

# NEAR-DECOMPOSABILITY AND THE TIME-SCALE RELATIVITY OF CAUSAL REPRESENTATIONS

NAFTALI WEINBERGER

ABSTRACT. A common strategy for simplifying complex systems involves partitioning them into subsystems whose behaviors are roughly independent of one another at shorter time-scales. Dynamic causal models (Iwasaki and Simon, 1994) explain how doing so reveals a system’s non-equilibrium causal relationships. Here I use these models to elucidate the idealizations and abstractions involved in representing a system at a time-scale. The models reveal that key features of causal representations – such as which variables are exogenous – may vary with the time-scale at which a system is considered. This has implications for debates regarding which systems can be understood causally.

A tell-tale feature of complex systems is that they exhibit different behaviors at different time-scales. Accordingly, we should expect there to be systematic differences between the causal relationships observed over longer and shorter time-scales. But how, precisely, do causal representations vary with the time-scale at which a system is considered? Answering this question is crucial for providing an account of causation that applies to complex systems. Moreover, questions about time-scale are relevant even in domains where the importance of time is less salient. As I will argue, even causal representations that do not explicitly refer to time may build in idealizing assumptions that are only warranted at particular scales. To understand causal representations in general, we need an account of the abstractions and idealizations involved in representing a system at a particular time-scale.

Here I develop one such account by appealing to the concept of *near-decomposability* (Simon, 1962). A common strategy for modeling complex systems involves partitioning them into subsystems whose behaviors are nearly independent of one another over shorter time-scales. Systems that can be partitioned this way are *near-decomposable*. The concept of near-decomposability has been central to philosophical discussions of complexity (Wimsatt, 1972; Bechtel and Richardson, 2010), and continues to ground discussions of the limits of causal or mechanistic explanations (e.g. Rathkopf, 2018). Nevertheless, there has been almost no philosophical discussion of how, specifically, decomposing a system facilitates causal representation. In what follows, I show how dynamic causal models (Iwasaki and Simon, 1994) partition near-decomposable systems into subsystems in order to represent the causal relationships in a system when it is away from equilibrium. I then argue that these models provide a basis for thinking about the abstractions and idealizations involved in considering a system at a particular time-scale.

The framework presented provides a tool for thinking about which types of systems can be represented causally. Two controversies in the philosophy of causation concern whether there can be causal relationships in 1) closed deterministic systems and 2) complex

systems. The worry about causation in closed deterministic systems is that the content of causal models is determined in part by their exogenous variables, and it is unclear how to interpret such variables in cases where the modeled variables are not being influenced by factors external to the system. The following discussion makes precise one sense in which whether a variable counts as exogenous depends on its relationship to a broader system and reveals that which variables count as exogenous can vary with time-scale. The issue for (certain) complex systems is that it is allegedly not possible to decompose such systems into modular parts. The following clarifies the relationship between causation and decomposition, and reveals that whether a system is decomposable is sensitive to the time-scale at which it is represented. Although the framework presented does not resolve these debates, it provides a representational framework for addressing them more systematically.

The paper is organized as follows. Section 1 reviews the notion of near decomposability. Section 2 presents Simon's causal ordering account, and section 3 explains how dynamic causal models generalize this account. Section 4 links dynamic causal models to near-decomposability. Section 5 offers an account of how causal representation vary with time-scale. Section 6 relates this account to debates about causation. Section 7 concludes.

## 1. NEAR DECOMPOSABILITY

A system is near-decomposable if it can be divided up into subsystems such that:

- (1) The short-run behavior of each of the component subsystems is approximately independent of the short-run behavior of the other components.
- (2) In the long run, the behavior of any one of the subsystems depends on the behavior of the other subsystems only in an aggregate way.

Simon (1962) illustrates near-decomposability with an example. Imagine a building that is partitioned into rooms, which in turn are partitioned into cubicles. The walls between the rooms are somewhat effective thermal insulators, and the partitions between the cubicles are poor insulators. The building is thermally insulated from its environment, and begins in a state in which each cubicle has a different temperature. In the analogy, the rooms are the subsystems. In the short run, each room reaches a local equilibrium temperature more-or-less independently of the others. In the long run, the temperature of any room depends on that of all the rooms, and the system converges on a common equilibrium temperature.

A few features of near-decomposability are worth highlighting. First, whether two subsystems count as independent varies with the time-scale at which the system is considered. That is, subsystems are not fully independent, but sufficiently independent when considered over a sufficiently short time-scale. Second, decomposing a system enables one to treat the subsystems as having local equilibrium states. Since the system is always tending towards the long-run equilibrium, there is no time at which the subsystems are entirely static. But in the short-run what happens within the subsystems is so much more important than what happens across them that one can model the subsystems as reaching constant values independently of one another. Third, while it is natural to focus on the

the *decomposability* of near-decomposable systems, the aggregate dependence of particular subsystems on the others is just as important.

Simon introduced near-decomposability as part of a theory of variable aggregation. When a system is near-decomposable, there are methods for representing each subsystem with a single course-grained variable. Wimsatt (1972) introduced near-decomposability into philosophical discussions of complexity. The concept also guides Bechtel and Richardson’s (1993/2010) work on experimental strategies for localization in complex systems. Near-decomposability continues to orient discussions of the limits of causal or mechanistic explanations. Notably, Rathkopf (2018) argues that certain network-based explanations are not causal/mechanistic on the basis that they describe non-decomposable systems.<sup>1</sup> As philosophers of science further consider more holistic forms of explanation, questions about the relationship between a system’s being near-decomposable and the possibility of making causal attributions about it will become increasingly pressing.

## 2. THE CAUSAL ORDERING METHOD

In addition to inventing the concept of near-decomposability, Simon also pioneered causal modeling methods, and saw these projects as closely connected.<sup>2</sup> The precise connection will not become clear until section 4, after introducing dynamic causal models in section 3. These models generalize Simon’s (1953) earlier account, to which we now turn.

Simon (1953) sought to determine the basis for the asymmetry of causation, given that scientific laws are represented using symmetric equations. For instance, the ideal gas law relates a gas’ equilibrium pressure ( $P$ ), temperature ( $T$ ), and volume ( $V$ ):

$$(1) PV = kT$$

In (1), which variables are on the right- or left-hand side is purely conventional, and the equation says nothing about which variables asymmetrically depend on which others. Simon’s insight is that even though a *particular* equation will not represent causal directionality, a *set* of equations can. For instance, consider a gas in a fixed-volume container immersed in a constant-temperature heat bath. The fact that the values of temperature and volume are determined *exogenously* – independently of the values of the other variables – can be represented with the following equations.

$$(2) T = c_2$$

$$(3) V = c_3$$

Given these equations one can solve for the values of  $T$  and  $V$ , and once these values are solved for equation (1) yields the value of  $P$ . Simon claims that the *causal ordering* of the variables – that is, the (partial) ordering of the variables such that effects come later than their causes – is determined by the ordering in which one solves for the variables. Since the equations for  $T$  and  $V$  need to be solved in order to solve for the value of  $P$ ,  $T$  and  $V$  are causes of  $P$ .

---

<sup>1</sup>See (Burnston, 2019) for a critical discussion of this and similar arguments

<sup>2</sup>Hoover (2015) explores this connection in the context of non-stationary time-series models from econometrics. See Weinberger (2019) for further discussion.




FIGURE 1. Bathtub Model:  $Q_{in}$ : Rate of flow in;  $Q_{out}$ : Rate of flow out;  $D$ : Depth;  $P$ : Pressure;  $K$ : Size of Drain

Simon’s method might look unfamiliar, but when there is a unique causal ordering it yields the structural equations that many philosophers know and love. The symmetric equations can be rewritten so that each variable appears on the left-hand side of a single equation and can be interpreted as an effect of the variables (if any) on the right-hand side. Such equations typically indicate how the effect variables would respond to *interventions* on their causes. But why should this work? To start, note that – like contemporary techniques – Simon’s method does not yield causal knowledge without causal assumptions. Facts about which variables are exogenous are causal facts. Simon is clarifying the types of assumptions that jointly yield a causal representation. His central idea is that facts about what causes what follow from facts about which sets of variables are governed by autonomous mechanisms. The autonomy of mechanisms is embedded in the fact that not every equation contains every variable, so the values of subsets of variables can be determined independently of those of the others. From this, the causal ordering follows.

### 3. DYNAMIC CAUSAL MODELS

I now turn to dynamic causal models, which generalize Simon’s method. I’ll keep technical details to a minimum. Interested readers may consult Simon and Rescher (1966), Iwasaki and Simon (1994), Dash (2003), and Weinberger (2019).

Iwasaki and Simon (1994) present an example in which water flows into a bathtub at a rate of  $Q_{in}$  and out at a rate of  $Q_{out}$  (figure 1). The short-run behavior of this system is simple and familiar. Increasing  $Q_{in}$  increases depth ( $D$ ), and the resulting increase in pressure ( $P$ ) at the bottom of the tub determines  $Q_{out}$  (which also depends on the size of the drain,  $K$ ). What makes this case interesting is the system’s longer-term dynamics. In some cases – and we’ll focus on these – the system reaches an equilibrium state in which  $Q_{in}$  equals  $Q_{out}$  and the other variables reach constant values.

We begin with a static *equilibrium* model for the system. Such a model would predict (e.g.) how the bathtub’s equilibrium depth responds to interventions changing  $Q_{in}$  or  $K$ . The model, given in figure 2, is counterintuitive.<sup>3</sup> It is crucial to remember that it concerns only the equilibrium values of the variables. For instance, one might suppose that  $Q_{in}$

<sup>3</sup>This model is derived from equations ((6)-(10)) below.




FIGURE 2. Bathtub Model at Equilibrium

influences  $Q_{out}$  via  $D$ . But even though a short term change in  $D$  (e.g. pouring a bucket of water into the tub) would influence  $Q_{out}$  and thus temporarily move the system away from equilibrium, the only thing determining  $Q_{out}$ 's equilibrium value is  $Q_{in}$ . Even given this explanation, there remain puzzles regarding whether  $Q_{out}$  and  $P$  can be independently manipulated. But since this model is a way-station for getting to the dynamic causal model, we need not address these here.

What makes this model an equilibrium one? Answering this requires us to further unpack the model's variables. Quantities such as depth are represented using a single variable. This is non-trivial. If we wanted to represent the change in depth over time, we would need distinct variables corresponding to distinct times.<sup>4</sup> Given that each quantity in the model is represented with a single variable, how are they temporally related? We should think of the variables as measured simultaneously (at equilibrium). The causal difference-making relationships are then explicated not temporally, but counterfactually:  $K$  causes  $P$  at  $t$  because  $K$  has value  $k$  and  $P = p$  and were  $K$  to have had some alternative value  $k'$ ,  $P$  would not equal  $p$ .

Since causal relationships (generally) take time, why are these represented as simultaneous? We need not understand the modeled relationships as genuinely simultaneous, but as indicating that the effect has had adequate time to respond to any changes in the values of its causes. There are two (complementary) ways one might think about this. One is that when considering the system at a longer time-scale, the time required for the cause to influence its effect is so small (relative to that time-scale) that we can treat it as instantaneous. A second is that the variables are measured at the same moment, but because the effect has fully adjusted to any changes in its causes' values and these values are no longer changing at equilibrium, the simultaneous values enable one to infer the diachronic relationship between the effect and its earlier causes.

Equilibrium models assume that the variables are at steady-state, and thus have had enough time to adjust to changes in the values of their causes. In *dynamic causal models*, not all of the variables have had time to reach steady-state. For such variables, we include both a variable and its time-derivative in the model. In the model we will presently consider, there is a time-derivative for  $D$ . Including a variable and its derivative in a model requires one to introduce additional equations:

---

<sup>4</sup>Proof: Since the values of a variable form a partition, a variable's taking on one value excludes its taking another. But the water being 1m at one minute does not exclude its being .5m at another. Thanks to Malcolm Forster for emphasizing this point.




FIGURE 3. Dynamic Bathtub Model

$$(4) \text{ Integrate}(D') = D$$

$$(5) D_0 = d$$

Formally, equation (4) indicates that integration and differentiation are inverse operations. Practically, integration takes the values of  $D'$  and  $D$  at a time and predicts  $D$ 's value at the next time step. It is because doing so requires the value of  $D$  at  $t$  that we need equation (5), which supplies an initial condition for the initial application of (4). Through derivatives and integration equations we incorporate time into a model, as we will see further.

Here are the forms of the symmetric equations from which the equilibrium model was derived:

$$(6) 0 = f_6(D, P)$$

$$(7) 0 = f_7(K, P, Q_{out})$$

$$(8) 0 = f_8(K)$$

$$(9) 0 = f_9(Q_{in})$$

$$(10) 0 = f_{10}(Q_{in}, Q_{out})$$

It was from equations (9) and (10) that we derived  $Q_{out}$  from  $Q_{in}$ , and  $Q_{in}$  determines  $Q_{out}$  (directly) only at equilibrium. When  $D$  is not at equilibrium, we replace (10) with:

$$(11) D' = f_{11}(Q_{in}, Q_{out})$$

This equation, combined with (4)-(9) yields the graph in figure 3.<sup>5</sup> One could further consider a graph with a time-derivative for every modeled variable (Iwasaki and Simon, 1994, p. 159), although this does not substantially change the representation.

The derived causal ordering is intuitive.  $D$  causes  $P$  which, along with  $K$ , determines  $Q_{out}$ . The variables linked by solid arrows should be understood as simultaneous (as interpreted above). The integration link indicates that the influence of  $Q_{in}$  and  $Q_{out}$  via depth on the rest of the system occurs over a longer time-scale than the other causal relationships, though more remains to be said about interpreting such links.

<sup>5</sup> $D_0$  is not represented in the model, but equation (5) affects the ordering by rendering  $D$  exogenous.




FIGURE 4. Rolled-Out Discrete Representation With the Derivatives Removed

#### 4. DECOMPOSING DYNAMIC SYSTEMS

We seem to have drifted far away from near-decomposability. It turns out, however, that when we reflect upon the temporal relations in the dynamic causal model, the introduction of the derivative-integration link pair constitutes a way of decomposing the variable set into nearly-independent subsystems. In particular, the dynamic model in figure 3 divides the system into the two (very unequal) subsystems  $\{Q_{in}\}$  and  $\{P, D, K, Q_{out}\}$ .<sup>6</sup> I will now spell out how dynamic models involve such a decomposition.

Although the dynamic model contains a cycle via the integration link, one could provide an equivalent “rolled out” acyclic representation in which variables connected by causal arrows are represented as synchronic and those connected by integration links are diachronic. To better understand the dynamic graph, it helps to compare it to a rolled out graph in which the derivatives and integration links are removed (figure 4). This graph gives the impression of providing a discrete representation of the causal relationships, with no fancy gadgets to confuse us. Contemplating what this graph does and does not capture elucidates the temporal relations in the dynamic graph.

A common reason for providing representations with time-indexed variables is to use the temporal ordering to infer facts about the causal ordering. Given the preponderance of causally-related simultaneous variables it should be clear that that is not what is going on here. The key information provided by the temporal indices is that synchronically- and diachronically-represented variables influence one another over comparatively shorter and longer time-scales. While in principle, one could specify the temporal units for which the graphs in figures 3 and 4 apply, the utility of the representation does not derive from its specifying the relations among a particular set of time-indexed variables, but in its making a qualitative distinction among interactions that are so quick that we can treat them as instantaneous (at a time-scale) and those we cannot.

<sup>6</sup>Readers having trouble thinking about  $\{Q_{in}\}$  as a subsystem are welcome to think about it as an aggregate variable for an additional bathtub system placed above our own, so that the output of that system is the input to our own. Additionally, one could also treat  $\{K\}$  as its own subsystem without any consequences for the discussion.

What do the derivatives add? When one distinguishes between variables that are at and away from equilibrium, the latter are history-dependent in a way that the former are not. Compare  $P^1$  to  $D^1$ . The value of  $P^1$  depends only on its cause at that time step ( $D^1$ ). In contrast, the value of  $D^1$  depends on  $D$ 's value at the previous time step. Variables away from equilibrium have a “memory” in a way that variables at equilibrium do not. Since a variable’s time-derivative predicts its subsequent values only in combination with its current value, time-derivatives enable one to capture this past-dependence.

I have been talking as if variables whose causes are at the same time-step are at equilibrium while others are not. But this is not quite right. Whenever some variables in a system are away from equilibrium, none of the variables can reach their long-run equilibrium values. It is more precise to say that the variables with simultaneous causes are at a local equilibrium, in that each has fully responded to the value of their causes at that time-step. So  $P^1$  and  $Q_{out}^1$  fully reflect the values of  $D^1$  and  $K^1$ . But because  $D$  reaches steady state more slowly than the others, it serves as a bottleneck hindering the whole system from reaching equilibrium.<sup>7</sup>

A final subtlety concerns the influence of  $Q_{in}$  and  $Q_{out}$  on  $D$  via  $D'$ . The influences of the flow variables on the others differ from other causal relationships not merely in their longer duration, but also in their influencing the system globally rather than locally. The causal significance of  $Q_{in}$  and  $Q_{out}$  is that their difference indicates how far the system is from equilibrium and thus matters for predicting its longer-term behavior. Accordingly, these variables influence the system in an aggregate way. One might worry about whether such aggregate influences are genuinely causal and one might suppose that a completely dynamical representation of the system (i.e. one obtaining at arbitrarily short time-scales) would not need to make assumptions about the system’s longer-term behaviors. I will not address this worry here. The key point is that if one is *not* considering the system down to arbitrarily small time-scales, one needs to model the fact that the bathtub’s current state in relation to  $Q_{out} - Q_{in}$  provides information about its future state. Whether the aggregate influences via the integration link are causal, they need to be considered in order to isolate the short-term local influences among the simultaneous variables.

Bringing these points together, dynamic causal models divide up a variables set into subsets of synchronic variables such that variables within each subset influence one another locally over shorter time-scales and reach a short-term equilibrium independently of variables outside of the subset. The variables within a subset depend on those outside of the subset, but only in an aggregate way. In other words, the temporal relations within and across the subsets correspond to those between and across the subsystems of a near-decomposable system, according to the definition in section 1. Dynamic causal models use derivative-integration link pairs to integrate both these shorter-term local and longer-term aggregate relationships into a single representation.

To summarize, we considered an example of a system away from equilibrium. In the example, a dynamic model was necessary to derive the set of causal relationships that we would intuitively ascribe to the variables in the system. We further saw that dynamic

---

<sup>7</sup>Shannon Nolen helpfully suggested this explication. A “bottleneck” is a rate-limiting factor.



causal models work by modeling the system as near-decomposable. We thus have a concrete illustration of how decomposing a system is necessary to represent the causal relationships among its variables away from equilibrium.

## 5. THE TIME-SCALE RELATIVITY OF CAUSAL REPRESENTATIONS

We now have a link between causal representations and near-decomposability. But one might still suppose that our discussion is limited to cases in which we are concerned with a system’s equilibrium behaviors. In this section, I propose that the tools presented here enable one to think more generally about how causal representations vary with the time-scale at which a system is considered.

Operationally, considering the system at a longer time-scale corresponds to sampling it at a slower rate and considering it at a shorter time-scale corresponds to sampling it more frequently. Equilibrium models may be thought of as those in which one considers the system at a longer time-scale at which all of the variables have had time to reach equilibrium, while in dynamic models not all of the variables have reached equilibrium. To be clear, equilibrium and dynamic models do not necessarily need be explicated in terms of time-scales – the difference between them concerns whether the variables are at steady-state. But for the wide range of systems with relatively stable equilibria and where perturbations away from equilibrium are relatively transient, equilibrium and dynamic causal models provide a basis for thinking about the way our causal representations of a system differ across longer and shorter time-scales.

Iwasaki and Simon (1994) provide two formal operators for deriving how a system’s representations change as one varies the time-scale at which it is considered: *equilibration* and *exogenization*. Applying equilibration to a variable in a dynamic model yields a model in which that variable has reached equilibrium.<sup>8</sup> One can think of this in terms of “zooming out” to a time-scale at which the time it takes for that variable to reach steady state is so small relative to that time-scale that we can treat it as instantaneous. In contrast, exogenization corresponds to “zooming in” and considering the system at a shorter time-scale. Doing so can have the following effects. First, variables that depend on other variables only minimally at the shorter time-scale can be treated as not depending on them at all. Second, variables changing extremely slowly at the shorter time-scale can be treated as constants. Both of these transform variables that are non-exogenous in the original representation to exogenous ones.<sup>9</sup>

I can now present my proposal for why dynamic models tell us something about causal representations more generally. Even causal representations not explicitly given as relative to a time-scale often rely on assumptions that are valid only at particular time-scales.

---

<sup>8</sup>For a description of equilibration, see Iwasaki and Simon (1994) or Dash (2003). In the bathtub model, equilibrating  $D$  involves replacing  $D'$  with 0 in equation (11), deleting equations (4) and (5), and solving for the causal ordering. The reader may verify that this yields the equilibrium causal model. Intriguingly, equilibration resembles what Wilson (2017) calls “Euler’s equilibrium-based strategy” (p. 77, see also pp. 63-70, 75-80)

<sup>9</sup>In dynamic causal models, the set of exogenous variables include not only constants, but also variables depending on their time-derivatives alone (i.e. those changing as a function of time).




FIGURE 5. Effects of Rainfall and Acres of Wheat Planted on Wheat Grown

Consider an example from Simon and Rescher (1966) in which the amount of wheat grown in a field depends on the amount of rainfall and the amount of acres planted (figure 5). *Rainfall* is about as uncontroversially exogenous as any variable – rainfall influences crop growth, not vice versa. But over a sufficiently large time-scale – say, a century – agriculture does influence climate. So the amount of wheat grown does influence rainfall, although the effect is minuscule at shorter time-scales. The basis for treating rainfall as exogenous is that if we are considering the relationship between rainfall and wheat growth over several years, we can ignore this longer-term influence. Formally, if began with a model with a cycle between *rainfall* and *wheat*, we could derive the shorter term model in figure 5 by exogenizing rainfall. The non-exogenized model would represent a broader system of which our model is a subsystem.

This simple example illustrates how causal representations can be implicitly time-scale relative. It is commonly remarked that the behaviors of systems (especially complex ones) vary with the time-scale at which they are considered. But, to my knowledge, no philosopher has offered a concrete illustration of how a system’s causal representations vary with time-scale or explained how the representations relate to one another. Equilibration and exogenization allow one to precisely specify the relationships between the representations of a system characterized at different time-scales.

Having formal operations for the way a system’s representations vary with time-scale is crucial for understanding the abstractions and idealizations involved in considering a system at a particular time-scale. To show that a model involves time-scale-dependent assumptions, one can illustrate how to derive it using equilibration or exogenization from models at different time-scales. A standard gloss on the distinction between abstractions and idealization is that while the former leave out details, the latter make false claims about the system. Equilibration involves abstraction in that it involves omitting details about the system’s shorter-term behavior away from equilibrium, and idealization in that a system is unlikely to be *exactly* at equilibrium and its interactions will not be truly instantaneous. Exogenization involves idealization, since the exogenized variables are not truly causally independent of the others. As usual, the line between abstraction and idealization is blurry, and nothing here rests on where we draw the line. What matters is that one can use these operations to spell out the assumptions involved in considering a system at a time-scale and to determine which features of our representations will depend on which assumptions.

Near-decomposability allows us to see why causal representations are time-scale relative. Causal attributions, on Simon’s account, rely on assumptions about which subsets of variables are independent. In near-decomposable systems, facts about variable independence are time-scale relative. At shorter time-scales we treat the subsystems as independent and view influences from other subsystems as exogenous. The subsystems are not independent at longer time-scales, and thus the set of variables considered exogenous changes. Causal representations are time-scale relative because they rely on facts which are themselves time-scale relative: independence and exogeneity.<sup>10</sup>

## 6. IMPLICATIONS

The framework presented provides a new tool for thinking about the content of causal models. For instance, Ismael (2016), following Pearl (2009), claims that the asymmetry of cause-effect relationships is relative to a choice of exogenous and endogenous variables. But what does this choice depend on? Using the present framework, we can see how a variable’s exogeneity depends on the time-scale at which the system is considered and its relationship to broader external systems (such as the broader climate system in the agriculture example). The question of whether, in general, exogenous variables correspond to exogenized systems has broad metaphysical implications. An affirmative answer to this question would support Hausman’s (1998) contention that there are no asymmetric causal relationships in closed deterministic systems. In contrast, (Frisch, 2014, p. 95) denies that causal models require exogenous variables implying the existence of a larger subsystem in order to be causally interpreted. Although the present discussion does not resolve this debate, it presents a way of making the relationship between exogenous variables and candidate broader systems concrete and reveals why in at least a wide range of cases the causal relations in a modeled system will depend on its relationship to a broader system.

The framework is also helpful for making headway in debates about whether causal (and/or mechanistic) explanations are appropriate in complex systems consisting of highly interdependent variables (e.g. Chemero and Silberstein, 2008; Rathkopf, 2018). Here the concern is that such systems are allegedly not decomposable and that decomposition is necessary for causal/mechanistic explanations. On the one hand, the foregoing confirms the link between decomposition and causal explanation by revealing that one must decompose the bathtub system to represent its non-equilibrium causal relationships. On the other hand, the discussion reveals that the variables in a system can *both* be highly interdependent in the long-run *and* nearly independent at shorter time-scales. These debates would benefit from more careful attention to the scale-relativity of representations, since this matters both for the decomposability of and the causal relationships that are attributed to a system.

---

<sup>10</sup>Jim Woodward’s symposium contribution also emphasizes independence and scale-relativity, although his independence corresponds to the relative invariance of coarse-grained variables to more fine-grained details of their realization. This provides a distinct and complementary basis for understanding the scale-relativity of causal representations.

## 7. CONCLUSION

The foregoing provides a proposal for thinking about how causal representations of a system vary with the time-scale at which it is considered. One aim of this discussion was to illustrate the philosophical payoff of having a formal representation of (a form of) time-scale dependence. Since the assumptions that here produce time-scale relativity involve idealizations, one might suppose that such relativity is *merely* a feature of the representation, and that in principle one could develop a representation that is not similarly scale-dependent. But I would argue that the types of idealizations described are inextricable from the foundational concepts of causal modeling. Assumptions about exogeneity and independence are the ingredients from which causal representations are built. If these concepts are themselves time-scale relative, then time-scale relativity is an indispensable feature of causal representations.

## ACKNOWLEDGEMENTS

Thanks to Jim Woodward, Lauren Ross and Ken Kendler for inviting me to participate in the PSA colloquium "Strategies for Dealing with Causal Complexity". Thank you to the Pittsburgh Center for Philosophy of Science for funding me during the year in which I wrote this paper and for the Humboldt foundation for funding me during the period when I revised it. The following people read earlier drafts of the paper and provided useful comments: Colin Allen, Janella Baxter, Tomasz Bigaj, Dan Burnston, Kevin Hoover, Chungyoung Lee, Edouard Machery, Antonella Tramacere, Bill Wimsatt, Jim Woodward, Liying Zhang, and two anonymous reviewers. Finally, I am indebted to Adam Edwards, Jonathan Livengood, Shannon Nolen, and Karen Zwier for years of fruitful discussions about causality and time.

## REFERENCES

- Bechtel, W. and R. C. Richardson (1993/2010). *Discovering complexity: Decomposition and localization as strategies in scientific research*. MIT press.
- Burnston, D. C. (2019). Getting over atomism: Functional decomposition in complex neural systems.
- Chemero, A. and M. Silberstein (2008). After the philosophy of mind: Replacing scholasticism with science. *Philosophy of science* 75(1), 1–27.
- Dash, D. (2003). *Caveats for Causal Reasoning with Equilibrium Models*. Ph. D. thesis, University of Pittsburgh.
- Frisch, M. (2014). *Causal reasoning in physics*. Cambridge University Press.
- Hausman, D. M. (1998). *Causal asymmetries*. Cambridge University Press.
- Hoover, K. D. (2015). The ontological status of shocks and trends in macroeconomics. *Synthese* 192(11), 3509–3532.
- Ismael, J. (2016). How do causes depend on us? the many faces of perspectivalism. *Synthese* 193(1), 245–267.
- Iwasaki, Y. and H. A. Simon (1994). Causality and model abstraction. *Artificial intelligence* 67(1), 143–194.

- Pearl, J. (2009). *Causality: models, reasoning and inference*. Cambridge university press.
- Rathkopf, C. (2018). Network representation and complex systems. *Synthese* 195(1), 55–78.
- Simon, H. (1953). Causal ordering and identifiability. In W. Hood and T. Koopmans (Eds.), *Studies in Econometric Method*. New York: John Wiley and Sons, pp. 49–74. New York: John Wiley and Sons.
- Simon, H. A. (1962). The architecture of complexity. *Proceedings of the American Philosophical Society* 106(6).
- Simon, H. A. and N. Rescher (1966). Cause and counterfactual. *Philosophy of Science* 33(4), 323–340.
- Weinberger, N. (2019). Reintroducing dynamics into static causal models. In S. Kleinberg (Ed.), *Time and Causality Across the Sciences*. Cambridge University Press.
- Wilson, M. (2017). *Physics Avoidance: and other essays in conceptual strategy*. Oxford University Press.
- Wimsatt, W. C. (1972). Complexity and organization. In *PSA: Proceedings of the biennial meeting of the Philosophy of Science Association*, Volume 1972, pp. 67–86. D. Reidel Publishing.