Restructuring Dynamic Causal Systems in Equilibrium

Denver Dash
Intel Research
3600 Juliette Lane, SC12-303,
Santa Clara, CA 95054, USA
denver.h.dash@intel.com

Abstract

In this paper I consider general obstacles to the recovery of a causal system from its probability distribution. I argue that most of the well-known problems with this task belong in the class of what I call *degenerate* causal systems. I then consider the task of discovering causality of dynamic systems that have passed through one or more equilibrium points, and show that these systems present a challenge to causal discovery that is fundamentally different from degeneracy. To make this comparison, I consider two operators that are used to transform causal models. The first is the well-known *Do* operator for modeling manipulation, and the second is the *Equilibration* operator for modeling a dynamic system that has achieved equilibrium. I consider a set of questions regarding the commutability of these operators i.e., whether or not an equilibrated-manipulated model is necessarily equal to the corresponding manipulated-equilibrated model, and I explore the implications of that commutability on the practice of causal discovery. I provide empirical results showing that (a) these two operators sometimes, but not always, commute, and (b) the manipulated-equilibrated model is the correct one under a common interpretation of manipulation on dynamic systems. I argue that these results have strong implications for causal discovery from equilibrium data.

1 Introduction

Causal Discovery refers to a special class of statistical analysis that seeks to infer, from a set of data, information about causal relations between variables. There has been much success on the topic of causal discovery in the past decade in Artificial Intelligence [Spirtes *et al.*, 2000; Verma and Pearl, 1991; Heckerman *et al.*, 1999; Tian and Pearl, 2001], building on structural-equation modelling techniques originating in early econometrics [cf., Simon, 1953; Wold, 1954].

There are, as one might expect, many difficulties with inferring reliable causal relationships from data. Latent common causes confounding relations between the observed variables, nonlinearity, acyclicity, and violations of faithfulness due to the cancelling of multiple causal paths are just a few. Identifying prospective pitfalls such as these is the critical first step to developing techniques to handle them in a principled way.

This paper exposes another obstacle to causal discovery that is likely prevalent and important, but is not currently being addressed by causal discovery research. I describe this event as *a violation of equilibration-manipulation commutability* (or *EMC violation*, for short), for reasons that I hope to make clear shortly. I show that EMC violation occurs in static systems, but when those systems have an underlying dynamics which have passed through some equilibrium points. I illustrate the existence of EMC violation by example. Then as further validation, I provide empirical results showing that EMC violation occurs in practice, and its occurrence depends on the time-scale at which the data is being collected relative to the important time-scales of the underlying dynamic systems. I argue that, since many real-world static systems are essentially equilibrium points of underlying dynamic systems, EMC violation is likely to be a common occurrence. I also argue that one can reduce the chance of an EMC violation when building causal models by taking care when choosing the set of variables to include in one's model.

In Section 2, I define some background concepts and explore known obstacles to causal discovery; in Section 3, I show a motivating example of a dynamic causal system going through equilibrium, I define the

EMC property and show why it is important; in Section 4 I show empirically how an EMC violating system can impact causal discovery in practice; in Section 5 I sketch two theorems that show sufficient conditions for systems to violate and obey EMC, and finally I conclude in Section 6.

2 Background Concepts

In this section I define a *causal system*, I explore known obstacles to causal discovery, I introduce the EMC questions and I demonstrate why these questions are important.

2.1 Causal Discovery

I define a *causal system* [c.f., Pearl, 2000] in terms of a set of structural equations:

Definition 1 (causal system) *A causal system over a set of variables* $\mathbf{V}$ *is a 4-tuple* $\langle \mathbf{U}, \mathbf{V}, \mathbf{E}, \phi \rangle$*, where* $\mathbf{U}$ *is a set of random variables that are determined outside the system ("exogenous variables"),* $\mathbf{V} = \{V_1, V_2, \ldots V_n\}$ *is a set of n variables determined by the system ("endogenous variables"),* $\mathbf{E}$ *is a set of n equations, and* $\phi : \mathbf{V} \rightarrow \mathbf{E}$ *is an onto mapping such that for every* $V_i \in \mathbf{V}$*,* $\phi(V_i)$ *can be written as* $V_i = f_i(\mathbf{Pa}_i, \mathbf{U}')$*, where* $\mathbf{Pa}_i \subseteq \mathbf{V} \setminus \{V_i\}$*,* $\mathbf{U}' \subseteq \mathbf{U}$*, and* f_i *is a function.*

A causal system defines a directed graph over variables in $\mathbf{V}$ as follows: For each V_i, let $\phi(V_i)$ be written as $V_i = f_i(\mathbf{Pa}_i, \mathbf{U}')$, and draw an arc from all variables $P_i^j \in \mathbf{Pa}_i \cup \mathbf{U}'$ to V_i. A graph constructed in this way is called a *causal graph*, and if P_i^j is a parent of V_i in this graph then P_i^j is a *cause* of V_i, and V_i is an *effect* of P_i^j. All Bayesian networks can be mapped onto a causal system [Druzdzel and Simon, 1993], but the converse is not true, e.g., causal systems can define cyclic graphs.

All randomness in a causal system is induced by the exogenous variables, which are assumed to be controlled by external forces and therefore are treated as random variables. To say that an equation $\phi(V_i)$ is *deterministic* means that $\mathbf{U}' = \emptyset$, in which case, V_i is a deterministic function of $\mathbf{Pa}_i$. Because the variables in $\mathbf{U}$ are random variables, and because in general the variables in $\mathbf{V}$ depend on $\mathbf{U}$, the causal system $S = \langle \mathbf{U}, \mathbf{V}, \mathbf{E}, \phi \rangle$ will define a probability distribution P over the set $\mathbf{V}$. A common assumption is to assume that each endogenous variable V_i in a causal system S depends on a single exogenous variable U_i and for all i, j, U_i is independent of U_j.

Causal systems such that all f_i functions are linear and all $U_i \in \mathbf{U}$ are normally distributed are called *linear structural equation models*, and for decades these have been widely used in econometrics and the social sciences to model causality [c.f., Simon, 1953; Wold, 1954].

Causal discovery or *causal inference* is the task of analyzing a probability distribution P, and possibly other background information I, to reconstruct the causal system S that generated P. In practice, however, even if P is known exactly, causal inference can do no better than identifying the *set* of causal models that define distributions identical to P that are consistent with I. Probability distributions which do not uniquely define a causal system are the most commonly observed obstacles to causal discovery. I call a causal system for which that is the case *degenerate*. Specific instances of features of causal systems that lead to degenerate probability distributions have been identified and are discussed in the next section.

2.2 Degenerate Causal Systems

Examining the conditional independence relations present in the probability distribution is a key method for causal discovery. Obviously, it is the presence of these relations that increases the specificity of the distribution and makes identification of causal relations possible. One of the most general problems one encounters when trying to perform causal inference from independence relations is a lack of *faithfulness*:

Definition 2 (faithfulness) *A probability distribution* $P(\mathbf{V})$ *over a set of variables* $\mathbf{V}$ *is* faithful *to a directed graph* G *over* $\mathbf{V}$ *if, for every conditional independence relation* $(V_1 \perp V_2 \mid \mathbf{V}')$ *in* P*, there exists a d-separation condition* $(V_1 \perp_d V_2 \mid \mathbf{V}')$ *in* G *and vice-versa*[1]*, for* V_1*,* $V_2 \in \mathbf{V}$ *and* $\mathbf{V}' \subset \mathbf{V}$*.*

If P is faithful to G then G is called a *perfect map* or *p-map* of P. P is called *causally faithful* to G when P is faithful to G, and G is a causal graph. Specific cases in which unfaithful distributions can be generated from real causal systems have been identified in Spirtes *et al.* [2000]. Two in particular are:

- **Determinism:** when a variable in a causal system depends deterministically on other variables. For example, in the causal graph with three variables $\{A, B, C\}$ such that: $A \rightarrow B \rightarrow C$ and $A \rightarrow C$, if C is a deterministic function of A, then C is independent of B given A although that d-separation condition does not exist in the causal graph.

[1] Some definitions of faithfulness do not require the converse.

- **Cancelling causal paths:** when two or more causal paths exactly cancel out. This event can make two or more variables non-correlated although they are causally connected. Although this is possible in principle, Spirtes *et al.* [2000] argue that its occurrence has Lebesque measure zero.

Other reasons for causal degeneracy are:

- **Statistical Indistinguishability**: when there exist other causal structures that have the same set of *adjacencies* and *v-structures*.
- **Lack of causal sufficiency.** A common cause $C \leftarrow A \rightarrow B$ will cause a dependence between C and B in the probability distribution over these variables. If A has been marginalized out of the distribution P, it becomes difficult to decide whether there is a direct causal arc between C and B given only P.
- **cyclic causality:** when a directed cycle exists in the causal graph. Although the physical relevance of these systems can be argued, they are not forbidden by definition, and the implications of their existence on independence relations is not fully explored.

One mitigating fact for all of these obstacles is that their occurrences are all detectable post-causal-discovery, at least sometimes: Determinism and cancelling causal paths will be detectable in the parameters of the model; statistical indistinguishability will be identifiable from the structure of the model; hidden common causes and cyclic causality can sometimes be detected: for example, when their presence causes many v-structures, the PC algorithm for causal discovery [Spirtes *et al.*, 2000] can produce bi-directed arcs or cycles, respectively.

Degenerate causal systems are on one hand problematic, but on the other hand are easy to understand. In the next section I introduce a qualitatively different type of obstacle to causal discovery which, in the author's opinion, is much less transparent and therefore more interesting than causal degeneracy. I call it "violation of *Equilibration-Manipulation Commutability*."

3 The EMC Property

When a causal system is based on a set of differential (or difference) equations, the probability distribution it specifies will not be static, but instead will be a function of time. The evolution of the probability distribution should be predictable. For example, consider the following discrete-time first-order difference system where the change, ΔX, in some variable X, is determined by a linear balance of factors:

$$\begin{aligned} X^0 &= x_0 && (1)\\ F_1^t &= \alpha_1 U_1^0 && (2)\\ F_2^t &= \alpha_2 X^t + U_2^0 && (3)\\ \Delta X^t &= \alpha_3 F_1^t + \alpha_4 F_2^t + U_x^0 && (4)\\ X^{t+1} &= X^t + \Delta X^t, && (5) \end{aligned}$$

where all α_i are constants. In this system, I have assumed that the exogenous variables $\{U_1^0, U_2^0, U_x^0\}$ are static throughout time (which is why they have a fixed $t = 0$ superscript). The causal graph for this system unrolled out to three time slices is shown in Figure 1-(a). The dotted boxes around F_1, F_2 and ΔX are used to denote the fact that the exogenous variables are static through time and are thus parents of those variables in each time slice. For conciseness I will use a shorthand graph, based on the notation of Iwasaki and Simon [1994], where arcs that occur through time are shown with dotted lines, instantaneous arcs are shown in solid, and static exogenous arcs shown dashed. The corresponding shorthand graph for our toy example is shown in Figure 1-(b). It should be emphasized that, although the shorthand version of this graph contains a directed cycle, it represents an acyclic dynamic graph. Sometimes I may also drop the exogenous variables

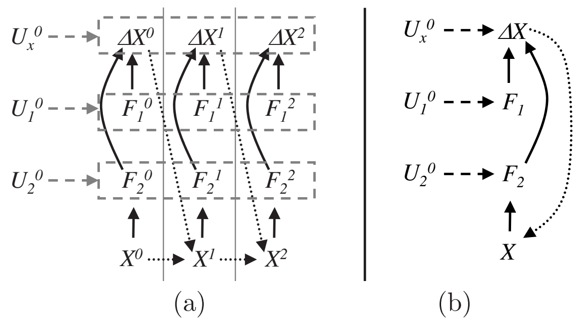

Figure 1: (a) A toy example dynamic causal graph based on difference equations and with static exogenous variables, and (b) the same graph in "shorthand" form.

from this graph to emphasize the endogenous causal relations.

If one considers what the probability distribution over the endogenous variables of this system at the nth time slice will look like, one needs only to expand this system out n slices and marginalize out all variables from previous time-slices to see the causal structure and to generate the probability distribution of the variables at the nth slice. The problem, however, is, if one waits long enough, it may very well be that this system achieves *equilibrium*, at which time there will be a qualitative change in the probability distribution. Specifically, at equilibrium, $\Delta X^t = 0$, so the system of equations reduces to:

$$
\begin{aligned}
F_1 &= \alpha_1 U_1^0 && (6)\\
F_2 &= -\alpha_3 F_1/\alpha_4 - U_x^0 && (7)\\
X &= F_2/\alpha_2 - U_2^0 && (8)
\end{aligned}
$$

In the distribution defined by this system, although U_2^0 is a *parent* of F_2 in the original causal system, it is marginally independent of F_2 in the equilibrium equation system. This fact is obvious by looking at the independence graph of this system shown in Figure 2, and it can also be easily derived from the equation system assuming independent exogenous variables.

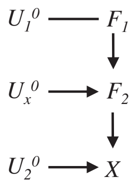

Figure 2: The independence graph of the equilibrium distribution defined by our toy causal system. Although U_2^0 is a parent of F_2 in the original dynamic system, it is marginally independent of F_2 in the equilibrium distribution.

This example illustrates a novel unanswered question associated with dynamic causal systems. Namely, if a causal system is a set of equations with some structure imposed upon it, and if, when a non-structural equation system goes through equilibrium, the equations go through a qualitative change, how should the causality of a system passing through an equilibrium point be modelled? That is, if some equations and variables are dropping out, how should the remaining equations be structured? In the above example, Equation 4 was originally used to determine ΔX; however in equilibrium ΔX has dropped out and Equation 4 has changed into Equation 7 and now "determines" F_2. In fact, if one were to learn a causal graph given the equilibrium distribution, obviously one would in general recover a totally different causal structure than would be recovered from the non-equilibrium system at some arbitrary time slice n. I show this fact empirically in Section 4.

It has been argued by Iwasaki and Simon [1994] that the causal relations governing a dynamic system can change as the time-scale of observation of the system is increased. In particular, they introduce the *Equilibration* operator that they argue produces the causal relations of a system in equilibrium given the dynamic (non-equilibrium) causal system. A detailed treatment of the equilibration operator is beyond the scope of this paper (see Iwasaki and Simon [1994] for more details), a sketch of the operator is as follows:

Definition 3 (Equilibration (sketch)) *Let $M = \langle \mathbf{U}, \mathbf{V}, \mathbf{E}, \phi \rangle$ be a causal model, and let $X \in \mathbf{V}$ be a dynamic variable in M. Equil(M, X) is a causal model $M' = \langle \mathbf{U}', \mathbf{V}', \mathbf{E}', \phi' \rangle$ that is defined by:*

1. *$\mathbf{V}'$ is equal to $\mathbf{V}$ with all of X's derivatives removed.*
2. *$\mathbf{E}'$ is equal to $\mathbf{E}$ with all integration equations removed, and*
3. *$\phi' : \mathbf{V}' \to \mathbf{E}'$ is an onto mapping.*

In general, such an operation may not define a unique mapping ϕ'; however, in the remainder of the paper I assume that ϕ' is unique and only present examples for which that is the case.

As we have done with our toy example, the equilibration operator formally makes the assumption of equilibrium which causes a modification to the equations of the system, then it recovers a structure consistent with the remaining set of equations. In many cases, there is a unique (independence) structure remaining (as in Figure 2). Iwasaki and Simon argue that under these circumstances that structure must be the causal structure of the system under equilibrium.

The *Do* operator, $Do(M, \mathbf{U} = \mathbf{u})$, is another operator of a causal system that transforms a causal model M to a new causal model M' where a subset of variables $\mathbf{U}$ in M' are fixed to specific values independent of the causes of $\mathbf{U}$. On the other hand, the *Equilibration* operator, *Equil*(M, X), transforms the model M with a dynamic (time-varying) variable X to a new causal model M' where X is static. This paper considers the relationship between these two operators. In particular I am interested in the what I call the *Equilibration Manipulation Commutability* property, or the *EMC* property for short:

Definition 4 (EMC Property) *Let $M(\mathbf{V})$ be a causal model over variables $\mathbf{V}$. M satisfies the Equilibration-Manipulation Commutability (EMC) property iff*

$$\textit{Equil}(\textit{Do}(M, \mathbf{U} = \mathbf{u}), X) = \textit{Do}(\textit{Equil}(M, X), \mathbf{U} = \mathbf{u}),$$

for all $\mathbf{U} \subseteq \mathbf{V}$ and all $X \in \mathbf{V}$.

In this paper, I consider the following set of questions (hereafter referred to as *the EMC questions*):

1. Does the EMC property hold for all dynamic causal models?

2. Does the EMC property hold for any dynamic causal models?
3. Under what conditions is the EMC property guaranteed to hold?
4. Under what conditions is the EMC property guaranteed to be violated?
5. In general, is it sensible to reason about causality in a dynamic system that has passed through some equilibrium points?

These questions are important for at least the following reason: Very often in practice a causal model is first built from either data or knowledge of equilibrium relationships, and then causal reasoning is performed on that model. This common approach takes path A in Figure 3. When a manipulation is performed on a system, however, the state of the system in general becomes "shocked" taking the system out of equilibrium, a situation which is modelled by path B in Figure 3. The validity of the common approach of taking path A thus hinges on the answers to the EMC Questions.

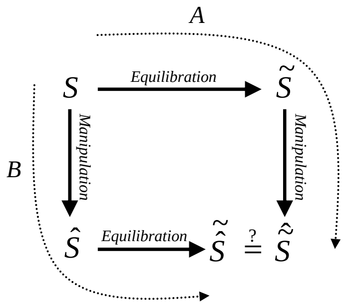

Figure 3: The EMC Questions consider under what conditions the *Do* operator commutes with the Equilibration operator operating on a dynamic causal model S.

This paper primarily answers EMC Questions 1, 2, and 5. The toy example I prove by example and by empirical tests that the answer to Question 1 is "No" and that of Question 2 is "Yes". These results in turn implies that the answer to Question 5 is "Sometimes". The answers to Questions 3 and 4 are addressed in Dash [2003], but I will sketch those results here as well.

4 Discovery from Data: Empirical Results

The previous section presented an example which implied that the answer to EMC Question 1 is "no". This section addresses the EMC Questions using empirical studies. I performed numerical simulations of a dynamic system to demonstrate that as the time scale was increased enough so that an equilibration could occur, the causal structure that was learned from data corresponds to the structure obtained by applying the *Equilibration* operator to the dynamic model. This fact is significant because it indicates that whenever a causal structure that is learned from equilibrium data is used for causal reasoning, then Path A of Figure 3 is being taken: if the EMC property does not hold for the model being used then subsequent causal reasoning will produce incorrect results.

Consider the causal system of five variables $\{Q_{in}, Q_{out}, D, K, P\}$ defined by Equations 9–13 below.

$$\begin{aligned} K &= K_0 && (9)\\ Q_{in} &= Q_0 && (10)\\ \dot{P} &= \alpha_2(\alpha_4 D - P) && (11)\\ \dot{Q}_{out} &= \alpha_3(\alpha_1 K P - Q_{out}) && (12)\\ \dot{D} &= \alpha_0(Q_{in} - Q_{out}) && (13) \end{aligned}$$

where $\dot{Q}_{out}$, $\dot{D}$ and $\dot{P}$ are the first time-derivatives of Q_{out}, D and P, respectively, and $\alpha_i : i \in \{0, 1, \ldots, 4\}$ are constants.

This system was taken from Iwasaki and Simon [1994]. To give some physical intuition, it roughly approximates a filling-bathtub where water is entering the bathtub from the faucet at a rate Q_{in} liters per second and is exiting the drain at a rate Q_{out} liters per second. The pressure of the water at the base of the drain is P, the depth of the water is D, and the diameter of the drain is K. In this system Q_{in} and K are exogenous and the remaining variables are dynamic. The causal graph of this system is shown in Figure 4.

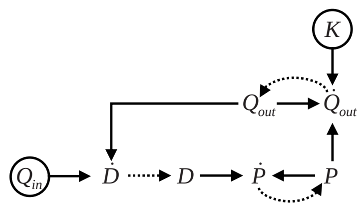

Figure 4: The dynamic causal graph S_0 of the bathtub system.

This system has three dynamic variables, and therefore three possible equilibrations, corresponding to the occurrence of $\dot{P} = 0$, $\dot{Q}_{out} = 0$, and $\dot{D} = 0$. When these conditions occur, Equations 11, 12, and 13 reduce to the equilibrium relations given by Equations 14–16, respectively:

$$\begin{aligned} P &= \rho g D && (14)\\ Q_{out} &= \alpha_1 K P && (15)\\ Q_{in} &= Q_{out} && (16) \end{aligned}$$

This system has many potential equilibrium causal orderings depending on the relative time scales of the three variables P, Q_{out} and D, and depending on the time scale at which the system is observed. If, for example, P and Q_{out} both reach equilibrium much sooner than D, and the system is observed before D has reached equilibrium but after P and Q_{out}, then Equations 11 and 12 get replaced by Equations 14 and 15, respectively. The causal ordering of this system is shown in Figure 5.

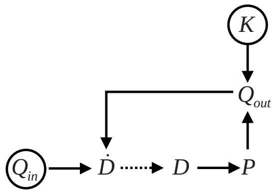

Figure 5: The causal ordering of the bathtub system when P and Q_{out} have been equilibrated but not D.

Because the system of Figure 4 involves three dynamic variables, there exist three important time-scales for this system, controlled by the inverse of the coefficients: $\tau_D \propto 1/\alpha_0$, $\tau_P \propto 1/\alpha_2$, and $\tau_Q \propto 1/\alpha_3$, for D, P and Q_{out} respectively. If $\tau_P \ll \tau_Q \ll \tau_D$, then there will exist four possible equilibrium causal structures learned from data depending on the time, τ, at which the data was observed. These four structures (over variables $\mathbf{V} = \{Q_{in}, Q_{out}, P, T, K\}$) are shown in Figure 6. At $\tau = 0$ each of the five variables in

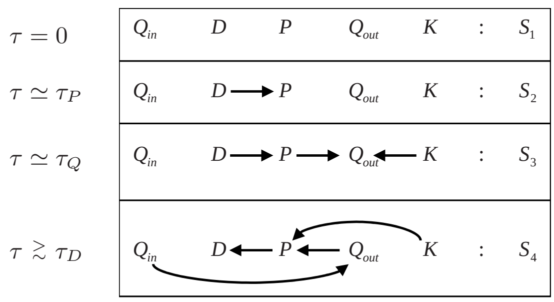

Figure 6: The bathtub system has four correct equilibrium structures depending on the time scale at which the system is observed.

$\mathbf{V}$ are given by their initial conditions and so are exogenous; in this case S_1 will be the structure learned from data. After enough time has passed for P to equilibrate ($\tau \simeq \tau_P$), then Equation 11 reduces to Equation 14, and the structure S_2 will result. After $\tau \simeq \tau_Q$, enough time has passed for Q_{out} to equilibrate, and Equation 12 reduces to Equation 15, resulting in the structure S_3. Finally, after $\tau > \tau_D$, enough time has passed for D to equilibrate and Equation 13 reduces to Equation 16, leading to a drastic restructuring of equations and resulting in model S_4.

I simulated learning over several time-scales for the filling-bathtub system. The following values for constants were used: $\alpha_1 = 1$, $\alpha_0 = 0.005$, $\alpha_2 = 0.05$, and $\alpha_3 = 0.01$. All variables were initialized from the uniform distribution over the interval $(0, 1)$. Independent Gaussian error terms with mean 0 were added to each derivative variable. The error terms for $\dot{D}$ and $\dot{Q}_{out}$ had standard deviation equal to 0.01, and $\dot{P}$ had standard deviation equal to 0.5. I assumed the bathtub was infinitely high (no bound on D was enforced), so given these constants, an equilibrium was guaranteed to exist.

A database of $N = 10000$ records was generated for each of the 29 time-scales given in the set $\mathcal{T} = \{0-10,$ 20, 30, 40, 50, 80, 100, 125, 150, 200, 250, 300, 500, 750, 1000, 1250, 1500, 1750, 2000$\}$, and for each of these databases the PC algorithm was run to retrieve a causal graph. A modified version of PC was used which forbade cycles or bi-directional arrows and randomized the order in which independencies were checked [Dash and Druzdzel, 1999]. Data for each variable took on a continuous range of values, and in all cases the Fisher's-z statistic was used to test for conditional independence using a significance level of $\alpha = 0.05$.

I restricted structure learning to the variables $\{D$, P, Q_{out}, Q_{in} and $K\}$, namely the variables relevant to the static analysis of this system. This was performed 50 times for each time scale, and the number of times the pattern corresponding to the graphs in Figure 6 were exactly recovered was counted.

The normalized results, showing the empirical probability of retrieving the four structures as a function of the time scale, are shown in Figure 7. For example, when the system is observed just one time step away from the initial conditions, Figure 7 shows that structure S_2 was learned around 45% of the time, structure S_1 was recovered around 6% of the time, structure S_3 was discovered less than 5% of the time, and some other structure was learned the remainder (about 44%) of the time. The time-scales $20 \leq \tau \leq 750$ are excluded from this figure—they produced empirical probabilities of 0 for all four structures. These results show convincingly that as the time-scale increases, the learned causal structure changes in the order predicted by the

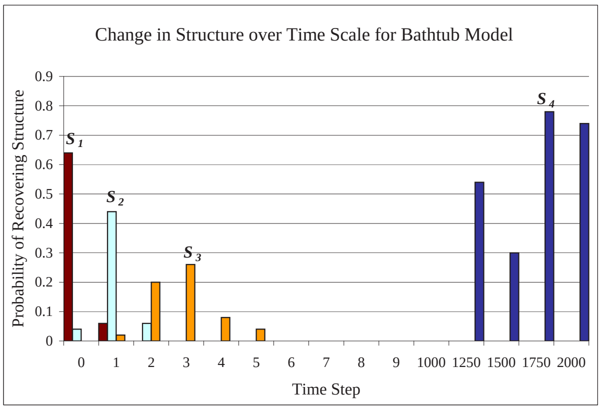

Figure 7: As the time step is varied, each of the four equilibrium structures can be recovered in sequence.

equilibration operator.

It is easy to verify that the EMC property holds when only P or Q_{out} are equilibrated and any of these five variables are manipulated (by verifying that manipulating any variable in S_1 or S_2 results in the same graph as manipulating them in Figure 4 then equilibrating). On the contrary, it is easy to verify that when D is equilibrated, the EMC property is violated: $Do(Equil(S_0, D), D)$ corresponds to S_4 with the arc from P to D removed; whereas $Equil(Do(S_0, D), D)$ (constructed by applying the Do operator to S_0 and then equilibrating all remaining dynamic variables) corresponds to S_3.

5 Theoretical Results: EMC Questions 3 and 4

The results from Section 4 show that in some cases the EMC property is preserved, while in others it is not. While it is beyond the scope of this paper to address the precise conditions when EMC will or will not be violated, I will briefly sketch in this section two results from Dash [2003] with proofs omitted. The first states conditions for which EMC is guaranteed to be violated, the second states conditions for which EMC is guaranteed to be satisfied.

These results involve the concept of a *feedback set*. A feedback set of a variable X in a causal model is the set of variables that are both ancestors and descendants of X in the shorthand causal graph. For example, in Figure 4, the feedback variables of D are $\dot{P}$, P, $\dot{Q}_{out}$, Q_{out} and $\dot{D}$. I let $\boldsymbol{Fb}(X)_M$ denote the set of feedback variables of X in model M.

For the following two theorems, we consider a dynamic causal model $M = \langle \mathbf{U}, \mathbf{V}, \mathbf{E}, \phi \rangle$ and let $M_{\tilde{v}} = \langle \mathbf{U}', \mathbf{V}', \mathbf{E}', \phi' \rangle$ denote the graph that results when a variable $V \in \mathbf{V}$ is equilibrated in M: $M_{\tilde{v}} = Equil(M, V)$. I assume that $M_{\tilde{v}}$ is unique.

Theorem 1 (EMC violation) *If both M and $M_{\tilde{v}}$ are recursive (have acyclic graphs) and there exists any $F \in \boldsymbol{Fb}(V)_M$ such that $F \in \mathbf{V}'$ then $Do(M_{\tilde{v}}, Y) \neq Equil(Do(M, Y), V)$ for any $Y \in \mathbf{V}$.*

For example, in Figure 6, the graph that results when D is equilibrated contains variables that are in the feedback set of D, so the bathtub system violates EMC when D is equilibrated.

Theorem 2 (EMC obeyance) *Let $\Delta^n V$ be the highest derivative (difference) of V in $\mathbf{V}$. If both M and $M_{\tilde{v}}$ are recursive (have acyclic graphs) and $V \in \mathbf{Pa}(\Delta^n V)$, then $Do(M_{\tilde{v}}, Y) = Equil(Do(M, Y), V)$.*

For example, in Figure 4, $\mathbf{Pa}(\dot{P}) = \{P\}$ and $\mathbf{Pa}(\dot{Q}_{out}) = \{Q_{out}\}$, so the bathtub system will obey EMC when either of these variables are equilibrated, as seen in Figure 6.

6 Conclusions

The results of this paper have important consequences for causal discovery. In particular, they emphasize the importance of considering the time-scale of the data being used for causal discovery. If data is recorded of a system for which some variables have achieved equilibrium, then learning a causal graph and using it to predict the effects of manipulating variables in the model amounts to taking path A in Figure 3; however, path B is the correct one to take: if the EMC property does not hold for that model, then incorrect inferences could result.

The fact that taking path A in Figure 3 produces predictions that differ from path B requires us, if we intend to perform causal reasoning with our model, to either ensure that we are taking path B or ensure that we are dealing with models that obey the EMC property. Currently, most work regarding the discovery or building of causal models takes path A and pays no regard to the EMC property. I hope that this work will bring attention to this fact and help to rectify it.

It is a valid question to ask why the EMC property is useful at all. That is, why treat an equation system that has passed through equilibrium as causal? The answer to that question lies in the extreme difficulty of knowing what the important time-scale of an unknown causal system might be. On top of that, to break a system down to its finest time-scale often involves modeling the system in intractable detail. For example, if it were necessary to model the microstates of a statistical ensemble of particles rather than using the macroscopic laws directly, then modeling the causality of any such system would be impossible.

The problem of identifying the relevant time-scales of a system is especially acute for the task of causal discovery (as opposed to building causal models from expert knowledge), because obviously, if one is trying to learn causal relations from data, it is likely that one is not privy to the details of the underlying dynamics of the system. The positive conclusion of this work is that, for systems that obey EMC, one does not need to consider the system on the shortest possible time scale for the resulting model to accurately reflect causality. The negative conclusion, however, is that at least some knowledge of temporal behavior of the system is likely necessary to ensure that the EMC condition is satisfied, and what knowledge is necessary and sufficient is not yet known.

Although this work raises important objections to some uses of causal reasoning with models learned from data, I believe that the great potential of causal modeling and causal discovery in artificial intelligence make it all the more important for these questions to be explored further and answered as forcefully as possible. The fact that equilibrium causal models are problematic for causal inference should not deter us from developing further the theory that can allow us to build and use them in practice.

7 Acknowledgements

This work was extracted from my PhD thesis [Dash, 2003] pursued at the Intelligent System Program at the University of Pittsburgh in partial collaboration with my advisor Marek Druzdzel. I would like to thank Richard Scheines, Nanny Wermuth and Gregory Cooper for their encouragement for this work.

References

Denver H. Dash and Marek J. Druzdzel. A hybrid anytime algorithm for the construction of causal models from sparse data. In *Proceedings of the Fifteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI–99)*, pages 142–149, San Francisco, CA, 1999. Morgan Kaufmann Publishers, Inc.

Denver Dash. *Caveats for Causal Reasoning.* PhD thesis, Intelligent Systems Program, University of Pittsburgh, Pittsburgh, PA, April 2003. http://etd.library.pitt.edu/ETD/available/etd-05072003-102145/.

Marek J. Druzdzel and Herbert A. Simon. Causality in Bayesian belief networks. In *Proceedings of the Ninth Annual Conference on Uncertainty in Artificial Intelligence (UAI–93)*, pages 3–11, San Francisco, CA, 1993. Morgan Kaufmann Publishers.

David Heckerman, Christopher Meek, and Gregory F. Cooper. A bayesian approach to causal discovery. In Clark Glymour and Gregory F. Cooper, editors, *Computation, Causation, and Discovery*, chapter four, pages 141–165. AAAI Press, Menlo Park, CA, 1999.

Yumi Iwasaki and Herbert A. Simon. Causality and model abstraction. *Artificial Intelligence*, 67(1):143–194, May 1994.

Judea Pearl. *Causality: Models, Reasoning, and Inference.* Cambridge University Press, Cambridge, UK, 2000.

Herbert A. Simon. Causal ordering and identifiability. In William C. Hood and Tjalling C. Koopmans, editors, *Studies in Econometric Method. Cowles Commission for Research in Economics. Monograph No. 14*, chapter III, pages 49–74. John Wiley & Sons, Inc., New York, NY, 1953.

Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, Prediction, and Search.* Springer Verlag, New York, 1993.

Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, Prediction, and Search.* The MIT Press, Cambridge, MA, second edition, 2000.

Jin Tian and Judea Pearl. Causal discovery from changes. In *Uncertainty in Artificial Intelligence: Proceedings of the Seventeenth Conference (UAI-2001)*, pages 512–521, San Francisco, CA, 2001. Morgan Kaufmann Publishers.

T.S. Verma and Judea Pearl. Equivalence and synthesis of causal models. In P.P. Bonissone, M. Henrion, L.N. Kanal, and J.F. Lemmer, editors, *Uncertainty in Artificial Intelligence 6*, pages 255 –269. Elsevier Science Publishing Company, Inc., New York, N. Y., 1991.

Herman Wold. Causality and econometrics. *Econometrica*, 22(2):162–177, April 1954.