

Causal Graphs and Biological Mechanisms*

Alexander Gebharter and Marie I. Kaiser¹

Abstract Modeling mechanisms is central to the biological sciences – for purposes of explanation, prediction, extrapolation, and manipulation. A closer look at the philosophical literature reveals that mechanisms are predominantly modeled in a purely qualitative way. That is, mechanistic models are conceived of as representing how certain entities and activities are spatially and temporally organized so that they bring about the behavior of the mechanism in question. Although this adequately characterizes how mechanisms are represented in biology textbooks, contemporary biological research practice shows the need for *quantitative, probabilistic* models of mechanisms, too. In this paper we argue that the formal framework of causal graph theory is well-suited to provide us with models of biological mechanisms that incorporate quantitative and probabilistic information. On the basis of an example from contemporary biological practice, namely feedback regulation of fatty acid biosynthesis in *Brassica napus*, we show that causal graph theoretical models can account for feedback as well as for the multi-level character of mechanisms. However, we do not claim that causal graph theoretical representations of mechanisms are advantageous in all respects and should replace common qualitative models. Rather, we endorse the more *balanced view* that causal graph theoretical models of mechanisms are useful for some purposes, while being insufficient for others.

3.1 Introduction

The search for mechanisms that underlie the phenomena under study is ubiquitous in many biological fields. Physiologists seek to find the mechanism for muscle contraction, cancer scientists try to discover the mechanisms that cause cell proliferation, and ecologists aim at elucidating the various mechanisms that bring about

* This is a draft paper. The final version of this paper is published under the following bibliographical data: Gebharter, A., & Kaiser, M. I. (2014). Causal graphs and biological mechanisms. In M. I. Kaiser, O. Scholz, D. Plenge, & A. Hüttemann (Eds.), *Explanation in the special sciences: The case of biology and history* (pp. 55-85). Synthese Library 367. Dordrecht: Springer. doi: [10.1007/978-94-007-7563-3_3](https://doi.org/10.1007/978-94-007-7563-3_3). The final publication is available at <http://link.springer.com/>.

¹ The authors are listed in alphabetical order.

the maintenance of species diversity – just to mention a few examples. In the last 15 years, the philosophical literature on mechanisms has dramatically increased. Among the major proponents of the “new mechanistic philosophy” (Skipper and Millstein 2005, 327) are Carl Craver (2007), William Bechtel (2006, 2008), Stuart Glennan (2002, 2005), Lindley Darden (2006, 2008), and Peter Machamer et al. (2000). According to the mechanist’s view, scientific practice consists in the discovery, representation, and manipulation of mechanisms. Scientific explanations are (exclusively or primarily) conceived as mechanistic explanations, that is, as descriptions of how the components of a mechanism work together to produce the phenomenon to be explained.²

Our primary interest in this paper is the modeling of biological mechanisms. How are, can, and should mechanisms be represented? Are certain kinds of models of mechanisms advantageous with regard to particular scientific purposes like explanation, understanding, prediction, or manipulation? Previous philosophical literature on this topic (e.g., Glennan 2005; Craver 2007; Bechtel 2008) regards mechanistic models as being primarily *qualitative* representations. According to the mechanist’s view, adequate models of mechanisms describe all and only those factors that contribute to bringing about the mechanism’s behavior of interest (i.e., the “constitutively relevant” factors; cf. Craver 2007, 139-159). These factors include the entities (or objects) that compose the mechanism, the activities (or operations or interactions) that these entities engage in, and the spatial and temporal organization of the entities and activities (i.e., how the entities are spatially distributed, which position shifts of entities take place, which activities initiate which other activities to what time, etc.). These qualitative models of biological mechanisms are typically depicted by diagrams (cf. Perini 2005), which scientists sometimes call “cartoon models” (Ganesan et al. 2009, 1621). Diagrams make it easier to understand how the steps of a mechanism together bring about the behavior in question. Hence, the representations of mechanisms that can be found in common biology textbooks are typically qualitative models.

However, biological research practice is much more diverse than what is depicted in biology textbooks. Whereas the models of mechanisms which are designed for textbooks aim at providing explanations and promote understanding, modeling strategies that are pursued in contemporary scientific practice, by contrast, serve multiple purposes. Besides offering explanation, models of mechanisms are also used, for instance, to make (quantitative or qualitative) predictions, to guide hypotheses building in scientific discovery, and to design manipulation experiments or even computer simulations. In some research contexts what will be needed are not purely qualitative models of mechanisms, but rather models that contain quantitative, probabilistic information. These models often have the virtue of being closer to the experiments and studies that are actually carried out in biological research practice. It is due to this closeness that probabilistic and quantita-

² Of course, one need not subscribe to all the details of the mechanistic view of science in order to acknowledge the importance of mechanisms to wide areas of biology.

tive models often allow for more usable predictions, in particular when it comes to predicting the probabilities of certain phenomena of interest under specific manipulations. Another advantage of models of mechanisms that combine qualitative with quantitative, probabilistic information might be that they allow for the integration of qualitative (e.g., molecular) studies and probabilistic (e.g., ecological or evolutionary) studies in a certain biological field. This is, for example, an urgent issue in epigenetics where the laboratory experiments performed by molecular epigeneticists and the observational studies and computer simulations conducted by ecologists and evolutionary biologists need to be brought together (cf. Baedke 2012).

With this paper we respond to the need of contemporary biology for models of mechanisms that include quantitative, probabilistic information. We argue that the formal framework of causal graph theory is well suited to provide us with probabilistic, (often) quantitative representations of biological mechanisms. We illustrate this claim with an example from actual biological research, namely feedback regulation of fatty acid biosynthesis in *Brassica napus*. Modeling this example allows us to show how causal graph theory is able to account for certain features of biological mechanisms that have been regarded as problematic (e.g., their multi-level character and the feedback relations that they frequently contain). However, besides the virtues our analysis of this case study also reveals which difficulties causal graph theoretical modeling strategies face when it comes to representing mechanisms. As a result, we argue for the *balanced view* that, even though causal graph theoretical models of mechanisms have advantages with respect to particular scientific purposes, they also have shortcomings with respect to other purposes.

We start with an introduction of the basic formal concepts of causal graph theory (Section 3.2). In Section 3.3 we present what can be regarded as the major characteristics of biological mechanisms, namely their multi-level character, their two kinds of components, and the spatial and temporal organization of their components. Section 3.4 deals with the case study that is central to our analysis: the mechanism for feedback inhibition of ACCase by 18:1-ACP in *Brassica napus*. In Section 3.5 we discuss how this mechanism (as well as one of its submechanisms) can be modeled by using causal graph theory. In doing so, we also address the possible objection that causal graph theory can account neither for the feedback relations that many biological mechanisms contain, nor for the fact that mechanisms are frequently organized in nested hierarchies. On basis of this analysis we can then specify, on the one hand, the virtues and, on the other hand, the shortcomings of modeling biological mechanisms within a causal graph framework (Section 3.6).

3.2 Causal Graph Theory

Causal graph theory is intended to model causality in a quite abstract and empirically meaningful way; it therefore provides principles which connect causal structures to empirical data. While causal structures are represented by graphs, empirical data is stored by means of probability distributions over sets of statistical variables. In this section we will introduce the basic formal concepts needed to investigate the question of whether a causal graph framework is capable of representing mechanisms. We start by giving some notational conventions and remarks concerning statistical variables and probability distributions (Section 3.2.1), before providing definitions for ‘probabilistic dependence’ and ‘probabilistic independence’ (Section 3.2.2). We introduce the concept of a causal graph (Section 3.2.3) and illustrate how such a causal graph, complemented by a probability distribution, becomes a causal model (Section 3.2.4).

3.2.1 Statistical Variables and Probability Distributions

A *statistical variable* X is a function that assigns exactly one of at least two mutually exclusive properties/possible values of X ($\text{val}(X)$ designates the set of X ’s possible values) to every individual in X ’s domain D_X . Statistical variables can be used in a way quite similar to predicate constants. ‘ $X(a) = x$ ’ (where ‘ a ’ is an individual constant), for instance, can be read as the token-level statement ‘individual a (e.g., a particular *Drosophila* fly) has property x (e.g., red eye color)’ and ‘ $X(u) = x$ ’ (where ‘ u ’ is an individual variable) as the type-level statement ‘having property x ’. Formulae like ‘ $X(u) = x$ ’ can be abbreviated as ‘ $X = x$ ’ or, even shorter, as ‘ x ’ whenever reference to individuals u is not needed. For the sake of simplicity we shall only use discrete variables, that is, variables X whose set of possible values $\text{val}(X)$ is finite. Continuous quantities can be captured by discrete variables whose values correspond to the accuracy of the used measurement methods.

Given a statistical variable X or a set of statistical variables X , then Pr is a *probability distribution* over X if and only if Pr is a function assigning a value $r_i \in [0,1]$ to every $x \in \text{val}(X)$, so that the sum of all assigned r_i equals 1. Since probability distributions should be capable of storing empirical data, we interpret probabilities as objective probabilities, that is, as inductively inferred limit-tendencies of observed frequencies.

3.2.2 Probabilistic Dependence and Independence Relations

Given a probability distribution Pr over variable set V , conditional probabilistic dependence between two variables X and Y can be defined in the following way:

(1) $DEP_{Pr}(X,Y|M)$ if and only if there are x , y , and m so that $Pr(x|y,m) \neq Pr(x|m)$, provided $Pr(y,m) > 0$.³

Read ' $DEP_{Pr}(X,Y|M)$ ' as ' X and Y are probabilistically dependent conditional on M '. According to definition (1) two variables X and Y are probabilistically dependent conditional on M if the probability of at least one value of one of these two variables is probabilistically sensitive to at least one value of the other variable in at least one context $M = m$. So 'probabilistic dependence' is a quite weak notion. 'Probabilistic independence', on the other hand, is a very strong notion. If two variables X and Y are probabilistically independent conditional on M , then there is not a single X -value x and not a single Y -value y so that x is probabilistically sensitive to y in any context $M = m$. Conditional probabilistic independence ($INDEP_{Pr}$) is defined as the negation of conditional probabilistic dependence:

(2) $INDEP_{Pr}(X,Y|M)$ if and only if for all x , y , and m : $Pr(x|y,m) = Pr(x|m)$, provided $Pr(y,m) > 0$.

Unconditional probabilistic dependence/ independence ($DEP_{Pr}(X,Y)$ / $INDEP_{Pr}(X,Y)$) turns out to be a special case of conditional probabilistic dependence/independence; it can be defined as conditional probabilistic dependence/independence given the empty context $M = \emptyset$:

(3) $DEP_{Pr}(X,Y)$ if and only if $DEP_{Pr}(X,Y|\emptyset)$, and
 (4) $INDEP_{Pr}(X,Y)$ if and only if $INDEP_{Pr}(X,Y|\emptyset)$.

3.2.3 Graphs and Causal Graphs

Let us turn to the concept of a causal graph. A *graph* G is an ordered pair $\langle V,E \rangle$, where V is a set of so-called vertices (which are statistical variables in causal graphs) while E is a set of so-called *edges*. Edges may be all kinds of arrows (e.g., ' \rightarrow ', ' $\dots>$ ', and ' \leftrightarrow ') or undirected links (' $—$ ') representing diverse binary relations among objects in V . Two variables in a graph's variable set V are called *ad-*

³ The condition $Pr(y,m) > 0$ is needed because $Pr(x|y,m)$ is defined as $Pr(x,y,m)/Pr(y,m)$ and division by 0 is undefined.

adjacent if and only if they are connected by an edge. A chain of $n \geq 1$ edges connecting two variables X and Y of a graph's variable set V is called a *path* between X and Y . A path of the form $X \rightarrow \dots \rightarrow Y$ is called a *directed path* from X to Y . Whenever a path contains a subpath of the form $X \rightarrow Z \leftarrow Y$, then Z is called a *collider* on this path; the path is called a *collider path* in that case. X is called an *ancestor* of Y if and only if there is a directed path from X to Y ; Y is called a *descendant* of X in that case. The set of all ancestors of a variable X is denoted by ' $Anc(X)$ ', while the set of all descendants of X is indicated by ' $Des(X)$ '. All X for which $X \rightarrow Y$ holds are called *parents* of Y ; the set of all parents of Y is referred to via ' $Pa(Y)$ '. All Y for which $X \rightarrow \dots \rightarrow Y$ holds are called *children* of X ; the set of all children of X is referred to via ' $Chi(X)$ '. Variables to which no arrowhead is pointing are called *exogenous* variables. Non-exogenous variables are called *endogenous* variables. A graph $G = \langle V, E \rangle$ containing a path of the form $X \rightarrow \dots \rightarrow X$ (with $X \in V$) is called a *cyclic graph*; an *acyclic graph* is a graph that is not a cyclic graph. A graph $G = \langle V, E \rangle$ is called a *directed graph* if E contains only directed edges.

A graph becomes a *causal graph* as soon as its edges are interpreted causally. We will interpret ' $X \rightarrow Y$ ' as ' X is a *direct cause* of Y in causal graph G '. X is a *cause* (i.e., a direct/indirect cause) of Y in G if and only if there is a causal chain $X \rightarrow \dots \rightarrow Y$ in G .

3.2.4 Bayesian Networks and Causal Models

A *directed acyclic graph* (DAG) $G = \langle V, E \rangle$ and a probability distribution Pr over G 's variable set V together become a so-called *Bayesian network* (BN) $\langle G, Pr \rangle$ if and only if G and Pr satisfy the *Markov condition*⁴ (MC). If G is an acyclic causal graph, then G and Pr become an *acyclic causal model* (CM) if and only if G and Pr satisfy the *causal Markov condition*⁵ (CMC):

(MC/CMC): $G = \langle V, E \rangle$ und Pr satisfy the (causal) Markov condition if and only if for all $X \in V$: $INDEP_{Pr}(X, V \setminus Des(X) | Pa(X))$.⁶

$V \setminus Des(X)$ is the set of all non-descendants of X . Note that ' $Des(X)$ ' and ' $Pa(X)$ ' in CMC refer to X 's effects and X 's direct causes, respectively, while ' $Des(X)$ ' and ' $Pa(X)$ ' are not causally interpreted at all in MC. The main idea behind CMC can

⁴ Cf. Glymour et al. 1991, 156.

⁵ Cf. Spirtes et al. 2000, 29.

⁶ In addition to MC/CMC, there are further principles of special interest when it comes to causal inference on the basis of empirical data (e.g., *causal sufficiency*, the *minimality condition*, and the *faithfulness condition*). For further details on these principles see, for example, Spirtes et al. (2000) or Williamson (2005).

be traced back to Reichenbach's *The Direction of Time* (1956).⁷ It captures the strong intuition that conditioning on all common causes as well as conditioning on intermediate causes breaks down the probabilistic influence between two formerly correlated variables X and Y . Or in other words: the direct causes of a variable X contain all the probabilistic information which can be found under the causes of event-types $X = x$; knowing the values of X 's parents screens X off from all of its indirect causes.

We illustrate how CMC works by providing some examples. CMC implies for the DAG in Figure 3.1, for instance, the following independence relations (as well as all probabilistic independence relations implied by them). These independence relations can directly be read off CMC applied to this DAG: $INDEP_{Pr}(X_1, X_4)$, $INDEP_{Pr}(X_2, \{X_3, X_4, X_6\} | X_1)$, $INDEP_{Pr}(X_3, \{X_2, X_4\} | X_1)$, $INDEP_{Pr}(X_4, \{X_1, X_2, X_3, X_5\})$, $INDEP_{Pr}(X_5, \{X_1, X_4, X_6\} | \{X_2, X_3\})$, and $INDEP_{Pr}(X_6, \{X_1, X_2, X_3\} | \{X_3, X_4\})$.




Fig. 3.1

It follows from MC/CMC that the equation $Pr(X_1, \dots, X_n) = \prod_i Pr(X_i | pa(X_i))$ ⁸ holds in every BN/acyclic CM $\langle V, E, Pr \rangle$ and, thus, that every BN/acyclic CM determines a fully defined probability distribution $Pr(X_1, \dots, X_n)$ over the variable set V of this BN/acyclic CM. Hence, BNs/acyclic CMs allow for probabilistic reasoning about events which can be described in terms of the variables in V . Because $Pr(X_1, \dots, X_n) = \prod_i Pr(X_i | pa(X_i))$ holds in acyclic CMs, the conditional probabilities $Pr(X_i | pa(X_i))$ —which are called X_i 's *parameters*—can represent the causal strengths of a variable X_i 's direct causes. Note that $Pr(X_1, \dots, X_n) = \prod_i Pr(X_i | pa(X_i))$ and thus MC/CMC do not hold in cyclic CMs, either. It is because of this that in cyclic CMs there are always some variables whose parameters are undefined (these are the variables lying on a cyclic directed path) and, thus, that also the causal strengths of their direct causes are undefined in such models.

⁷ See also Williamson (2010).

⁸ Note that ' $pa(X_i)$ ' stands for X_i 's parents taking certain values, while ' $Pa(X_i)$ ' stands for X_i 's parents, i.e., the variables which are X_i 's direct predecessors in the corresponding graph.

3.3 Biological Mechanisms

Before we can assess the strengths and shortcomings of causal graph theoretical models of biological mechanisms we need to know what the main features of biological mechanisms are. In the last 15 years, philosophical interest in mechanisms has significantly increased. Those who endorse the mechanistic account place the concept of a mechanism at the heart of their philosophical analysis of scientific practice. They regard models of mechanisms as being involved in almost all scientific activities, let it be explanation, discovery, prediction, generalization, or intervention. There are still controversies in the debate with regard to how the notion of a mechanism should be specified, for instance, to which ontological category the components of a mechanism belong (Machamer et al. 2000; Tabery 2004; Torres 2008), whether the regular occurrence of the mechanism's behavior is a necessary condition (Bogen 2005; Craver and Kaiser 2013), or whether the concept of a mechanism can be extended such that it also accounts for the behavior of complex systems (Bechtel and Abrahamsen 2010, 2011) or for historical processes (Glennan 2010; see also Glennan's paper in this volume). Despite these differences there are also many points of accord. In what follows we will briefly present what are regarded as the major characteristics of biological mechanisms in the debate.

To begin with, a mechanism is always a mechanism *for* a certain behavior (Glennan 2002), for instance, the mechanism *for* protein synthesis or the mechanism *for* cell division. This is crucial because only those factors (i.e., entities and activities/interactions) that contribute to producing the specific behavior of the mechanism are said to be *components* of this mechanism.⁹ An important consequence is that, although, for example, protein synthesis is the behavior of a cell not all parts of the cell are also components of the mechanism for protein synthesis. Some parts of the cell (e.g., the centrosome and the cytoskeleton) are causally irrelevant for synthesizing proteins and thus do not count as components of the mechanism for protein synthesis.¹⁰ In other words, the decomposition of a mechanism into its components depends on how the behavior of the mechanism is characterized (Kauffman 1970; Craver and Darden 2001).

A second major characteristic of mechanisms is their *multi-level character*. The notion 'multi-level character' refers to two distinct but related features of mechanisms: first, it appeals to the part-whole relation that exists between a mechanism and its components. This part-whole relation gives rise to the ontological claim

⁹ Craver calls these factors "constitutively relevant" and specifies this notion by his criterion of "mutual manipulability" (2007, 139-159).

¹⁰ However, the parts of the cell that are not components of the mechanism for protein synthesis may be components of other mechanisms. For instance, the centrosome and the cytoskeleton are components of the mechanism for cell division.

that the mechanism as a whole is located on a higher level of organization¹¹ than the entities and activities/interactions that compose the mechanism. For instance, the mechanism for muscle contraction is said to be located on a higher level than the calcium ions, the sarcoplasmic reticulum, the myosin and actin molecules etc. that interact with each other in a certain way (or that perform certain activities) in order to bring about the behavior of the mechanism as a whole (i.e., the contraction of the muscle fiber). Second, what is also meant by ‘multi-level character’ is the fact that many mechanisms (in particular, in the biological realm) occur in nested hierarchies. Many mechanisms have components that are themselves (lower-level) mechanisms; and many mechanisms themselves constitute a component in a higher-level mechanism. For instance, the calcium pump that actively transports the calcium ions from the cytosol back into the sarcoplasmic reticulum is a part of the mechanism for muscle contraction. However, the calcium pump is also a mechanism on its own, namely a mechanism for active transport of calcium ions. As such, it has its own components (e.g., A-, N-, and P-domain, transmembrane domain, calcium ions, ATP, etc.) with their own organization. Furthermore, the mechanism for muscle contraction constitutes itself a part in a higher-level mechanism, for instance, in the mechanism for crawling by peristalsis, a behavior that is exhibited, for example, by earthworms.

The third feature of mechanisms concerns their components. It is the one with respect to which there exists least conformity. The proponents of the mechanistic view concur that mechanisms consist of components, but they use different terminologies to classify the components and some of them assign the components to different ontological kinds (whereas others are just not interested in metaphysical issues). For instance, Machamer et al. (2000) endorse the dualistic thesis that mechanisms are composed of entities *and* activities, which they conceive as two distinct ontological kinds. By contrast, Glennan (1996, 2002) characterizes mechanisms in a monist fashion, that is, as being constituted exclusively by entities that interact with each other and thereby change their properties. Other mechanists do not take a stand on this ontological dispute, but nevertheless draw the distinction between the spatial components of a mechanism and “what the spatial components are doing” or “the changes in which the spatial components are involved”. Moreover, these authors adopt a different terminology to describe this difference. Bechtel (2006, 2008), for example, speaks of component parts and component operations (or functions). We think that it is not necessary (although legitimate) to become engaged in the ontological dispute about whether mechanisms consist of components that belong to one or to two distinct ontological kinds. One can avoid this dispute and yet argue that the two concepts – let it be entities and activities, entities and interactions, component parts and component operations, or whatever

¹¹ We leave it open whether the notion of a level of organization must be spelled out in a mechanistic way, as for example Craver claims (2007, 184-195). Alternatively, one could try to offer an account of levels, according to which levels are not only defined in local explanatory contexts, but rather globally. In this spirit, for instance, Wimsatt takes levels to be local maxima of regularity and predictability (1976, 1994, and 2007).

one likes – are *descriptively adequate*, that is, useful for representational purposes. When biologists represent mechanisms, they typically distinguish between the object itself (e.g., ribosome) and what the object is doing or the interactions in which the object is involved (e.g., binding, moving along the mRNA, releasing polypeptide). Thus, one should account for this difference when one models biological mechanisms. This, however, leaves open the ontological question of whether activities can be reduced to property changes of entities¹² or not. In sum, the third feature of mechanisms is that they are represented as having two kinds of components, entities and activities (or operations or interactions).

A fourth major characteristic of mechanisms is the importance of the *spatial and temporal organization* of their components for the functioning of the mechanism. Only if the components of a mechanism are organized in a specific way, the mechanism as a whole brings about the behavior in question. It is important to note that mechanisms are organized in a spatial as well as in a temporal manner. The spatial organization refers to the fact that certain entities are localized in certain regions of the mechanism, move from one region to another, and perform different activities in different regions. For instance, it is significant to the functioning of the mechanism of photosynthesis that the transport of electrons through the thylakoid membrane causes the transport of protons from the chloroplast stroma into the thylakoid lumen, and that the resulting chemiosmotic potential is used for ATP synthesis by transporting the protons back into the stroma again. The temporal organization means that a mechanism is temporally divided into certain stages which have characteristic rates and durations as well as a particular order. Earlier stages give rise to latter stages so that there exists a “productive continuity” (Machamer et al. 2000, 3) between the stages of a mechanism. In other words, the activities or interactions are “orchestrated” (Bechtel 2006, 33) such that they produce the phenomenon of interest. Consider the mechanism of photosynthesis again. This mechanism is also characterized by a specific sequence of activities. The first step is the absorption of a photon (by the photosystem II). This causes the excitation of an electron, which is followed by the transport of this electron down the electron transport chain. This transport brings about the transport of protons, and so on.

At this point one could discuss further features of mechanisms, like the fact that most mechanisms produce a certain behavior in a regular way (given certain conditions) or that the components of mechanisms might be connected by a special kind of causal relations, namely “productive causal relations” (Bogen 2008). However, these characteristics of mechanisms are far more controversial than the ones we have mentioned so far. This is why we do not take them for granted here. In what follows we examine the question of whether causal graph theoretical models of biological mechanisms are able to capture the major characteristics of mechanisms that we have presented in this section, namely the multi-level charac-

¹² Or, in the case of an activity that involves two entities, two events in which the change of one property of one object causes the property change of another entity.

ter of mechanisms, their two kinds of components, and the spatial and temporal organization of their components. We do this by means of an extended analysis of an example from recent biological research. As announced before, the result of our analysis will be that causal graph theory succeeds with regard to some respects while failing with regard to others (Section 3.5). But before, we give a short introduction to the case study that we are concerned with (Section 3.4).

3.4 Feedback Inhibition of ACCase by 18:1-ACP in *Brassica napus*

Feedback inhibition is a common mode of metabolic control. Generally speaking, in feedback inhibition a product P produced late in a reaction pathway inhibits an enzyme E that acts earlier in the pathway and that transforms the substrate S into an intermediate product IP₁. Figure 3.2 illustrates this general connection.




Fig. 3.2 The General Mechanism for Feedback Inhibition

The figure shows that the substrate S is transformed in several steps into the product P (via the intermediate products IP₁, ..., IP_n). As P accumulates, it slows down and finally switches off its own synthesis by inhibiting the regulatory enzyme E that often catalyzes the first committed step of the pathway. That way, feedback inhibition prevents the cell from wasting resources by synthesizing more P than necessary. Because enzyme activity can be rapidly changed by allosteric modulators, feedback inhibition of regulatory enzymes provides almost instantaneous control of the flux through the pathway.

Many instances of this general mechanism of feedback inhibition can be found in nature. In this paper, we focus on an example from contemporary botanical research, namely on the feedback regulation of fatty acid biosynthesis in *Brassica napus*. This particular feedback inhibition mechanism has only recently been identified by Andre et al. (2012).¹³ Fatty acid biosynthesis is a crucial process for both plants and animals, providing the cell with components for membrane biogenesis and repair and with energy reserves in specialized cells (such as epidermal cells or

¹³ However, empirical work on similar regulation mechanisms, for instance, in tobacco suspension cells (Shintani and Ohlrogge 1995) and in *Escherichia coli* (Heath and Rock 1995; Davis and Cronan 2001), has been carried out before.

the cells of oil seeds). Since the need for fatty acids varies not only with the cell type, but also depends on the stage of development, time of the day, or rate of growth, fatty acid biosynthesis must be closely regulated to meet these changes. Although the biochemistry of plant acid biosynthesis has been extensively studied,¹⁴ comparatively little is known about its regulation and control (Ohlrogge and Jaworski 1997). Hence, a major challenge for contemporary plant fatty acid research is to provide an understanding of how plants regulate differential fatty acid synthesis. Andre et al. (2012) take up this challenge with regard to canola (*Brassica napus*). Knowing the mechanism of how fatty acid biosynthesis in plants is regulated is important, not least because it may give rise to the design of strategies for increasing fatty acid synthesis in plants (cf. Tan et al. 2011). This is particularly significant in light of the economic potential of genetically manipulated oil crops for improved nutritional quality or as renewable sources of petrochemical substitutes.¹⁵

The main aim of the experimental studies conducted by Andre et al. (2012) was to discover the feedback system that regulates the biosynthesis of fatty acids in the plastids of *Brassica napus*. In other words, they sought after characterizing both the “enzymatic target” of the inhibition, that is, the enzyme E that is inhibited by the product P, and the “feedback signal” of the inhibition, that is, the product P that inhibits the enzyme E (see Figure 3.2). In order to do so, Andre et al. carried out several *in vitro* experiments, for instance, they tested the effect that Tween 80 (which contains predominantly oleic acid (18:1)) has on the rate of fatty acid synthesis as compared to the effect other Tweens have (2012, Table 1). The major results of their studies are twofold: first, they provide evidence for the hypothesis that plastidic acetyl-CoA carboxylase (in short, ACCase) is the target for feedback inhibition (i.e., the regulatory enzyme that is inhibited). ACCase catalyzes the transformation of acetyl-CoA into malonyl-CoA. This finding is in line with other studies that have revealed ACCase as a major regulatory point for plant fatty acid synthesis (cf. Ohlrogge and Jaworski 1997, 115-118). Second, their experiments indicate that the 18:1-acyl carrier protein (in short, 18:1-ACP) is the feedback signal, that is, the inhibitor of ACCase. On basis of these findings they proposed the mechanism for feedback inhibition of fatty acid synthesis in *Brassica napus* that is illustrated in Figure 3.3.

¹⁴ For an overview about lipid biosynthesis see, for instance, Ohlrogge and Browse (1995).

¹⁵ Canola (*Brassica napus*) is the third largest source of vegetable oil supply. It is of high nutritional value (because of its high concentrations of unsaturated C18 fatty acids and a low level of erucic acid) and a suitable source for biodiesel fuels as well as for raw materials in industry (Tan et al. 2011).




Fig. 3.3 Mechanism for Feedback Inhibition of Fatty Acid Synthesis in *Brassica napus* (Reproduced from Andre et al. 2012)

The mechanism for feedback inhibition that takes place in the plastid (depicted in the upper, inner box) can be characterized as an instance of the general mechanism presented in Figure 3.2. The enzyme ACCase (E) converts the substrate acetyl-CoA (S) into the intermediate product malonyl-CoA (IP₁), which is then transformed into the product 18:1-ACP (P). If the concentration of 18:1-ACP increases, more and more 18:1-ACP molecules bind to ACCase molecules and inhibit them. This, in turn, slows down and finally switches off the synthesis of further 18:1-ACP.

3.5 Modeling the Mechanism for Feedback Inhibition

The mechanism presented in the previous section can be characterized as bringing about a certain phenomenon, namely *the regulation of the synthesis of 18:1-ACP* (which is a fatty acid). One way to characterize this phenomenon in more detail is to specify it quantitatively: the concentration of 18:1-ACP is regulated such that it very likely does not exceed a certain upper bound b (i.e., that the probability for a concentration of 18:1-ACP lower than b is greater than a certain defined probability threshold r). Figure 3.4 shows an illustration.




Fig. 3.4 The Explanandum Phenomenon of 18:1-ACP Regulation [The dots stand for the 18:1-ACP concentrations ($C_{18:1-ACP}$) measured over time (t).¹⁶ More than r (95% in this example) of 18:1-ACP concentrations measured so far are lower than b .]

3.5.1 A Causal Graph Theoretical Model of the Mechanism for Feedback Inhibition

How can the mechanism that brings about the regulation of fatty acid synthesis (more precisely, the regulation of the synthesis of 18:1-ACP) be represented within a causal graph framework? At first, we need to introduce a variable P , standing for the concentration of the product 18:1-ACP.¹⁷ P shall be a discrete variable fine-grained enough to correspond to the given measurement accuracy. The phenomenon may then be described as $Pr(p < b) > r$.

Furthermore, the concentration of the substrate acetyl-CoA (from now on represented by variable S) is causally relevant for the 18:1-ACP concentration P : the higher the concentration of acetyl-CoA is, the *higher* will be the probability for higher 18:1-ACP concentrations. Another factor that is causally relevant for the 18:1-ACP concentration is the concentration of the regulatory enzyme ACCase. Here we have to distinguish between active enzymes and enzymes which bind the product 18:1-ACP (at the effector interaction site). We represent the former by the variable E_{active} and the latter by the variable $E_{P-bound}$. While the concentration of active enzymes is causally relevant to the concentration of the product 18:1-ACP (the higher E_{active} 's value, the *higher* the 18:1-ACP concentration), the 18:1-ACP concentration is causally relevant to the concentration of P-bound enzymes (the higher P 's value, the *higher* $E_{P-bound}$'s value) which is, again, causally relevant to the concentration of active enzymes (the higher $E_{P-bound}$'s value, the *lower* E_{active} 's

¹⁶ To be precise, the empirical data that biologists actually gather are not concentrations. Rather, they measure, for instance, optical densities (in spectrophotometric studies), and then draw inferences from the density values about the concentrations.

¹⁷ Note that variables are always represented by italic letters. The italic ' P ', for example, stands for a variable describing the concentration of the product 18:1-ACP, while the non-italic ' P ' stands for the concentration of the product 18:1-ACP itself.

value) and so on. The negative causal influence of $E_{P-bound}$ on E_{active} represents the fact that the binding of 18:1-ACP molecules to active ACCases causes the inhibition of the ACCases (i.e., the ACCases becoming inactive), and the negative causal influence of E_{active} on S stands for the fact that many active enzymes decrease the amount of the ACCases. According to these considerations, we may illustrate the mechanism by the causal graph depicted in Figure 3.5.




Fig. 3.5 Static Cyclic CM of the Mechanism for Feedback Inhibition [S and E_{active} are direct causes of P . P is a direct cause of $E_{P-bound}$ which is a direct cause of E_{active} which is, again, a direct cause of S and P , and so on. Direct causal influences are represented by arrows. A plus ('+') above an arrow stands for a positive causal influence (i.e., high cause values lead to *high* effect values) and a minus ('-') stands for a negative causal influence (i.e., high cause values lead to *low* effect values).]¹⁸

To get a causal model, we have to supplement the causal graph depicted in Figure 3.5 with a measured probability distribution Pr over variable set $V = \{S, P, E_{active}, E_{P-bound}\}$. This probability distribution will imply that the probability of $p < b$ will be greater than r (this is the phenomenon the mechanism brings about). The probabilities Pr will correspond to the positive/negative causal influences as described above. So the probability for high P -values, for example, will be high given high S - and E_{active} -values, and low given low S - or E_{active} -values. This corresponds to the facts that many 18:1-ACPs are produced when there are many substrates and active enzymes, and that few 18:1-ACPs are produced when there are few substrates or active enzymes available, respectively.

The probabilities Pr are interpreted as inductively inferred limit-tendencies of the observed frequencies of the diverse concentrations, as they are found under *normal conditions*. These *normal conditions* can be captured by adding a context $C = c$. This context $C = c$ is simply an instantiation of a variable or a set of variables which stand for the typical experimental setup and are not (or only slightly)

¹⁸ One might object that this causal graph theoretical model is inadequate because it contains two variables that are analytically dependent, namely $E_{P-bound}$ and E_{active} . We do not think that this is the case. $E_{P-bound}$ and E_{active} are *analytically independent* variables because there is a temporal distance between the binding of P to E and the inactivation of E (i.e., the conformational change of the substrate binding site). In other words, the binding of P to E and the inactivation of E are not the same processes occurring at the same time, but rather the former causes the latter. This is also why there exists a submechanism that specifies this causal relation.

changed during measuring or manipulating S , P , E_{active} , or $E_{P-bound}$. With regard to our case study the context $C = c$ will include a certain temperature (or range of tolerable temperatures), a particular level (or tolerable range) of salinity, and a certain pH value (or range of tolerable pH values). The conditional probabilities along the causal arrows should correspond to the causal strengths of the variables' direct causes under normal conditions (i.e., in context $C = c$).

Here we can observe the *first* problem of our causal model: while the parameters of a causal model are uniquely defined in an acyclic CM, this is not the case in cyclic CMs. This is a problem when it comes to explaining or predicting certain phenomena. We typically explain or predict a variable X 's taking value x by means of this variable's direct or indirect causes and its parameters or the parameters of the variables lying between X and its indirect causes. So we explain or predict $X = x$ by reference to X 's causes and only to X 's causes and not to X 's effects. But in our cyclic CM some variable's causes are also their effects. P , for example, is a cause *and* an effect of $E_{P-bound}$. So conditioning on P does not correspond to the probabilistic influence of $E_{P-bound}$'s direct causes alone, but rather to a mixture of the probabilistic influences one gets from $E_{P-bound}$'s direct causes and some of its effects. In other words: conditioning on P does not give us the probabilistic influence of P on $E_{P-bound}$ transported only over path $P \rightarrow E_{P-bound}$, but the mixed probabilistic influence of P transported over $P \rightarrow E_{P-bound}$ and $E_{P-bound} \rightarrow E_{active} \rightarrow P$. A *second* problem of our causal model is that it does not capture the dynamic aspect of mechanisms – it does not show how the parts of the mechanism described influence each other over a period of time. A *third* deficit of our causal model is that it does not represent any hierarchical organization, that is, it does not account for the fact that mechanisms are often embedded in higher-level mechanisms and have parts that are (sub)mechanisms themselves (see Section 3.3). The above model just describes the causal relations that are responsible for bringing about the behavior of the mechanism, that is, it refers only to causes at one and the same ontological level and therefore (even if the problem one would not exist) does not, strictly speaking, allow for *interlevel* explanation/prediction. In order to cope with these three problems, in the next two subsections we expand our causal model that represents the mechanism for feedback inhibition of fatty acid synthesis in *Brassica napus*.

3.5.2 Dynamic Causal Models

The first two problems discussed in the last section can be solved by unrolling the causal model over a period of time and thereby constructing a dynamic CM.¹⁹ In doing so, we quite plausibly presuppose that causal influences need some time to

¹⁹ Similar considerations can already be found in the first (but not in the second) edition of Spirtes et al. (2000).

spread and do not occur instantaneously. We get a dynamic CM if we add time indices to the variables of our system $V = \{S, P, E_{active}, E_{P-bound}\}$, representing the mechanism's diverse stages. By presupposing that causal influences need some time to take place, we can generate the dynamic CM whose causal graph is depicted in Figure 3.6 (for five stages) on the basis of our static CM in Section 3.5.1.




Fig. 3.6 The Causal Graph of a Five Stage dynamic CM representing the Mechanism for Feedback Inhibition in *Brassica napus*

The dashed arrows transport probabilistic influences (the substrate concentration S_i , for instance, is always probabilistically relevant to the substrate concentration at the next stage) in exactly the same way as their non-dashed counterparts. The only difference is that we interpret continuous arrows as direct causal connections while we want to leave it open whether the dashed arrows represent such causal connections. Dashed arrows could, for example, also be interpreted as analytic dependencies²⁰: the later S-concentration may, for instance, analytically depend on the earlier S-concentration. The variables of the five stages together with the continuous and the dashed arrows constitute the dynamic CM's causal graph.

The corresponding static CM's topological structure can be read off from the dynamic CM. One just has to abstract from the diverse stages of the dynamic CM and look at the continuous arrows: there has to be an arrow from S to P , from P to $E_{P-bound}$, from $E_{P-bound}$ to E_{active} , and from E_{active} to S and to P in the corresponding static CM, and these all have to be causal arrows in this static CM.

Note that the time intervals between two stages of a dynamic CM should be suitably chosen. On the one hand, if they are too small, then the causal influence may not have enough time to spread from the cause to the effect variable and correlations between causes and effects will get lost. On the other hand, these intervals should not be too large, either. This may lead to violations of very basic causal intuitions. To give an example, suppose the causal model in Figure 3.6 shows the correct causal structure of the mechanism for feedback inhibition of fatty acid synthesis in *Brassica napus*. Then S is an indirect but not a direct cause of $E_{P-bound}$. S 's causal influence on $E_{P-bound}$ is mediated via P . But if the interval between two

²⁰ 'Analytic dependence' is a notion that captures a wide range of non-causal dependences, e.g., conceptual dependence, definitional dependence, and dependence which is due to a part-whole relation.

stages were too large, say, for example, it were chosen such that stage 3 in the dynamic CM in Figure 3.6 would be the next stage after stage 1, then S and $E_{P-bound}$ would be correlated and this correlation would not break down under conditionalization on the intermediate cause P . Thus, conditioning on an effect's direct causes would not screen it off from its indirect causes.

Dynamic CMs have some advantages over static CMs. First of all, they are acyclic CMs and, thus, we can use the same methods as in BNs to compute the probabilities we are interested in. Furthermore, CMC holds and the causal model's parameters are defined. So we know the causal strengths of a variable's causes and we can thus use dynamic CMs to explain certain phenomena which can be described by means of endogenous variables. So the *first* problem discussed in Section 3.5.1 can be solved: we can generate explanations and predictions by referring to the causes of the event of interest and to the probabilistic influence of these causes on this event. In addition, we can predict the probabilities of certain effects of interventions. We can, for example, predict the probability of certain P-concentrations at stage 5 given certain S - and E_{active} -concentrations at stage 1 when we change the concentration of S in a certain way at stage 3 via manipulation. The *second* problem can also be solved: the dynamic CM tells us how the parts of the mechanism described influence each other over a period of time and we can thus also make predictions about what will (most likely) happen at later stages of the mechanism when we manipulate certain variables at earlier stages of the mechanism. Another nice feature of dynamic CMs is, provided the time intervals between the diverse stages of the mechanism are suitably chosen, that standard methods can be used for causal discovery because CMC holds for dynamic CMs. Causal discovery is still a serious problem for cyclic CMs and there are only a few algorithms which, in general, do not lead to very detailed causal information (cf. Richardson 1996; Spirtes 1995). The third problem, however, still remains: our dynamic CM captures only causal information at one and the same ontological level, and thus does not allow for interlevel mechanistic explanation, manipulation, and prediction.

3.5.3 Hierarchically Ordered Causal Models

There are at least two possibilities to represent the hierarchic organization of mechanisms within causal graph theory, that is, to solve the *third* problem that we mentioned at the end of Section 3.5.1. Each of these approaches has its own merits and deficits. One of these possibilities is developed in detail in Casini et al. (2011). Casini et al. provide a quite powerful formalism. They propose to start to represent a mechanism's top-level by a causally interpreted BN. Such a BN's variable set V may then contain some so-called *network variables*. These are variables whose values are BNs themselves. Network variables (or, more precisely, the BNs which are their possible values) are intended to represent the possible states

(e.g., “functioning” and “malfunctioning”) of a mechanism’s submechanisms. These BNs’ variable sets may then themselves contain network variables which stand for the possible states of a submechanism’s submechanisms and so on. To connect the diverse levels of the mechanism represented by such BNs, Casini et al. suggest an additional modeling assumption: the *recursive causal Markov condition* (RCMC). Whenever this condition holds, then Casini et al.’s formalism allows for probabilistic reasoning across the diverse levels of the represented mechanism.

In this paper, we can discuss Casini et al.’s (2011) approach only very briefly. Though their formalism is definitely powerful, their crucial modeling assumption RCMC is quite controversial. First of all, it is neither obvious that RCMC holds in general, nor is it clear how one could distinguish cases in which it holds from cases in which it does not. Secondly, RCMC leads to contra-intuitive consequences. We have the strong intuition that learning information about a mechanism’s micro-structure should at least sometimes lead to better (or at least different) predictions of the phenomena this mechanism will bring about. This should be the case, for example, when the macro-variable describing the possible states of the mechanism is described in a quite coarse-grained way, while more and more knowledge about the mechanism’s micro-structure is collected. But, according to RCMC, a mechanism’s micro-variables are probabilistically screened off from its macro-variables whenever the state of the submechanism represented by a network variable is known. A third deficit of Casini et al.’s approach is that it does not provide any information about how a submechanism’s micro-structure is connected to the macro-structure of the overlying mechanism, that is, how exactly changes of some of the submechanism’s micro-variables’ values influence the mechanism’s macro-variables due to probabilistic influences transported over its causal micro-structure. Such information is crucial when it comes to the question of how macro-phenomena can be controlled by manipulating some of their underlying mechanisms’ micro-variables.

In what follows we sketch an alternative approach for representing the hierarchic structure of mechanisms which avoids these problems. According to our approach, the submechanisms that a particular mechanism contains are, at least in most cases, adequately represented not via network variables, as Casini et al. (2011) propose, but via *causal arrows*. We will illustrate this claim on basis of the case study that we have already introduced, namely the mechanism for feedback inhibition of fatty acid synthesis in *Brassica napus*. This mechanism can be modeled within a causal graph framework as described in Section 3.5.1. An example for a submechanism of this mechanism is the mechanism for allosteric inhibition. This submechanism specifies the causal arrow between the variables $E_{P-bound}$ and E_{active} (see Figure 3.7). That is, it describes how exactly the binding of the product 18:1-ACP (i.e., P) to the regulatory enzyme ACCase (i.e., E_{active}) causes the inhibition or inactivation of ACCase with the effect that ACCase cannot bind the substrate acetyl-CoA (i.e., S) and convert it into 18:1-ACP anymore. In other words,

this submechanism discloses why it is the case that the higher the concentration of 18:1-ACP, the lower the concentration of active ACCase.



Fig. 3.7 Static CM of the Phenomenon that is brought about by the Submechanism for Allosteric Inhibition

But how can such a submechanism be modeled within a causal graph framework, and how can it be related to the mechanism for feedback inhibition of which it is a part? In order to assess these questions we need to go into more scientific details. Unfortunately, the biochemical submechanism that explains how the binding of 18:1-ACP to the enzyme ACCase ($E_{P-bound}$) causes the inhibition of ACCase (E_{active}) in *Brassica napus* has not been discovered yet (Andre et al. 2012). The same is true for the biochemical inhibition mechanisms in other species, for instance, in *Escherichia coli* (Heath and Rock 1995; Davis and Cronan 2001). However, in order to get an idea of how the model of the submechanism might look like, we will consider a different but analogous example, in which extensive molecular and structural studies have been carried out to unravel the biochemical mechanism of inhibition. In their recent work Ganesan et al. (2009) investigated a different feedback system, namely the allosteric inhibition of the enzyme serine protease (more precisely, of hepatocyte growth factor activator, in short ‘HGFA’) by an antibody (Ab40). Their goal was to unravel the molecular details of this inhibition mechanism. That is, they aimed at characterizing the molecular interactions and conformational changes that are caused by the binding of Ab40 (in general terms: of product P) to the effector interaction site of the enzyme HGFA (in general terms, to enzyme E) and that bring about the inhibition or deactivation of HGFA. Their work is very useful for our analysis because, on an abstract level, Ganesan et al. (2009) were interested in discovering the same submechanism as the one we singled out above, namely the submechanism that explains how the binding of P to E causes the inhibition of E, in other words, why it is the case that the higher $E_{P-bound}$ ’s value, the lower E_{active} ’s value.

The exact route by which the amino acids that compose E transmit the allosteric effect, that is, by which intermediate steps the binding of P to the remote effector interaction site of E causes the altered catalytic activity of E, is in general very poorly known (Sot et al. 2009). However, the structural and kinetic studies that Ganesan et al. (2009) performed produce some relief. One of their main results is that the binding of Ab40 (i.e., P) to the effector interaction site of HGFA (i.e., E) is accompanied by a major structural change (called the “allosteric switch”; Ganesan et al. 2009, 1620), namely the movement of a certain part of the enzyme, the 99-loop, from the competent into the non-competent conformation. This, in turn, obstructs the binding of the substrate to the enzyme E; more precisely, it causes a steric clash between the P2-Leu and the S2 subsite of E and the loss of stabilizing

interactions between P4-Lys and the S4 subsite of E. The diagram in Figure 3.8 provides a general illustration of these changes (while leaving out most of the molecular details).




Fig. 3.8 Qualitative Model of the Mechanism for Allosteric Inhibition of HGFA by Fab40²¹ (Reproduced from Ganesan et al. 2009)

The molecular interactions could be described in far more details.²² However, the foregoing description suffices for our purposes. How can this submechanism for allosteric inhibition of HGFA by Ab40 be modeled in a causal graph framework? We propose to model the submechanism with a static CM containing the variables and causal topology depicted in Figure 3.9.



Fig. 3.9 Static CM of the Submechanism for Allosteric Inhibition of HGFA by Fab40

The first thing to note is that B , 99-loop , $S2$, and $S4$ are binary (and, thus, qualitative) variables. This is because the studies that were carried out in order to reveal the details of the submechanism are not quantitative, but qualitative studies: Ganesan et al. (2009) primarily performed structural analyses, for instance, by using crystallographic methods. B can take one of the two values “bindings between functional groups of Ab40 and the effector interaction site of HGFA are estab-

²¹ Fab40 is a special type of Ab40.

²² For instance, Ganesan et al.’s (2009) structural analyses show that the steric clash between the P2-part of S and the S2 subsite of E is due to a conformational shift of the hydroxyl side of Ser99 (which is an amino acid of E). In the competent conformation the hydrophobic S2 pocket is ideally shaped to recognize Leu (i.e., an amino acid that is part of the substrate) as a P2 residue. By contrast, in the non-competent conformation there is no steric fit between P2 and S2.

lished” and “bindings between functional groups of Ab40 and the effector interaction site of HGFA are not established”. *99-loop* can take one of the two values “being in the competent state” and “being in the non-competent state”. *S2* can take one of the two values “having an ideally shaped hydrophobic pocket to recognize Leu” and “having a deformed pocket so that P2-Leu cannot be recognized”. *S4* can take one of the two values “being able to perform stabilizing interactions to P4-Lys” and “being unable to perform stabilizing interactions to P4-Lys”. This model describes that if bindings between functional groups of Ab40 and the effector interaction site of HGFA are established, then the probability is high that *99-loop* is in its competent state, which is why the probability is high that *S2* has an ideally shaped hydrophobic pocket to recognize Leu and *S4* is able to perform stabilizing interactions to P4-Lys. On the higher level, we would say that if P (Ab40) binds to E (HGFA), this submechanism brings about the behavior that E (HGFA) is inactive (which means, on the lower level, that the two amino acids P2-Leu and P4-Lys of the substrate cannot bind to the substrate binding sites *S2* and *S4* of the enzyme (HGFA)).

We are aware of the fact that it is very unlikely that the biochemical submechanism for the inhibition of ACCase by 18:1-ACP in *Brassica napus* looks exactly like the submechanism for the inhibition of HGFA by Ab40, which we just described. There are too many molecular differences between the two enzymes and the two inhibitory products. However, for the sake of the argument suppose that also in the case of the inhibition of ACCase the binding of 18:1-ACP causes the movement of some part of the enzyme *X* from a competent state into a non-competent state. Suppose further that this allosteric switch brings about certain molecular and conformational changes in two substrate binding sites *S2* and *S4* of the enzyme ACCase, which prevent the substrate to bind to the enzyme. A static CM of this hypothetical submechanism would look like the one in Figure 3.10.




Fig. 3.10 Static CM of the Hypothetical Submechanism for Allosteric Inhibition of ACCase by 18:1-ACP²³

²³ The corresponding possible values of the variables are the following: *B* can take one of the two values “bindings between functional groups of 18:1-ACP and the effector interaction site of ACCase are established” and “bindings between functional groups of 18:1-ACP and the effector interaction site of ACCase are not established”. *X* can take one of the two values “being in the competent state” and “being in the non-competent state”. *S2* and *S4* can take one of the two values “having an ideal conformation that allows its binding to a certain part of 18:1-ACP” and “having a deformed conformation that inhibits its binding to a certain part of 18:1-ACP”.

On this basis we can now tackle the crucial question of how the model of the mechanism for feedback inhibition, which we developed in Section 3.5.1, and the model of one of its submechanisms, namely of the biochemical mechanism of allosteric inhibition, can be related within a causal graph framework. We propose to model the hierarchic order of this multi-level mechanism by means of a *hierarchic static causal model* with the topological structure depicted in Figure 3.11.




Fig. 3.11 Hierarchic Static CM of the Mechanism for Feedback Inhibition and of one of its Submechanisms, namely the Biochemical Mechanism for Allosteric Inhibition


The two-headed arrows between $E_{P-bound}$ and B as well as between $S2$ and E_{active} and $S4$ and E_{active} which connect the two levels of the two mechanisms do not stand for causal, but rather for constitutive relevance relations, for instance, in the sense of Craver (2007). Hence they transport probabilistic dependencies and the effects of manipulations in the same way as direct causal loops in static CMs. Note that the causal arrow $E_{P-bound} \rightarrow E_{active}$ in our original static CM disappeared in the hierarchic causal model. It is replaced by the underlying mechanism of this causal arrow, that is, by a causal structure whose input and output variables are connected to $E_{P-bound} \rightarrow E_{active}$, respectively, via constitutive relevance relations. Also note that it is not clear how the submechanism represented by $E_{P-bound} \rightarrow E_{active}$ could be analyzed in Casini et al.'s (2011) approach. They would need to add a network variable N between $E_{P-bound} \rightarrow E_{active}$ ($E_{P-bound} \rightarrow N \rightarrow E_{active}$). But then and because there is no intermediate (macro-level) cause N between $E_{P-bound}$ and E_{active} , it is unclear what this network variable N should represent at the mechanism's macro-level.

Our hierarchic static CM can be used for mechanistic reasoning²⁴ across diverse levels. In contrast to Casini et al.'s (2011) models, our model also tells us

²⁴ The main difference between mechanistic reasoning and causal reasoning is that mechanistic reasoning makes use not only of causal, but also of constitutive relevance relations. In other

how exactly probabilistic influence between macro-variables is transported over the underlying mechanism's causal micro-structure and how exactly (i.e., over which causal and/or constitutive relevance paths) manipulations of micro-variables influence certain macro-variables. For example, if we manipulate $S4$, this will change E_{active} and $S2$ because $S4$ and $S2$ are constitutively relevant for E_{active} . Since X is a direct cause of $S4$, changing $S4$ will, on the other hand, not have a direct influence on X 's value. But changing $S4$ will nevertheless have a quite indirect influence on X : a change of $S4$'s value will have an influence on E_{active} 's value at the macro-level, which influences its macro-level effect $E_{P-bound}$. Since B is constitutively relevant for $E_{P-bound}$, $E_{P-bound}$ -changes will lead to B -changes which will, since B is a direct cause of X at the micro-level, lead to certain X -changes.

Though such hierarchic models as the one depicted in Figure 3.11 can be used for probabilistic reasoning across a multi-level mechanism's diverse levels, they cannot generally be used for explanation and prediction. The reason is the same as in the case of static CMs, as illustrated in Section 3.5.1: a certain $E_{P-bound}$ -value, for example, can be explained or predicted only via reference to $E_{P-bound}$'s causes, e.g., P . But in our hierarchic static CM, P does influence $E_{P-bound}$ not only as a cause, but also as an effect: P influences $E_{P-bound}$ not only over $P \rightarrow E_{P-bound}$, but also over $P \leftarrow E_{active} \leftarrow \dots \leftarrow S2 \leftarrow X \leftarrow B \leftarrow \dots \leftarrow E_{P-bound}$ and $P \leftarrow E_{active} \leftarrow \dots \leftarrow S4 \leftarrow X \leftarrow B \leftarrow \dots \leftarrow E_{P-bound}$. So the probabilistic influence of P on $E_{P-bound}$ does not correspond to P 's causal influence on $E_{P-bound}$ alone. We can solve this problem by rolling out our hierarchic model over time as we have already done for our original static CM in Section 3.5.2. Figure 3.12 is an illustration of the result of this procedure.



words, mechanistic reasoning contains not only *intralevel* reasoning, but also *interlevel* reasoning.

Fig. 3.12 Hierarchic Dynamic Causal Model of the Mechanism for Feedback Inhibition and the Biochemical Mechanism for Allosteric Inhibition

Note that, while causal influences need some time to spread, value changes produced by constitutive relevance relations occur instantaneously. Because of this, the two-headed dashed arrows representing such constitutive relevance relations only connect variables at one and the same stage. This also corresponds to the fact that one cannot change one of two constitutively dependent variables without changing the other. Note also that the causal arrows from $E_{P-bound}$ to E_{active} disappeared in the hierarchic dynamic CM. This is because these arrows represented a submechanism at work which is explicated in more detail in the hierarchic dynamic CM – the hierarchic dynamic CM tells us exactly (and, in contrast to our original dynamic CM developed in Section 3.5.2, in a mechanistic way)²⁵ how $E_{P-bound}$ influences E_{active} and thus finally solves problem three, too: hierarchic dynamic CMs allow for probabilistic interlevel explanation and prediction of certain E_{active} -values. Certain E_{active} -values, for instance, can be mechanistically explained or predicted by certain $E_{P-bound}$ -values: $E_{P-bound}$ at stage 1 has some influence on its constitutive part B at stage 1. B at stage 1 causes X at the micro-level at stage 1.5 which causes $S2$ and $S4$ at the micro-level at stage 2, and, since $S2$ and $S4$ are constitutively relevant for E_{active} , they have a direct probabilistic influence on E_{active} at stage 2.

One could object that, since the two-headed dashed arrows in our hierarchic dynamic CM transport the influences of interventions in both directions, CMC does not hold in such models and, hence, they should have the same problems as static CMs when it comes to explanation and prediction. The first point of such an objection is definitely true: CMC does not hold for hierarchic dynamic CMs. However, this does not lead to the suspected consequence. The problem for explanation and prediction in static CMs was that the probabilities one gets when conditioning on some variables also provide some information which can only be achieved if one also knew these variables effects (in other words, probabilistic information is transported not only over cause paths, but also over effect paths). But the events that we want to explain do not occur *because* some of their effects occurred (i.e., because they had a probabilistic influence on them), and events we want to predict cannot be predicted via reference to some of their effects (which have not occurred yet). However, this problem does not arise for hierarchic dynamic CMs. In a hierarchic dynamic CM cycles appear only due to constitutive relevance relations within certain stages and, thus, conditioning on a variable's

²⁵ Note also that Casini et al.'s (2011) approach does not allow for mechanistic reasoning in this sense. In their approach the question of how two or more macro-variables (e.g., $E_{P-bound}$ and E_{active} in our example) influence each other can only be answered by causal connections at the macro-level. In our approach, on the other hand, we can explain such an influence by reference to the underlying mechanism – we can tell a story about how $E_{P-bound}$ influences E_{active} by demonstrating how $E_{P-bound}$'s constitutively relevant parts causally influence E_{active} 's constitutively relevant parts at the micro-level.

causes does *only* provide probabilistic information about this variable's values transported over cause or constitutive relevance paths. It *never* provides probabilistic information transported over an effect path.

3.6 Merits and Limits of Causal Graph Theoretical Models

On basis of the preceding analysis we can now approach the question of whether causal graph theory is suited for modeling biological mechanisms and what the advantages and shortcomings of representing mechanisms within a causal graph framework are. In the previous literature the concern has been raised that, even if it is possible to provide causal graph theoretical models of biological mechanisms, they are deficient because they fail to comprise some important kinds of information. In this line, for instance, Weber (2012) argues that because causal graph theoretical models only encompass sets of variables and relations of causal dependence, they fail to include information about the structure of biological entities (such as information about the DNA double helix topology and the movements undergone by a replicating DNA molecule) and about their spatiotemporal organization. However, claims like these remain on a quite general level. Our goal in this section is to use the results of our analysis of the case study in the previous section in order to assess and to specify these claims. We do so by pointing out which kinds of information about biological mechanisms cannot or can only insufficiently be represented within a causal graph framework, and what are the reasons for these failures. In addition to revealing the *limitations* of causal graph theoretical models of mechanisms, we also highlight the *virtues* they have with respect to certain scientific purposes.

To begin with, recall the major characteristics of biological mechanisms that we identified in Section 3.3. First, mechanisms possess a multi-level character, which means, on the one hand, that there exists a part-whole relation between the mechanism and its components and, on the other hand, that mechanisms frequently occur in nested hierarchies. Second, mechanisms are represented as having two different kinds of components: entities (having particular properties) and activities (or interactions, operations, etc.). Finally, a mechanism brings about a specific behavior only if its components are spatially and temporally organized in a certain way. Can all these three features of biological mechanisms adequately be represented by causal graph theoretical models?

Consider first the *multi-level character* of mechanisms. As we have shown in the previous section, the fact that many mechanisms occur in nested hierarchies (i.e., that they are embedded in higher-level mechanisms and have components that are themselves submechanisms) can be represented in at least two ways. On the one hand, one can represent a mechanism's submechanisms by so-called network variables, as, for instance, Casini et al. (2011) do. We, on the other hand, think that there are good reasons for representing such submechanisms by causal

arrows between variables X and Y . In our approach one can generate a hierarchic causal model by replacing such a causal arrow by another causal structure. This causal structure should be on a lower ontological level than X and Y , it should contain at least one constitutively relevant part of X and at least one of Y , and there should be at least one causal path going from the former to the latter at the micro-level. Such hierarchic models allow, in contrast to purely qualitative models, for probabilistic mechanistic reasoning *across different levels*. Hierarchic dynamic CMs do even allow for probabilistic mechanistic interlevel explanation and prediction. Contrary to Casini et al.'s models, they can also provide detailed information about how certain causal influences at the macro-level are realized by their underlying causal influences propagated at the micro-level. This is important when it comes to questions about how certain manipulations of macro- or micro-variables influence certain other macro- or micro-variables of interest and how a mechanism's causal micro-structure is connected to its macro-structure.

Let us now turn to the second feature of mechanisms. Do causal graph theoretical models succeed in representing mechanisms as being composed of two different kinds of components, namely *entities and activities* (or operations, interactions, etc.)? It is quite clear that causal models represent entities. Precisely speaking, the individuals in the domains D_{X_1}, \dots, D_{X_n} of the causal model's variables X_1, \dots, X_n represent the entities that are components of the mechanism. Furthermore, the variables X_1, \dots, X_n taking certain values represent different properties or different behaviors of these entities. But can causal graph theoretical models represent activities, too?

A convenient first step towards an answer to this question seems to be to scrutinize the activities that are involved in our case study. Examples of activities that are part of the mechanism for feedback inhibition of fatty acid synthesis in *Brassica napus* are: the *binding* of 18:1-ACP (P) to ACCase (E), the *transformation* of acetyl-CoA (S) into 18:1-ACP (P) (via the intermediate product malonyl-CoA), and the *inhibition* of ACCase (E) by 18:1-ACP (P) (see description of Figure 3.5). The submechanism that brings about the activity of the inhibition of ACCase by 18:1-ACP is, in turn, composed of the following micro-activities: the *establishment* of a certain kind of binding between a functional group of 18:1-ACP and the effector interaction site of ACCase, the *shifting* the conformation of a particular part of ACCase, the *deformation* of the conformation of the S2-part of the substrate binding site of ACCase, etc. (see description of Figure 3.9). What all these activities have in common is that they are temporally extended *processes* that involve some kind of *change*. Correspondingly, Machamer et al. have characterized activities as being "the producers of change" (2000, 3). It should be noted that not all activities must involve interactions between two or more distinct entities.²⁶ There might also be activities (so-called "non-interactive activities" (Tabery 2004,

²⁶ According to Glennan (2002, 344), an interaction is an occasion on which a change in a property of one component of the mechanism brings about a change in a property of another component.

9; Torres 2008, 246), like the shifting of the conformation of a particular part of ACCase) that involve only one entity (i.e., the particular part of ACCase) and a change of its properties (i.e., from the property “being in a competent state” to “being in a non-competent state”).²⁷ In any case, activities involve the change of properties. In principle, the variables of a causal graph theoretical model could just be chosen in such a way that the different values they can take represent different processes or changes of properties. However, such a choice of variables would completely be at odds with experimental practice in biology. In most cases it is difficult or even impossible to measure entire processes by just measuring once. Rather, what biologists do, for instance to collect empirical data about the inhibition of ACCase by 18:1-ACP, is that they measure the concentration of the product (which is an indicator of ACCase’s activity, and thus, also of its inhibition) to different times. Against this background it would be inadequate to choose the variable in such a way that one of its values represents the entire process/activity of inhibition of ACCase by 18:1-ACP. The option of representing activities simply by variables taking certain values can also be ruled out by the following argumentation: if activities were represented by variables taking certain values, then activities would neither involve changes, nor be productive – they would rather occur due to other productive causal relations. Since activities are productive and involve changes, they must be represented differently.

We think that there are two ways in a causal graph theoretical model by which the activities that compose a mechanism can be captured: they can either be represented by *causal arrows* between variables. For instance, the causal arrow between S and P in Figure 3.5 represents the activity “transformation of acetyl-CoA into 18:1-ACP”. This is the option that matches the neat picture that several authors seem to have in mind: in a causal model the variables represent the entities (and their possible properties) and the arrows represent the activities. However, our analysis shows that things are not that neat. There is a second, equally adequate way to represent activities in causal graph theoretical models, namely representing them by the *change of the value of a variable*. For instance, the activity “shifting the conformation of a particular part of ACCase” is represented in Figure 3.9 by the variable X , changing its value from “being in a competent state” to “being in a non-competent state”.

A related view of static CMs, which we have to give up, is the neat view that the different variables in static CMs always represent the possible properties and activities of *distinct* entities. The flexibility of the choice of variables allows that one static CM contains variables that represent different possible properties (and activities) of *the same* entity. For instance, in our static CM depicted in Figure 3.5 the variables $E_{P-bound}$ and E_{active} both refer to the concentrations of enzymes, but describe different properties of these enzymes, namely “being bound to P” and

²⁷ As mentioned in Section 3.3, we leave it open whether activities can be reduced to state transformations via property changes or whether there is something lost by this reduction (such as the *productive nature* of activities).

“being active”. In other words, in causal graph theoretical models the boundaries between different entities and between entities and activities often become *fuzzier* than in qualitative models. This fuzziness may have the disadvantage of impeding the understanding of how a mechanism brings about a certain phenomenon – when one looks at a static CM or at a dynamic CM, one does not recognize at first sight what the entities are and which activities they perform.

To conclude, we think that it is possible to represent mechanisms as being composed of entities and activities in a causal graph framework. However, what one does not get are neat static CMs in which each variable represents a distinct entity and the arrows represent activities. This might be disadvantageous for some purposes, but not for others.

Finally, how do things stand with the third main feature of mechanisms, namely with the *spatial and temporal organization* of their components? How much and which structural and spatial information one actually represents simply depends on one’s choice of variables. In our case study, for instance, the causal graph theoretical model depicted in Figure 3.11 contains structural as well as spatial information: the variable S_2 , for example, refers to a particular entity, namely the S_2 -part of the substrate binding site of ACCase, and to the two possible *structural* properties that this entity can exhibit, namely “having an ideal conformation that allows its binding to a certain part of 18:1-ACP” and “having a deformed conformation that inhibits its binding to a certain part of 18:1-ACP”.²⁸ A different example is the variable $E_{P-bound}$ which represents the concentration of those regulatory enzymes (ACCases) that are bound to, that is, *spatially* connected to, the product 18:1-ACP. Hence, it is possible to include certain crucial structural and spatial information about the components of a mechanism into a causal graph theoretical model – one just has to choose variables that refer to structural and spatial properties.

Information about the temporal organization can be captured by and read off from the causal arrows of dynamic CMs: in the example we discussed in Section 3.5.2, for instance, S at stage 1 causes P at stage 2, which causes $E_{P-bound}$ at stage 3. So *at first* S interacts with P , and *then* P interacts with $E_{P-bound}$ etc. However, even if there are no in-principle reasons for why it is impossible to include all the details of the spatial and temporal organization of a mechanism’s components into a causal graph theoretical model, this does not preclude that there may be heuristic reasons for doing so. For instance, including all the relevant spatial, structural, and dynamic information might give rise to a causal model that includes too many different variables, so that it is unmanageable and thus not useful.

In sum, causal graph theoretical models *can* account for the three main features of mechanisms. However, they do so in a quite abstract way, which is why they are far worse than purely qualitative models with respect to the purpose of providing understanding. Qualitative models tell us in a very intelligible way how the

²⁸ Of course, these two properties could be and, in fact, are specified in more detail in biological practice. We give this general and brief characterization just for heuristic reasons.

components of a mechanism interact to bring about the phenomenon of interest. They make, contrary to probabilistic causal models, clear distinctions between the macro- and the micro-level (i.e., between mechanisms and their submechanisms) and between distinct entities and activities (or operations, interactions, etc.). Purely qualitative models of mechanisms can also be used to explain certain behaviors of systems by revealing how the components of a mechanism bring about the behavior in question. These qualitative models are, however, limited. They fail when it comes to explain why certain systems frequently (but not always) bring about certain behaviors. In other words, they fail when it comes to explaining probabilistic phenomena like the phenomenon described in Section 3.5. Moreover, they do not allow for probabilistic prediction and (interlevel) manipulation. But knowing how we can bring about a particular phenomenon with high probability is a crucial investigative strategy in the biological sciences. Finally, purely qualitative models fail to integrate qualitative information with quantitative, probabilistic information. The latter is an important task in certain research areas like epigenetics where laboratory molecular experiments need to be brought together with ecological or evolutionary observational studies and computer simulations.

3.7 Conclusion

In this paper we have shown how the formal framework of causal graph theory can be used to model biological mechanisms in a probabilistic and quantitative way. Our analysis of the mechanism for feedback regulation of fatty acid biosynthesis in *Brassica napus* revealed that causal graph theoretical models can be extended such that they can also account for more complex forms of organization of the components of a mechanism (like feedback) as well as for the fact that mechanisms are frequently organized into nested hierarchies. We argued that, because causal graph theoretical models are not purely qualitative, but rather include probabilistic and quantitative information, they are useful in the context of causal discovery – in particular if one wants to make quantitative, probabilistic predictions or conduct manipulations. What is more, since causal graph theoretical models allow us to represent different levels of mechanisms in the same model (e.g., a mechanism, one of its submechanisms, and the relations between them) they enable us to carry out *interlevel* mechanistic manipulation and prediction, too.

However, our analysis of the case study did not only disclose advantages of representing biological mechanisms within a causal graph framework. Rather, it gave rise to the more *balanced view* that probabilistic, quantitative models of mechanisms – although there are clear merits with respect to some purposes – also have shortcomings with respect to other purposes. Accordingly, our analysis revealed that causal graph theoretical models have the resources to represent the three main features of biological mechanisms, namely their multi-level character, their two kinds of components, and the spatial and temporal organization of their

components. However, it also became clear that in some respects probabilistic, quantitative models of mechanisms are insufficient (e.g., because the boundaries of entities and between entities and activities become fuzzy and because the amount of structural/spatial and dynamical information that can be represented is limited) which makes them inadequate for some purposes (in particular for providing understanding). With this analysis we hope to have shed some light on the merits and limitations of modeling biological mechanisms within a causal graph framework and to have provided some interesting prospect for future philosophical work.

DRAFT

Acknowledgements We would like to thank the members of the research group “Causation and Explanation”, the participants of the colloquia at the University of Cologne and at the University of Düsseldorf, and the members of the Lake Geneva Biological Interest Group at the University of Geneva for their helpful comments on earlier drafts. This project was made possible by the funding provided by the Deutsche Forschungsgemeinschaft (DFG).

References

- Andre C, Haslam RP, Shanklin J (2012) Feedback regulation of plastidic acetyl-CoA carboxylase by 18.1-acyl carrier protein in *Brassica napus*. *PNAS* 109(25):10107-10112
- Baedke J (2012) Causal explanation beyond the gene: manipulation and causality in epigenetics. *Theoria* 74:153-174
- Bechtel W (2006) *Discovering Cell Mechanisms. The Creation of Modern Cell Biology*. Cambridge University Press, Cambridge
- Bechtel W (2008) *Mental Mechanisms. Philosophical Perspectives on Cognitive Neuroscience*. Taylor and Francis Group, New York/London
- Bechtel W, Abrahamsen A (2010) Dynamic mechanistic explanation: Computational modeling of circadian rhythms as an exemplar for cognitive science. *Studies in History and Philosophy of Science Part A* 41(3):321-333
- Bechtel W, Abrahamsen A (2011) Complex biological mechanisms: Cyclic, oscillatory, and autonomous. In: Hooker CA (ed) *Philosophy of complex systems. Handbook of the philosophy of science, Volume 10*. Elsevier, New York, pp 257-285
- Bogen J (2008) Causally Productive Activities. *Studies in History and Philosophy of Science Part A* 39(1):112-123
- Casini L, Illari MP, Russo F, Williamson J (2011) Models for Prediction, Explanation, and Control: Recursive Bayesian Networks. *Theoria* 70:5-33
- Craver CF (2007) *Explaining the Brain. Mechanisms and the Mosaic Unity of Neuroscience*. Clarendon Press, Oxford
- Craver C, Darden L (2001) *Discovering Mechanisms in Neurobiology: The Case of Spatial Memory*. In: Machamer P, Grush R, McLaughlin P (eds) *Theory and Method in Neuroscience*. University of Pittsburgh Press, Pittsburgh, pp 112-137
- Craver C, Kaiser MI (2013) Mechanisms and Laws: Clarifying the Debate. In: Chao H-K, Chen S-T, and Millstein RL (eds) *Mechanism and Causality in Biology and Economics*. Springer
- Darden L (2006) *Reasoning in Biological Discoveries*. Cambridge University Press, New York
- Darden L (2008) Thinking Again About Mechanisms. *Philosophy of Science* 75(5):958-969
- Davis MS, Cronan JE Jr (2001) Inhibition of *Escherichia coli* Acetyl Coenzyme A Carboxylase by Acyl-Acyl Carrier protein. *Journal of Bacteriology* 183(4):1499-1503
- Ganesan R, Eigenbrot C, Wu Y, Liang W-C, Shia S, Lipari MT, Kirchhofer D (2009) Unraveling the Allosteric Mechanism of Serine protease Inhibition by an Antibody. *Structure* 17:1614-1624
- Glennan SS (1996) Mechanism and the Nature of Causation. *Erkenntnis* 44:49-71
- Glennan SS (2002) Rethinking Mechanistic Explanation. *Philosophy of Science* 69:342-353
- Glennan SS (2005) Modeling Mechanisms. *Studies in the History and Philosophy of Biological and Biomedical Sciences* 36(2):443-464
- Glennan SS (2010) Ephemeral Mechanisms and Historical Explanation. *Erkenntnis* 72(2):251-266
- Glymour C, Spirtes P, Scheines R (1991) Causal Inference. *Erkenntnis* 35:151-189
- Heath RJ, Rock CO (1995) Regulation of Malonyl-CoA Metabolism by Acyl-Acyl Carrier protein and β -Ketoacyl-Acyl Carrier Protein Synthases in *Escherichia coli*. *The Journal of Biological Chemistry* 270(26):15531-15538
- Kauffman SA (1970) Articulation of Parts Explanation in Biology and the Rational Search for them. *Boston Studies in the Philosophy of Science* 8:257-272
- Machamer P, Darden L, Carver CF (2000) Thinking About Mechanisms. *Philosophy of Science* 67:1-25
- Ohlrogge JB, Browse JG (1995) Lipid Biosynthesis. *The Plant Cell* 7:957-970
- Ohlrogge JB, Browse JG (1997) Regulation of Fatty Acid Synthesis. *Ann. Rev. Plant Physiol. Plant Mol. Biol.* 48:109-136

- Perini L. (2005) Explanation in Two Dimensions: Diagrams and Biological Explanation. *Biology and Philosophy* 20:257-269
- Reichenbach H (1956) *The Direction of Time*. University of California Press, Berkeley
- Richardson T (1996) A Discovery Algorithm for Directed Cyclic Graphs. In: Horvitz E, Jensen F (eds) *Proceedings of the 12th Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann, Portland, pp 454-461
- Shintani DK, Ohlrogge J (1995) Feedback inhibition of fatty acid synthesis in tobacco suspension cells. *The Plant Journal* 7(4):577-587
- Skipper RA, Millstein RL (2005) Thinking about evolutionary mechanisms: natural selection. *Studies in the History and Philosophy of Biological and Biomedical Sciences* 36(2):327-347
- de Sol A, Tsai C-J, Ma B, Nussinov R (2009) The origin of allosteric functional modulation: multiple pre-existing pathways. *Structure* 17:1042-1050
- Spirtes P (1995) Directed cyclic graphical representations of feedback models. *Proceedings of the 11th conference on uncertainty in Artificial Intelligence*. Morgan Kaufmann, San Francisco, pp 491-498
- Spirtes P, Glymour C, Scheines R (2000) *Causation, Prediction, and Search*, 2nd edition. MIT Press, Cambridge
- Tabery JG (2004) Synthesizing Activities and Interactions in the Concept of a Mechanism. *Philosophy of Science* 71:1-15
- Tan H, Yang X, Zhang F, Zheng X, Qu C, Mu J, Fu F, Li J, Guan R, Zhang H, Wang G, Zou J (2011) Enhanced Seed Oil Production in Canola by Conditional Expression of Brassica napus LEAFY COTYLEDON1 and LEC1-LIKE in Developing Seeds. *Plant Physiology* 156:1577-1588
- Torres PJ (2008) A Modified Conception of Mechanisms. *Erkenntnis* 71(2):233-251
- Weber M (2012) Experiment in Biology. In: Zalta EN (ed) *The Stanford Encyclopedia of Philosophy* (Spring 2012 Edition), URL = <<http://plato.stanford.edu/archives/spr2012/entries/biology-experiment/>>
- Williamson J (2010) Probabilistic Theories. In: Beebe H, Hitchcock C, Menzies P (eds): *The Oxford Handbook of Causation*. Oxford University Press, Oxford, pp 185-212
- Williamson Jon (2005) *Bayesian Nets and Causality*. Oxford University Press, Oxford
- Wimsatt WC (1976) Reductionism, Levels of Organization, and the Mind-Body Problem. In Globus GG (ed) *Consciousness and the Brain*. Plenum Press, New York/London, pp 205-267
- Wimsatt WC (1994) The Ontology of Complex Systems: Levels, Perspectives, and Causal Thickets. *Canadian Journal of Philosophy (Suppl.)* 20:207-274
- Wimsatt WC (2007) *Re-Engineering Philosophy for Limited Beings: Piecewise Approximations to Reality*. Harvard University Press, Cambridge