_n$ converges in distribution to $\tilde{\mathbf{E}}^*$ (see Remark 6.4 and Corollary 13.19 in [42]), that is, the distribution of $\tilde{\mathbf{E}}_n$ converges weakly to the distribution of $\tilde{\mathbf{E}}^*$, that is, $\text{w-lim}_{n \rightarrow \infty} \mathbb{P}^{\tilde{\mathbf{E}}_n} = \mathbb{P}^{\tilde{\mathbf{E}}^*}$.³¹ Similarly, we have $\text{w-lim}_{n \rightarrow \infty} \mathbb{P}^{(\mathbf{E}_j)_n} = \mathbb{P}^{\mathbf{E}_j^*}$ for every $j \in \mathcal{J}$. Applying Theorem 2.8 in [4] gives that

$$\mathbb{P}^{\tilde{\mathbf{E}}^*} = \text{w-lim}_{n \rightarrow \infty} \mathbb{P}^{\tilde{\mathbf{E}}_n} = \text{w-lim}_{n \rightarrow \infty} \prod_{j \in \mathcal{J}} \mathbb{P}^{(\mathbf{E}_j)_n} = \prod_{j \in \mathcal{J}} \mathbb{P}^{\mathbf{E}_j^*}.$$

³¹ Let $\mathbb{P}, \mathbb{P}_1, \mathbb{P}_2, \dots$ be probability distributions over \mathbb{R}^d , then \mathbb{P}_n converges weakly to \mathbb{P} , denoted by $\text{w-lim}_{n \rightarrow \infty} \mathbb{P}_n = \mathbb{P}$, if $\lim_{n \rightarrow \infty} \mathbb{P}_n(U) = \mathbb{P}(U)$ for all measurable sets U in \mathbb{R}^d with $\mathbb{P}(\partial U) = 0$, where ∂U is the boundary of U , that is, the closure of U minus the interior of U .

We conclude that the family of random variables $(\mathbf{E}_j^*)_{j \in \mathcal{J}}$ is independent. \square

Proof of Theorem 4.11. Let \mathbf{X} be an equilibrating solution and let \mathbf{E} converge a.s. to the random variable \mathbf{E}^* . From Lemma 4.7 it follows that

$$\mathbf{X}^* = \mathbf{f}(\bar{\mathbf{X}}^{(n)*}, \mathbf{E}^*) = \mathbf{f}(\bar{\mathbf{l}}(\mathbf{X}^*), \mathbf{E}^*) = \mathbf{f}^*(\mathbf{X}^*, \mathbf{E}^*) \quad \text{a.s.},$$

where we used in the second equality that $\bar{\mathbf{l}}(\mathbf{X}^*) = \bar{\mathbf{X}}^{(n)*}$, since for all $i \in \mathcal{I}$ we have that $\bar{X}_i^{(n_i)*}$ is a.s. equal to $(X_i^*, 0, \dots, 0)$ by Proposition 4.3. \square

Proof of Proposition 4.15. Suppose that the equilibrated SDCM $\mathcal{M}_{\mathcal{R}}$ has a solution \mathbf{X}^* . Then the stochastic process $\mathbf{X} : T \times \Omega \rightarrow \mathcal{X}$ defined by $\mathbf{X}_t(\omega) := \mathbf{X}^*(\omega)$ is a solution of \mathcal{R} that equilibrates to \mathbf{X}^* . \square

Proof of Proposition 4.16. By definition, the graph of the equilibrated model $\mathcal{M}_{\mathcal{R}}$ has nodes $\mathcal{I} \subseteq \bar{\mathcal{I}}^{(n)}$. For every $i \in \mathcal{I}$, a functional parent of i in $\mathcal{M}_{\mathcal{R}}$ is a functional parent in \mathcal{R} , since for all $\mathbf{e} \in \mathcal{E}$ and for all $\mathbf{x} \in \mathcal{X}$ we have

$$x_i = f_i^*(\mathbf{x}, \mathbf{e}) \implies x_i = f_i(\bar{\mathbf{l}}(\mathbf{x}), \mathbf{e}).$$

Note there are no integrated parents of i in $\mathcal{M}_{\mathcal{R}}$. \square

Proof of Theorem 4.18. This follows directly from Definitions 3.1, 3.7 and 4.9. One can easily check that

$$\begin{aligned} (\mathcal{M}_{\mathcal{R}})_{\text{do}(I, \mathbf{K}_I^*)} &= \langle \mathcal{I}, I \cup \mathcal{J}, \mathcal{X}, \mathcal{X}_I \times \mathcal{E}, \widetilde{\mathbf{f}}^*, (\mathbf{K}_I^*, \mathbf{E}^*) \rangle \\ &= \langle \mathcal{I}, I \cup \mathcal{J}, \mathcal{X}, \mathcal{X}_I \times \mathcal{E}, \widetilde{\mathbf{f}}^*, (\mathbf{K}_I, \mathbf{E})^* \rangle \\ &= \mathcal{M}_{\mathcal{R}_{\text{do}(I, \mathbf{K}_I)}} \end{aligned}$$

where the intervened and equilibrated dynamic causal mechanism

$$\widetilde{\mathbf{f}}^* = \widetilde{\mathbf{f}}^* : \mathcal{X} \times (\mathcal{X}_I \times \mathcal{E}) \rightarrow \mathcal{X}$$

is given by

$$\widetilde{f}_i^*(\mathbf{x}, (\mathbf{e}_I, \mathbf{e}_{\mathcal{J}})) := \begin{cases} f_i(\bar{\mathbf{l}}(\mathbf{x}), \mathbf{e}_{\mathcal{J}}) & i \in \mathcal{I} \setminus I \\ e_i & i \in I. \end{cases}$$

\square

Proof of Corollary 4.22. The statement follows immediately from Theorem 1 of Fisher and Fuller [27] followed by application of Proposition 4.4.

Theorem 1 of Fisher and Fuller [27] states that under the stated condition, there exists an invertible diagonal stabilization matrix $\Lambda \in \mathbb{R}^{d \times d}$ such that $-\Lambda^{-1}A$ is Hurwitz.³²

Note first that by construction, $\mathcal{M}_{\mathcal{R}_{\mathcal{M}, \Lambda}} = \mathcal{M}$. The SDCM $\mathcal{R}_{\mathcal{M}, \Lambda}$ satisfies Assumption 1- \mathcal{I} , that is, it can be written in the form of the equations in Proposition 4.4 with $I = \mathcal{I}$, where $B_{II'} = -\Lambda$ and $B_{II} = B$, and hence $B_{II'}^{-1}(-B_{II} + \mathbb{I}_I) = -\Lambda^{-1}A$ is Hurwitz. The statements now follow from Proposition 4.4. \square

Proof of Corollary 4.23. The statement follows from Theorem 2.1 of Locatelli and Schiavoni [47] followed by application of Proposition 4.4 and Theorems 4.11 and 4.18.

Theorem 2.1 of Locatelli and Schiavoni [47] states that for every matrix $A \in \mathbb{R}^{d \times d}$ that satisfies for all subsets $I \subseteq \mathcal{I}$ the condition $\det(A_{II}) \det(\text{diag}(A_{II})) > 0$, there exists a diagonal matrix $D \in \mathbb{R}^{d \times d}$ such

32 A simple counterexample of a system that cannot be stabilized in this way is given by taking the matrix

$$B = \begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix},$$

for which $\Lambda^{-1}(B - \mathbb{I})$ is not Hurwitz for any diagonal invertible matrix Λ .

that the matrix $D_{II}A_{II}$ is Hurwitz for all $I \subseteq \mathcal{I}$. In particular, observe that this matrix D is invertible, since D_{II} is invertible for every $I \subseteq \mathcal{I}$ (note $\det(D_{II}) \neq 0$ due to $\det(D_{II}A_{II}) \neq 0$).

Let $\Lambda \in \mathbb{R}^{d \times d}$ be an invertible diagonal matrix such that $-\Lambda_{II}^{-1}A_{II}$ is Hurwitz for every $I \subseteq \mathcal{I}$. Note first that by construction, $\mathcal{M}_{\mathcal{R}_{\mathcal{M},\Lambda}} = \mathcal{M}$. Now let $\text{do}(J, \mathbf{K}_J)$ be a stochastic perfect intervention for some subset $J \subseteq \mathcal{I}$ and \mathbf{K}_J some stochastic process that is constant in time. The intervened SDCM $(\mathcal{R}_{\mathcal{M},\Lambda})_{\text{do}(J,\mathbf{K}_J)}$ satisfies Assumption 1-I for $I := \mathcal{I} \setminus J$, that is, it can be written in the form of the equations in Proposition 4.4, where $B_{II'} = -\Lambda_{II}$, $B_{JJ} = -\mathbb{I}_{JJ}$, $B_{JI} = \mathbf{0}_{JI}$ the zero matrix and $\Gamma_{JJ}\mathbf{e} = \mathbf{K}_J$. Moreover,

$$B_{II'}^{-1}(B_{IJ}B_{JJ}^{-1}B_{JI} - B_{II} + \mathbb{I}_I) = -\Lambda_{II}^{-1}(\mathbb{I}_I - B_{II}) = -\Lambda_{II}^{-1}A_{II},$$

which is Hurwitz, from which we conclude that every solution \mathbf{X} of $(\mathcal{R}_{\mathcal{M},\Lambda})_{\text{do}(J,\mathbf{K}_J)}$ is an equilibrating solution. Hence, from Theorem 4.11 it follows that for every solution \mathbf{X} of $(\mathcal{R}_{\mathcal{M},\Lambda})_{\text{do}(J,\mathbf{K}_J)}$, its limit \mathbf{X}^* is a solution of the equilibrated model

$$\mathcal{M}_{((\mathcal{R}_{\mathcal{M},\Lambda})_{\text{do}(J,\mathbf{K}_J)})} = (\mathcal{M}_{\mathcal{R}_{\mathcal{M},\Lambda}})_{\text{do}(J,\mathbf{K}_J)} = \mathcal{M}_{\text{do}(J,\mathbf{K}_J)}$$

where we made use of Theorem 4.18. Note that \mathbf{E} is assumed constant (in time), and hence $\mathcal{R}_{\mathcal{M},\Lambda}$ is steady; in addition, \mathbf{K}_J is assumed to be constant. The solutions of $\mathcal{M}_{\text{do}(J,\mathbf{K}_J)}$ are a.s. unique, because they satisfy the equations $\mathbf{X}_I^* = A_{II}^{-1}(B_{IJ}\mathbf{X}_J^* + \Gamma_{IJ}\mathbf{E})$ and $\mathbf{X}_J^* = \mathbf{K}_J$ almost surely. \square

References

- [1] O. O. Aalen. Dynamic modelling and causality. *Scandinavian Actuarial Journal*, 1987(3–4):177–190, 1987.
- [2] U. M. Ascher and L. R. Petzold. *Computer Methods for Ordinary Differential Equations and Differential-Algebraic Equations*. Society for Industrial and Applied Mathematics, Philadelphia, USA, 1st edition, 1998.
- [3] S. Bauer, N. S. Gorbach, D. Miladinovic, and J. M. Buhmann. Efficient and flexible inference for stochastic systems. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 6988–6998. Curran Associates, Inc., 2017. URL <http://papers.nips.cc/paper/7274-efficient-and-flexible-inference-for-stochastic-systems.pdf>.
- [4] P. Billingsley. *Convergence of Probability Measures*. Wiley Series in Probability and Statistics: Probability and Statistics. John Wiley & Sons Inc., New York, USA, 2nd edition, 1999.
- [5] T. Blom and J. M. Mooij. Causality and independence in perfectly adapted dynamical systems. *arXiv.org preprint*, arXiv:2101.11885 [cs.AI], 2021. URL <https://arxiv.org/abs/2101.11885>.
- [6] T. Blom, S. Bongers, and J. M. Mooij. Beyond structural causal models: Causal constraints models. In R. P. Adams and V. Gogate, editors, *Proceedings of the 35th Conference on Uncertainty in Artificial Intelligence (UAI-19)*. AUAI Press, July 2019. URL <http://auai.org/uai2019/proceedings/papers/205.pdf>.
- [7] T. Blom, M. M. van Diepen, and J. M. Mooij. Conditional independences and causal relations implied by sets of equations. *Journal of Machine Learning Research*, 22(178):1–62, 2021. URL <http://jmlr.org/papers/v22/20-863.html>.
- [8] K. Bollen. *Structural Equations with Latent Variables*. John Wiley & Sons, New York, USA, 1989.
- [9] S. Bongers, P. Forré, J. Peters, and J. M. Mooij. Foundations of structural causal models with cycles and latent variables. *Annals of Statistics*, 49(5):2885–2915, 2021.
- [10] H. Bunke. *Gewöhnliche Differentialgleichungen mit Zufälligen parametern*. Akademie-Verlag, Berlin, DE, 1972.
- [11] E. A. Coddington and N. Levinson. *Theory of Ordinary Differential Equations*. International series in pure and applied mathematics. McGraw-Hill, New York, USA, 1955.
- [12] D. Commenges and A. Gégout-Petit. A general dynamical statistical model with causal interpretation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 71(3):719–736, 2009.
- [13] G. F. Cooper. A simple constraint-based algorithm for efficiently mining observational databases for causal relationships. *Data Mining and Knowledge Discovery*, 1(2):203–224, 1997.
- [14] P. Dagum, A. Galper, and E. Horvitz. Dynamic network models for forecasting. *Uncertainty in Artificial Intelligence*, pages 41–48, 1992.
- [15] D. Dash. Restructuring dynamic causal systems in equilibrium. In R. Cowell and Z. Ghahramani, editors, *Proceedings of the Tenth International Workshop on Artificial Intelligence and Statistics*, pages 81–88. The Society for Artificial Intelligence and Statistics, Jan. 2005. URL <http://www.gatsby.ucl.ac.uk/aistats/fullpapers/264.pdf>.
- [16] V. Didelez. *Graphical models for event history analysis based on local independence*. PhD thesis, Universität Dortmund, 2000.
- [17] V. Didelez. Graphical models for composable finite Markov processes. *Scandinavian Journal of Statistics*, 34(1):169–185, 2007.

- [18] V. Didelez. Graphical models for marked point processes based on local independence. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(1):245–264, 01 2008.
- [19] V. Didelez. Causal reasoning for events in continuous time: A decision–theoretic approach. In R. Silva, I. Shpitser, R. J. Evans, J. Peters, and T. Claassen, editors, *Proceedings of the UAI 2015 Workshop on Advances in Causal Inference (UAI15-ACI)*, volume 1504 of *CEUR Workshop Proceedings*, pages 40–45. CEUR-WS.org, 2015. URL http://ceur-ws.org/Vol-1504/uai2015aci_paper3.pdf.
- [20] O. Duncan. *Introduction to Structural Equation Models*. Academic Press, New York, 1975.
- [21] F. Eberhardt. Direct causes and the trouble with soft interventions. *Erkenntnis*, 79(4):755–777, Aug. 2014.
- [22] F. Eberhardt and R. Scheines. Interventions and causal inference. *Philosophy of Science*, 74(5):981–995, 2007.
- [23] M. Eichler. Granger causality and path diagrams for multivariate time series. *Journal of Econometrics*, 137(2):334–353, 2007.
- [24] M. Eichler and V. Didelez. Causal reasoning in graphical time series models. In R. Parr and L. van der Gaag, editors, *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence (UAI-07)*, pages 109–116, Corvallis, Oregon, 07 2007. AUAI Press. URL <https://dslpitt.org/papers/07/p109-eichler.pdf>.
- [25] F. M. Fisher. A correspondence principle for simultaneous equation models. *Econometrica*, 38(1):73–92, 1970.
- [26] F. M. Fisher. A simple proof of the Fisher-Fuller theorem. *Mathematical Proceedings of the Cambridge Philosophical Society*, 71(3):523–525, 1972.
- [27] M. E. Fisher and A. Fuller. On the stabilization of matrices and the convergence of linear iterative processes. *Mathematical Proceedings of the Cambridge Philosophical Society*, 54:417–425, 1958.
- [28] J.-P. Florens and D. Fougere. Noncausality in continuous time. *Econometrica*, 64(5):1195–1212, 1996.
- [29] P. Forré and J. M. Mooij. Causal calculus in the presence of cycles, latent confounders and selection bias. In R. P. Adams and V. Gogate, editors, *Proceedings of the 35th Conference on Uncertainty in Artificial Intelligence (UAI-19)*. AUAI Press, July 2019. URL <http://auai.org/uai2019/proceedings/papers/15.pdf>.
- [30] K. Friston, L. Harrison, and W. Penny. Dynamic causal modelling. *NeuroImage*, 19(4):1273–1302, 2003.
- [31] Z. Ghahramani. Learning dynamic Bayesian networks. In G. M. Giles C.L., editor, *Adaptive Processing of Sequences and Data Structures*, volume 1387 of *NN 1997. Lecture Notes in Computer Science*, pages 168–197, Berlin, Heidelberg, 1998. Springer.
- [32] A. Goldberger and O. Duncan. *Structural Equation Models in the Social Sciences*. Seminar Press, New York, 1973.
- [33] C. W. J. Granger. Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 37(3):424–438, 1969.
- [34] T. Haavelmo. The statistical implications of a system of simultaneous equations. *Econometrica*, 11(1):1–12, 1943.
- [35] J. Halpern. Axiomatizing causal reasoning. In G. Cooper and S. Moral, editors, *Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence (UAI-98)*, pages 202–210, San Francisco, CA, USA, 7 1998. Morgan Kaufmann. URL <https://dslpitt.org/uai/papers/98/p202-halpern.pdf>.
- [36] X. Han and P. E. Kloeden. *Random Ordinary Differential Equations and Their Numerical Solution*, volume 85 of *Probability Theory and Stochastic Modelling*. Springer, Singapore, 2017.
- [37] N. Hansen and A. Sokol. Causal interpretation of stochastic differential equations. *Electronic Journal of Probability*, 19, 2014.
- [38] Y. Iwasaki and H. A. Simon. Causality and model abstraction. *Artificial Intelligence*, 67(1):143–194, 1994.
- [39] D. J. J. Peters and B. Schölkopf. Causal inference on time series using restricted structural equation models. In *Advances in Neural Information Processing Systems 26*, pages 154–162, 2013. URL <https://papers.nips.cc/paper/5063-causal-inference-on-time-series-using-restricted-structural-equation-models.pdf>.
- [40] A. Jentzen and P. E. Kloeden. *Taylor Approximations for Stochastic Partial Differential Equations*, volume 83 of *PCBMS-NSF regional conference series in applied mathematics*. Society for Industrial and Applied Mathematics, Philadelphia, USA, 2011. URL <https://epubs.siam.org/doi/abs/10.1137/1.9781611972016>.
- [41] H. K. Khalil. *Nonlinear systems*. Prentice Hall, Upper Saddle River, NJ, USA, 2nd edition, 1996.
- [42] A. Klenke. *Probability Theory*. Universitext. Springer, London, GB, 2nd edition, 2014.
- [43] P. Kloeden and E. Platen. *Numerical Solution of Stochastic Differential Equations*, volume 23 of *Stochastic Modelling and Applied Probability*. Springer, Berlin, DE, 1992.
- [44] K. B. Korb, L. R. Hope, A. E. Nicholson, and K. Axnick. Varieties of causal intervention. In C. Zhang, H. W. Guesgen, and W.-K. Yeap, editors, *PRICAI 2004: Trends in Artificial Intelligence*, pages 322–331, Berlin, DE, 2004. Springer. ISBN 978–540-28633-2.
- [45] G. Lacerda, P. L. Spirtes, J. Ramsey, and P. O. Hoyer. Discovering cyclic causal models by independent components analysis. In D. McAllester and P. Myllymaki, editors, *Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence (UAI-08)*, pages 366–374, Corvallis, Oregon, USA, July 2008. AUAI Press. URL <https://dslpitt.org/papers/08/p366-lacerda.pdf>.
- [46] J. Liu, Z. Long, R. Wang, J. Sun, and B. Dong. RODE-Net: Learning ordinary differential equations with randomness from data. *arXiv.org preprint*, arXiv:2006.02377 [math.NA], 2020. URL <https://arxiv.org/abs/2006.02377>.
- [47] A. Locatelli and N. Schiavoni. A necessary and sufficient condition for the stabilisation of a matrix and its principal submatrices. *Linear Algebra and its Applications*, 436(7):2311–2314, 2012.

- [48] H. Lütkepohl. *New Introduction to Multiple Time Series Analysis*. Springer, Berlin, Heidelberg, 2005.
- [49] S. Mani. *A Bayesian Local Causal Discovery Framework*. PhD thesis, University of Pittsburg, March 2006. URL <http://d-scholarship.pitt.edu/10181/>.
- [50] S. W. Mogenssen, D. Malinsky, and N. R. Hansen. Causal learning for partially observed stochastic dynamical systems. In A. Globerson and R. Silva, editors, *Proceedings of the Thirty-Fourth conference on Uncertainty in Artificial Intelligence (UAI-18)*. AUAI Press, Aug. 2018. URL <http://auai.org/uai2018/proceedings/papers/142.pdf>.
- [51] J. M. Mooij and T. Claassen. Constraint-based causal discovery using partial ancestral graphs in the presence of cycles. In J. Peters and D. Sontag, editors, *Proceedings of the 36th Conference on Uncertainty in Artificial Intelligence (UAI-20)*, volume 124, pages 1159–1168. PMLR, 8 2020. URL <http://proceedings.mlr.press/v124/m-mooij20a/m-mooij20a.pdf>.
- [52] J. M. Mooij, D. Janzing, and B. Schölkopf. From ordinary differential equations to structural causal models: the deterministic case. In A. Nicholson and P. Smyth, editors, *Proceedings of the 29th Conference on Uncertainty in Artificial Intelligence (UAI-13)*, pages 440–448, Corvallis, Oregon, USA, Aug. 2013. AUAI Press. URL <http://auai.org/uai2013/prints/papers/24.pdf>.
- [53] J. M. Mooij, S. Magliacane, and T. Claassen. Joint causal inference from multiple contexts. *Journal of Machine Learning Research*, 21(99):1–108, 2020. URL <http://jmlr.org/papers/v21/17-123.html>.
- [54] R. Neal. On deducing conditional independence from d -separation in causal graphs with feedback. *Journal of Artificial Intelligence Research*, 12:87–91, 2000. URL <https://www.jair.org/index.php/jair/article/download/10250/24409>.
- [55] T. Neckel and F. Rupp. *Random Differential Equations in Scientific Computing*. Versita, De Gruyter publishing group, Warsaw, Dec. 2013.
- [56] J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, New York, USA, 2nd edition, 2009.
- [57] J. Peters, D. Janzing, and B. Schölkopf. *Elements of Causal Inference: Foundations and Learning Algorithms*. MIT Press, Cambridge, MA, USA, 2017.
- [58] J. Peters, S. Bauer, and N. Pfister. Causal models for dynamical systems. *arXiv.org preprint*, arXiv:2001.06208 [stat.ME], 2020. URL <https://arxiv.org/abs/2001.06208>.
- [59] N. Pfister, S. Bauer, and J. Peters. Learning stable and predictive structures in kinetic systems. *Proceedings of the National Academy of Sciences*, 116(51):25405–25411, 2019.
- [60] T. S. Richardson and J. M. Robins. ACE bounds; SEMs with equilibrium conditions. *Statistical Science*, 29(3):363–366, 2014.
- [61] P. K. Rubenstein, S. Bongers, B. Schölkopf, and J. M. Mooij. From deterministic ODEs to dynamic structural causal models. In A. Globerson and R. Silva, editors, *Proceedings of the 34th Conference on Uncertainty in Artificial Intelligence (UAI-18)*. AUAI Press, Aug. 2018. URL <http://auai.org/uai2018/proceedings/papers/43.pdf>.
- [62] T. Schweder. Composable Markov processes. *Journal of Applied Probability*, 7(2):400–410, 1970.
- [63] C. Sims. Macroeconomics and reality. *Econometrica*, 48(1):1–48, 1980.
- [64] K. Sobczyk. *Stochastic Differential Equations*, volume 40 of *Mathematics and its Applications*. Springer, Dordrecht, NL, 1st edition, 1991.
- [65] T. Soong. *Random Differential Equations in Science and Engineering*, volume 103 of *Mathematics in Science and Engineering*. Academic Press, New York, USA, 1973.
- [66] P. Spirtes, C. Meek, and T. S. Richardson. An algorithm for causal inference in the presence of latent variables and selection bias. In C. Glymour and G. F. Cooper, editors, *Computation, Causation and Discovery*, chapter 6, pages 211–252. The MIT Press, 1999.
- [67] P. Spirtes, C. Glymour, and R. Scheines. *Causation, Prediction, and Search*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, Massachusetts, 2nd edition, 2000. ISBN 0-262-19440-6.
- [68] M. Voortman, D. Dash, and M. J. Druzdzel. Learning why things change: The difference-based causality learner. In P. Grunwald and P. Spirtes, editors, *Proceedings of the Twenty-Sixth Conference on Uncertainty in Artificial Intelligence (UAI-10)*, pages 641–650, Corvallis, Oregon, USA, July 2010. AUAI Press. URL <https://dslpitt.org/papers/10/p641-voortman.pdf>.
- [69] H. White. Time-series estimation of the effects of natural experiments. *Journal of Econometrics*, 135(1):527–566, 2006.
- [70] J. Woodward. *Making Things Happen: A Theory of Causal Explanation*. Oxford Studies In Philosophy Of Science. Oxford University Press, New York, 2003.
- [71] S. Wright. Correlation and causation. *Journal of Agricultural Research*, 20:557–585, 1921.
- [72] J. Zhang. On the completeness of orientation rules for causal discovery in the presence of latent confounders and selection bias. *Artificial Intelligence*, 172(16-17):1873–1896, 2008.