
Beyond Structural Causal Models: Causal Constraints Models

Tineke Blom
Informatics Institute
University of Amsterdam
The Netherlands

Stephan Bongers
Informatics Institute
University of Amsterdam
The Netherlands

Joris M. Mooij
Informatics Institute
University of Amsterdam
The Netherlands

Abstract

Structural Causal Models (SCMs) provide a popular causal modeling framework. In this work, we show that SCMs are not flexible enough to give a complete causal representation of dynamical systems at equilibrium. Instead, we propose a generalization of the notion of an SCM, that we call Causal Constraints Model (CCM), and prove that CCMs do capture the causal semantics of such systems. We show how CCMs can be constructed from differential equations and initial conditions and we illustrate our ideas further on a simple but ubiquitous (bio)chemical reaction. Our framework also allows to model functional laws, such as the ideal gas law, in a sensible and intuitive way.

1 INTRODUCTION

Real-world processes are often complex and time-evolving. The dynamics of such systems can be modeled by (random) differential equations, which offer a fine-grained description of how the variables in the system change over time. A coarser but more tractable approach are Structural Causal Models (SCMs), which provide a modeling framework that is used in many fields such as biology, the social sciences, and economy (Pearl, 2000). Although SCMs have been successfully applied to certain static systems, a pressing concern is whether SCMs are able to completely model the causal semantics of the stationary behavior of a dynamical system. In this work, we prove that generally SCMs are not flexible enough to completely model dynamical systems at equilibrium.

We generalize the notion of SCMs and introduce a novel type of causal model, that we call Causal Constraints Models (CCMs). We prove that they give a complete de-

scription of the causal semantics of dynamical systems at equilibrium and show how a CCM can be derived from differential equations and initial conditions. We further motivate our approach by pointing out that CCMs, contrary to SCMs, correctly describe the causal semantics of functional laws (e.g. the ideal gas law), which describe relations between variables that are invariant under all interventions. We illustrate the benefits of CCMs on a simple but ubiquitous (bio)chemical reaction.

Causal models that arise from studying the behavior of dynamical systems have received much attention over the years. Fisher (1970), Mogensen et al. (2018), Rubenstein et al. (2018), Sokol and Hansen (2014), and Voortman et al. (2010) consider causal relations in systems that can be modelled by (stochastic) differential equations that are not in equilibrium. In contrast, we consider the stationary behaviour of dynamical systems, which does not require us to model the system's dependence on time. Bongers and Mooij (2018), Hyttinen et al. (2012), Lacerda et al. (2008), Mooij et al. (2011), and Mooij et al. (2013) show how cyclic SCMs may arise from studying the stationary behavior of certain dynamical time-series or differential equations, and how in some cases cyclic SCMs can be learned from equilibrium data. SCMs are well-understood and have recently been extended to also include the cyclic case (Bongers et al., 2018; Forré and Mooij, 2017). The drawback of the extension in Forré and Mooij (2017), with respect to modeling equilibria of dynamical systems, is that it requires the model to have a globally compatible solution under any intervention, which dynamical systems do not, in general, possess. Another modeling approach for dynamical systems at equilibrium is to construct a, possibly cyclic, SCM from the differential equations as Mooij et al. (2013) and Bongers and Mooij (2018) do. In this work, we show that these approaches to model the causal semantics of the stationary behavior in dynamical systems cannot accommodate the dependence of the equilibria on the initial conditions of the system.

In previous work, researchers have come across subtleties regarding the relation between the causal semantics and conditional independence properties of dynamical systems at equilibrium (Dash, 2005; Iwasaki and Simon, 1994; Lacerda et al., 2008). Previously, researchers have made additional assumptions about the underlying dynamical system to circumvent these. Although Rubenstein et al. (2017) and Bongers and Mooij (2018) do not make such restrictions, the price that one pays is that either one must limit the interventions that can be modeled or the equilibrium is no longer uniquely specified and one is limited to modeling the fixed points of the system. To the best of our knowledge, Causal Constraints Models are the first models that can completely capture the causal semantics of the stationary behavior of dynamical systems in general.

A disadvantage of CCMs over SCMs is that they do not yet possess the intuitive graphical interpretation that SCMs have. We consider representations of the independence structure of CCMs outside the scope of this work.

1.1 STRUCTURAL CAUSAL MODELS

A statistical model over random variables, taking value in a measurable space \mathcal{X} , usually is a pair $(\mathcal{X}, \mathbb{P}^{\mathcal{X}})$ where $\mathbb{P}^{\mathcal{X}}$ is a (parametrized) family of probability distributions on \mathcal{X} . A causal model on the other hand, can be thought of as a family of statistical models, one for each (perfect) intervention,

$$\bar{\mathbb{P}}^{\mathcal{X}} = \left(\mathbb{P}_{\text{do}(I, \xi_I)}^{\mathcal{X}} : I \in \mathcal{P}(\mathcal{I}), \xi_I \in \mathcal{X}_I \right), \quad (1)$$

where \mathcal{I} is an index set and $\mathcal{P}(\mathcal{I})$ denotes the power set of \mathcal{I} (i.e. the set of all subsets of \mathcal{I}). I represents the intervention target and ξ_I a tuple of intervention values. The null intervention $\text{do}(\emptyset)$ for $I = \emptyset$ corresponds to the observed system.

SCMs are a special type of causal models that are specified by structural equations. Our formal treatment of SCMs mostly follows Bongers et al. (2018) and Pearl (2000). For the purposes of this paper, we deviate from the usual definition of SCMs by not assuming independence of exogenous variables and by not requiring acyclicity (i.e. recursiveness).

Definition 1. Let \mathcal{I} and \mathcal{J} be index sets. A *Structural Causal Model (SCM)* \mathcal{M} is a triple $(\mathcal{X}, F, \mathbf{E})$, with:

- a product of standard measurable spaces $\mathcal{X} = \prod_{i \in \mathcal{I}} \mathcal{X}_i$ (domains of endogenous variables),
- a tuple of exogenous random variables $\mathbf{E} = (E_j)_{j \in \mathcal{J}}$ taking value in a product of standard measurable spaces $\mathcal{E} = \prod_{j \in \mathcal{J}} \mathcal{E}_j$,

- a family F of measurable functions:¹

$$f_i : \mathcal{X}_{\text{pa}(i) \cap \mathcal{I}} \times \mathcal{E}_{\text{pa}(i) \cap \mathcal{J}} \rightarrow \mathcal{X}_i, \quad \forall i \in \mathcal{I}.$$

Note that a cyclic structural causal model does not need to imply a unique joint distribution $\mathbb{P}_{\text{do}(\emptyset)}^{\mathcal{X}}$ on the space of endogenous variables in the observed system, although acyclic SCMs do (Bongers et al., 2018). When there exists a unique solution $\mathbf{x} \in \mathcal{X}$ to the *structural equations*

$$x_i = f_i(\mathbf{x}_{\text{pa}(i) \cap \mathcal{I}}, \mathbf{e}_{\text{pa}(i) \cap \mathcal{J}}), \quad \forall i \in \mathcal{I}$$

for almost all $\mathbf{e} \in \mathcal{E}$, we say that the model is uniquely solvable.

Definition 2. We say that a random variable $\mathbf{X} = (X_i)_{i \in \mathcal{I}}$ is a *solution* to an SCM $\mathcal{M} = (\mathcal{X}, F, \mathbf{E})$ if

$$X_i = f_i(\mathbf{X}_{\text{pa}(i) \cap \mathcal{I}}, \mathbf{E}_{\text{pa}(i) \cap \mathcal{J}}) \quad \text{a.s.}, \quad \forall i \in \mathcal{I}.$$

An SCM may have a unique (up to zero sets) solution, multiple solutions, or there may not exist any solution at all.

There are many types of interventions, corresponding to different experimental procedures, that can be modeled in an SCM. For the remainder of this work, we consider *perfect* (also known as “surgical” or “atomic”) interventions that force variables to take on a specific value through some external force acting on the system.

Definition 3. A *perfect intervention* $\text{do}(I, \xi_I)$ with target $I \subseteq \mathcal{I}$ and value $\xi_I \in \mathcal{X}_I$ on an SCM $\mathcal{M} = (\mathcal{X}, F, \mathbf{E})$ maps it to the *intervened* SCM $\mathcal{M}_{\text{do}(I, \xi_I)} = (\mathcal{X}, \tilde{F}, \mathbf{E})$ with \tilde{F} the family of measurable functions:

$$\begin{aligned} & \tilde{f}_i(\mathbf{x}_{\text{pa}(i) \cap \mathcal{I}}, \mathbf{e}_{\text{pa}(i) \cap \mathcal{J}}) \\ &= \begin{cases} \xi_i & i \in I, \\ f_i(\mathbf{x}_{\text{pa}(i) \cap \mathcal{I}}, \mathbf{e}_{\text{pa}(i) \cap \mathcal{J}}) & i \in \mathcal{I} \setminus I. \end{cases} \end{aligned}$$

Note that the solvability of an SCM may change after a perfect intervention, e.g. a uniquely solvable SCM may no longer be so after certain interventions.

1.2 DYNAMICAL SYSTEMS

We consider dynamical systems \mathcal{D} describing $p = |\mathcal{I}|$ (random) variables $\mathbf{X}(t)$ taking value in $\mathcal{X} = \mathbb{R}^p$. They consist of a set of coupled first-order ordinary differential equations (ODEs) where the initial conditions $\mathbf{X}(0)$ are determined by exogenous random variables $\mathbf{E} = (E_i)_{i \in \mathcal{I}}$ taking value in $\mathcal{E} = \mathbb{R}^p$. That is,

$$\begin{aligned} \dot{X}_i(t) &= f_i(\mathbf{X}(t)), & \forall i \in \mathcal{I}, \\ X_i(0) &= E_i, & \forall i \in \mathcal{I}, \end{aligned}$$

¹ $\text{pa}(i) \subseteq \mathcal{I} \cup \mathcal{J}$ denotes a subset of indexes that are sufficient to determine the values of f_i .

where the f_i are locally Lipschitz continuous functions.² Throughout this paper, we will assume for any dynamical system we encounter that for \mathbb{P}^E -almost every $e \in \mathbb{R}^p$ the initial value problem with $\mathbf{X}(0) = e$ has a unique solution $\mathbf{X}(t, e)$ for all $t \geq 0$, given by

$$\mathbf{X}(t, e) = \mathbf{X}(0, e) + \int_0^t \mathbf{f}(\mathbf{X}(s, e)) ds. \quad (2)$$

This solution $\mathbf{X}(t, e)$ can be trivially extended to \mathcal{E} and it is measurable in e for all t (Han and Kloeden, 2017).

A *fixed point* (or equilibrium point) of \mathcal{D} is a point $\mathbf{x}^* \in \mathbb{R}^p$ for which $\mathbf{f}(\mathbf{x}^*) = \mathbf{0}$. For $e \in \mathbb{R}^p$, the dynamical system converges to an equilibrium $\mathbf{X}^*(e) \in \mathbb{R}^p$ if

$$\lim_{t \rightarrow \infty} \mathbf{X}(t, e) = \mathbf{X}^*(e). \quad (3)$$

If for \mathbb{P}^E -almost every e the limit in equation (3) exists, then we say that \mathcal{D} converges to the equilibrium solution $\mathbf{X}^* = \lim_{t \rightarrow \infty} \mathbf{X}(t, \mathbf{E})$.

Interventions on dynamical systems can be modeled in different ways. One could for example fix the value of targeted values at one time-point. Alternatively, one could fix the trajectory of the targeted values as in Rubenstein et al. (2018). Here, we follow Mooij et al. (2013) and define interventions as operations that fix the value of the targeted variables to a constant (for all time).

Definition 4. A *perfect intervention* $\text{do}(I, \xi_I)$ where $I \subseteq \mathcal{I}$ and $\xi_I \in \mathcal{X}_I$ results in the *intervened* dynamical system $\mathcal{D}_{\text{do}(I, \xi_I)}$ specified by

$$\begin{aligned} \dot{X}_i(t) &= 0, & X_i(0) &= \xi_i, & \forall i \in I, \\ \dot{X}_i(t) &= f_i(\mathbf{X}(t)), & X_i(0) &= E_i, & \forall i \in \mathcal{I} \setminus I. \end{aligned}$$

We say that a causal model \mathcal{M} *completely* captures the causal semantics of the stationary behaviour of a dynamical system \mathcal{D} if for all $I \subseteq \mathcal{I}$ and all $\xi_I \in \mathcal{X}_I$: the equilibrium solutions of $\mathcal{D}_{\text{do}(I, \xi_I)}$ coincide with the solutions of $\mathcal{M}_{\text{do}(I, \xi_I)}$ (up to \mathbb{P}^E -null sets).

The construction of SCMs from dynamical systems in Mooij et al. (2013) relies on the fact that for systems that converge to a fixed point independent of initial conditions (i.e. globally asymptotically stable systems), the fixed point directly gives a complete description of its stationary behavior. A much weaker stability assumption is (*global*) *semistability* (Bhat and Bernstein, 1999; Campbell and Rose, 1979), where solutions of a system converge to a stable equilibrium determined by initial conditions. Our definition follows Haddad et al. (2010).

²If the dynamics depends on (random) parameters, they can be modeled as additional endogenous variables with vanishing time derivatives and initial conditions corresponding to the (random) parameters. Therefore, without loss of generality, we may assume that the functions f_i only depend on \mathbf{X} .

Definition 5. Let \mathcal{D} be a dynamical system and $\mathcal{U} \subseteq \mathbb{R}^p$ an invariant subset (i.e. if $\mathbf{x}(0) \in \mathcal{U}$ then $\mathbf{x}(t) \in \mathcal{U}$ for all $t \geq 0$). A fixed point $\mathbf{x}^* \in \mathcal{U}$ is Lyapunov stable with respect to \mathcal{U} if for all $\mathbf{x}(0) \in \mathcal{U}$: for all $\epsilon > 0$ there exists $\delta > 0$ such that if $\|\mathbf{x}(0) - \mathbf{x}^*\| < \delta$ then for all $t \geq 0$, $\|\mathbf{x}(t) - \mathbf{x}^*\| < \epsilon$. It is *semistable* w.r.t. \mathcal{U} if, additionally, there exists a relatively open subset³ \mathcal{N} of \mathcal{U} that contains \mathbf{x}^* such that $\mathbf{x}(t)$ converges to a Lyapunov stable fixed point for all $\mathbf{x}(0) \in \mathcal{N}$. If $\mathcal{N} = \mathcal{U}$ then \mathbf{x}^* is globally semistable w.r.t. \mathcal{U} . Finally, we say that \mathcal{D} is globally semistable w.r.t. \mathcal{U} if all its fixed points are globally semistable w.r.t. \mathcal{U} .

Definition 6. A dynamical system \mathcal{D} is *structurally semistable* if for all $I \subseteq \mathcal{I}$ there exists $\mathcal{U} \subseteq \mathbb{R}^p$ with $\mathbb{P}^{E_{\mathcal{I} \setminus I}}(\mathcal{U}_{\mathcal{I} \setminus I}) = 1$ such that: $\mathcal{D}_{\text{do}(I, \xi_I)}$ is globally semistable w.r.t. \mathcal{U} (for any $\xi_I \in \mathcal{X}_I$).


Whether a dynamical system converges to a certain fixed point depends on initial conditions. This dependence can often be described by constants of motion, and there exists a vast literature on how and when these can be derived from differential equations. The notion of semistability is appropriate in many real-world applications in chemical kinetics, environmental, and economic systems (Haddad et al., 2010). For chemical reaction networks, there exist convenient criteria on the network structure that guarantee global semistability (Chellaboina et al., 2009), and for mechanical systems semistability characterizes the motion of rigid bodies subject to damping (Bhat and Bernstein, 1999).

2 DYNAMICAL SYSTEMS AS SCMs

We consider SCM representations of the equilibria in a chemical reaction and conclude that, generally, SCMs are not flexible enough to completely capture the causal semantics of stationary behaviour in dynamical systems.

2.1 BASIC ENZYME REACTION

The basic enzyme reaction is a well-known example of a (bio)chemical reaction network. It describes a system where a substrate S reacts with an enzyme E to form a complex C which is then converted into a product P and the enzyme (Murray, 2002). In the open enzyme reaction a constant influx of substrate and an efflux of product are added (Belgacem and Gouzé, 2012). The process can be presented by the following reaction graph,



³ \mathcal{N} is a relatively open subset of \mathcal{U} if there is an open set $\mathcal{N}' \subseteq \mathbb{R}^p$ such that $\mathcal{N} = \mathcal{N}' \cap \mathcal{U}$.

and $\mathbf{k} = [k_0, k_{-1}, k_1, k_2, k_3]$ strictly positive parameters.

Differential equations for the concentrations of each molecule in the system can be obtained by application of the law of mass-action, which states that the rate of a reaction is proportional to the product of the concentration of the reactants (Murray, 2002), yielding:

$$\dot{S}(t) = k_0 - k_1 S(t) E(t) + k_{-1} C(t), \quad (4)$$

$$\dot{E}(t) = -k_1 S(t) E(t) + (k_{-1} + k_2) C(t), \quad (5)$$

$$\dot{C}(t) = k_1 S(t) E(t) - (k_{-1} + k_2) C(t), \quad (6)$$

$$\dot{P}(t) = k_2 C(t) - k_3 P(t), \quad (7)$$

$$(S(0), E(0), C(0), P(0)) = (s_0, e_0, c_0, p_0). \quad (8)$$

We simulated the system in (4) to (7) with random initial conditions and also under interventions on S and E . Figures 1a to 1c show how the time-trajectories of the concentrations depend on initial conditions in different interventional settings.

2.1.1 EQUILIBRIUM SOLUTIONS

By explicit calculation one can verify that given strictly positive initial conditions, the dynamical system converges to an equilibrium (S^*, C^*, E^*, P^*) if it exists, for any perfect intervention (one can also check that the system is structurally semistable).⁴ The equilibria can be found by deriving constraints on solutions of the system:

- At equilibrium the system is at rest and all time derivatives (in the equations of motion) must vanish. The equation of motion of each variable then results in a constraint that is invariant under all interventions that do not target that variable. For example, equation (4) yields the equilibrium equation

$$k_0 - k_1 S^* E^* + k_{-1} C^* = 0,$$

which constrains the equilibrium state unless S is targeted by an intervention.

- Symmetries or (linear) dependencies between the time derivatives lead to conservation laws (i.e. constants of motion), which are relations between variables that are time-invariant but that are typically invariant under fewer interventions than constraints of the first type. For example, since $\dot{C}(t) + \dot{E}(t) = 0$ for all t , we have that

$$C(t) + E(t) = c_0 + e_0, \quad \forall t, \quad (9)$$

unless C , E or both C and E are targeted by an intervention.

⁴See Belgacem and Gouzé (2012) and supplementary material for details.

- A system may contain (derived) variables whose time-derivative does not depend on itself. Since

$$\dot{S}(t) - \dot{E}(t) = k_0 - k_2 C(t), \quad (10)$$

the variable C cannot be ‘freely manipulated’, in the sense that $S(t) - E(t)$ does not converge to equilibrium under interventions $\text{do}(C = \xi_C)$ when $\xi_C \neq \frac{k_0}{k_2}$. For $\xi_C = \frac{k_0}{k_2}$ a new constant of motion is introduced so that $S(t) - E(t) = s_0 - e_0$ unless S , E or both S and E are targeted by an intervention.

It can be shown, through explicit calculation, that for any perfect intervention these constraints have no solution when the dynamical system does not converge to an equilibrium and they have a unique solution when the system does converge to an equilibrium. A complete causal description of the system can be found in Table 2 in the supplementary material. Table 1 and Figure 1 illustrate the rich causal semantics of this system (e.g. an intervention on S makes C^* dependent on the initial conditions, while an intervention on E makes S^* independent of the initial conditions).

2.2 SCM REPRESENTATION

Globally asymptotically stable dynamical systems converge to a unique fixed point and Mooij et al. (2013) show how SCMs can then be constructed from ordinary differential equations. For the basic enzyme reaction (which is not globally asymptotically stable) their construction method would yield the structural equations:

$$S^* = \frac{k_0 + k_{-1} C^*}{k_1 E^*}, \quad (11)$$

$$E^* = \frac{(k_{-1} + k_2) C^*}{k_1 S^*}, \quad (12)$$

$$C^* = \frac{k_1 S^* E^*}{k_{-1} + k_2}, \quad (13)$$


$$P^* = \frac{k_2}{k_3} C^*. \quad (14)$$

While this SCM represents the causal semantics of the system’s fixed points, it would be underspecified as an SCM for the stationary behavior of the basic enzyme reaction. Indeed, this SCM has multiple solutions, corresponding to different possible initial conditions of the dynamical system and it does not contain any information on which of its solutions is realized. Theorem 1 shows that a complete SCM representation of the stationary behavior in the basic enzyme reaction does not exist.

Theorem 1. *The causal semantics of the stationary behaviour of the basic enzyme reaction, and its dependence on initial states, cannot be completely represented by an SCM with endogenous variables S^*, E^*, C^*, P^* .*

Proof. The system converges to an equilibrium under the intervention $\text{do}(E^* = e, C^* = c, P^* = p)$.⁵ Setting $\dot{S} =$

⁵See supplementary material A2.



(a) S and P converge to an equilibrium that depends on initial conditions in the observed system.

(b) C , E , and P converge to an equilibrium that depends on the initial conditions after an intervention on S .

(c) S , C , and P converge to an equilibrium that is independent of the initial conditions after an intervention on E .

Figure 1: Temporal dependence of concentrations in the basic enzyme reaction in (4) to (7) with random initial conditions and $\mathbf{k} = [0.4, 0.3, 1.0, 1.1, 0.5]$. Other choices for the rate parameters give qualitatively similar results.

Table 1: Equilibrium solutions to the intervened dynamical system of the basic enzyme reaction in (4) to (7) under various interventions, where $y = \frac{1}{2}\sqrt{(e_0 - s_0)^2 + 4\frac{k_0(k_{-1}+k_2)}{k_1k_2}}$.

I	S^*	C^*	E^*
\emptyset	$\frac{k_0+k_{-1}\frac{k_0}{k_2}}{k_1(e_0+c_0-\frac{k_0}{k_2})}$	$\frac{k_0}{k_2}$	$e_0 + c_0 - \frac{k_0}{k_2}$
$S = \xi_s$	ξ_s	$\frac{k_1\xi_s(e_0+c_0)}{k_{-1}+k_2+k_1\xi_s}$	$\frac{(k_{-1}+k_2)(e_0+c_0)}{k_{-1}+k_2+k_1\xi_s}$
$C = \frac{k_0}{k_2}$	$\frac{(e_0-s_0)}{2} + y$	$\frac{k_0}{k_2}$	$\frac{-(e_0-s_0)}{2} + y$
$E = \xi_e$	$\frac{k_0+k_{-1}\frac{k_0}{k_2}}{k_1\xi_e}$	$\frac{k_0}{k_2}$	ξ_e

0 in (4) and solving for S , we find that $S^* = \frac{k_0+k_{-1}c}{k_1e}$ and therefore any SCM that models the effect of this intervention correctly must have a structural equation that is equivalent to equation (11). Analogously, considering the remaining three interventions on three out of four variables, we find that an SCM that correctly models the effects of those interventions must have structural equations for E^* , C^* and P^* that are equivalent to the structural equations (12) to (14), respectively. Table 1 shows that the system converges to an equilibrium that depends on the initial conditions c_0 and e_0 under the null intervention. This equilibrium is a solution of the structural equations in (11) to (14). However, these structural equations do not depend on initial conditions and admit other solutions as well. Therefore they do not completely represent the stationary behaviour of the system. \square

3 CAUSAL CONSTRAINTS MODELS

We introduce Causal Constraints Models (CCMs) and prove that they completely capture the causal semantics of the stationary behaviour of dynamical systems.

SCMs are specified by structural equations which con-

strain its solutions unless the corresponding variable is targeted by an intervention. CCMs are specified by *causal constraints*: relations between variables that constrain the solutions of the model under explicitly specified intervention targets.

Definition 7. Let \mathcal{I} , \mathcal{J} and \mathcal{K} be index sets. A *Causal Constraints Model (CCM)* is a triple $(\mathcal{X}, \Phi, \mathbf{E})$, with:

- \mathcal{X} and \mathbf{E} as in Definition 1 (domain of endogenous variables and tuple of exogenous random variables respectively),
- a set $\Phi = \{\phi_k : k \in \mathcal{K}\}$ of causal constraints, each of which is a triple $\phi_k = (f_k, c_k, A_k)$ where,
 - $f_k : \mathcal{X}_{\text{pa}(k) \cap \mathcal{I}} \times \mathcal{E}_{\text{pa}(k) \cap \mathcal{J}} \rightarrow \mathcal{Y}_k$ is a measurable function, \mathcal{Y}_k a standard measurable space and $\text{pa}(k) \subseteq \mathcal{I} \cup \mathcal{J}$,
 - $c_k \in \mathcal{Y}_k$ is a constant,
 - $A_k \subseteq \mathcal{P}(\mathcal{I})$ specifies the set of intervention targets under which ϕ_k is *active*.

The following example illustrates the interpretation of causal constraints in CCMs.

Example 1. Consider the price, supply, and demand of a certain product, denoted by P , S , and D respectively, related by the following causal constraint:

$$(f, c, A) = (S - D, 0, \{\emptyset, \{D\}, \{S\}, \{D, S\}\}). \quad (15)$$

The constraint $S - D = 0$ is active in the observational setting because $\emptyset \in A$. It is also active when either D , S or both D and S are targeted by an intervention. The constraint becomes inactive after an intervention on P . In other words, supply equals demand unless the price of the product is intervened upon (e.g. price-fixing). \triangle

3.1 CCM SOLUTIONS

We define a solution of a CCM in complete analogy with the definition of a solution of an SCM.

Definition 8. Let $\mathcal{M} = (\mathcal{X}, \Phi, \mathbf{E})$ be a CCM and let $\Phi^\emptyset := \{\phi_k = (f_k, c_k, A_k) \in \Phi : \emptyset \in A_k\}$. A random variable \mathbf{X} taking value in \mathcal{X} is a *solution* of \mathcal{M} if

$$f_k(\mathbf{X}_{\text{pa}(k) \cap \mathcal{I}}, \mathbf{E}_{\text{pa}(k) \cap \mathcal{J}}) = c_k \text{ a.s.}, \forall (f_k, c_k, A_k) \in \Phi^\emptyset.$$

Similar to SCMs, a CCM has either no solution, or it has a solution and all its solutions may either induce a unique or multiple distributions.

3.2 CCM INTERVENTIONS

Interventions on SCMs act on its structural equations. Analogously, an intervention on a CCM acts on its causal constraints. Roughly speaking, the activation sets of the causal constraints in the model are updated and additional causal constraints describe the intervention.

Definition 9. Let $\mathcal{M} = (\mathcal{X}, \Phi, \mathbf{E})$ be a CCM and let $I \subseteq \mathcal{I}$ be the intervention target and $\xi_I \in \mathcal{X}_I$ the target value. The *intervened CCM* is given by $\mathcal{M}_{\text{do}(I, \xi_I)} = (\mathcal{X}, \tilde{\Phi}, \mathbf{E})$ where:

- for each $i \in I$ we add a causal constraint describing the intervened value of the targets, $(x_i, \xi_i, \mathcal{P}(\mathcal{I} \setminus \{i\})) \in \tilde{\Phi}$,
- for each causal constraint $(f, c, A) \in \Phi$ we get a modified causal constraint $(f, c, A_{\text{do}(I)}) \in \tilde{\Phi}$ if $A_{\text{do}(I)} \neq \emptyset$, where

$$A_{\text{do}(I)} = \{A_i \setminus J : A_i \in A, J \subseteq I \subseteq A_i\}.$$

Definition 9 says that for any $A_i \in A$, and for any combination of two subsequent interventions such that $I_1 \cup I_2 = A_i$, the constraint will be active. So after I_1 (which needs to be a subset of A_i), any I_2 that adds the remaining elements $A_i \setminus I_1$ (plus possibly any elements that were already in I_1) will activate the constraint.

Example. The effect of different interventions on a set $A_{\text{do}(\emptyset)} = \{\emptyset, \{1, 2\}, \{2, 3\}\}$:

$$\begin{aligned} A_{\text{do}(1)} &= \{\{2\}, \{1, 2\}\}, \\ A_{\text{do}(2)} &= \{\{1\}, \{1, 2\}, \{3\}, \{2, 3\}\}, \\ A_{\text{do}(\{1, 2\})} &= A_{\text{do}(1)\text{do}(2)} = A_{\text{do}(2)\text{do}(1)} \\ &= \{\emptyset, \{1\}, \{2\}, \{1, 2\}\}, \\ A_{\text{do}(\{1, 2, 3\})} &= \emptyset. \end{aligned}$$

Lemma 1 shows that the effect of multiple interventions on a CCM depends neither on whether the interventions are performed simultaneously or sequentially nor on the order in which they are performed.

Lemma 1. *Let \mathcal{M} be a CCM for variables indexed by \mathcal{I} and let $I, J \subseteq \mathcal{I}$ be two disjoint sets of intervention*

targets with intervention values $\xi_I \in \mathcal{X}_I$ and $\xi_J \in \mathcal{X}_J$ respectively. Then

$$\begin{aligned} (\mathcal{M}_{\text{do}(I, \xi_I)})_{\text{do}(J, \xi_J)} &= (\mathcal{M}_{\text{do}(J, \xi_J)})_{\text{do}(I, \xi_I)} \\ &= \mathcal{M}_{\text{do}(I \cup J, \xi_{I \cup J})}. \end{aligned}$$

Proof. The result follows directly from Definition 9. \square

The following example illustrates interventions on a CCM.

Example 1 (Continued). Suppose that the supply of a product, if it is not targeted by an intervention, is determined by a function f_S , which takes as input the price of the product P and an exogenous random variable E (e.g. cost of production). The system for price, supply, and demand can be represented by an (underspecified) CCM $\mathcal{M} = (\mathbb{R}^3, \Phi, \mathbf{E})$, where Φ consists of two causal constraints:

$$\begin{aligned} (S - D, \quad 0, \quad \{\emptyset, \{D\}, \{S\}, \{D, S\}\}), \\ (S - f_S(P, E), \quad 0, \quad \{\emptyset, \{D\}, \{P\}, \{D, P\}\}). \end{aligned}$$

After an intervention on P we get $\mathcal{M}_{\text{do}(P, \xi_P)} = (\mathbb{R}^3, \tilde{\Phi}, \mathbf{E})$, where the updated set of causal constraints is given by

$$\begin{aligned} (S - f_S(P, E), \quad 0, \quad \{\emptyset, \{D\}, \{P\}, \{D, P\}\}), \\ (P, \quad \xi_P, \quad \{\emptyset, \{D\}, \{S\}, \{D, S\}\}). \end{aligned}$$

Note that after an intervention on P , there would be no intervention under which the causal constraint $(S - D, 0, \{\emptyset, \{D\}, \{S\}, \{D, S\}\})$ is still active (not even for the null intervention), so it is discarded from $\tilde{\Phi}$. \triangle

3.3 FROM SCM TO CCM

Structural equations in SCMs are constraints that are active as long as their corresponding variables are not targeted by interventions. This can be used to demonstrate how, for real-valued SCMs, an equivalent CCM with the same solutions under interventions can be constructed.⁶

Proposition 1. Let $\mathcal{M}^{\text{SCM}} = (\mathbb{R}^p, F, \mathbf{E})$ be a real-valued SCM and $\mathcal{I} = \{1, \dots, p\}$ an index set. The CCM $\mathcal{M}^{\text{CCM}} = (\mathbb{R}^p, \Phi, \mathbf{E})$ with causal constraints Φ :

$$(f_j(\mathbf{x}_{\text{pa}(j)}, \mathbf{e}_{\text{pa}(j)}) - x_j, \quad 0, \quad A_j = \mathcal{P}(\mathcal{I} \setminus \{j\})), \forall j \in \mathcal{I},$$

has the same solutions as \mathcal{M}^{SCM} under any intervention.

Proof. The result follows from Definitions 8 and 9. \square

⁶The general case, where variables take value in a standard measurable space, requires an additive structure on the variable domains with a zero-element.

3.4 EQUILIBRIUM CAUSAL MODELS

We have seen that SCMs may fail to completely capture the causal semantics of stationary behaviour in dynamical systems. Here we prove that CCMs can always completely represent such causal semantics.

Theorem 2. *Let \mathcal{D} be a dynamical system such that for all $I \subseteq \mathcal{I}$ and all $\xi_I \in \mathcal{X}_I$, $\mathcal{D}_{\text{do}(I, \xi_I)}$ has a unique solution of the form (2). Then there exists a CCM $\mathcal{M}(\mathcal{D})$ such that for all $I \subseteq \mathcal{I}$ and all $\xi_I \in \mathcal{X}_I$:*

- the equilibrium solutions of $\mathcal{D}_{\text{do}(I, \xi_I)}$ coincide with the solutions of $(\mathcal{M}(\mathcal{D}))_{\text{do}(I, \xi_I)}$,
- the following diagram commutes:

$$\begin{array}{ccc} \mathcal{D} & \xrightarrow{\quad} & \mathcal{M}(\mathcal{D}) \\ \downarrow & & \downarrow \\ \mathcal{D}_{\text{do}(I, \xi_I)} & \xrightarrow{\quad} & (\mathcal{M}(\mathcal{D}))_{\text{do}(I, \xi_I)}. \end{array}$$

Proof. By assumption, the intervened system $\mathcal{D}_{\text{do}(I, \xi_I)}$ has a unique solution $\mathbf{X}_t(\xi_I, e_{\mathcal{I} \setminus I}) := \mathbf{X}(t, (\xi_I, e_{\mathcal{I} \setminus I}))$ which is measurable in $(\xi_I, e_{\mathcal{I} \setminus I})$ for all t . For $I \subseteq \mathcal{I}$, let $\mathcal{C}_I := \{(\xi_I, e_{\mathcal{I} \setminus I}) \in \mathbb{R}^{|I|} \times \mathbb{R}^{|\mathcal{I} \setminus I|} : \mathbf{X}_t(\xi_I, e_{\mathcal{I} \setminus I}) \text{ converges for } t \rightarrow \infty\}$. Consider the measurable function $\mathbf{g}^I : \mathbb{R}^{\mathcal{I}} \times \mathbb{R}^{\mathcal{I} \setminus I} \rightarrow \mathbb{R}^{\mathcal{I}}$ defined by

$$\begin{aligned} \mathbf{g}^I(\mathbf{x}, e_{\mathcal{I} \setminus I}) &:= \mathbf{X}^*((\mathbf{x}_I, e_{\mathcal{I} \setminus I})) \mathbf{1}_{\mathcal{C}_I}((\mathbf{x}_I, e_{\mathcal{I} \setminus I})) \\ &\quad + (\mathbf{x} + \mathbf{1})(\mathbf{1} - \mathbf{1}_{\mathcal{C}_I}((\mathbf{x}_I, e_{\mathcal{I} \setminus I}))) - \mathbf{x}. \end{aligned}$$

The constraint $\mathbf{g}^I(\mathbf{x}, e_{\mathcal{I} \setminus I}) = \mathbf{0}$ gives a contradiction if and only if $(\mathbf{x}_I, e_{\mathcal{I} \setminus I}) \notin \mathcal{C}_I$, and reduces to the equation $\mathbf{x} = \mathbf{X}^*((\mathbf{x}_I, e_{\mathcal{I} \setminus I}))$ otherwise. Therefore, the equilibrium solutions of $\mathcal{D}_{\text{do}(I, \xi_I)}$ coincide with the solutions of the equation $\mathbf{g}^I(\mathbf{x}, e_{\mathcal{I} \setminus I}) = \mathbf{0}$. The CCM $\mathcal{M}(\mathcal{D}) := (\mathcal{X}, \Phi, \mathbf{E})$ with $\Phi = \{(\mathbf{g}^I, \mathbf{0}, A^I = \{I\}) : I \subseteq \mathcal{I}\}$ satisfies the properties of the theorem by construction. \square

Theorem 2 proves that a CCM representation always exists that completely characterizes the causal semantics of a dynamical system at equilibrium. Although we construct a CCM in the proof of the theorem, it does not give a parsimonious representation of the system.⁷ In the next section, we will outline an intuitive and more convenient construction method in the context of ODEs.

4 FROM ODE TO CCM

We consider how and when parsimonious CCM representations can be derived from ODEs and initial conditions in a dynamical system. We demonstrate how causal constraints completely capture the stationary behavior of the basic enzyme reaction and how, unlike SCMs, they are able to correctly represent non-convergence.

⁷Interestingly, the CCM construction in the proof of Theorem 2 can be applied to dynamical systems at finite time t .

4.1 CAUSAL CONSTRAINTS FROM DIFFERENTIAL EQUATIONS

When modeling the stationary behavior of a system of ODEs, setting the time-derivatives equal to zero constrains the solution space of the equilibrium model to the fixed points of the system. A CCM allows us to interpret such constraints as *causal* by explicitly specifying under which interventions they put constraints on the equilibrium solutions of the system.

Example 2. For the basic enzyme reaction, some of the causal constraints are obtained by setting the time derivatives of the four variables of the system in equations (4) to (7) to zero. The resulting equations constrain the solutions of the system as long as the corresponding variables are not targeted by an intervention. This leads to the causal constraints in equations (16) to (19) below,

$$(k_0 + k_{-1}C^* - k_1S^*E^*, \quad 0, \quad \mathcal{P}(\mathcal{I} \setminus \{S\})), \quad (16)$$

$$(k_1S^*E^* - (k_{-1} + k_2)C^*, \quad 0, \quad \mathcal{P}(\mathcal{I} \setminus \{C\})), \quad (17)$$

$$(-k_1S^*E^* + (k_{-1} + k_2)C^*, \quad 0, \quad \mathcal{P}(\mathcal{I} \setminus \{E\})), \quad (18)$$

$$(k_2C^* - k_3P^*, \quad 0, \quad \mathcal{P}(\mathcal{I} \setminus \{P\})), \quad (19)$$

with \mathcal{I} an index set for (S, C, E, P) . At this stage, the CCM is equivalent to the underspecified SCM of the dynamical system (see also section 2.3). In the next section we will proceed by adding more causal constraints. \triangle

Example 3. The Lotka-Volterra model (Murray, 2002) is a set of differential equations that is often used to describe the dynamics of a system where prey (e.g. deer) and predators (e.g. wolves), X_1 and X_2 , interact. The dynamics of the biological model are given by

$$\dot{X}_1 = X_1(t)(k_{11} - k_{12}X_2(t)), \quad (20)$$

$$\dot{X}_2 = -X_2(t)(k_{22} - k_{21}X_1(t)), \quad (21)$$

with initial values $X_1(0) > 0, X_2(0) > 0$ and strictly positive rate parameters. The system has two fixed points $(X_1^*, X_2^*) = (0, 0)$ and $(X_1^*, X_2^*) = (k_{22}/k_{21}, k_{11}/k_{12})$, which can be represented either by causal constraints,

$$(X_1^*(k_{11} - k_{12}X_2^*), 0, \{\emptyset, \{2\}\}), \quad (22)$$

$$(X_2^*(k_{22} - k_{21}X_1^*), 0, \{\emptyset, \{1\}\}), \quad (23)$$

or (equivalently) by structural equations:

$$X_1^* = X_1^* + X_1^*(k_{11} - k_{12}X_2^*),$$

$$X_2^* = X_2^* - X_2^*(k_{22} - k_{21}X_1^*).$$

These (structural) equations do not describe the stable steady state behavior of the model, because the system displays undamped oscillations around the positive fixed point, as was pointed out by Mooij et al. (2013) and Murray (2002). In the next section we proceed by adding additional relevant constraints to the CCM. \triangle

4.2 CAUSAL CONSTRAINTS FROM CONSTANTS OF MOTION

For dynamical systems that admit a constant of motion (i.e. a conserved quantity), the trajectories of its solutions are confined to a space that is constrained by its initial conditions. Hence, the solutions for the equilibrium must be similarly constrained. In a CCM we interpret these constraints as causal by specifying under which interventions they constrain the solution space.

Example 2 (Continued). For the basic enzyme reaction, we include the conservation law that results from the linear dependence between the time derivative of the free enzyme E and the complex C in equation (9). Since this relation holds as long as the ‘cycle’ between C and E is not broken, we obtain the following causal constraint

$$(C^* + E^* - (c_0 + e_0), \quad 0, \quad \mathcal{P}(\mathcal{I} \setminus \{C, E\})). \quad (24)$$

Another conservation law appeared after intervention on the variable C . The resulting conservation law $S(t) - E(t) = s_0 - e_0$ applies as long as the ‘cycle’ between S and E is not broken by another intervention on the system. This leads to the final causal constraint:

$$(S^* - E^* - (s_0 - e_0), \quad 0, \quad \{\{C\}, \{C, P\}\}). \quad (25)$$

Let Φ be the set of causal constraints in (16) to (19) and (24) to (25). In Section 2.1.1 we showed that the active constraints in Φ have a unique solution under any intervention. If $\mathbf{E} = (s_0, e_0, c_0, p_0)$ is a set of exogenous random variables then the CCM $\mathcal{M} = (\mathbb{R}_{>0}^4, \Phi, \mathbf{E})$ completely captures the stationary behaviour of the basic enzyme reaction. \triangle

Remark 1. Interestingly, if we treat C as a latent endogenous variable that cannot be intervened upon, the equilibrium to which the dynamics of the basic enzyme reaction converges can be described by the following marginal CCM (see supplementary material for details):

$$\begin{aligned} \frac{k_0 + k_{-1} \frac{k_0}{k_2}}{k_1 E^*} - S^*, & \quad 0, & \quad \mathcal{P}(\mathcal{I} \setminus \{S\}), \\ \frac{(k_{-1} + k_2)(c_0 + e_0)}{k_{-1} + k_2 + k_1 S^*} - E^*, & \quad 0, & \quad \mathcal{P}(\mathcal{I} \setminus \{E\}), \\ \frac{k_2}{k_3} \frac{k_1 S^* E^*}{k_{-1} + k_2} - P^*, & \quad 0, & \quad \mathcal{P}(\mathcal{I} \setminus \{P\}), \end{aligned}$$

where \mathcal{I} is an index set for $\{S, E, P\}$. From Proposition 1 it can be seen that there exists an equivalent SCM that does completely capture the causal semantics of S, E , and P , as long as one does not intervene on C .

Example 3 (Continued). The Lotka-Volterra model provides an example of a system that admits a non-linear conservation law:

$$\begin{aligned} k_{21} X_1 + k_{22} \log(X_1) - k_{12} X_2 + k_{11} \log(X_2) = & \quad (26) \\ - k_{21} X_1(0) + k_{22} \log(X_1(0)) - k_{12} X_2(0) + k_{11} \log(X_2(0)), \end{aligned}$$

which represents a constraint that is only active in the observational setting. If the system would converge to an equilibrium (X_1^*, X_2^*) the causal constraints derived from the differential equations should hold simultaneously. These constraints are only satisfied when the system starts out in one of the fixed points (e.g. $(X_1(0), X_2(0)) = (k_{22}/k_{21}, k_{11}/k_{12})$). Otherwise the dynamical system exhibits steady-state oscillations and the set of causal constraints has no solution.

A complete causal description can be obtained by adding the following two causal constraints:

$$(X_1^* - X_1(0) \mathbf{1}_{\{k_{11} - k_{12} X_2^* \geq 0\}}, \quad 0, \quad \{\{2\}\}), \quad (27)$$

$$(X_2^* - X_2(0) \mathbf{1}_{\{k_{22} - k_{21} X_1^* \leq 0\}}, \quad 0, \quad \{\{1\}\}). \quad (28)$$

Addition of the causal constraint in equation (27) ensures that after an intervention on the amount of predators X_2 : a) the prey X_1 goes extinct when there are too many predators b) the model has no solution if there are too few predators and c) the amount of prey is constant if the amount of predators is exactly right. The causal constraint in equation (28) can be interpreted similarly. Together, the causal constraints in equations (22), (23), (26), (27), and (28) capture the stationary behavior of the predator-prey model.⁸ The SCM on the other hand has the fixed points of the system as a solution and does not predict the non-convergent behavior. \triangle

4.3 CONSTRUCTING CCMs

Causal constraints (or structural equations) derived from differential equations result in a causal description of the fixed points in a system. For structurally semistable systems the addition of causal constraints derived from constants of motion results in a complete causal description of the system’s stationary behavior when the constraints specify the equilibria in terms of initial conditions.

Theorem 3. *Let \mathcal{D} be a dynamical system that converges to a fixed point if it has at least one. Let \mathcal{M} be a CCM constructed from the ODEs and constants of motion in \mathcal{D} for which all solutions, if they exist, are unique up to $\mathbb{P}^{\mathbf{E}}$ -zero sets. \mathcal{D} converges to an equilibrium \mathbf{X}^* if and only if \mathbf{X}^* is a solution of \mathcal{M} .*

Proof. First assume that \mathcal{D} has a fixed point, so that \mathcal{D} converges to an equilibrium $\mathbf{X}^*(e)$ for almost every $e \in \mathbb{R}^p$. We have that a) $\mathbf{X}^*(e)$ satisfies the constants of motion in the dynamical system and b) for $\mathbf{X}^*(e)$ the time-derivatives appearing in the ODEs are equal to zero. Hence if \mathcal{D} converges to \mathbf{X}^* then \mathbf{X}^* is a solution of \mathcal{M} . Since \mathcal{M} has no more than one solution (up to zero sets),

⁸This can be verified by explicitly calculating the solutions of the model under all interventions.

the reverse statement is also true. Now assume that \mathcal{D} has no fixed point. In that case \mathcal{M} has no solutions, and \mathcal{D} cannot converge to an equilibrium. \square

Corollary 1. *Let \mathcal{D} be structurally semistable and \mathcal{M} a CCM constructed from the ODEs and constants of motion in \mathcal{D} for which under any intervention, all solutions, if they exist, are unique up to \mathbb{P}^E -zero sets. Then for all $I \subseteq \mathcal{I}$ and $\xi_I \in \mathbb{R}^{|\mathcal{I}|}$: $\mathcal{D}_{\text{do}(I, \xi_I)}$ converges to an equilibrium $X^*(I, \xi_I)$ iff $X^*(I, \xi_I)$ is a solution of $\mathcal{M}_{\text{do}(I, \xi_I)}$.*

Proof. If $\mathcal{M}_{\text{do}(I, \xi_I)}$ has a solution then $\mathcal{D}_{\text{do}(I, \xi_I)}$ has a fixed point with $x_I^* = \xi_I$ and it converges because \mathcal{D} is structurally semistable. If $\mathcal{M}_{\text{do}(I, \xi_I)}$ has no solution then $\mathcal{D}_{\text{do}(I, \xi_I)}$ does not converge to a fixed point. The result follows from Theorem 3 and Definition 9. \square

The basic enzyme reaction in Example 2 is structurally semistable, while the Lotka-Volterra model in Example 3 is not. Corollary 1 tells us that for structurally semistable systems, if a CCM constructed from ODEs and constants of motions has at most one solution under any intervention, then the CCM completely captures the causal semantics of the stationary behaviour of the system.

5 FUNCTIONAL LAWS

CCMs can also represent *functional laws*, which are relations between variables that are invariant under *all* interventions. Causal constraints allow one to explicitly state under which interventions a constraint is active. Therefore a CCM never admits a solution that violates the functional law, where an SCM would.

Example 4. It is well-known that the pressure P and temperature T for N particles of an ideal gas in a fixed volume V are related by the ideal gas law. In absence of any knowledge about the environment, this system can be represented by the (underspecified) CCM $\mathcal{M} = (\mathbb{R}^2, \{(PV - Nk_B T, 0, \mathcal{P}(\mathcal{I})\}, \mathbb{P}_\emptyset)$, where k_B is Boltzmann’s constant, and \mathcal{I} is an index set for the variables (P, T) in the system. If we were to describe the same system using an SCM, then we would need two copies of this causal constraint as structural equations:

$$P = \frac{Nk_B T}{V}, \quad T = \frac{PV}{Nk_B}.$$

Indeed, considering interventions on one of the variables leaves no choice for the structural equation of the other one. Furthermore, a simultaneous intervention on P and T always has a solution in the SCM representation, even when this means that the ideal gas law is violated. The CCM representation typically does not have a solution under such an intervention (unless the target values satisfy the ideal gas law constraint). Therefore, the CCM representation of functional laws like the ideal gas law

is more parsimonious and more natural than any SCM representation can be. \triangle

A functional law can be any relation that is invariant under all interventions. For example, a transformation of a (set of) variables to another (set of) variables describing the same system can also be modeled as a functional law.

Example 5. Let \mathcal{I} be an index set of (T, V, O) . Suppose that the viscosity T of a salad dressing, consisting of a certain amount of oil O and a certain amount of vinegar V is determined by a causal constraint $\phi = (f, 0, \mathcal{P}(\mathcal{I} \setminus \{T\}))$ where f is a function depending on the amount of oil and vinegar. By adding causal constraints

$$\begin{aligned} (O_r - O/(O + V), & \quad 0, & \quad \mathcal{P}(\mathcal{I}), \\ (V_r - V/(O + V), & \quad 0, & \quad \mathcal{P}(\mathcal{I}), \end{aligned}$$

a CCM allows us to have the relative amounts of oil and vinegar O_r and V_r in the model without running into logical contradictions. \triangle

6 CONCLUSION

While Structural Causal Models (SCMs) form a very popular modeling framework in many applied sciences, we have shown that they are neither powerful enough to model the rich equilibrium behavior of simple dynamical systems such as the basic enzyme reaction, nor simple functional laws of nature like the ideal gas law. This raises the question whether the common starting point in causal discovery—that the data-generating process can be modeled with an SCM—is tenable in certain application domains, for example, for biochemical systems.

We believe that the examples presented in this paper form a compelling motivation to extend the common causal modeling framework to potentially broaden the impact of causal modeling in dynamical systems. In this work, we introduced Causal Constraints Models (CCMs). We showed how they can be ‘constructed’ from differential equations and initial conditions and proved that they can completely capture the causal semantics of functional laws and stationary behavior in dynamical systems.

One intuitively appealing aspect of SCMs is their graphical interpretation. In contrast, CCMs are not equipped with graphical representations yet. In future work, we plan to investigate graphical representations of the conditional independence structure of CCMs. This will allow us to better understand the causal interpretation of the results of existing causal discovery algorithms.

Acknowledgements This work was supported by the ERC under the European Union’s Horizon 2020 research and innovation programme (grant agreement 639466) and by NWO (VIDI grant 639.072.410).

References

- Belgacem, I. and Gouzé, J.-L. (2012). “Global Stability of Full Open Reversible Michaelis-Menten Reactions”. *IFAC Proceedings Volumes* 45.15, pp. 591–596.
- Bhat, S. P. and Bernstein, D. S. (1999). “Lyapunov analysis of semistability”. *Proceedings of the American Control Conference* 3.
- Bongers, S. and Mooij, J. M. (2018). “From Random Differential Equations to Structural Causal Models: the stochastic case”. *arXiv.org preprint arXiv:1803.08784v2 [cs.AI]*.
- Bongers, S., Peters, J., Schölkopf, B., and Mooij, J. M. (2018). “Theoretical Aspects of Cyclic Structural Causal Models”. *arXiv.org preprint arXiv:1611.06221v2 [stat.ME]*.
- Campbell, S. and Rose, N. (1979). “Singular Perturbation of Autonomous Linear Systems”. *SIAM Journal of Math. Analysis* 10, pp. 542–551.
- Chellaboina, V., Bhat, S. P., Haddad, W. P., and Bernstein, D. S. (2009). “Modeling and analysis of mass-action kinetics”. *IEEE Control Systems* 29 (4).
- Dash, D. (2005). “Restructuring Dynamic Causal Systems in Equilibrium”. *Proceedings of the Tenth International Workshop on Artificial Intelligence and Statistics (AISTATS 2005)*.
- Fisher, F. M. (1970). “A Correspondence Principle for Simultaneous Equation Models”. *Econometrica* 38.1, pp. 73–92.
- Forré, P. and Mooij, J. M. (2017). “Markov Properties for Graphical Models with Cycles and Latent Variables”. *arXiv.org preprint arXiv:1710.08775 [math.ST]*.
- Haddad, W. M., Chellaboina, V., and Hui, Q. (2010). *Nonnegative and Compartmental Dynamical Systems*. Princeton University Press.
- Han, X. and Kloeden, P. E. (2017). *Random Ordinary Differential Equations and Their Numerical Solution*. Springer Singapore.
- Hytinen, A., Eberhardt, F., and Hoyer, P. O. (2012). “Learning Linear Cyclic Causal Models with Latent Variables”. *Journal of Machine Learning Research* 13.1, pp. 3387–3439.
- Iwasaki, Y. and Simon, H. A. (1994). “Causality and model abstraction”. *Artificial Intelligence* 67.1, pp. 143–194.
- Lacerda, G., Spites, P., Ramsey, J., and Hoyer, P. O. (2008). “Discovering Cyclic Causal Models by Independent Components Analysis”. *Proceedings of the Twenty-Fourth Conference on Uncertainty in Artificial Intelligence*.
- Mogensen, S. W., Malinsky, D., and Hansen, N. R. (2018). “Causal Learning for Partially Observed Stochastic Dynamical Systems”. *Proceedings of the 34th Annual Conference on Uncertainty in Artificial Intelligence (UAI-18)*.
- Mooij, J. M., Janzing, D., Heskes, T., and Schölkopf, B. (2011). “On causal discovery with cyclic additive noise models”. *Advances in Neural Information Processing Systems*, pp. 639–647.
- Mooij, J. M., Janzing, D., and Schölkopf, B. (2013). “From Ordinary Differential Equations to Structural Causal Models: the deterministic case”. *Proceedings of the 29th Annual Conference on Uncertainty in Artificial Intelligence (UAI-13)*, pp. 440–448.
- Murray, J. D. (2002). *Mathematical Biology I: An Introduction*. Third. Springer-Verlag New York.
- Pearl, J. (2000). *Causality : models, reasoning, and inference*. Cambridge University Press.
- Rubenstein, P. K., Bongers, S., Schölkopf, B., and Mooij, J. M. (2018). “From Deterministic ODEs to Dynamic Structural Causal Models”. *Proceedings of the 34th Annual Conference on Uncertainty in Artificial Intelligence (UAI-18)*.
- Rubenstein*, P. K., Weichwald*, S., Bongers, S., Mooij, J. M., Janzing, D., Grosse-Wentrup, M., and Schölkopf, B. (2017). “Causal Consistency of Structural Equation Models”. *Proceedings of the 33rd Conference on Uncertainty in Artificial Intelligence (UAI-17)*. *equal contribution.
- Sokol, A. and Hansen, N. R. (2014). “Causal Interpretation of Stochastic Differential Equations”. *Electronic Journal of Probability* 19, pp. 1–24.
- Voortman, M., Dash, D., and Druzdzel, M. J. (2010). “Learning Why Things Change: The Difference-Based Causality Learner”. *Proceedings of the Twenty-Sixth Annual Conference on Uncertainty in Artificial Intelligence (UAI)*.