Cascade versus Mechanism: The Diversity of Causal Structure in Science

Lauren N. Ross

According to mainstream philosophical views causal explanation in biology and neuroscience is mechanistic. As the term 'mechanism' gets regular use in these fields it is unsurprising that philosophers consider it important to scientific explanation. What is surprising is that they consider it the only causal term of importance. This paper provides an analysis of a new causal concept—it examines the cascade concept in science and the causal structure it refers to. I argue that this concept is importantly different from the notion of mechanism and that this difference matters for our understanding of causation and explanation in science.

1. Introduction

According to mainstream philosophical views causal explanation in biology and neuroscience is mechanistic. As the term 'mechanism' gets regular use in these fields it is unsurprising that philosophers consider it important to scientific explanation. What is surprising is that they consider it the only causal term of importance. Although scientists rely on a wide variety of causal concepts in their explanations, most philosophers assume that all of these concepts can be well understood with the single notion of 'mechanism'. For example, Craver ([2007], p. 3) claims that while scientists 'say that they discover systems and pathways in the flow of information, and molecular cascades, mediators, and modulators […] The term mechanism could do the same work'. Similarly, Darden and Craver ([2001], p. 113) state that scientists 'often speak of 'systems' and 'cascades' to describe what we call, also consistently with the field's language, 'mechanisms'. This 'new mechanist' position claims that 'most or all the phenomena found in nature depend on mechanisms' and that it is the 'chief business' of science to examine them (Glennan [2017], p. 1). This expansive notion of mechanism is dominant in the literature and it has held this position for at least the past two decades. Continued support for this view is evident by outright claims in the literature,[1] the common assumption that explanations are either non-causal or mechanistic,[2] and an enormous research programme that interprets all causal

[1] In addition to the earlier quotes, see (Robins and Craver [2009], p. 42; Glennan [2017], p. 1).

[2] For this assumption, see (Brigandt [2013]; Levy and Bechtel [2013]; Silberstein and Chemero [2013]; Reutlinger and Andersen [2016]; Lange [2018]; Huneman [2018]; Saatsi and Reutlinger [2018]; Bich and Bechtel [2022]). A further (and more extreme) example of this are claims that legitimate

systems in science as mechanisms, even when they are associated with other causal terms.[3]

A significant challenge for these mainstream views is that scientists cite a variety of causal terms that are not all well understood with the mechanist programme. To clarify this, consider that mechanisms are described as causal structures with distinct features. While debate surrounds which exact features they have, three features that are often viewed as characteristic of mechanisms are (1) constitutive relations, (2) fine-grained detail, and (3) mechanical interactions. These features are also present in the machine analogy, which is often associated with the mechanism concept. However, if mechanisms are causal structures with specific features—whatever these features are—it should be an open question whether other causal structures exist with different features. This paper explores this question by examining the cascade concept in science. It is argued that the cascade concept is importantly different from the notion of mechanism and that this difference matters for our understanding of causation and explanation in science.

This paper provides an analysis of the cascade concept in science and the causal structure it refers to.[4] I examine the main features of this causal structure, analogies it is associated with, and strategies used to study it. While scientific work supports distinguishing the cascade and mechanism concepts, this analysis is not merely descriptive. Instead, it provides a theoretical framework for how these concepts should be understood. This framework matters for our assessment of the causal structure of the world, how we study this structure, use it to produce particular outcomes, and communicate about it to others. Before proceeding with this analysis, two clarifications are in order. First, I do not suggest that scientists always use these causal terms in the distinct ways indicated in this analysis, but that they often do and should use them in this way. This reveals normative features of this work and an important way that philosophy can contribute to science, namely, by making suggestions for how these concepts should be understood and used. Second, my

---

explanations necessarily describe mechanisms (Kaplan and Craver [2011]; Kaplan and Bechtel [2011]).

[3] For example, the mechanistic framework is used to analyse biological pathways (Thagard [2003]; Bogen and Machamer [2010]; Fagan [2013]; Levy and Bechtel [2013]) and cascades (Craver [2007]; Brigandt [2013]).

[4] I use causal 'concept' and causal 'structure' somewhat interchangeably, as the concept is used to refer to the structure in the world.

analysis of these concepts articulates clear ways that they differ, but leaves space for some structures in science to be borderline. The presence of such cases should not prevent us from articulating useful categories that distinguish causal structures in the majority of cases, even if the distinction can (in rare cases) be blurred.

One way this project differs from mechanist analyses is that it examines causal concepts in a broader set of scientific fields, beyond just biology and neuroscience. This is possible because the cascade concept refers to a causal pattern found in many sciences, including physics, chemistry, biology, psychology, and economics. While the mechanism concept is also likely to have broad applicability, most analyses focus on its role in the biological sciences. This point prompts a final clarification. In particular, while this paper argues for a distinction across these causal concepts, it does not deny that mechanism plays an important role in scientific explanation and reasoning. What is denied are the stronger claims that mechanism is the only causal concept of importance for scientific explanation, that all other causal terms can be interpreted as mechanisms, and that all causal explanation fits the mechanist pattern. There have been many constructive and convincing criticisms of this expansive mechanistic view (Skipper and Millstein [2005]; Woodward [2013]; Dupré [2013]; Skillings [2015]; Ross [2021a]). This analysis is in agreement with these other projects. It aims to build on them and clarify a way forward in understanding causal structure, causal concepts, and causal explanation in biology and science more generally.

2. Causation and Mechanism in Science

In the current philosophical literature, it is common to use 'mechanism' to capture a type of causal structure without specifying an account of causation.[5] To say that there is a 'mechanism' for some effect suggests that, at minimum, there is (i) some set of causal factors that interact to produce this effect, as opposed to a single cause. However, (i) does not specify how this causation among factors should be understood, whether in terms of a regularity, conserved quantity, counterfactual, law-based,

[5] In earlier work, the notion of 'mechanism' was associated with Salmon's ([1984]) causal mechanical model, a type of process theory of causation. Current work no longer associates 'mechanism' with this work and different accounts of causation are used by different mechanist philosophers.

powers-based, or any other account of causation.[6] This is reflected by the fact that different philosophical accounts of mechanism rely on different accounts of causation.[7] This captures an important distinction between 'causation' and 'mechanism'—mechanism depends on an account of causation and further specifies a type of causal structure.

While having (i) many interacting causes is a minimal feature of mechanisms, we often view 'mechanism' as communicating more than this. Philosophical accounts of mechanism are in the business of capturing what more this is. Although there is not complete consensus on how to understand mechanism, a number of promising characteristics have emerged. These characteristics draw on mechanical and machine-like views of causal systems, which are often said to have originated with the work of Descartes (Craver and Tabery [2015]). A first characteristic feature of mechanisms is that they are constitutive, in the sense of containing causal 'parts' that make up the 'whole' mechanism. In this manner, scientific mechanisms are often analogized to machines, such as a car engine or watch mechanism, in which a discrete system is comprised of component parts (Levy [2014]). This constitutive feature promotes a reductive and hierarchical picture, in which a mechanism's parts are at a 'lower-level' relative to the 'higher-level' behaviour produced. Second, mechanisms are expected to capture fine-grained causal detail, as opposed to highly abstract relations. This is related to the sense in which mechanisms should provide some information about 'how' a cause produces its effect, as opposed to omitting, abstracting from, or 'black-boxing' these details (Craver [2007]; Glennan [2017]). We see this in the fact that mono-causal relationships and sparse causal networks are not considered mechanisms, although they may be instantiated by mechanistic details.[8] Third, causal relations in mechanisms are often described with mechanical language that emphasizes force, action, and motion. This involves stating that a cause 'bends', 'pushes', or 'pulls' another factor, as opposed to just saying that it 'causes' some effect. This third feature supports the second because specifying 'how' X causes Y and provides more information than just saying 'that' X causes Y.

---

[6] One exception to this is found in the work of Glennan ([1996]), who uses mechanism to provide an account of causation.

[7] As Craver and Tabery ([2015]) state, 'Mechanists have disagreed with one another about how to understand the cause in causal mechanism'.

[8] While this feature is supported in the scientific and philosophical literature, it remains debated in the literature (as will be discussed more later).

In addition to the fact that mechanisms are associated with these features and the machine analogy, they are often studied with particular methods. As many philosophers have indicated, mechanisms are studied with decomposition and localization (Wimsatt [1974]; Bechtel and Richardson [2010]). These methods involve specifying some an explanatory target and then 'drilling down' to identify the lower-level component parts that produce this target. In addition to supporting the reductive and hierarchical nature of the mechanism concept, these methods reveal two addition features. First, mechanisms are relative to and defined by single effects, as seen in the first step of applying these methods. Second, mechanisms are viewed as relatively discrete systems with boundaries that are determined by methodological and pragmatic considerations (Craver [2009]; Bechtel [2015]).

This account of mechanism successfully captures particular causal structures in biology and neuroscience. However, it should not be viewed as a comprehensive account of all causal structure in these domains. The mechanism concept is better understood as one of a variety of types of causal structure. One way to explore this claim is to examine the causal terms that scientists use in their descriptions and explanations of natural phenomena. These include causal terms such as 'cascade' and 'pathway', which philosophers have traditionally interpreted as mechanisms. This paper provides an analysis of the 'cascade' concept and structure and argues that it is importantly different from the notion of mechanism. While cascades meet criterion (i), they lack standard features of mechanisms and have their own distinct characteristics, analogies, and investigative strategies. After introducing and examining the cascade concept in the next two sections, I explore how this concept differs from mechanism, and why this matters for philosophical accounts of causal explanation.

3. Cascade Concept: Introduction

The causal world is complex. Scientists encounter systems with many causes, different causes, and causes that interact in varied and dynamic ways. Sometimes, these systems fit particular patterns that repeatedly present in many different contexts. In these cases, it can be useful to describe these patterns, while giving them a name or 'linguistic label' that captures their features and implications (Genter and Smith [2012], p. 130). This has happened with the cascade concept—this concept refers to a

unique causal pattern that occurs in a wide variety of domains. Biologists discuss the blood coagulation cascade, the complement cascade, and cell signalling cascades, to name a few. Psychologists study developmental cascades, ecologists examine trophic cascades, and economists cite cascading failures. In physics and chemistry, we find collision cascades, oxidative cascades, cascade showers, and cascade liquefaction. Similar to mechanisms, cascades meet criterion (i) as they contain multiple causes that work together to produce an outcome, as opposed to a single cause. However, when compared to mechanisms, cascades have their own distinct set of characteristics, analogies, and causal investigative strategies.

What exactly is a cascade? I am going to suggest that cascades are causal systems that involve (i) an initial trigger, (ii) sequential amplification, and (iii) stable progression from start to finish. A key feature of these processes is that they involve amplifying steps that convert a small signal into a huge, explosive effect. This is supported by the fact that these systems are analogized to avalanches, the snowball effect, and the ripple effect, which all involve a small cause triggering a large outcome (Deyo [2002]; Hougie [2004]; Dodge et al. [2009]; Masten and Cicchetti [2010]). Other common analogies are cascade circuits in electronics with amplifiers arranged in series and natural cascades such as waterfalls, in which a narrow stream of water progressively amplifies in speed and distribution of spread (Macfarlane [1964], [1966]). The stable amplification in these systems leads scientists to refer to them as 'explosions' that run to their 'inevitable conclusion' (Deyo [2002], p. 23). These three features help explain how cascades can be central to the functioning of various physiological and electronic systems and how, in other cases, they produce widespread, catastrophic damage that 'ripples' uncontrollably through a system.

Before examining these three features in detail, it will help to briefly describe an early and influential use of the cascade concept in science. One of the first uses of this concept in modern biology was in reference to blood coagulation. Initiated by some disruption or tear in a blood vessel, coagulation serves to stop blood loss by forming a thick, fibrous clot at the site of injury. In the early 1960s, it was generally accepted that enzymes were involved in clot formation, but it was unclear exactly how this took place. According to a dominant theory, one main enzymatic 'molecular system' was responsible for clot formation, although this theory faced a number of issues. A first issue was that there appeared to be a variety of different enzymes that when deficient,

could each individually cause clotting disorders, as opposed to a single enzyme. Second, it was unclear how a single enzyme system could link the 'minute physical stimulus' of vessel damage with the 'final enzymatic explosion' of clot formation (Macfarlane [1964], p. 499). How. How a single molecular system could produce the many-fold increase in product given a small stimulus remained to be explained.

The field was soon transformed by two similar but independent publications suggesting the 'cascade' and 'waterfall' models of blood coagulation (Macfarlane [1964]; Davie [2003]). Instead of a single enzymatic system these theories postulated a 'multiple-factor theory' involving an eight-step sequence of enzyme activations, in which one enzyme activated another, which activated another and so on, in series. This new theory suggested that blood coagulation involved many factors, that these factors that were initiated in succession, and that the final enzyme was amplified relative to the starting material. When these amplified steps were arranged in succession, they could account for the huge explosion of final clotting product, which explained an important and previously mysterious aspect of the clotting process.[9]

A key feature of this new model was amplification—not only at each step, but when organized in succession, a huge overall amplification from start to finish. In fact, MacFarlane—creator of the cascade model—estimated that if each enzyme produced ten times its own weight in product, that there would be a million-fold overall gain in final clotting material, which is 'a figure close to modern observations' (Hougie [2004], p. 1335). MacFarlane stated that he chose the term 'cascade' to capture this amplification feature by relying on an analogy to electrical circuits, in which amplifiers were arranged in series. In this electrical context, 'cascade' is understood 'as 'a succession of stages in which each stage derives from, or acts, sometimes cumulatively, upon the product or output of the preceding' (Macfarlane [1966], p. 591). MacFarlane ([1966], pp. 596–97) also emphasized amplification by describing blood coagulation a 'biological amplifier' and by comparing this system to a photomultiplier[10]:

---

[9] This sequential amplification allowed for 'an explosive generation of thrombin; [as] a more gradual generation of thrombin is haemostatically ineffective' (Macfarlane [1964], p. 499). This is now referred to as the thrombin 'burst', highlighting the explosive amplification of final product.

[10] Later work on blood coagulation would show a preference for 'cascade' over 'waterfall', as the former was thought to better capture the amplification feature (Hougie [2004], p. 1231).

> It is probably that such an enzyme sequence would develop a progressive increase in activity from stage to stage. This recalls the principle of the photomultiplier in which one electron striking the first dynode releases several more electrons each one of which releases yet more on striking the second dynode and so on, and it was this analogy which prompted the use of the term 'cascade' in the present context. In this concept, the clotting factors thus form a device by which a small stimulus is transmitted as a rising wave of activity to culminate in the large effect of fibrin formation.

What was previously criticized as an overly complex, multi-part enzyme sequence, was now viewed as highly functional and finely tailored to the goals of the system. Soon after publication of the cascade model of blood coagulation this 'relatively simple pattern' was applied to other biological systems including visual processing, hormone signalling, and complement initiation in immunology (Macfarlane [1964], p. 498; Wald [1965]). Discussions of these systems cited the cascade model of blood coagulation, referring to 'the same physiological necessity for amplifying a weak signal' (Bowness [1966], p. 1370). As the cascade notion captured 'a composite of all bare facts in a nutshell' it provided 'an easy to understand and general concept' (Hougie [2004], p. 1231). Furthermore, this notion can be 'applied to other systems' and remains used in science today, which attests to 'the great utility of the concept' (Hougie [2004], p. 1232).

4. Cascade Concept: Features

The cascade concept is not just applicable to systems in biology—it is found in many scientific domains including physics, chemistry, psychology, ecology, and economics. Getting a better understanding of this concept requires appreciating its three main features: (i) an initial trigger, (ii) sequential amplification, and (iii) stable progression from start to finish.

4.1. Trigger

A first feature of cascades is that they involve an initial trigger. The notion of a trigger captures a single main cause, which takes on binary values, and reliably produces some outcome of interest. These features are compatible with the trigger's namesake, the trigger of a firearm (or explosive), which is represented as a single, binary switch that reliably discharges ammunition. Causal triggers are viewed as an important 'first' or 'starting' cause of an outcome or sequence of steps. These factors mark the

beginning of a cascade even when other upstream and nearby causes exist. Triggers are distinguished from other causes on the basis of being located upstream of the causal steps in question and producing their effects with a high likelihood. Causal triggers typically take on binary 'switch-like' values—they are either initiated or not with no in-between. Similarly, a firearm's trigger is either pulled or not—we do not talk about a trigger being initiated more-or-less, but whether it has been initiated at all (or not). Even if the triggering cause can take on a continuous set of values, scientists will identify some threshold that distinguishes the 'on' values from the 'off'. For example, in the blood coagulation cascade scientists state that 'Below the threshold, the system will not 'fire', but above the threshold, explosive propagation will ensue' (Jetsy and Beltrami [2005], p. 2465). This binary feature of causal triggers is also supported by fact that they are often referred to as 'switches' in various scientific contexts (Di Ventura et al. [2006]; Mer [2018]).

An important feature of causal triggers is that they have significant control over their effects, often producing them with a high likelihood. Due to this, scientists claim that these causes have a 'strong' influence and 'kick-start' the causal process (Paine [1980]; Little et al. [2020]). This leads causal triggers to take on significant causal responsibility—they are viewed as having more explanatory power and as being 'better predictors' of the downstream outcome than other causes (Herren and McMahon [2018]).This is due to the fact that these are single causes that do not share causal power with other factors, they reliably produce some downstream effect, and, in the case of cascades, they produce a large effect of some kind.[11] Also, notice that we also see this trigger feature in the causal systems that scientific cascades are analogized to. An avalanche is trigged by a small amount of snowfall and a ripple is triggered by a drop of water.

4.2. Sequential amplification

The second (and perhaps most important) feature of cascades is sequential amplification. While cascades involve a sequence of amplifying steps it will be helpful to first examine amplification at a single causal step. Amplification refers to a

[11] Of course, the trigger relies on a set of background conditions, but these conditions are likely considered important, common, or typically fixed, such that in them, releasing the trigger reliably leads to the effect.

situation in which a small amount of cause produces an amplified or large amount of effect. At the scale of a single causal step two types of amplification should be distinguished. Given some causal factor amplification can involve (1) an increased amount of a single effect, single-product amplification, or (2) an increased number of different effect types, multi-product amplification.[12] An example of single product amplification is found in the blood coagulation cascade, in which an enzyme 'a' produces an increased amount of the single product enzyme 'b'. An example of multi-product amplification is a natural disaster, such as an earthquake, in which a single insult produces a variety of different downstream effects, such as fires, flooding, landslides, and so on. In multi-product amplification a cause influences many different effects,[13] without necessarily increasing their individual values.[14] Both single-product and multi-product amplification capture one-to-many cause–effect scenarios, in which every single unit of cause produces many units of effect.

Amplification is often represented as a continuous property as opposed to being all-or-none. In many sciences, there is interest in quantifying amplification 'degree' especially for single-product cases. An example from electronics is the 'gain' of an electrical circuit, which refers to 'the degree to which a signal has been strengthened' from input to output (Saggio [2014], p. 342). In this case, the gain or 'amplification factor' is defined as the ratio of output power (or voltage or current) to input. A second example is the notion of a 'multiplier' in economics, which captures the 'magnitude of change' in an effect, with respect to its cause. For example, if autonomous spending causes an increase in gross domestic product (GDP), 'The multiplier is the ratio of the change in GDP to the change in autonomous spending—that is, the change in GDP

---

[12] Causal amplification is a distinction among causation that should be added to a list of other important and common distinctions (Woodward [2010]; Ross [2021b], [unpublished]; Ross and Woodward [forthcoming]). Similarly, the opposing notion of 'dampening' should also be added, in which a large cause produces a small effect.

[13] Consider the example of a disaster cascade to make this clearer. An earthquake directly causes three things: the start of fires, the collapse of bridges, and the destruction of city water pipes. What matters here is that one cause produced changes in three different things—not that it (also) increased or decreased the value of these things. We could define the variables so that the earthquake decreased the number of functioning bridges and so on, but this isn't the point. The disaster cascade involves amplification on the second definition (multi-product), because one causes influences many different effects.

[14] Multi-product amplification is related to a type of causal specificity, which can be called 'variable specificity', in which a single cause is capable of producing many different types of effects (Woodward [2010]; Ross [2021b]). A straightforward example from genetics is pleiotropy, in which one gene produces many different outcomes.

divided by the change in autonomous spending that caused the change' (Lipsey and Crystal [2015], p. 410). Another example of amplification degree is R_0 or the 'basic reproduction number' in epidemiology (Cha et al. [2008]). This captures the number of secondary infections caused by a primary infected agent, which represents disease transmissibility. A final example, is the 'multiplication coefficient' of branch chain reactions in physics and chemistry (Baskakov [2007]; Soustelle [2011]). This coefficient captures an increase in reactive species of the product relative to the substrate.

While amplification can take place at a single causal step, cascades involve a sequence of these amplifying steps. This produces increased amplification because when amplifiers are organized in succession the overall 'gain' of the system equals the product of the gains at each step. For example, if one unit of enzyme 'a' produces ten units of enzyme 'b' and one unit of enzyme 'b' produces ten units of enzyme 'c', the overall process has a gain of one hundred (producing one hundred units of 'c' for every single unit of 'a'). Similarly, if one insult produces five distinct outcomes, which each produce five other distinct outcomes, the effects balloon-out, multiplying twenty-five fold. The systems that scientific cascades are analogized to also have this sequential amplification feature. The steady increase of a building avalanche, expanding circles in a ripple effect, and increasing product in an enzyme cascade all involve a sequence of amplifying steps.

4.3. Stable progression

The third feature of cascades is stable progression. This feature refers to the fact that once initiated, cascades involve a strength that extends through their sequence of causal steps. In other words, changes in the causal trigger 'propel' ensuing steps, such that once a cascade is triggered, its sequence of steps unfold with high probability (Dodge et al. [2009], p. 23).

This stable progression is related to the sense in which cascades 'gain momentum' as they move through a series of steps—the growing amount of reactive product at each step can make downstream steps more likely as there are more reactive factors available to produce the next outcome. This is sometimes referred to as a chain-reaction, in which the effect of one reaction propels another to take place, and so on, in a self-sustaining manner. One implication of this stable progression feature is that

once initiated, cascades can be very difficult to stop. They have the potential to 'run-away' or 'become uncontrolled' (Bloomfield and Stephens [1996], pp. 168, 171). This is very similar to descriptions of snowball effect, in which a small snow fall triggers a growing, unavoidable avalanche. As Stein ([1990]) states, 'cascade refers to a process that once started, proceeds stepwise to its full, seemingly inexorable, conclusion […] the danger of a cascade is that it can be inappropriately triggered […] once triggered, it is virtually impossible to stop'.

The three features discussed in this section start to reveal why we are so interested in cascades. We are interested in cascades because they are powerful causal systems. Their power is related to their (a) amplification and (b) stable progression. Cascades produce a huge, expansive outcome and they do so in a way that is difficult to stop. It almost seems that once triggered, they are destined to run to completion. The power of these causal systems is very useful and also very dangerous. These features allow us to purposely set off a causal process that will have a huge outcome in a way that is nearly impossible to stop. And, for the same reasons, we want to make sure that we do not accidentally set them off in a way that will result in a disastrous outcome. Put another way, 'What makes the study of cascading failures so important is the fact that an actual catastrophe, such as infrastructure collapse, global epidemic or a financial meltdown may happen seemingly without warning, starting from a very small failure' (Smolyak et al. [2020]). We have a strong interest in identifying cascades and their triggers because of their explosive effects and nearly unstoppable nature.

4.4 Examples

It will help to illustrate the three features of cascades with a number of examples. These examples will show that the cascade concept has broad applicability in science, that it involves similar features and analogies in different contexts, and that a characteristic diagram is used to represent cascades in many different scientific fields. This characteristic diagram contains a fan-out, one-to-many, and branching structure, which captures sequential amplification in these systems.

A first set of examples are enzyme cascades, which are found in biological systems and various technologies. In these cases, one unit of an enzyme 'a' produces multiple units of enzyme 'b', each unit of 'b' produces multiple units of enzyme 'c', and so

on.[15] Diagrams of these systems depict a fan-out structure, as shown in Figure 1(a), in which a small cause leads to a product that grows progressively larger at each step. Amplification in these cascades is typically directed toward some explicit purpose or goal. For example, these cascades can trigger an immediate burst of life-saving clot, in our visual system they can convert minute visual inputs into a larger, more manageable signal for our brain to process, and they can allow trace amounts of hormone to produce large-scale effects on the body. With respect to various technologies, such as polymerase chain reaction (PCR), cascades allow for amplification of trace amounts of key substances so that we can more easily identify and study them.

A second set of cascade examples are 'branching chain reactions' in physics and chemistry. These involve nuclear or chemical reactions in which an initial substrate creates multiple products and each product creates many more, and so on, in a 'branching', tree-like structure. A paradigmatic example of this is nuclear fission where an initial 'triggering' event—such as the collision of a subatomic particle with an atom—produces an explosive, sequential amplification of reactive products. For example, in the context of uranium substrate, 'Fission of a single $^{235}_{92}$U nucleus produces, on average, three neutrons, each of which can cause another fission if it strikes $^{235}_{92}$U nucleus before escaping from the uranium containing substance' (Moore and Pearson [1981], p. 408). This 'cascade of nuclear fissions' is illustrated in Figure 1, which captures the characteristic fan-out, one-to-many, and branching structure used to represent cascades (Spyrou and Mittig [2017]). Similar diagrams are used to illustrate 'cascade reactions' or 'branching chain' reactions in chemistry and 'collision cascades' in physics (Thompson [2002]; Baskakov [2007]).[16] The high degree of

[15] For example, in hormonal regulation 'a single hormone molecule (e.g. epinephrine) can result in the production of millions of product molecules (e.g. glucose)' (MacDonald [2004], p. 94).

[16] In these and other cascades, the steps (and their products) are referred to in terms of generations (first, second, third, and so on), similar to biological reproduction or fission processes. This, in fact, is what nuclear fission is named after.

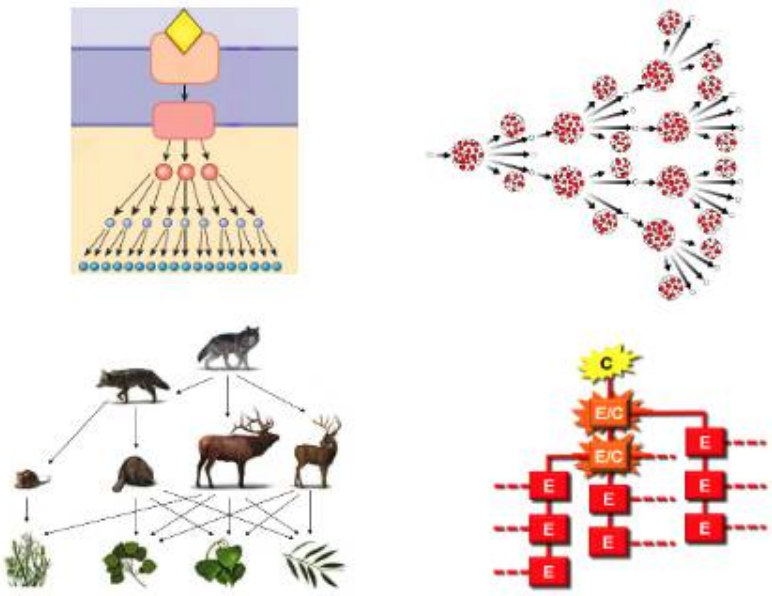

Figure 1. Cascade examples and diagrams. These include (a) cell-signalling cascades in biology (Silverthorn [2013], p. 197), (b) cascading-reactions of nuclear fission in physics (Fox [2020]), (c) trophic cascades in ecology (Freeman et al. [2017], p. 1105), and (d) cascading disasters in socio-ecological systems (Pescaroli and Alexander [2016], p. 176).

amplification in these chemical and physics processes is exploited in the creation of explosives (such as fission bombs) and signal detection methods.

A third example are trophic cascades in ecology, in which changes to a 'key predator' triggers large-scale effects in an ecosystem.[17] In these cases 'The loss of a single species can have a devastating ripple effect across an ecosystem', producing 'ripple-effect extinctions' in other species (Laurance and Vasconcelos [2009], p. 37; Holley [2017], p. 894). A well-known example of this is the elimination of the grey wolf population in Yellowstone National Park in the 1920s after excessive hunting (Boyce [2018]). As shown in Figure 1, alteration to this single wolf species triggered many effects that cascaded through the ecosystem. The removal of the wolf community triggered a surge in the elk population, this decimated the willow tree population, which strained and decreased beaver communities, which caused a decreased water table and ultimately reduced various fish populations. Removal of the wolf also diminished various scavenger species as they normally fed on leftover

[17] For the first use of 'trophic cascade' in ecology, see (Paine [1980]). Further scientific and philosophical discussion can be found in (Ripple et al. [2016]; Millstein [2019]).

wolf kills. This cascade meets the second type of amplification—this involves an increase in the number of effects that a single cause influences, not an increase in the value (or amount) of each of these effects. So if changing the wolf directly causes a change in five species (elk, coyote, bear, raven, and eagle), and changes in each of these produces changes in many more (including beaver, fish, amphibian, and song birds) then this second definition of amplification is met even if the number of individuals in a given species reduce in number.[18] The amplification in these systems interests ecologists because they capture how some species exert a 'disproportionate influence' on the ecosystem relative to their abundance (Power et al. [1996]; Cottee-Jones and Whittaker [2012]). This is related to the notions of a 'keystone' and 'umbrella' species in ecology, in which a single species exerts a stabilizing force (such as a keystone in an arch) or a protective influence over many species (under their 'umbrella') in the ecosystem (Roberge and Angelstam [2004]). Just as removal of these species can produce wide-spread damage, reintroducing them can re-stabilize the ecosystem through the same distributed causal connections.

A fourth set of examples are cascading disasters, in which a man-made accident or natural disaster triggers a sequence of expanding and damaging effects. An example of this is the 2011 Tohoku earthquake in Japan, which triggered a set of secondary insults (a tsunami, power loss, collapse of roadways and railways), which each produced many tertiary effects (failure of nuclear power plants, rupture of dams), leading to quaternary outcomes (radiation damage), and so on. The sequential amplification feature of cascading disasters is represented in Figure 1, which depicts expanding effects after an initial stimulus. In the context of these disasters, it is suggested that 'a vital factor of the definition of cascades' is the involvement of elements that 'amplify each other' and produce an 'amplification of damage'

[18] It may seem puzzling that this cascade decreases the number of individual within downstream species, when the notion of amplification seems to involve an increase in something. It is important to clarify, that the second type of amplification involves an increase in the number of outcomes influenced, not an increase in the value (or amount) of each outcome. So, if changes in one species cause changes in many other species, this is all that is needed for the second amplification type—it doesn't need to be the case that the many other species all increase in number. The amplification here is in terms of number of effects or variables that the single cause influences—not that it also increases their individual values. These two types of amplification—single and multi-product—are related to the notions of value specificity and variable specificity (Ross [2021b]). In particular, a cause can have influence over many values (or states) of a single effect versus few of its values (value specificity). Alternatively, a cause can have many different effects as opposed to a few (variable specificity).

(Pescaroli and Alexander [2015], p. 59). In these cases, the potential for a cascading failure is associated with the interdependency of components in a system, such that the failure of one inevitably leads to the failure of another. In particular, it is stated that as systems become 'more interdependent, they become more vulnerable to large-scale cascading disruptions' (Pescaroli and Alexander [2016], p. 180).

As a final example consider the cascade-like spread of infectious diseases, such as COVID-19, that is captured in epidemiology with the basic reproduction number (or R_0). This spread is illustrated with branching diagrams that depict how one person can transmit this disease to multiple people, and each of these individuals can transmit it to many more and so on, in a set of steps that amplify disease occurrence. If this keeps happening disease incidence will balloon out, as a large number of cases will be triggered by single patient. This is a cascade—in this case, disease incidence dangerously amplifies as it spreads through population. If you want to clearly communicate the threat and manner of spread of COVID-19 through the population, you should not call this transmission a mechanism—you should call it a cascade. The cascade concept better captures the amplified spread of this disease through the population, which it what we aim to draw attention to with these descriptions and illustrations of disease transmission.[19]

5. Cascade versus Mechanism

How exactly do cascades and mechanisms differ from each other? One answer to this question involves citing the different features and analogies that these causal structures are associated with, which has been discussed in this analysis. Mechanisms involve constitutive relations, cascades do not. Cascades are analogized to the snowball effect, mechanisms are analogized to machines, and so on. While there are many differences between these causal structures, it will help to explicitly discuss a few of these in more detail.

5.1. Differences

A first difference between these concepts is that cascades lack the constitutive, part-whole feature of mechanisms. Mechanisms have a hierarchical structure, in which

[19] Other examples not discussed in detail are collision cascades, oxidative cascades, and social cascades.

their lower-level parts produce some higher-level outcome of the whole mechanism.[20] Cascades, on the other hand, are level-agnostic with respect to their causes and effects. Cascades can have causes and effects at the same level, higher-level causes that produce lower-level effects, and lower-level causes that produce higher-level effects.[21] Additionally, the hierarchical nature of mechanisms contributes to the view that they are 'discrete' systems with lower-level parts that are contained within higher-levels. Cascades are not confined in this way—they can spill over from one system to another as they ripple through some domain. With cascades the emphasis is on causal influence that moves upstream-to-downstream, without being confined to levels or always moving bottom-up. Cascades have causal influence that is better understood as 'relationships between distinct' factors, as opposed to 'a whole and its parts' (Craver [2007], p. 178).

Second, cascades and mechanisms are studied with different investigative strategies. Recall that mechanisms are studied with decomposition and localization, which involves fixing an explanatory target and then drilling down to identify its lower-level causal parts (Bechtel and Richardson [2010]). Cascades are not studied in this way for a variety of reasons. First, as indicated above, cascades do not always have lower-level causes that one can drill down to find. Causal factors in cascades can extend across many different levels and scales. Second, the first step of fixing an explanatory target is ill-suited to cascades because—unlike mechanisms—they are not always relative to single outcomes. We see this in cascades with multi-product amplification, in which many effects are produced. Another difference is that when cascades are studied their entire set of effects are not always known, but part of what scientists want to uncover. Instead of starting with an effect and drilling down, a better approach for cascades can involve starting with the trigger and expanding out to

[20] As Craver ([2007], p. 165) states, 'The property or activity at a higher level of mechanisms is the behavior of the mechanism as a whole (the explanandum phenomenon); the parts of the mechanism and their activities are at a lower level'. This is also seen in (Wimsatt [1974], p. 686).

[21] Enzymatic cascades are an example of the first type as the cause and effect are both enzymes. An example of higher-level causes producing lower-level effects are energy cascades involved in turbulence, which involve the transfer of energy from 'large scales of motion to the small scales' (Richardson [1922], p. 66; see also Wilson [2021]). Another example are traumatic experiences that alter gene expression, referred to as 'downward cascades' (Masten and Cicchetti [2010], p. 492). Finally, a cascade example with a lower-level cause that produces a higher-level effect is a hormone cascade, in which a hormone trigger produces some system-level behaviour. Another example is a pharmacological intervention that alters behaviour, sometimes called an 'upward cascade' (Masten and Cicchetti [2010], p. 492).

identify what effects are produced. This requires a system-wide outlook, which contrasts with the more local focus on discrete mechanisms.

Finally, studying cascades involves other challenges that are not necessarily present in mechanisms. As it is difficult to stop a cascade once initiated there is often significant focus on experimental manipulations of the causal trigger. This is related to the significance of the causal trigger for prediction, prevention, and control. For example, consider developmental cascades in psychology where triggers in early childhood can cascade into many other outcomes. In this case 'the evidence in prevention science indicates a higher return on investment in early childhood interventions' (Masten and Cicchetti [2010], p. 491)—in other words, targeting upstream portions of the cascade provides more control over downstream outcomes (compared to targeting intermediate or later factors). The amplification and stable progression of cascades have motivated other means of controlling these systems. In an effort to prevent catastrophic large-scale avalanches, scientists will purposefully trigger smaller 'controlled' avalanches to reduce excessive snowfall build-up. Similar strategies have been suggested for COVID-19, by exposing 'individuals whose probability of developing serious health conditions is low' under 'controlled supervision' to reach heard immunity faster (Klement et al. [unpublished]). While the cascade's trigger receives significant attention for the purposes of prediction, prevention, and control, there are also efforts to identify intermediate causes and background factors that can be manipulated to control cascades even when the trigger has been initiated.

Third, unlike mechanisms, cascades are not expected to contain mechanical or other types of fine-grained causal detail. Various philosophers suggest that identifying a 'complete' mechanism requires specifying the 'activities' present in causal connections 'without leaving gaps, or using filler terms to 'stand for some-process-we-know-not' (Craver [2007], pp. 113–14). In particular, models that 'black box' or abstract from information about 'working components' are not sufficiently mechanistic, as 'mechanisms are the working components revealed by opening black boxes' (Craver [2007], p. 113). Overly sparse models fail to meet mechanist standards as one 'loses sight of how the nodes work to produce, underlie, or maintain the phenomenon' (Craver and Darden [2013], p. 91). Clarifying 'how' causes produce

their effects is very much the business of mechanisms and this is provided by specifying 'mechanical' or 'activity' information about causal connections.

While mechanisms are expected to contain fine-grained and mechanical details, cascades are not. Cascades capture 'that' a cause amplifies its effect, but they need not detail 'how' this amplification takes place. Appreciating this distinction is important for two main reasons. First, when cascades are explanatory their lower-level mechanistic detail is often unnecessary for the explanation. In these cases, it is the presence of one-to-many amplification that does the explanatory work, not the details of how this amplification is produced. Perturbing the grey wolf population in Yellowstone National Park results in a trophic cascade because of their position in a set of interconnections in the ecosystem. Disrupting the raven or song bird populations would not have this effect, because they lack a similar interconnected position. Once these higher-level connections in the ecosystem are identified, including lower-level mechanistic information is possible, but unnecessary for the explanation. In the grey wolf trophic cascade, you do not need to know the causal mechanical details of how the Elk consumes the willow tree and other plant species, you just need to know that it consumes many species, as opposed to few.

Distinguishing these causal structures is also important because, in some cases cascade information explains, while mechanistic information is completely unexplanatory. In order to see this, consider that the same cascade can be instantiated by different causal mechanical 'how' details. As an example, consider the cascade-like spread of COVID-19 and other infectious diseases. Notice that we use the same branching diagram to represent and explain the amplified spread of COVID-19 in various contexts, even though the mechanisms of spread differ drastically across these contexts. COVID-19 can be spread through different mechanisms of airborne transmission (for example, talking, coughing, sneezing), contact with contaminated objects (for example, door handles, shared utensils), and so on. In these cases, the cascade structure is helpful because it communicates that the spread is one-to-many as opposed to one-to-one. This is important because it means the disease is more easily transmitted, harder to control, and that early containment is crucial. All of these are implications of the shared higher-level cascade structure and not lower-level mechanisms (which differ greatly across these systems). If you want to explain COVID transmission across these contexts you cannot appeal to lower-level

mechanisms because they differ across these cases. However, we can appeal to the cascade causal structure because it is shared across these cases and it captures the main features of this causal spread.

Fourth, while mechanisms are discrete systems in a hierarchical sense, the beginning and end of cascades are more precisely defined. Mechanistic philosophers have claimed that the start and finish of mechanisms are dictated by pragmatic, contextual factors, without representing natural divisions in the world (Craver [2009]; Bechtel [2015]). The bounds of a mechanism can change depending on goals, interests, and explanatory targets. With cascades, no matter what pragmatic and contextual considerations are involved, we know that their beginning and end have objective features. The cascade begins with a small causal trigger and ends with a large downstream effect.

There is a lot more that can be said about how cascades and mechanisms differ. One way to discuss this is to respond to potential objections that could be raised in response to this analysis.

5.2. Potential objections

This paper suggests that cascades and mechanisms are distinct causal structures—that they are associated with distinct features, analogies, and causal investigative strategies. This differs from common philosophical views that all causal concepts and causal structures in science are well understood as mechanisms. In completing this analysis, it will help to address some potential objections.

A first objection states that the discussion in this paper is merely semantic or terminological. This could be used to imply that this analysis is trivial and inconsequential. Who cares whether a causal structure is referred to by one name or another? Who cares whether we call a causal system ‘cascade’ versus ‘mechanism’? While there are important arguments for how these structures are named (which I discuss soon), this detracts from deeper questions that motivate this analysis and that are viewed as foundational in philosophy. These include: What types of causal structure are present in the world? And are these structures explanatory? Mainstream responses to these questions cite mechanisms, so we are obligated to ask what mechanisms are and what causal system or systems they refer to. This paper answers these questions by specifying two types of causal structure (whatever you want to call

them) that are identified in science and that figure in scientific explanations. This provides a novel and scientifically informed answer to these questions—questions that are not trivial or inconsequential, but that are viewed as foundational in philosophy.

A second objection accepts my distinction between these causal concepts and structures, but claims that mechanisms are more explanatory than cascades. This is often motivated by a reductive mechanism position, which views the fine-grained, lower-level detail of mechanisms as providing enhanced explanatory power. A first response to this objection (outlined in the last subsection) asserts that for many explanatory targets mechanistic details are unnecessary, while cascade information explains. In these cases, it is the presence of one-to-many amplification that does the explanatory work, not the details of how this amplification is produced. In these cases, the mechanistic details can be included, but they do not enhance or improve the explanation.

There is a stronger response to this objection. In some cases, mechanism information is unexplanatory in the sense of detracting from the explanation, while the cascade structure explains. This occurs when a group of systems share a cascade structure, but have different mechanisms instantiating this structure. This is seen in explanations of the cascade-like spread of COVID in different locations, despite the fact that the causal mechanical details of transmission vary from one context to another. Other examples include physiological cascades (such as visual phototransduction) in which there is the same initial trigger and amplified effect, but the enzymes processing the amplification can differ across instances of the cascade. The reason why these systems exhibit a similar outcome is not explained by lower-level mechanisms, which differ across systems. It is explained by the one-to-many cascade structure that is shared across all cases. We can also make this point in a broader way. Why do trophic cascades, enzyme cascades, and cascading reactions in physics all produce large amounts of some effect? It is not because they all share lower-level mechanisms—their mechanistic and micro-structural details differ enormously. The explanation for this is that all of these systems share the same higher-level cascade structure, in which a triggering cause initiates some sequence of one-to-many causal amplification steps.

A third objection to this analysis rejects these distinctions and claims that what I call a ‘cascade’ is really a mechanism. This represents an expansive mechanistic

position, as it proposes that all (or most) causal structures are well understood as mechanisms. This is a common claim in the literature—it is associated with views that mechanisms can be abstract and that they can be linear (or etiological) without containing the constitutive feature sometimes attributed to them (Levy and Bechtel [2013]; Craver and Tabery [2015]). On this view mechanism encompasses the cascade structure and cascade is viewed as type of mechanism.

Here is how I think we should respond to this objection. As philosophers we are free to define mechanisms however we like—this is often done according to various standards that we find appropriate. We can define mechanism in a broad way suggested by this objection or in a narrow way as I have done in my analysis. If we choose to define mechanism in this broad way, this comes with various disadvantages. In order to see what these are consider the following. As scientists have explored the world they have identified various causal structures, which they distinguish from one another. Some causal systems have constitutive relations, other are more linear, some have amplifying effects, others do not, and so on. Part of what this expansive definition suggests is that all of these causal structures are well understood as mechanisms. However, if all of these structures are mechanisms—and they have varying features such that none are characteristic of unique to the mechanism concept—then mechanism just means 'complex causal structure' and we might as well replace it with this. This is a problem because we often think that 'mechanism' communicates more than this. Thus, a first problem with the broad definition of mechanism is that we lose out on 'mechanism' meaning anything more than 'complex causal structure' and we usually think that it conveys more than this.

A second problem is that this broad notion of mechanism fails to accurately capture how scientists use causal terminology. This noteworthy because scientific use of this term was an original motivation for many mechanistic accounts of causal explanation (Machamer et al. [2000], p. 2). What is sometimes overlooked is that scientists often expect descriptions of mechanisms to contain significant causal detail. This is seen in discussions of a drug's mechanism of action and mechanisms of enzyme catalysis. These mechanisms are expected to provide 'a comprehensive understanding of the entire sequence of events' and 'detailed knowledge of the causal and temporal relationships among all the steps leading to a specific effect' (Hutchinson [2007], p. 7; see also Ankley et al. [2010], p. 731). This expectation of significant detail often

leads scientists to admit that they do not yet have the complete mechanism of interest although they continue to work toward this.

With the cascade concept scientists often explicitly acknowledge the focus on a sparser causal structure. In early work on the cascade model of blood coagulation, MacFarlane distinguished this model from a more detailed mechanistic understanding. He claimed that the cascade model captured 'the gross ways in which [components] interact [rather] than in details of their individual structure' and that this work is distinct from 'later detailed investigation of its underlying mechanism in terms of molecular structure' (Macfarlane [1966], p. 592). Cascade models are similar to other abstract causal systems that only capture 'selected key events' and that have 'gaps and black boxes in which mechanistic details are either unknown or not needed' (Hutchinson [2007], p. 1; see also Ankley et al. [2010], p. 732). The fact that scientists reserve the term 'mechanism' for highly detailed systems puts pressure on claims that it could 'do the same work' as concepts such as 'cascade', which refer to systems that are less detailed and have other unique features.

Third, if we adopt this broad notion of mechanism we lose out on the utility of analogy and analogical reasoning. One advantage of the mechanism and cascade concepts is that they are 'linguistic labels' that originate from causal systems we are familiar with in everyday life (Gentner and Smith [2012]). Drawing on knowledge of familiar causal structures is a useful way to make complex causal systems in science more cognitively accessible. This allows scientists to highlight features of causal systems that are important for their functioning, behaviour, and for the questions of interest. Use of these concepts necessarily depends on background knowledge of the audience—this means that these concepts can mislead if used inappropriately or facilitate understanding when cited at the right time. If the goal is to communicate the steps of a developmental pathway or the sequential growth of a cascade, calling these 'mechanisms' can lead to incorrect conclusions about the system. Use of 'mechanism' can lead audiences to assume these systems have properties that they lack (reductive character, lower-levels components, and mechanical interactions) and to overlook properties that they have (higher-level causal connections, sequential amplification, flow of material, and so on). Appreciating the role of these terms in communicating causal information is important for the general project of science communication, how scientists convey their work to various audiences (experts, lay persons, and so on),

and how philosophers formulate and justify their views of the causal structure of the world.

Appreciating these causal terms and the distinct structures they refer to also promotes analogical reasoning and problem-solving. These analogies do not just pick out similar structures, but also the unique implications that follow from them. Enzyme cascades gain momentum and become difficult to stop just like the snowball effect. In considering how to controls these enzyme systems, it can be helpful to formulate the problem in terms of avalanches and other ordinary life cascades. These can lead to the conclusion that if one wants to prevent or stop a cascade outcome, early interventions more likely to be successful—these circumvent the increasing momentum and sequential amplification features that cause them to 'run-away' uncontrollably. Similar advantages are present for other causal terms such as pathways and mechanisms. Various scientific pathways—such as metabolic, developmental, and anatomical pathways—are analogized to roadways, highways, and city streets (Ross [2021a]). In many of these scientific systems researchers conceive of 'traffic' build-up along the pathway and they identify remedies for this by thinking in terms of a freeway traffic model. The lessons do not follow from mechanism because this concept lacks the unique features present in cascades and pathways, which lead to the behaviours in these systems.

My responses to this final objection and the expansive notion of mechanism are pointing to a similar worry. This worry is that if mechanism means everything, it starts to mean nothing. Indeed, on some accounts 'it seems that mechanisms just are whatever explains whatever happens' and 'there is a serious danger of vacuity in [these] treatments of the topic' (Dupré [2013]). There are substantial concerns about whether this notion of mechanism can capture the diversity of causal structure in science and whether it does justice to the mechanism concept itself. It should be uncontroversial that 'If the concept of a mechanism is to do any work, we must surely have some sense of what isn't a mechanism' (Dupré [2013]). Along these lines, if 'mechanism' is more than just a synonym for 'causal structure' (which would trivialize the mechanism concept) it should be clear what causal structures do and do not count as mechanisms.

Our world contains different types of causal systems—this is a fact that we cannot deny. Some causal relationships are oriented linearly, others branch out; some causes

are fast others are slow; some causes are strong others are weak; some are stable, while others are sensitive. We care about these distinctions, in science and everyday life, because causal structures with different features have different implications. They provide different types of control over outcomes in the world and these are more or less useful given our interests. Scientists are going to study and appreciate these distinctions no matter what they are called. This is harder to do when all causal structures are referred to as mechanisms and 'mechanism' loses it meaning. This is easier to do when we accept and appreciate different causal terms, analogies they are associated with, and the distinct causal structures they refer to.

6. Conclusion

A considerable amount of philosophical work has focused on the causal structure of the world and how this structure figures in explanation. Many accounts interpret this structure as mechanistic and claim that genuine causal explanations appeal to mechanisms. If mechanism just means 'complex causal structure' then these accounts are trivially true. On this reading, causal explanations appeal to mechanisms, of course, but nothing new has been said. However, if mechanism has a non-trivial meaning and refers to a causal structure with particular features—constitutive relations, mechanical interactions, or others—we have good reason to conclude that it alone will fail to capture the diversity of causal structure in science. While all complex causal structures contain causal relations, there is incredible variation when it comes to their other features.

This paper suggests a new way of understanding causal structure and causal explanation in science, but it also opens many questions that have yet to be addressed. For example, there is much more to say about the cascade concept in particular, including how it may differ across contexts, how it relates interconnections within a system (and the concepts of redundancy, modularity, and so on), and the role of background conditions in supporting cascade-like reactions. It should be clear how this structure relates to 'the butterfly effect', positive and negative feedback loops, and systems that involve causal dampening, as opposed to amplification. While other work has argued for the distinctiveness of the pathway concept (Ross [2021a]), it will be illuminating to explore whether other non-mechanistic causal structures exist. It seems likely that pathways and cascades only scratch the surface of the diverse types

of causal systems in science. Capturing this diversity is necessary for an accurate conception of the causal structure of the world, for a full picture of the methods and reasoning we use in identifying this structure, and a realistic starting point in communicating this structure to various audiences.

Acknowledgements

I would like to thank James Woodward, William Bechtel, Lindley Darden, and two anonymous reviewers for helpful feedback on this paper. I would also like to acknowledge the Munich Center for Mathematical Philosophy at Ludwig Maximilian University of Munich and the National Science Foundation (Grant [1945647]) for supporting this research.

Department of Logic and Philosophy of Science
University of California, Irvine
California, USA
rossl@uci.edu

References

Ankley, G. T., Bennett, R. S., Erickson, R. J., Hoff, D. J., Hornung, M. W., Johnson, R. D., Mount, D. R., Nichols, J. W., Russom, C. L., Schmieder, P. K., Serrrano, J. A., Tietge, J. E. and Villeneuve, D. L. [2010]: 'Adverse Outcome Pathways: A Conceptual Framework to Support Ecotoxicology Research and Risk Assessment', *Environmental Toxicology and Chemistry*, **29**, pp. 730–41.

Baskakov, I. [2007]: 'Branched Chain Mechanism of Polymerization and Ultrastructure of Prion Protein Amyloid Fibrils', *The FEBS Journal*, **274**, pp. 3756–65.

Bechtel, W. [2015]: 'Can Mechanistic Explanation Be Reconciled with Scale-Free Constitution and Dynamics?', *Studies in History and Philosophy of Biological and Biomedical Sciences*, **53**, pp. 84–93.

Bechtel, W. and Richardson, R. C. [2010]: *Discovering Complexity*, Cambridge, MA: MIT Press.

Bich, L. and Bechtel, W. [2022]: 'Organization Needs Organization: Understanding Integrated Control in Living Organisms', *Studies in History and Philosophy of Science*.

Bloomfield, M. M. and Stephens, L. J. [1996]: *Chemistry and the Living Organism*, New York: John Wiley.

Bogen, J. and Machamer, P. [2010]: *Mechanistic Information and Causal Continuity*, Oxford: Oxford University Press.

Bowness, J. M. [1966]: 'Epinephrine: Cascade Reactions and Glycogenolytic Effect', *Science*, **152**, pp. 1370–71.

Boyce, M. S. [2018]: 'Wolves for Yellowstone: Dynamics in Time and Space', *Journal of Mammalogy*, **99**, pp. 1021–31.

Brigandt, I. [2013]: 'Systems Biology and the Integration of Mechanistic Explanation and Mathematical Explanation', *Studies in History and Philosophy of Biological and Biomedical Sciences*, **44**, pp. 477–92.

Cha, M., Mislove, A., Adams, B. and Gummadi, K. P. [2008]: 'Characterizing Social Cascades in Flickr', in Association for Computing Machinery (*ed.*), *Proceedings of the First Workshop on Online Social Networks*, New York: ACM, pp. 13–18.

Cottee-Jones, H. E. W. and Whittaker, R. J. [2012]: 'The Keystone Species Concept: A Critical Appraisal', **4**, *Frontiers of Biogeography*, pp. 117–27.

Craver, C. and Darden, L. [2013]: *In Search of Mechanisms*, Chicago, IL: University of Chicago Press.

Craver, C. and Tabery, J. [2015]: 'Mechanisms in Science', in E. N. Zalta (*ed.*), *Stanford Encylopedia of Philosophy*, available at <plato.stanford.edu/archives/win2015/entries/science-mechanisms/>.

Craver, C. F. [2007]: *Explaining the Brain*, Oxford: Oxford University Press.

Craver, C. F. [2009]: 'Mechanisms and Natural Kinds', *Philosophical Psychology*, **22**, pp. 575–94.

Darden, L. and Craver, C. [2001]: 'Discovering Mechanisms in Neurobiology', P. K. Machamer, R. Grush and P. McLaughlin (*eds*), *Theory and Method in Neuroscience*, Pittsburgh, PA: University of Pittsburgh Press, pp. 112–37.

Davie, E. W. [2003]: 'A Brief Historical Review of the Waterfall/Cascade of Blood Coagulation', *Journal of Biological Chemistry*, **278**, pp. 50819–32.

Deyo, R. A. [2002]: 'Cascade Effects of Medical Technology', *Annual Review of Public Health*, **23**, pp. 23–44.

Di Ventura, B., Lemerle, C., Michalodimitrakis, K. and Serrano, L. [2006]: 'From in vivo to in silico Biology and Back', *Nature*, **443**, pp. 1–7.

Dodge, K. A., Malone, P. S., Lansford, J. E., Miller, S., Pettit, G. S., Bates, J. E., Collins, A., Schulenberg, J. E. and Maslowsky, J. [2009]: *A Dynamic Cascade Model of the Development of Substance-Use Onset*, Oxford: Wiley-Blackwell.

Dupre, J. [2013]: 'Living Causes', *Aristotelian Society* Supplementary Volume, **87**, pp. 19–37.

Fagan, M. B. [2013]: *Philosophy of Stem Cell Biology*, New York: Palgrave Macmillan.

Fox, M. [2020]: 'Nuclear Fission', *AccessScience*, McGraw Hill, available at <doi.org/10.1036/1097-8542.458400>.

Freeman, S., Quilin, K., Allison, L., Black, M., Taylor, E., Podgorski, G., and Carmichael, J. [2017]: *Biological Science*, Harlow: Pearson.

Gentner, D. and Smith, L. [2012]: 'Analogical Reasoning', in V. Ramachandran (*ed.*), Encyclopedia of Human Behavior, Oxford: Elsevier, pp. 130–36.

Glennan, S. [2017]: *The New Mechanical Philosophy*, Oxford: Oxford University Press.

Glennan, S. S. [1996]: 'Mechanisms and the Nature of Causation', *Erkenntinis*, **44**, pp. 49–71.

Herren, C. M. and McMahon, K. D. [2018]: 'Keystone Taxa Predict Compositional Change in Microbial Communities', *Environmental Microbiology*, pp. 2207–17.

Holley, D. [2017]: *General Biology II: Organisms and Ecology*, Indianapolis, IN: Dog Ear Publishing.

Hougie, C. [2004]: 'The Waterfall-Cascade and Autoprothrombin Hypotheses of Blood Coagulation: Personal Reflections from an Observer', *Journal of Thrombosis and Haemostasis*, **2**, pp. 1225–33.

Huneman, P. [2018]: 'Diversifying the Picture of Explanations in Biological Sciences: Ways of Combining Topology with Mechanisms', *Synthese*, **195**, pp. 115–46.

Hutchinson, T. [2007]: 'Intelligent Testing Strategies in Ecotoxicology: Mode of Action Approach for Specifically Acting Chemicals', Technical Report TR 102, ECETOC, available at <www.ecetoc.org/publication/tr-102-intelligent-testing-strategies-in-ecotoxicology-mode-of-action-approach-for-specifically-acting-chemicals/>.

Jetsy, J. and Beltrami, E. [2005]: 'Positive Feedbacks of Coagulation', *Arteriosclerosis, Thrombosis, and Vascular Biology*, **25**, pp. 2463–69.

Kaplan, D. M. and Bechtel, W. [2011]: 'Dynamical Models: An Alternative or Complement to Mechanistic Explanations?', *Topics in Cognitive Science*, **3**, pp. 438–44.

Kaplan, D. M. and Craver, C. F. [2011]: 'The Explanatory Force of Dynamical and Mathematical Models in Neuroscience: A Mechanistic Perspective', *Philosophy of Science*, **78**, pp. 601–27.

Klement, E., Klement, A., Chinitz, D., Harel, A., Fattal, E. and Klausner, Z. [unpublished]: 'Controlled Avalanche: A Regulated Voluntary Exposure Approach for Addressing Covid19', available at <www.medrxiv.org/content/10.1101/2020.04.12.20062687v1>.

Lange, M. [2018]: *Because without Cause: Non-causal Explanations in Science and Mathematics*, Oxford: Oxford University Press.

Laurance, W. F. and Vasconcelos, H. L. [2009]: 'Deforestation and Fragmentation in the Amazon', in K. Del Claro, P. S. Oliveira and V. Rico-Gray (*eds*), *Tropical Biology and Conservation Management*, Vol. 2, Oxford: EOLSS Publications pp. 23–39.

Levy, A. [2014]: 'Machine-Likeness and Explanation by Decomposition', *Philosopher's Imprint*, **14**, pp. 1–15.

Levy, A. and Bechtel, W. [2013]: 'Abstraction and the Organization of Mechanisms', *Philosophy of Science*, **80**, pp. 241–61.

Lipsey, R. and Crystal, A. [2015]: *Economics*, Oxford: Oxford University Press.

Little, J. C., Garcia-Garcia, E., Sul, A. and Kalderon, D. [2020]: 'Drosophila Hedgehog Can Act as a Morphogen in the Absence of Regulated Ci Processing', *eLife*, available at <doi.org/10.7554/eLife.61083>.

MacDonald, J. A. [2004]: 'Signal Transduction Pathways and the Control of Cellular Responses to External Stimuli', in K. B. Storey (*ed.*), *Functional Metabolism: Regulation and Adaptation*, Hoboken, NJ: Wiley-Liss.

Macfarlane, R. G. [1964]: 'An Enzyme Cascade in the Blood Clotting Mechanism, and Its Funciton as a Biochemical Amplifier', *Nature*, **202**, pp. 498–99.

Macfarlane, R. G. [1966]: 'The Basis of the Cascade Hypothesis of Blood Clotting', *Thrombosis et Diathesis Haemorrhagica*, **15**, pp. 591–602.

Machamer, P., Darden, L. and Craver, C. F. [2000]: 'Thinking about Mechanisms', *Philosophy of Science*, **67**, pp. 1–25.

Masten, A. S. and Cicchetti, D. [2010]: 'Developmental Cascades', *Development and Psychopathology*, **22**, pp. 491–95.

Mer, G. [2018]: 'Mayo Researchers Find Off/On Switch for DNA Repair Protein', *News Network*, available at <newsnetwork.mayoclinic.org/discussion/mayo-researchers-find-off-on-switch-for-dna-repair-protein/>.

Millstein, R. L. [2019]: 'Types of Experiments and Causal Process Tracing: What Happened on the Kaibab Plateau in the 1920s', *Studies in History and Philosophy of Science Part A*, **78**, pp. 98–104.

Moore, J. W. and Pearson, R. G. [1981]: *Kinetics and Mechanism*, New York: Wiley.

Paine, R. T. [1980]: 'Food Webs: Linkage, Interaction Strength, and Community Infrastructure', *Journal of Animal Ecology*, **49**, pp. 666–85.

Pescaroli, G. and Alexander, D. [2015]: 'A Definition of Cascading Disasters and Cascading Effects: Going beyond the Toppling Dominos Metaphor', *Planet@Risk*, **3**, pp. 58–67.

Pescaroli, G. and Alexander, D. [2016]: 'Critical Infrastructure, Panarchies and the Vulnerability Paths of Cascading Disasters', *Natural Hazards*, **82**, pp. 175–92.

Power, M., Tilman, D., Estes, j. A., Menge, B. A., Bond, W. J., Daily, G., Castilla, J. C., Lubchenco, J., Paine, R. T. and Mills, L. S. [1996]: 'Challenges in the Quest for Keystones', *BioScience*, **46**, pp. 609–20.

Reutlinger, A. and Andersen, H. [2016]: 'Abstract versus Causal Explanations?', *International Studies in the Philosophy of Science*, **30**, pp. 129–46.

Richardson, L. F. [1922]: *Weather Prediction by Numerical Process*, Cambridge: Cambridge University Press.

Ripple, W. J., Estes, j. A., Schmitz, O. J., Constatnt, V., Kaylor, M. J., Lenz, A., Motley, J. L., Self, K. E., Taylor, D. S. and Wolf, C. [2016]: 'What Is a Trophic Cascade?', *Trends in Ecology and Evolution*, **31**, pp. 842–49.

Roberge, J.-M. and Angelstam, P. [2004]: 'Usefulness of the Umbrella Species Concept as a Conservation Tool', *Conservation Biology*, **18**, pp. 76–85.

Robins, S. K. and Craver, C. [2009]: *Biological Clocks: Explaining with Models of Mechanisms*, Oxford: Oxford University Press.

Ross, L. and Woodward, J. [forthcoming]: 'Irreversible (One-hit) and Reversible (Sustaining) Causation', *Philosophy of Science*, available at <doi.org/10.1017/psa.2022.70>.

Ross, L. N. [unpublished]: 'Explanation in Contexts of Causal Complexity: Lessons from Psychiatric Genetics'.

Ross, L. N. [2021a]: 'Causal Concepts in Biology: How Pathways Differ from Mechanisms and Why It Matters', *British Journal for the Philosophy of Science*, **72**, pp. 131–58.

Ross, L. N. [2021b]: 'Causes with Material Continuity', **36**, *Biology and Philosophy*, available at <doi.org/10.1007/s10539-021-09826-x>.

Saatsi, J. and Reutlinger, A. [2018]: *Introduction: Scientific Explanations beyond Causation*, Oxford: Oxford University Press.

Saggio, G. [2014]: *Principles of Analogy Electronics*, Boca Raton, FL: CRC Press.

Salmon, W. [1984]: *Scientific Explanation and the Causal Structure of the World*, Princeton, NJ: Princeton University Press.

Silberstein, M. and Chemero, A. [2013]: 'Constraints on Localization and Decomposition as Explanatory Strategies in the Biological Sciences', *Philosophy of Science*, **80**, pp. 958–70.

Skillings, D. J. [2015]: 'Mechanistic Explanation of Biological Processes', *Philosophy of Science*, **82**, pp. 1139–51.

Skipper, R. A. and Millstein, R. L. [2005]: 'Thinking about Evolutionary Mechanisms: Natural Selection', *Studies in History and Philosophy of Biological and Biomedical Sciences*, **36**, pp. 327–47.

Smolyak, A., Levy, O., Vodenska, I., Buldyrev, S. and Havlin, S. [2020]: Mitigation of Cascading Failures in Complex Networks. Nature: Scientific Reports.

Soustelle, M. [2011]: *An Introduction to Chemical Kinetics*, London: John Wiley.

Spyrou, A. and Mittig, W. [2017]: 'The Science behind the First Nuclear Chain Reaction, which Ushered in the Atomic Age 75 Years Ago', *Smithsonian Magazine*, available at <www.smithsonianmag.com/innovation/the-science-behind-first-nuclear-chain-reaction-180967375/>.

Stein, H. F. [1990]: *American Medicine as Culture*, New York: Taylor and Francis.

Thagard, P. [2003]: 'Pathways to Biomedical Discovery', *Philosophy of Science*, **70**, pp. 235–54.

Thompson, M. W. [2002]: 'Atomic Collision Cascades in Solids', *Vacuum*, **66**, pp. 99–114.

Silverthorn, D. U. [2013]: *Human Physiology*, Boston, MA: Pearson.

Wald, G. [1965]: 'Visual Excitation and Blood Clotting', *Science*, **150**, pp. 1028–30.

Wilson, M. [2021]: *Imitation of Rigor*, Oxford: Oxford University Press.

Wimsatt, W. C. [1974]: 'Reductive Explanation: A Functional Account', *Philosophy of Science*, **1974**, pp. 671–710.

Woodward, J. [2010]: 'Causation in biology: Stability, specificity, and the choice of levels of explanation', *Biology and Philosophy*, **25**, pp. 287–318.

Woodward, J. [2013]: 'Mechanistic explanation: Its scope and limits', *Proceedings of the Aristotelian Society*, **87**, pp. 39–65.