# Modelling Cyclic Causal Structures*

Alexander Gebharter & Bert Leuridan[1]

**Abstract.** Many causal systems studied by sciences such as biology, pharmacology, and economics feature causal cycles. Most accounts of causal modelling currently on the market are, however, explicitly designed to study acyclic structures. This chapter focuses on causal cycles and the challenges such cycles pose for causal modelling. First we distinguish between different types of causal cycles. Then we introduce causal models and discuss a selection of general challenges for cyclic models when it comes to representation, prediction, and causal discovery. Finally, we zoom in on a concrete case from biochemistry: the PI3K/mTOR signalling network that plays a crucial role in regulating the cell cycle. Using this case



* This is a draft paper. Please do not cite or quote without permission. The final version oft he paper is published under the fokllowing bibliographical data: Gebharter, A., & Leuridan, B. (2024). Modelling cyclic causal structures. In P. M. K. Illari, & F. Russo (Eds.), *Routledge handbook of causality and causal methods* (pp. 269-280). Routledge. doi:10.4324/9781003528937-30

we discuss some possible pitfalls for the application of causal modelling tools to complex biological cases.

**Notes on Contributors**

Alexander Gebharter is an assistant professor and co-director of the Center for Philosophy, Science, and Policy at the Marche Polytechnic University, Italy. He is the author of Causal Nets, Interventionism, and Mechanisms (Synthese Library, 2017).

Bert Leuridan is Professor in Philosophy of Science at the University of Antwerp, Belgium. His research interests focus on causation, mechanisms, scientific laws, explanation, experimentation, evidence and statistics. He has applied his views to, or drawn inspiration from, several special sciences. Recently he has turned to the issue of policy in relation to evidence and causation.



# 1 Introduction

Many of the systems studied by the sciences are not as nice and simple as we would like them to be. Systems studied by biology, medicine, and

economics, for example, often involve causal cycles. Feedback inhibition, for instance, is a typical mechanism for metabolic control in biological organisms. A product is produced at a certain point during a reaction pathway that then inhibits a regulatory enzyme which was involved in the production of that product earlier in the pathway. Another example would be the interaction of a population of rabbits and foxes in a certain region. Changes in the size of the rabbit population have effects on the size of the population of foxes and vice versa. Finally, the more demand there is for a certain good, the higher the price for said good typically becomes (provided supply is held fixed). However, the good becoming more expensive over time will also have a negative effect on how attractive it is for customers, which in turn might decrease the demand. (For a cyclic graph on price and demand and on predator-prey dynamics see, for example, Mella 2008.)

Systems like these come with specific challenges due to their cyclic nature. Some of these challenges concern the question of how they can be modelled. In this chapter we will see that the answer to this question depends on the specific types of causal cycles involved.[2] We introduce different kinds of causal cycles in section 2. In section 3 we introduce causal models and discuss a selection of general challenges for cyclic models when it comes to representation, prediction, and causal discovery. For introductory and

illustration purposes we operate with simple toy examples in this section. In section 4 we then turn to a concrete case from biochemistry, the PI3K/mTOR signalling network. This signalling network plays a crucial role in regulating the cell cycle: cell growth, cell survival and metabolism (for instance of glucose), cell death etc. Using this case we highlight some of the specific challenges for the application of causal modelling techniques to real and complex cyclic structures.

The questions we will discuss are methodological rather than metaphysical: How can we represent cyclic causal systems (representation)? How can predictions be derived starting from cyclic causal models (prediction)? How can cyclic causal models be derived from empirical data (causal discovery)?[3] This is not to say that metaphysical issues are completely absent. A key question, after all, is: What is causality?[4] In response to this, we will adopt a general interventionist view on causation (cf. Woodward 2003). We do this not because we think it is the only viable theory of causation, nor because we think it suffers no problems whatsoever (see, e.g., Baumgartner 2009 2012; Gebharter 2017: sec. 5.3; Strevens 2007). (In fact, Maziarz rightly claims in Chapter CROSSREF-chapter-Maziarz, this volume, that every monistic account of causality faces some criticism and counterexamples.) Our choice is rather motivated by the fact that interventionist accounts of causation are

typically formulated in a language that is familiar to practicing scientists and because they were developed with the formal framework of causal models in mind.[5]

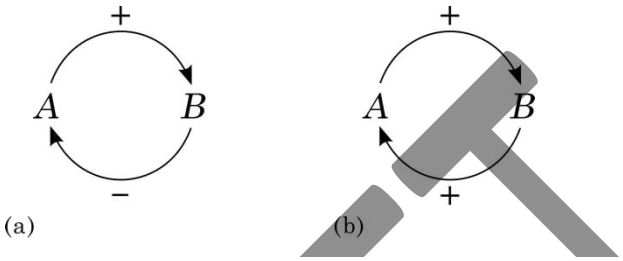


*Figure 1: Negative feedback (a) and positive feedback (b). Contingent feedback consists in fluctuations between (a) and (b) over time.*

From now on, we describe causes and effects in terms of variables (instead of, for instance, letters referring to single events). According to an interventionist understanding, a variable $X_i$ is causally relevant w.r.t. another variable $X_j$ iff changes brought about by wiggling $X_i$ in the right way[6] make $X_j$ (or its probability distribution) wiggle as well. Interventionist intuitions can also be represented in terms of causal Bayesian networks or structural equation models (cf. Gebharter & Schurz 2014; Pearl 2000; Spirtes, Glymour, & Scheines 2000), though these more technical frameworks usually treat causation as a basic concept and abstain from providing an explicit definition of causation in terms of interventions (see, e.g., Gebharter 2017; Glymour 2004; Schurz & Gebharter 2016). However, we bracket most of these technicalities in this chapter and rather try to keep it as informal as possible.

(For more information about the technicalities, see the literature referred to in this paragraph.)

## 2 Types of causal cycles

Before we come to the question of how to best represent cyclic causal structures in section 3, it will be useful to briefly highlight and discuss different types of causal cycles. Clarke, Leuridan, and Williamson (2014) distinguish between the following three types (we add explications; see Figure 1 for a graphical illustration):

- *Negative feedback:* $X_i$ has a positive effect (directly or indirectly) on $X_j$ which in turn has a negative effect (also directly or indirectly) on $X_i$.
- *Positive feedback:* $X_i$ has a positive effect (directly or indirectly) on $X_j$ which in turn has a positive effect (also directly or indirectly) on $X_i$.
- *Contingent feedback:* $X_i$ has a positive effect (directly or indirectly) on $X_j$ which in turn sometimes has a negative and sometimes has a positive effect (also directly or indirectly) on $X_i$.

Concrete examples of each type will be given below.

Since we describe causes and effects in terms of variables, we interpret all three types of feedback as about causal cycles between *types of events* rather than between particular events. (We leave it open whether cases of *particular events* causing themselves exist.)

All three types of causal cycles are closely connected to the notion of a system's *equilibrium state*. An equilibrium state is a state of balance the system approaches in the absence of varying external influences.[7] After shocking the system from the outside, it typically returns to its equilibrium state. Here is a simple (acyclic) example: Assume you are pouring blue colored ink into a basin filled with clear water. At first, the water will be intensely blue close to the exact spot where you poured the ink into the basin, but clear everywhere else. Slowly the color will spread. Larger areas of water become blue too, but the more the ink spreads, the lighter the color becomes. In the end you will find equally light blue water at every spot in the basin. This is the system's equilibrium state.

Or take the predator-prey system briefly mentioned in section 1: Assume the size of the rabbit population has a positive causal effect on the size of the population of foxes (the availability of rabbits makes a fox's life easier), but that the latter has a negative effect on the former (the more foxes there are in the region, the harder it becomes for a rabbit to survive). Now the system may

have natural breaking points. If, for example, the ratio of foxes compared to rabbits is too high, then all rabbits will be eaten by foxes and the population of rabbits will become 0. Shortly after this, also the population of foxes will become 0 due to missing prey (assuming rabbits are the only food source for foxes). Let us further assume that the sizes of the two populations are such that no such natural breaking points are reached. Now the population sizes of rabbits and foxes may push each other such that their magnitudes reach a state of balance that does not change drastically anymore and stably oscillates around specific numbers (in absence of external disturbance factors). This is the system's equilibrium state.[8]

A *negative feedback* cycle typically pushes a system towards its equilibrium state. The negative influence of the size of the fox population on the size of the rabbit population, for example, keeps the latter within its proper bounds (see Figure 2(a)). The same goes for many kinds of metabolic cycles and homeostatic pathways. A *positive feedback* cycle, on the other hand, typically pushes the system away from its equilibrium state. A prime example would be cancer (see Figure 2(b)). Apoptosis (programmed cell death) is a natural process that leads to the death of abnormal or mutated cells. It depends on a complex web of negative feedback cycles which push the organism towards the equilibrium state of having close to no mutated cells. If apoptosis

is drastically diminished, damaged or obsolete cells are not removed and ever more cells get damaged (tumours). Conversely, in cases of hyperapoptosis too many cells die, leading to medical conditions such as neuronal degeneration and diabetes. In both cases the organism is, thus, pushed away from its equilibrium state. A *contingent feedback* cycle sometimes pushes the system towards and sometimes away from its equilibrium state. Cancer might serve as an example here as well. This time, however, we are looking at a somewhat different case. In healthy individuals, apoptosis pushes the number of mutated cells towards zero. But once this mechanism gets damaged enough, the number of mutated cells increases and is drawn away from equilibrium. After the treatment of the patient, things might change and the rate of mutated cells might decrease again for some time. But then cancer might take the upper hand again and so on. Thus, contingent feedback essentially consists of positive and negative feedback cycles interacting with each other and competing for dominance in the long run. Thus, which type of feedback a system involves depends on the specific time at which we choose to describe it as well as on the presence or absence of additional external factors (see again Clarke et al.

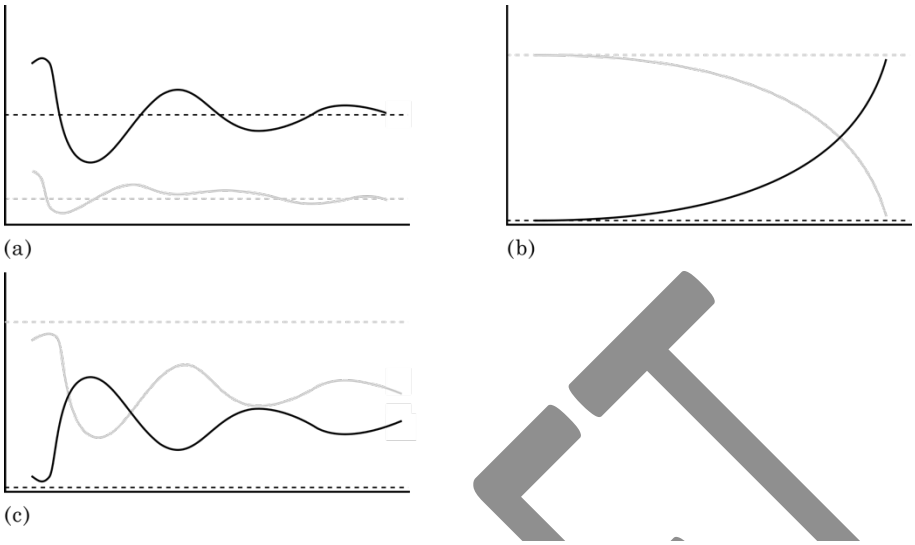


*Figure 2: (a) solid black line: number of rabbits; solid grey line: number of foxes. (b) and (c) solid black lines: number of damaged cells; solid grey lines: apoptosis. (b) shows the untreated case where cancer rapidly spreads and overthrows apoptosis. (c) shows the case where cancer is treated. The diagram leaves it open whether treatment ultimately succeeds.*

2014: 1657, on the notion of granularity). Finally, some systems (e.g., some chaotic ones) do not have equilibrium states at all.

Before turning to causal models, we should address one possible worry. The word "equilibrium" can cause semantic unease as three different ideas (and seemingly opposing notions) are intertwined here. First, the equilibrium involved in homeostasis in biological contexts should neither be interpreted as static, nor as the total absence of order. Turner (2019) describes homeostasis as "persistence of a living system in a state of specified and dynamic disequilibrium" (2019: 3) and as "the active striving of living systems towards a persistent and specified orderliness" (ibid.). Second, the persistent state of a

given living system can be strongly different from that of its surroundings. In this sense they are often called "far-from-equilibrium" (see Chapter CROSSREF-chapter-Bechtel-and-Bollhagen, this volume, by Bechtel and Bollhagen; see also Leuridan & Lodewyckx 2021). Third, the orderliness of a given homeostatic system can be more visible or less visible, depending on granularity (cf. supra). Here is a simple example[9]: Normal body temperature in humans is in equilibrium around 37°C[10] but varies throughout the day, both absent external factors (by 0.25°C to 0.50°C) and due to external factors (causing fever or hypothermia), thus exhibiting a specified and dynamic disequilibrium, and is typically strongly different from the environment's temperature (far-from-equilibrium).

## 3 Causal models and cyclicity

As already mentioned before, we represent types of events by variables. Once a set **V** of variables $X_1,...,X_n$ of interest is chosen, arrows will represent causal dependencies among these variables. The structure resulting from all these arrows is called a causal graph. If $X_i \rightarrow X_j$, then $X_i$ is called a causal parent of $X_j$ and $X_j$ is called a causal child of $X_i$. If two variables $X_i$ and $X_j$ are connected by a chain of directed arrows $X_i \rightarrow ... \rightarrow X_j$, $X_i$ is called a causal ancestor of $X_j$ and $X_j$ a causal descendant of $X_i$. If a path $\pi$ features a substructure of the form

$\rightarrow X_i \leftarrow$, then $X_i$ is called a collider on $\pi$ (since the arrows collide on $X_i$). The semantics determining the causal relations forming such causal structures given an interventionist understanding of causation is as follows: $X_i$ is directly causally relevant for $X_j$ iff there are interventions on $X_i$ changing $X_j$'s value (or probability distribution) if the values of all other variables in **V** are held fixed by additional interventions. A variable $X_i$ connected to another variable $X_j$ via a chain of arrows $\pi : X_i \rightarrow ... \rightarrow X_j$ is causally relevant for $X_j$ iff there are interventions on $X_i$ that change $X_j$'s value (or probability distribution) if the values of all other variables not lying on $\pi$ are held fixed by additional interventions. An intervention on a variable $X_i$ is an event that changes $X_i$'s value directly, meaning that it has an influence on other variables in **V** (if it has one at all) only through directed causal chains running through $X_i$.[11]

In addition to an interventionist interpretation, one can formulate several conditions that, if satisfied, link causal structure to probabilities. One of these is the causal Markov condition (Pearl 2000: 19; Spirtes et al. 2000: 29), which can be stated in different ways. For our endeavor in this chapter, its global version is more useful, since it also applies to cyclic systems (Pearl & Dechter 1996; Spirtes 1995). Informally, the *global causal Markov condition* is satisfied iff every (conditional or unconditional) probabilistic dependence between any two (non-intersecting) sets of variables in **V** can be explained by

some suitable causal connection among these variables. A causal path is suitable to do this job iff (i) every collider on this path is conditioned on (or has a causal descendant that is conditioned on) and (ii) all other causal variables on this path are allowed to vary. (For details why these conditions adequately characterize causal relations see, e.g., Gebharter 2017: sec. 4.2.)

To arrive at a full-fledged causal model that can be used for causal inference, we need to specify the causal strengths of the specific causal links. One can do so by supplementing the causal graph with a probability distribution or by assigning an equation to every variable $X_i$ that determines $X_i$'s value on the basis of its causal parents (and possible latent factors represented by error terms). To keep things simple and since this strategy is more general, we go for probability distributions in this chapter. The causal graph linked together with a probability distribution as specified by the global causal Markov condition allows us now to use the model to generate inferences about what would happen if we would learn the values of some of the variables in **V** by observation or if they were brought about by intervention (Pearl 1988 2000). This allows one to test and assess complex causal hypotheses and sometimes even to abduce the existence of yet undiscovered causal background mechanisms (see, e.g., Koch, Eisinger, & Gebharter 2017). If the causal graph and the probability distribution meet additional conditions

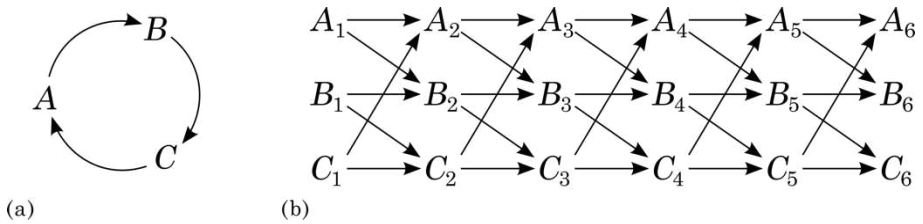


*Figure 3: Simple static (a) and dynamic cyclic causal model (b)*

such as *faithfulness*[12], causal structure can even be systematically learned on the basis of (passive) observation (Spirtes et al. 2000). These features of causal models constitute some of the specific advantages causal models have over other, non-causal modelling approaches. They allow for distinguishing the effect of observing that a particular variable $X$ has taken value $x$ from the effect of setting $X$ to $x$ by a intervention – which is typically not captured by sets of ordinary (i.e., non-causal) equations (cf. Pearl & Mackenzie 2018, introduction). In principle they can also be used even without a deeper understanding of the system's dynamics (for example in the form of differential equations).[13]

So far we have not touched the issue of cyclic causality in this section. There are basically two ways to represent cyclic causal structures, a static and a dynamic one (cf. Clarke et al. 2014; Gebharter & Schurz 2016). Assume we are interested in the variable set $\mathbf{V} = \{A,B,C\}$ and that this system forms a simple causal cycle. In particular, assume that $A$ is directly causally relevant

for *B*, that *B* is directly causally relevant for *C*, and that *C* is directly causally relevant for *A*. The two possibilities of how to model this system are shown in Figure 3.

The model in Figure 3(a) is static; it abstracts away from its development over time, while the one in (b) is dynamic, as it rolls out the causal cycle over time. The model in (b) was rolled out over 6 time steps, but in principle it can cover as many time steps as one wants. Both representations are compatible with interventionist causation: Fixing the value of *A* in (a) by an intervention may lead to changes in its causal descendants and likewise for *B* and for *C*. Representation (b) is a bit more complex: The repeating pattern of arrows between consecutive stages accounts for (i) *A* directly causing *B*, *B* directly causing *C*, and *C* directly causing *A* as well as (ii) the fact that the state of each variable at any time *i* is directly relevant for its state at the next step $i + 1$. To this end, it is important to note that arrows do not skip temporal stages. This means, in turn, that fixing the values of all variables at any stage will screen all earlier stages off from later ones.[14]

Let us start with having a closer look at *static* models. First of all, we can observe that the model's probability distribution can be somewhat tricky to choose (Clarke et al. 2014; Gebharter & Schurz 2016). The problem is that while the causal cycles within the system work over time, the probability

distribution over $\mathbf{V} = \{A,B,C\}$ can change. The specific probabilities observed at a specific time are only guaranteed to correctly represent the causal strengths within the system at that particular time. But we want to create a static causal model that can be used for inference regardless of the specific moment the system is in. There is a straightforward solution to this problem only for cyclic systems featuring negative feedback cycles powerful enough to push these systems towards their equilibrium states and without important oscillations around that equilibrium state (see again Clarke et al. 2014: 1657, on granularity). In such cases, we can take the equilibrium distribution in order to specify the model's probability distribution (cf. Pearl & Dechter 1996; Spirtes 1995). A system's equilibrium distribution is the probability distribution that, once reached by the system in the long run, does not change anymore over time (in the absence of external disturbances).[15] This is a severe limitation of static causal models.

In principle, prediction under intervention in static models works exactly as in ordinary acyclic causal models (cf. Pearl 2000: sec. 1.3). In a first step, one fixes the value of the variable $X_i$ on which one wants to intervene. In a second step, one deletes all the arrows pointing at $X_i$ from one's graph. In a third step, one applies the global causal Markov condition to the graph that resulted from deleting the arrows into $X_i$, which may result in further

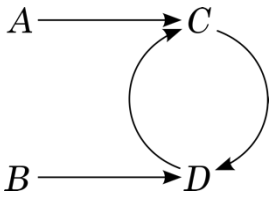


*Figure 4: Simple cyclic structure producing a unique probabilistic footprint*

independence constraints on one's probability distribution. In a fourth step, one computes the probabilities of interest based on this possibly further constrained probability distribution.

Finally, also causal discovery based on observational data is possible with static causal models. The core observation here is that certain cyclic structures leave unique probabilistic footprints. As Richardson (1996) shows, the causal structure depicted in Figure 4, for example, implies that $A$ and $B$ are independent, become dependent conditional on $C$ or on $D$, but remain independent conditional on both $\{C,D\}$. Furthermore, this is the only causal structure compatible with these dependencies and independencies.[16]

Let us next have a look at *dynamic* models. One of the main problems they share with static models is the choice of the probability distribution. It is, again, best identified with the system's equilibrium distribution. One then uses these probabilities to specify the probabilities for every variable $X_i$ conditional on its direct predecessors (parents) in the graph at step $i - 1$. Note that these conditional probabilities will be identical for any step, regardless of how we

choose *i*. This guarantees that regardless of how we specify the prior probability distributions over the exogenous variables at stage 1, the model will approach equilibrium in the long run (provided we keep system-external influences constant). As before, this will only work for systems displaying negative feedback loops that trump the effects of possibly involved positive feedback loops. Systems not reaching equilibrium will not produce a suitable distribution to specify these probabilities.

Another limitation of dynamic models is that the time difference between consecutive stages *i* and *i*+1 needs to be chosen very carefully. More specifically, the *measurement timescale* must match the *causal timescale*. The reason is that causes typically require some time to produce their effects. If time intervals are chosen too small, this will result in measured independencies between variables at some stage *i* and variables at the next stage *i* + 1 though these variables are causally connected. If these intervals are chosen too large, on the other hand, then fixing all of our variables' values at a stage *i* will not screen off earlier from later stages anymore. To guarantee that these things do not happen, it does not suffice that we take much care in choosing time intervals between stages. Also the system of interest must be nice enough. In particular, all the causal influences we are interested in need to spread more or less equally fast. This is a severe limitation for dynamic causal models.

Also in dynamic causal models computing the effects of interventions and hypothetical experiments is straightforward. Since the rolling-out of the model over time always results in an acyclic structure—compare the static model in Figure 3(a) to its dynamic counterpart in (b)—the procedure is identical to computing the effects of interventions in ordinary acyclic causal models (see Pearl 2000: sec. 1.3 and above). Such interventions are typically also used when it comes to causal discovery. One applies a shock particularly to one variable at one specific stage and then measures how this induces changes to other variables in the network over time. However, Danks and Plis (2013) show that the causal structure can often at least partially be recovered even when undersampling, i.e., even if the measurement timescale is slower than the causal timescale. It can be shown that whenever we infer $X_i$ to be a direct cause of $Y_j$ (where $i < j$) on the basis of the measurement timescale, then $X_i$ is a direct or indirect cause of $Y_j$ given the true causal timescale.[17] This result is based on technical results about how variables can be marginalized out from richer causal structures.[18]



## 4 Case: the PI3K/mTOR signalling network

Up till now, we have kept our examples as simple as possible. In scientific practice, however, it is very common to find highly complex

networks. PI3K/mTOR may serve as an example (see the recent review by Ghomlaghi, Hart, Hoang, Shin, & Nguyen 2021). To repeat, PI3K/mTOR is a signalling network that plays a crucial role in regulating the cell cycle: cell growth, cell survival and metabolism (for instance of glucose), cell death etc. It is a remarkably complex structure involving a multitude of positive and negative feedback loops, feed-forward loops, double-negative feedback and mutual inhibition (Ghomlaghi et al. 2021: sec. 2 and 3). Together, these result in network behaviour that is both flexible (having multiple stable output states) and robust in response to extra- and intra-cellular perturbations (Ghomlaghi et al. 2021: sec. 1 and 3). Many aspects of signal transduction and regulation within the PI3K/mTOR network are well-known; for others, little or no detailed mechanistic knowledge is available yet (Ghomlaghi et al. 2021: sec. 2 and 3). The workings of some of the network's loops are dependent on specific network conditions (e.g., concentrations of certain network components, mutual inhibition mediated by protein-protein competition; Ghomlaghi et al. 2021: sec. 3.3). Several similar feedback mechanisms seem to co-exist redundantly, perhaps with an eye to more robust control and/or fine-tuning (Ghomlaghi et al. 2021: sec. 3.1). Moreover, these signalling pathways do in turn interact with other networks (e.g., the cell-cycle signalling network) and with nutrient sources; they engage in "crosstalk" so as to

integrate information from a variety of sources (Ghomlaghi et al. 2021: sec. 4). Non-linear and dynamic behaviour typically results from such complex interactions, including oscillatory behaviour, switchlike and biphasic responses (passim).

Ghomlaghi et al. (2021) argue that increasing our understanding of the PI3K/mTOR signalling network is of great importance since it plays a crucial role in cellular homeostasis (read: dynamic equilibrium) and it is one of the most frequently deregulated pathways in cancer contexts (read: cancer may result when the signalling pathway is pushed away from equilibrium). Their review is full of references to current and future research into cancer therapy. However, they add, the pathway's behaviours are too complex for intuitive prediction and reasoning. In order to increase predictive accuracy, a combination of computational modelling, accurate and comprehensive biological knowledge and an eye for the role of the cellular context is needed (Ghomlaghi et al. 2021: sec. 5). They conclude in the following optimistic vein:

> As research continues, the combination of this knowledge [of the pathway itself, its interactions and the role of cellular context] with computational modelling will one day enable us

> to make incredibly specific and accurate predictions about therapeutic perturbations, down to the individual patient level. (Ghomlaghi et al. 2021: 13)

We submit that the cyclic causal modelling approach briefly outlined in section 3, with its own graphic, probabilistic and axiomatic underpinnings, may offer a useful, complementary tool for this computational endeavour. By "complementary" we mean in addition to other computational tools. At a very abstract level, all the feedback cycles discussed by Ghomlaghi et al. (2021) display negative, positive, or contingent feedback. It is at least worth trying to see how far one gets using the approach we have discussed. Unfortunately, neither Ghomlaghi et al., nor—as far as we could see—any of their sources, refer to the causal modelling literature.

Of course, when applying causal modelling tools to complex biological cases, one should not turn a blind eye to the possible pitfalls of doing so.[19] These tools start from a number of assumptions (such as the global Markov condition and the faithfulness condition). We have also stressed the role of granularity and appropriate time scales. Finally, Ghomlaghi et al. (2021) rightly stress the indispensability of substantive biological knowledge.

In the following, we focus on one of the main challenges the PI3K/mTOR signalling network seems to pose for causal modelling approaches due to its high degree of complexity.[20] To build an adequate causal model of this system, either by hand or on the basis of causal search algorithms, one typically uses the system's equilibrium distribution. The problem is that, as Ghomlaghi et al. (2021) emphasize, the PI3K/mTOR signalling pathway guarantees the flexibility required to do its job by displaying bi-stability and oscillation due to sophisticated submechanisms and crosstalk with neighboring signalling pathways. Bi-stability means that the system can reach several different equilibrium states, depending on whether required threshold-concentrations are reached. The system displaying oscillatory behaviour means that its inner cyclic workings together with its wider environment push it towards one of its equilibrium states at one time, but to another at the next, rendering the PI3K/mTOR pathway a contingent cycle. As a consequence, we will not be able to choose one single adequate distribution in order to build a causal model of the whole system.

Does this mean that causal modelling approaches ultimately fail when it comes to the application to more complex real world cases such as the PI3K/mTOR signalling network? Not necessarily. We believe that cases such as this show a clear limitation. However, we can still try a patchwork

approach. If we can keep relevant system-internal or -external factors steady enough so that the target system is pushed towards one of its equilibrium states, we will be able to use the resulting equilibrium distribution in order to complete our model. This will not give us a causal model for the whole system with all its nuanced behaviours, but we can repeat this strategy for different circumstances pushing the system towards different equilibrium states, which will give us different smaller-scale causal models, each of which can serve as a device for causal inference under specific circumstances.

This patchwork causal modelling approach of course requires substantive biological knowledge. One needs sufficient knowledge, for instance, about systeminternal and -external factors in order to be able to keep them steady enough. But this does not set causal modelling apart from other computational approaches (witness the repeated message in this sense by Ghomlaghi et al. 2021). In some sense, this is also not too different from what we do when building causal models for ordinary non-cyclic systems. We keep possible confounders fixed (or as fixed as possible) and assume that relevant background variables do not vary too drastically, since such variations can have a huge impact on the probability distribution associated with the causal structure we are actually interested in. So also ordinary models in practice never give us a complete model of the whole system of interest, but rely on

explicit and often also on unknown background conditions. They are adequate tools for causal inference only if these conditions are met.

There is, however, also a difference in complications between modelling cyclic and acyclic causal systems: While a simple acyclic model can in principle be expanded to a more extensive model by adding some of the background factors to one's model and allowing them to vary, this cannot be done with cyclic models of systems displaying dominant contingent feedback such as the PI3K/mTOR pathway. Once we add the variables we needed to keep constant in order to avoid oscillation, the system will not produce a suitable equilibrium distribution anymore.

# 5 Conclusion

Since a couple of decades there is a very large literature on causal modelling for acyclic causal structures. There is also a growing literature on cyclic causal modelling. The latter can be very useful for practicing scientists, especially in tandem with other computational tools. It goes without saying that cyclic causal modelling techniques should be applied in a cautious way, keeping in mind the various assumptions underlying them. That being said, it's better to only throw away the bathwater and not the baby.

## Key messages

- Many causal systems studied by sciences such as biology, pharmacology, and economics feature causal cycles.
- Three types of causal cycles can be distinguished: negative feedback, positive feedback and contingent feedback, all of which are closely connected to the notion of a system's *equilibrium state*.
- A causal model consists of a causal graph and a way to specify the causal strengths of the specific causal links (for instance a probability distribution or a set of equations).
- Most accounts of causal modelling currently on the market are explicitly designed to study acyclic structures, but there exist different ways to generalize the causal modelling framework to cyclic causal structures.
- These causal modelling tools can be applied to complex biological cases, yet one should not turn a blind eye to the possible pitfalls of doing so.

## Key readings

- Clarke, B., Leuridan, B., & Williamson, J. (2014) "Modelling mechanisms with causal cycles," *Synthese* 191(8): 1651–81.
- Gebharter, A. (2017) *Causal nets, interventionism, and mechanisms*. Cham: Springer.
- Pearl, J. (2000) *Causality* (1st ed.). Cambridge: Cambridge University Press.
- Pearl, J., & Mackenzie, D. (2018) *The Book of Why*. New York: basic Books.
- Spirtes, P., Glymour, C., & Scheines, R. (2000) *Causation, prediction, and search* (2nd ed.). Cambridge, MA: MIT Press.

## References

Baumgartner, M. (2009) "Interdefining causation and intervention," *Dialectica* 63(2): 175–94.

Baumgartner, M. (2012) "The logical form of interventionism," *Philosophia* 40(4): 751–61.

Clarke, B., Leuridan, B., & Williamson, J. (2014) "Modelling mechanisms with causal cycles," *Synthese* 191(8): 1651–81.

Danks, D., & Plis, S. (2013) "Learning causal structure from undersampled time series," *JMLR: Workshop and conference proceedings* 1: 1-10. doi: 10.1184/R1/ 6492101.v1

Gebharter, A. (2017) *Causal nets, interventionism, and mechanisms*. Cham: Springer.

Gebharter, A. & Hüttemann, A. (2023) "Causal bayes nets and token-causation: Closing the gap between token-level and type-level," *Erkenntnis*.

Gebharter, A. & Schurz, G. (2014) "How Occam's razor provides a neat definition of direct causation," in J. M. Mooij, D. Janzing, J. Peters, T. Claassen, & A. Hyttinen (Eds.), *Proceedings of the UAI workshop Causal Inference: Learning and Prediction.* Aachen. Retrieved from http://ceur-ws.org/Vol-1274/uai2014ci_paper1.pdf

Gebharter, A., & Schurz, G. (2016) "A modeling approach for mechanisms featuring causal cycles," *Philosophy of Science* 83(5): 934–45.

Ghomlaghi, M., Hart, A., Hoang, N., Shin, S., & Nguyen, L. K. (2021). "Feedback, crosstalk and competition: Ingredients for emergent non-linear behaviour in the PI3K/mTOR signalling network," *International Journal of Molecular Sciences* 22(6944): 1-20.

Glymour, C. (2004) “Critical notice,” *British Journal for the Philosophy of Science* 55(4): 779–90.

Kaiser, M. I. (2016) “On the limits of causal modeling: Spatially-structurally complex biological phenomena,” *Philosophy of Science* 83(5): 921–33.

Koch, D., Eisinger, R. S., & Gebharter, A. (2017) “A causal Bayesian network model of disease progression mechanisms in chronic myeloid leukemia,” *Journal of Theoretical Biology* 433: 94–105.

Leuridan, B., & Lodewyckx, T. (2021) “Diachronic causal constitutive relations,” *Synthese* 198: 9035–65.

Mella, P. (2008) “Systems thinking: The art of understanding the dynamics of systems,” *International journal of learning* 15.

Papineau, D. (2022) “The statistical nature of causation,” *The Monist* 105(2): 247–75.

Pearl, J. (1988) *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. San Mateo, CA: Morgan Kaufmann.

Pearl, J. (2000) *Causality* (1st ed.). Cambridge: Cambridge University Press.

Pearl, J., & Dechter, R. (1996) “Identifying independencies in causal graphs with feedback,” in *UAI’96: Proceedings of the Twelfth International Conference on Uncertainty in Artificial Intelligence* (pp. 420–426). San Francisco, CA: Morgan Kaufmann.

Pearl, J., & Mackenzie, D. (2018) *The Book of Why*. New York: basic Books.

Richardson, T. (1996) "A discovery algorithm for directed cyclic graphs," in *UAI'96: Proceedings of the Twelfth International Conference on Uncertainty in Artificial Intelligence* (pp. 454–461). San Francisco, CA: Morgan Kaufmann.

Richardson, T., & Spirtes, P. (2002) "Ancestral graph Markov models," *Annals of Statistics* 30(4): 962–1030.

Schurz, G., & Gebharter, A. (2016) "Causality as a theoretical concept: Explanatory warrant and empirical content of the theory of causal nets," *Synthese* 193(4): 1073–103.

Spirtes, P. (1995) "Directed cyclic graphical representations of feedback models," in P. Besnard & S. Hanks (Eds.), *Proceedings of the 11th Conference on Uncertainty in Artificial Intelligence* (pp. 491–498). San Francisco: arXiv.org.

Spirtes, P., Glymour, C., & Scheines, R. (2000) *Causation, prediction, and search* (2nd ed.). Cambridge, MA: MIT Press.

Strevens, M. (2007) "Review of Woodward, Making Things Happen," *Philosophy and Phenomenological Research* 74(1): 233–49.

Turner, J. S. (2019) "Homeostasis as a fundamental principle for a coherent theory of brains," *Philosophical Transactions of the Royal Society B: Biological Sciences* 374(1774): 20180373.

Weber, M. (2016) "On the incompatibility of dynamical biological mechanisms and causal graphs," *Philosophy of Science* 83(5): 959–71.

Woodward, J. F. (2003) *Making things happen*. Oxford: Oxford University Press.

---

[1] Names of authors are listed alphabetically. We would like to thank William Bechtel, Sebastiaan Tieleman, Phyllis Illari and Federica Russo for their helpful comments on this chapter.

[2] See also Chapter CROSSREF-chapter-Bechtel&Bollhagen, this volume, for a complementary approach to the issue of causal cycles in biology.

[3] See also Chapter CROSSREF-chapter-Kleinberg, this volume, by Samantha Kleinberg for a discussion of the use (or lack of use) of causal models in general and Chapter CROSSREF-chapter-MonetaTieleman, this volume, by Alessio Moneta and Sebastiaan Tieleman on causal modelling and complexity in macroeconomics.

[4] See also Chapter CROSSREF-chapter-Andersen, this volume, by Holly Andersen or (Gebharter & Hüttemann 2023; Papineau 2022) on the metaphysics required by causal modelling methods.

[5] Other chapters in this volume that also adopt an interventionist interpretation of causality are Chapters CROSSREF-chapter-Runhardt and CROSSREF-chapter-Suarez.

[6] For a more detailed account—in terms of intervention variables and interventions—of what counts as wiggling $X_i$ *in the right way*, see (Woodward 2003, ch. 3).

[7] Note that a system's equilibrium state is typically defined only *ceteris paribus*, meaning that certain external disturbance factors are assumed to stay equal or close enough to equal.

[8] This stable oscillation counts as an equilibrium state only if you treat it at a coarsegrained level and over a relatively long span of time; see also Clarke et al. 2014: 1657, on the notion of granularity.

[9] See https://medlineplus.gov/ency/article/001982.htm.

[10] Mean body temperature varies among normal individuals; this we will ignore here.

[11] For details and additional motivation of these notions see, for example, (Woodward 2003).

[12] For details, see (Spirtes et al. 2000: 31).

[13] This is not to say that such a deeper understanding does not have added value.

[14] $A$ screens $B$ and $C$ off each other iff $Pr(b|a,c) = Pr(b|a)$ or $Pr(a,c) = 0$, for all $A$-, $B$-, and $C$-values $a$, $b$, and $c$, respectively.

[15] Note that only systems that have an equilibrium state can develop an equilibrium distribution. Note further that systems displaying positive or contingent feedback outweighing the effects of possible negative feedback cycles will not reach their equilibrium state. Thus, identifying our model's probability distribution with an equilibrium distribution would not provide us with an adequate representation of these systems' causal dynamics.

[16] This holds only under the assumption of faithfulness and in the absence of latent common causes.

[17] This result assumes, again, faithfulness as well as the absence of latent common causes.

[18] For the full approach, see (Richardson & Spirtes 2002); for a simplified and less technical approach, see (Gebharter & Schurz 2016).

[19] See also Chapter CROSSREF-Chapter-Andersen, this volume.

[20] For more problems that can come with a high degree of complexity over and above the one we discuss in this chapter see, for example, (Kaiser 2016; Weber 2016).