Causal concepts in biology: How pathways differ from mechanisms and why it matters

Lauren N. Ross

In the last two decades few topics in philosophy of science have received as much attention as mechanistic explanation. A significant motivation for these accounts is that scientists frequently use the term "mechanism" in their explanations of biological phenomena. While scientists appeal to a variety of causal concepts in their explanations, many philosophers argue or assume that all of these concepts are well understood with the single notion of mechanism (Robins and Craver 2009; Craver 2007). This reveals a significant problem with mainstream mechanistic accounts– although philosophers use the term "mechanism" interchangeably with other causal concepts, this is not something that scientists always do. This paper analyses two causal concepts in biology–the notions of "mechanism" and "pathway"–and how they figure in biological explanation. I argue that these concepts have unique features, that they are associated with distinct strategies of causal investigation, and that they figure in importantly different types of explanation.

1 Introduction. For nearly two decades few topics in philosophy of science have received as much attention as mechanistic explanation. The beginnings of this "new mechanist" philosophy are often associated with a paper by Machamer et al. (2000), which outlines the general view and remains one of the most cited publications in *Philosophy of Science*. While various accounts of mechanistic explanation exist,[1] many of them describe mechanisms as organized sets of entities and activities that underlie and produce some phenomenon of interest. This explanatory pattern involves explaining some outcome by appealing to the mechanism that produces it. While this basic picture is thought to well-represent explanation in many domains, it has been most extensively examined and applied to the biological sciences where it has led to a research program that has quite literally "exploded" (Bechtel and Richardson 2010, xlviii). Not only does this explanatory framework receive significant attention in the philosophical literature, but many view it as fundamental to any understanding of explanation in biology and as the "dominant view of explanation in the philosophy of science at present" (Kaplan and Craver 2011, 606).

One motivation for these accounts is that scientists frequently use the term "mechanism" in their explanations of biological phenomena (Machamer, Darden, and Craver 2000, 2), (Bechtel and Richardson 2010, xvii) (Wimsatt 1976, 671). Biologists, of course, use a variety of causal concepts in their explanations, including concepts like pathways, cascades, triggers, and processes. Despite this variety, mainstream philosophical views interpret *all* of these concepts with the notion of mechanism. For example, Robins and Craver (2009) state that although scientists appeal to terms like "cascades, pathways, systems, and substrates...[w]e use the term mechanism for all of these" (Robins and Craver 2009, 42). Similarly, Craver claims that while scientists "say that they discover "systems and pathways in the flow of information, and molecular cascades, mediators, and modulators...[t]he term mechanism could do the same work" (Craver 2007, 3). These claims

[1]For some of these accounts, see (Glennan 1996; Craver 2007; Bechtel and Richardson 2010). For an excellent overview of different philosophical projects connected with the "mechanism" concept, see (Andersen 2014a; Andersen 2014b).

receive widespread support in the philosophical literature. This is evidenced by the fact that numerous philosophical projects analyze the appeal to various causal concepts (e.g. pathways, cascades, processes, etc.) as instances of mechanistic explanation.[2] Further evidence of this is seen in discussions of non-mechanistic explanation, where the goal is to find explanations that are non-causal as it is often assumed that all causal explanations (and those causal concepts figuring in them) are mechanistic (Kaplan and Craver 2011).

This mechanistic approach faces a significant problem. Although philosophers use the notion of a mechanism interchangeably with other causal concepts, this is not something that scientists always do. Consider the notion of a pathway, which commonly figures in biological explanation. Examples of this concept include gene expression pathways, cell signaling pathways, metabolic pathways, anatomical pathways, developmental pathways, and ecological pathways. Scientists often distinguish pathways from mechanisms. They claim that a single pathway can be instantiated by different mechanisms, that distinct pathways can have similar mechanisms, and that pathways can be discovered without any knowledge of the mechanisms that underlie them. When introducing scientific material they often claim to discuss both the mechanisms and pathways relevant to some topic and they appear to consistently use each concept in some situations, while not in others. These points raise a number of puzzles for the dominant mechanistic program. If all or most of the causal concepts in biology are well interpreted with the notion of mechanism, why do scientists often distinguish mechanisms from these other concepts? Why do they use a variety of causal terms if the single notion of mechanism would suffice? Finally, what explains their seemingly consistent use of particular causal concepts in some situations, while not in others? These puzzles suggest that it is worth exploring how scientists distinguish various causal concepts in biology and how these distinctions matter for understanding causal explanation in this domain.

This paper examines how the mechanism and pathway concepts are used in the biology and how they figure in biological explanation. I argue that these concepts: (a) have unique features, (b) that they are associated with distinct strategies of causal investigation, and (c) that they figure in importantly different types of explanation.[3] One important caveat of this project is that I argue for these claims in the context of the biological sciences, which I take to include neuroscience and medicine. I do not extend these claims to other fields such as the causal modeling literature (in general), economics, or political science, where the mechanism concept is often used in a more flexible way.[4] Furthermore, even in the context of the biological sciences I am not going to claim that scientists always, in every case, use these concepts in the ways I indicate. Surely at least one scientist uses one of these concepts in a different manner. Surely at least one counterexample exists. What I am claiming is that scientists *often* use these concepts in the ways I discuss and that it makes sense that they do. In relation to this, a main theme throughout my analysis is

[2]For example, the philosophical account of mechanism is used to analyze biological phenomena including metabolic pathways (Bechtel 2011; Bogen and Machamer 2010; Bechtel and Levy 2013; Bechtel and Abrahamsen 2005; Bechtel and Richardson 2010; Thagard 2003), developmental pathways (Fagan 2013; Tabery 2014), gene-expression pathways (Bickle 2006), and signaling cascades (Brigandt 2013).

[3]My interest in this project has been importantly influenced by conversations with Ken Schaffner and his discussions of the pathway concept, which he suggests is unlikely to be well-accommodated by the new mechanist paradigm (Schaffner 2016). My analysis explores this suggestion by providing a novel characterization of the "pathway" concept, how it differs from the notion of "mechanism," and how these differences matter for causal investigation and explanation in biology.

[4]I should also exclude from my focus highly mathematical areas of biology where the "pathway" concept is not always causal, but sometimes captures correlations (as is the case for particular network models).

that when scientists refer to biological systems as "mechanisms" and "pathways" they often do so because these systems share features with structures in ordinary life that we associate with these concepts. In this sense, scientists are analogizing these systems to structures in ordinary life that we are familiar with. This is a strategy for making complicated features of complex biological systems more cognitively accessible. While the presence of such analogies in scientific discourse and reasoning is not a new observation, its relevance to the mechanism concept (and other causal concepts in biology) has been surprisingly unexplored.[5] I explore this suggestion and argue for the aforementioned claims (a, b, c) in the rest of this paper, which is structured as follows. In section 2, I briefly discuss the mechanism concept, its features, the strategies of causal investigation that it is associated with, and how it figures in biological explanation. Section three introduces the pathway concept, its main features, and a common causal investigative strategy that it is associated with. This section compares the pathway concept to the mechanism concept, in order to clarify how they differ. In section 4, I explore (three examples of) one type of pathway explanation, where pathway information is explanatory and mechanistic information is not. The final section returns to the topic of analogy in scientific reasoning and contains some concluding remarks.

2 Mechanisms: The basics. Biologists frequently appeal to mechanisms in their explanations and descriptions of biological phenomena. They discuss the mechanisms of gene regulation, DNA synthesis, nerve firing, muscle contraction, visual processing, and so on. When they use the mechanism concept they often suggest that some biological phenomenon can be understood as a kind of machine or mechanical system–such as a car engine or clock–in the sense of having particular features. This machine analogy encourages thinking of biological phenomena as having component parts, which are spatially organized, and that causally interact to produce some behavior of the system. A key feature of this explanatory pattern (which I discuss in more detail later) is that it involves explaining some outcome by appealing to its causal parts. The system-level behavior serves as the effect or explanatory target (i.e. explanandum), while the interacting mechanical parts are what explain this behavior (i.e. explanans).

Three features of this mechanism concept should be highlighted. First, mechanisms are often characterized as having a constitutive makeup, in the sense of involving particular systems with higher level behaviors that can be decomposed into lower-level causal parts. This feature is exploited in efforts to discover mechanisms through common investigative strategies of "decomposition and localization," which are considered the "central heuristics" of mechanism discovery (Wimsatt 1974; Bechtel and Richardson 2010; Bechtel and Levy 2013). This strategy involves a process where scientists identify a system or behavior of interest and then "drill down" to identify the system's parts, their location, and how they interact to produce the behavior in question. This can be understood as starting with an effect and searching "backward" or causally upstream to identify the causal mechanism that produces it. Scientists also approach discovery from the opposite direction. They can start their inquiry with the causal components of a mechanism and search "forward" or causally downstream to identify its effects. In other words, the former involves starting with an effect and looking backward to search for its causes, while the latter involves starting with causes and looking forward to identify their effects. In both cases a set of factors are circumscribed as the causal parts that makeup the mechanism where these factors are determined on the basis of their causal relevance to the system and, in particular, their causal relevance to some effect of

[5]For discussion of analogy and analogical reasoning in science, see: (Hempel 1965, 434), (Hesse 1966), (Lewis 1986, 220), (Nersessian 2002).

interest. This contributes to our conception of mechanisms as discrete singular causal entities in the same way that we talk about particular car engines or clock mechanisms as single, distinct causal systems (Bechtel and Richardson 2010, 35). These causal systems have boundaries and they can be discussed as individual units that need not be connected to other causal systems in the world (Andersen 2014a, 276).

A second feature of the mechanism concept is that it is used to refer to causal systems that are described in significant amounts of causal detail as opposed to abstracting from such information. Consider the "mechanism of enzyme catalysis" where an enzyme catalyses (or speeds up) the chemical conversion of an upstream substrate into a downstream product. Scientists refer to these enzymes as "molecular machines," because they perform these conversions in multi-subunit complexes, which are understood as having many causally interacting parts (Spirin 2002, 153). These parts and their interactions are represented in "reaction mechanism" diagrams, which include components such as the enzyme itself, its substrate, and various cofactors and regulators that alter its functionality. Scientists expect complete descriptions of these mechanisms to contain large amounts of causal information. Consider the following quote:

> "An understanding of the complete mechanism of action of a purified enzyme requires identification of all substrates, cofactors, products, and regulators. Moreover, it requires a knowledge of (1) the temporal sequence in which enzyme-bound reaction intermediates form, (2) the structure of each intermediate and each transition state, (3) the rates of interconversion between intermediates, (4) the structural relationship of the enzymes to each intermediate, and (5) the energy contributed by all reacting and interacting groups to intermediate complexes and transition states. *As yet, there is probably no enzyme for which we have an understanding that meets all these requirements*" (Lehninger and Cox 2008, 205, emphasis added).

As this quote suggests, scientists expect descriptions of these mechanisms to contain a large degree of causal information–so much information, in fact, that it has not been acquired it in our best scientific understanding of these systems. This same sentiment is present in scientists's discussion of the "mechanism of action" for particular drugs. They claim that these mechanisms must involve "a complete and detailed understanding of each and every step in the sequence of events that leads to..[an]...outcome," "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, 1,7) (Ankley 2010, 731). This mechanistic understanding is contrasted with other approaches, which only capture "selected key events" and have "gaps and black boxes in which mechanistic details are either unknown or not needed" (Hutchinson 2007, 1) (Ankley 2010, 732).[6] The expectation that mechanisms contain significant causal detail is expressed by many philosophical accounts of mechanism (Machamer, Darden, and Craver 2000; Darden 2006; Craver 2007; Craver and Darden 2013). This feature of mechanisms is associated with our interest in understanding how they work and our assumption that this often involves identifying more and more information about their causal components, organization, and so on. Furthermore, acquiring such information is useful for various reasons–it suggests different potential targets to use in changing a final outcome, it provides more information on how a mechanism might

[6](Ankley, Bennett, Erickson, Hoff, Hornung, Johnson, Mount, Nichols, Russom, Schmieder, Serrrano, Tietge, and Villeneuve 2010, 731)

break, and it can lead to the identification of causal relationships that are more invariant or stable across different contexts. Additionally, the characterization of mechanisms as single, individual causal systems with boundaries, makes the expectation of significant detail more natural, than it would be if they were understood as having various causal connections to other systems in the world.

A third feature of the mechanism concept is that it often involves an emphasis on the "force," "action," and "motion" involved in causal relationships. This emphasis is evident in how we discuss machines in ordinary life–machines have parts, such as pulleys, levers, hammers, and gears, which actively do things. We do not simply say that these parts "cause" various outcomes in each system, we say that they "push," "pull," "bend," and "compress" some downstream component. Mechanism descriptions in biology involve a similar emphasis–scientists say that a cofactor "activates" an enzyme, which then "binds" to a substrate, before "splicing" off a chemical moiety, and "attaching" it to another molecule. The fact that the mechanism concept has this feature is should be somewhat unsurprising, because the term "mechanism" literally draws on mechanics or the branch of science and mathematics concerned with "motion and the forces producing motion" (OUP 2012, 449). What is the significance of this feature? Emphasizing the force or action of causal relationships serves a few functions in biological (and other) contexts. First, it helps to satisfy our interest in understanding "how" a mechanism works–adding force or motion terms adds something more than just saying that X causes Y. Second, these terms also function to fill in space between cause and effect variables, which can suggest closer physical proximity and satisfy our interest in getting more detail about the mechanism of interest. Causal terms involving force and motion appear to fill in black boxes and suggest that we know more about some causal process than merely saying "that" X causes Y.

Thus, in the biological sciences "mechanism" is often used to refer to causal systems that have a constitutive character, that are represented in significant, fine-grained detail, and that contain an emphasis on the "force," "action", or "motion" of causal relations. This concept is associated with the causal investigative strategies of decomposition and localization and it is involved in an explanatory pattern where some outcome is explained by appealing to the causal mechanism that produces it.

3 The pathway concept. The pathway concept is commonly found in the biological sciences. Biologists refer to gene expression pathways, metabolic pathways, developmental pathways, anatomical pathways, and ecological pathways, just to name a few. In all of these cases the notion of a pathway refers to a sequence of causal steps that string together an upstream cause to a set of causal intermediates to some downstream outcome. For example, gene expression pathways track causal connections from genes, to their downstream products, to a final phenotype (Figure 1). Metabolic pathways capture a sequence of steps in the conversion of some initial metabolic substrate into a final downstream product (Figure 2). Developmental pathways depict a step-wise set of changes in the development of some early precursor system (e.g. cell, tissue, or organism) into a later final state (Figure 3). Anatomic pathways capture physical routes, which outline causal paths such as lymphatic pathways, blood vessels, and nerve tracts (Figure 4). Finally, ecological pathways track causal links of predator-prey relationships where these are represented in food chains that make up larger food webs (Figure 5 (Smith and Smith 2012, 325)).

When biologists use the pathway concept they often imply that some system can be understood in terms of causal routes or roadways, which capture interconnected paths that track the flow or

movement of some entity through a system. In these cases they analogize a biological system to our ordinary life conception of roadways, highways, and city streets. However, what exactly are the features of this pathway concept? How is it used in causal investigation and explanation in biology? Finally, how does it differ from the mechanism concept, if it does at all?

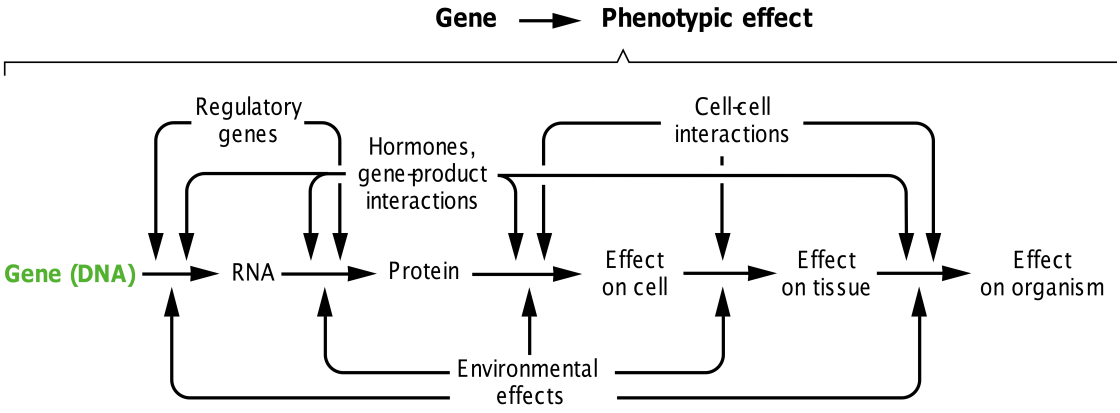

Figure 1: Cell-signaling pathway: An abstract model.

3.1 Main features. The pathway concept, as it is commonly used in biology, has at least four main features: it captures a (i) sequence of steps, where these steps (ii) track the flow of some entity through a system, (iii) abstract from significant causal detail, and (iv) emphasize the "connection" aspect of causal relationships. A (i) first feature of the pathway concept, which the above cases make clear, is that it captures a sequence of causal steps in some process. This sequence captures a fixed order of causal relationships that reflects which outcomes need to occur before and after others in the unfolding of a causal process. For example, consider the first three steps of the glycolytic pathway, which represents the biochemical conversion of glucose into pyruvate (fig 2). In these steps, glucose is first converted into glucose-6-phosphate (G6P), which is converted into fructose-6-phosphate (F6P), which is converted into fructose-1,6,-bisphosphate (F-1,6-BP). The glycolytic pathway captures the sense in which these steps need to take place in a particular order–glucose cannot be directly converted into F-1,6-BP without first forming the G6P and F6P intermediates, and it must go through both intermediates without omitting one or reversing their order. Furthermore, the pathway does not just capture a fixed order of entities, but a fixed order of *causally related* entities. The glycolytic pathway captures a *causal* chain in the sense that every downstream product depends on an upstream substrate. In other words, the upstream substrate is at least one causally relevant factor in the production of the most immediately downstream product. Other cases of the pathway concept in biology involve this same fixed sequence of causal relations, but differ in terms of the causal relata along the pathway. Biologists sometimes refer to these causal chains as this as "domino causality," because like a sequence of falling dominos, the effect at one step becomes a cause of the next (Grotzer and Basca 2003). While these cases depict pathways as linear sequences, they can also have branching and circular arrangements (and they can be nested in complex causal networks or interconnected sets of pathways).

It might seem that this first feature does not capture a true difference between the pathway and mechanism concepts, because both can be understood in terms of sequential causal steps. There

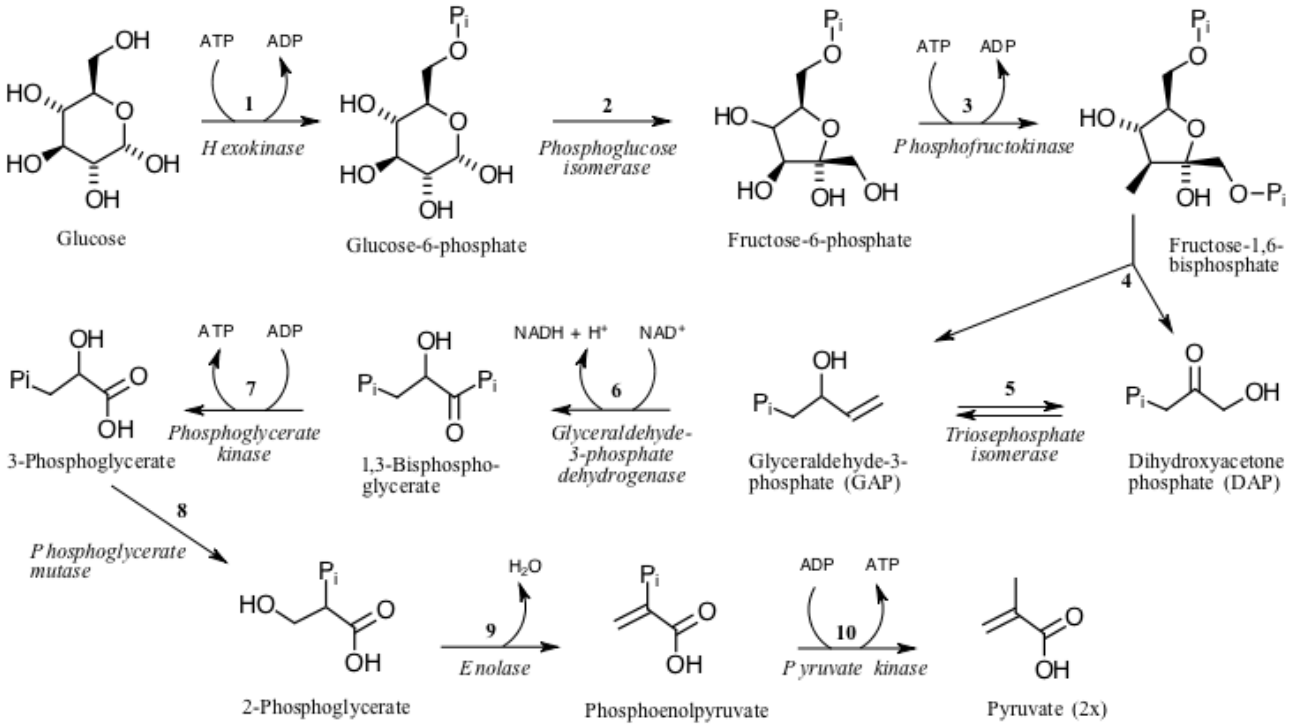

Figure 2: Glycolysis: A metabolic pathway.

is more to say about the causal sequences captured by pathways that reveals how this concept differs from the notion of mechanism. In particular, one difference, and a (ii) second key feature of the pathway concept, is that it represents causal sequences that capture the "flow" of some entity through a system. For example, metabolic pathways trace the flow of chemical substances through stepwise changes, stem cell pathways capture the flow of cells through developmental sequences, anatomical pathways such as blood vessels trace the flow of blood through the body, and ecological pathways trace the flow of energy through ecosystems. This notion of "flow" refers to something that is carried over from one causal step to the next–it involves the permanence or continuity of something that moves or travels along causal connections.[7] Use of the "pathway" concept in these cases is not just a happy coincidence–scientists use the pathway concept to refer to these biological systems, because they have features that are similar to pathways we discuss in everyday life. One of these shared features is the notion of "flow." This notion is common to ordinary life examples of pathways, such as how cars move along freeways and how water courses along pipes in a plumbing system. The fact that scientists explicitly point out this similarity is seen in the language they use to describe these system. They refer to the "flux" of chemicals along "metabolic roads and byways" (Lehninger and Cox 2008, 528), (Pardee 1994, 375), they claim

[7]To be clear, I am not suggesting that the pathway concept supports connected process or mark transmission accounts of causation. Instead, pathways represent causal relationships that have the additional feature of capturing the flow of some entity, while not all causal relationships have this feature (as is the case in absence causation examples). Additionally, while these cases involve the flow or movement of some material through a system, other pathway examples involve the flow of immaterial phenomena such as the flow of information in cell-signaling pathways.

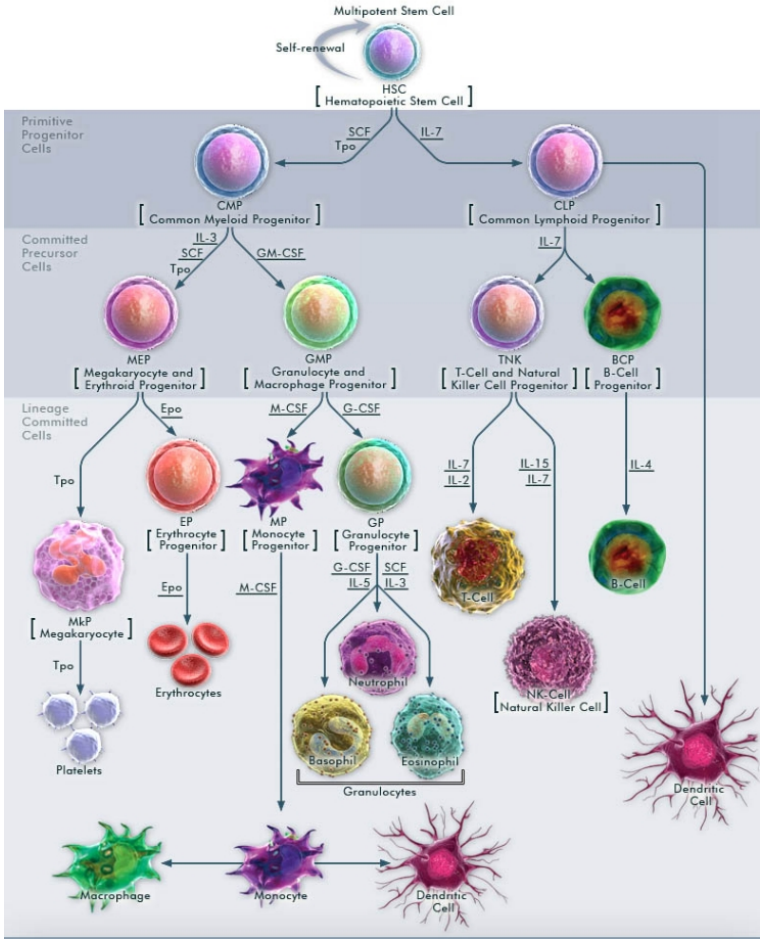

Figure 3: Stem cell pathways.

that lymphatic pathways are "avenues" and "routes" along which "lymph is transported" (Meyers, Charnsangavej, and Oliphant 2005, 4), (Richter and Feyerabend 2004, 6), that food chains are "energy channels that propagate matter and energy...linked by predators"' (Moore and de Ruiter 2012, 225) and literally that ecological pathways are "plumbing of sorts–through which matter and energy flow within ecosystems" (Caswell 2005, viii). The pathway concept and language of "flow" are not simply meaningless, colorful metaphors–they are pointing out objective, physical features of these systems that reveal how they operate in the world and how we can best study, discover, and understand them. For example, flow through these systems is often experimentally studied with tracer and tagging techniques that exploit the physical flow through them.[8] Furthermore, "blockages" in these pathways, as in the case of "inborn errors of metabolism," can be understood through analogical reasoning with ordinary life traffic-freeway examples.[9] Instead of being a trivial

[8]These techniques tag material, which is sent into the pathway, so that pathway steps, flow rates, and interconnections can be discovered.

[9]These blockages result in a pathologic build-up of material upstream of the blockage or accident. Successful therapeutic measures involve rerouting this material along "bypass" routes, shunting it to some

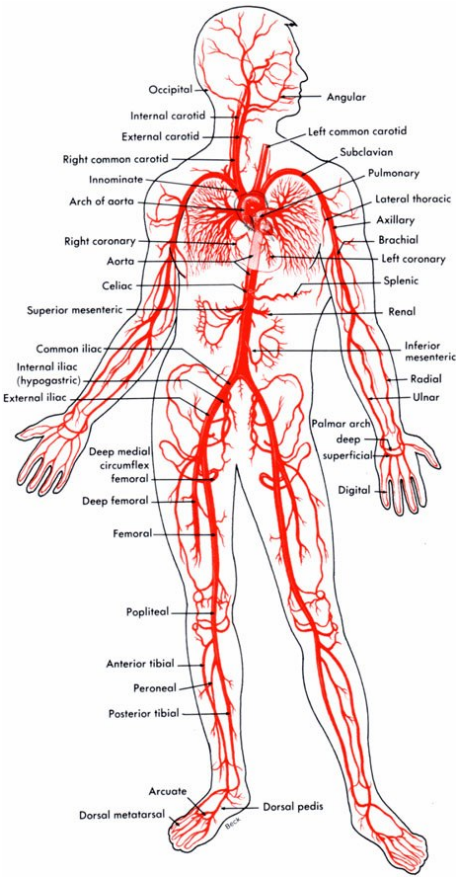

Figure 4: Anatomical (vascular) pathways: Blood vessels.

turn-of-phrase, the pathway concept is used to highlight features of these biological systems that matter for how they are scientifically studied and for how they are used to control outcomes in the world.

When mechanisms are discussed in biological contexts, there is not usually an emphasis on the "flow" of some entity through the mechanism. Mechanisms contain parts that interact to produce a final effect, but there is usually not an emphasis on the movement of something across or along these parts. Consider an objection to this claim. One might suggest that "causal influence" flows through mechanisms, where this refers to the propagation of causal force through a set of a intermediates. Of course, there is a sense in which all causal relationships involve the "flow" of causal influence, yet something more is present in these pathway cases that is not found in all causal relationships, viz. the movement of some *further* entity (e.g. metabolites, cells, blood, or energy rich substances), besides causal influence. This is also supported by the fact that, in biological contexts where "pathways" are identified, there are numerous causal relationships transmitting "causal influence," but only some of these relationships are highlighted and represented as pathways–viz. those that trace the movement of some particular entity of interest.

non-pathologic product, reducing influx into the blocked pathway, and supplying the needed material downstream of the blockade.

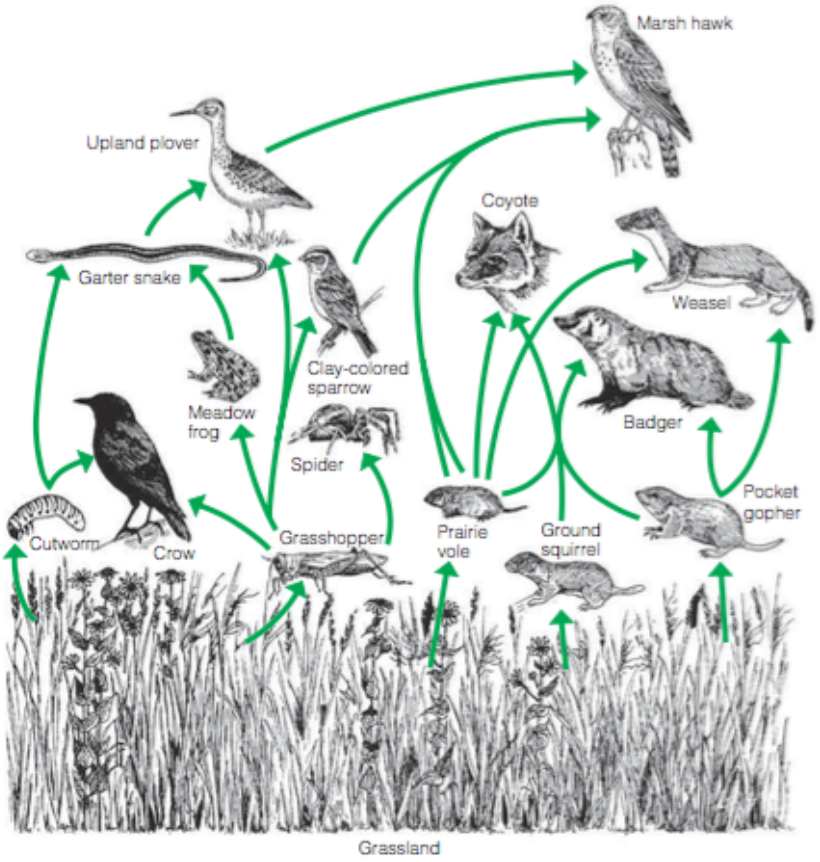

Figure 5: Ecological pathways: Food chains within a food web (Smith and Smith 2012, 325).

A third feature of the pathway concept–and one that clearly differs from the notion of a mechanism–is that they represent causal sequences that abstract from significant amounts of causal detail. One way that pathways abstract from detail is that they only represent causal factors that capture the flow of some entity through system and not the entirely of factors that support or are causally relevant to this flow. For example, metabolic pathways represent the flow of metabolites and not the many other factors that can control or regulate this flow, such as enzymes, cofactors, temperature, pH, etc. This is similar to how road maps represent freeways and city streets without also depicting traffic lights, police officers, or road blocks, which can regulate or alter the flow of traffic along these routes. These pathway diagrams abstract from this type of information. A second way that pathways abstract from detail is by representing complex processes with an economy of causal steps. This is easily seen in the case of developmental pathways, which capture the development of living cells, tissues, and organisms in a limited number of stages. For example, the entire life cycle of many organisms is represented in anywhere from 4 to 12 main steps (Mahadeo and Parent 2006, 116), which could each be further divided into numerous causal links. Biologists make explicit reference to this difference between the mechanism and pathway concepts. While they emphasize the need for detail in mechanism cases, they explicitly state that pathways are "not intended...to be exhaustive descriptions" (McClanahan and Branch 2008, 5). This is also seen in the fact that they claim that "complete pathways" have been identified, while "complete

mechanisms" are still beyond our reach. While scientists admit that no "complete" mechanism of enzyme catalysis has yet to be uncovered due to the immense detail that this requires, they claim that the "whole pathway" of glycolysis was discovered in the 1930s (Lehninger and Cox 2008, 528). This, and other discussions of the identification of "complete" pathways in ecology reveals the lower standard of causal detail that they are expected to meet (Wentsel, Beyer, Forbes, Maund, and Pastorok 2008, 217).

I am suggesting that one clear difference between these concepts is that pathways abstract from significant causal detail, while mechanisms include it. Consider a few objections to this position. First, one might claim that significant causal detail is not a necessary or characteristic feature of mechanisms and that some mechanisms do, in fact, abstract from such information. While some mechanistic philosophers subscribe to this "abstract mechanism view," most understand mechanisms as highly detailed.[10] Philosophers who promote this former view are likely to claim that pathways are easily accommodated by philosophical accounts of mechanism, because there is nothing problematic about the notion of an "abstract mechanism" and that this is exactly what pathways are. One main problem with this response is that biologists simply do not use the mechanism concept in this way–they consistently use "mechanism" to refer to causal structures that contain significant amounts of detail and they rarely apply it to structures with sparse causal information. These pathway cases and other examples of abstract causal relationships (such as monocausal models of disease) are rarely if every referred to as mechanisms and they are consistently viewed by scientists as devoid of mechanistic information. Furthermore, frequent use of the "mechanism" concept in biology was one of the original motivations for constructing a mechanistic account of explanation–to misinterpret the use of this term conflicts with one of the very motivations for constructing such an account of explanation and it detracts from the potential strength it might have. If we want a philosophical account of mechanism that is representative of how biologists use this concept, it should accommodate the fact that they often use it in cases where they expect and want significant, fine-grained causal detail. If not, we should admit that our philosophical analysis is not representative of this scientific concept. Of course, we *can* formulate a philosophical account of mechanism that is divorced from the scientific use of the this term, but, if we do this, we should be upfront about it and refrain from claiming that biological terminology and reasoning support this view. I do not think that we should do this, because it would obscure the reasons scientists have for using these concepts and it would prevent us from accommodating the role of analogy in these cases.

A second objection to my analysis retains the view that mechanisms involve significant detail and simply claims that pathways are early-stage mechanism sketches or schemata, which have yet to be filled in with sufficient detail (Craver 2007, 113-114) This is a standard interpretation of these pathway examples in the literature (Craver and Darden 2013, 91). As these interpretations suggest that increases in causal detail track increases in explanatory power, it is claimed that the lack of detail in pathways results in the explanatory deficiency of this causal concept. In fact, these mechanistic accounts associate the pathway concept with the "vice of chainology" where "[o]ne becomes fascinated by nodes in a causal chain but loses sight of how the nodes work to produce, underlie, or maintain the phenomenon" (Craver and Darden 2013, 91). These pathways are causal structures that are "incomplete" and reflect a "shallowness" of understanding (Craver and Darden 2013, 91-92). This interpretation represents a misconception of the pathway concept in biology

[10]Proponents of the former position include (Bechtel and Levy 2013) and the latter (Machamer, Darden, and Craver 2000; Darden 2006; Craver 2007; Kaplan and Craver 2011).

and its role in explanation. This view seems to suggest that scientists aim to fill pathways in with detail, but this conflicts with their explicit statement to the contrary–that such pathways are "not intended...to be exhaustive descriptions" (McClanahan and Branch 2008, 5). Scientists consider pathways to be "whole" and "complete" when they contain far less detail than "complete mechanisms" (Lehninger and Cox 2008, 528) (Wentsel, Beyer, Forbes, Maund, and Pastorok 2008, 217). These points suggest that pathways are not properly viewed as precursor, incomplete mechanism sketches, but as a "complete," albeit distinct, causal concept in their own right. Furthermore, pathways are cited in explanations, without being viewed as explanatorily deficient–I explore this further in section 4, where I consider the role of pathways in biological explanation.[11]

A fourth and final feature of the pathway concept is that it emphasizes the "connection" involved in causal relationships as opposed to the "force," "action," and "motion," which is emphasized in the causal relationships in mechanisms. Where mechanisms involve specifying "how" X causes Y, pathways involve simply capturing "that" X causes Y–in particular, the goal with the pathway concept is to show, of some set of entities in a system, what is causally connected to what and not the fine-grained details of "how" they are connected. This feature is alluded to when biologists refer to ecological food webs as "connectance webs" and "wiring diagrams" that involve "showing which species are connected to which" (Caswell 2005, vii,vii) This connection feature plays a significant role in clarifying the causal investigative strategies that this concept is associated with and the role in plays in explanation, which are topics I turn to now.

3.2 Investigative strategy. Recall that the mechanism concept is associated with the causal investigative strategies of decomposition and localization. These strategies involve "drilling down" or decomposing a system into its lower-level parts. Before this strategy can be implemented it requires a first step that mechanists refer to as identifying a "locus of control," which involves specifying some (a) system and (b) effect of interest (Bechtel and Richardson 2010, 35). All causal components of the mechanism are identified on the basis of these specifications–they are included or omitted from the mechanism on the basis of whether they are found in (a) and whether they causally contribute to (b). In this sense, mechanisms are circumscribed on the basis of which parts causally interact to produce to the effect of interest. This effect-relative approach leads to the identification of a single, discrete causal system–a set of causal parts that are responsible for the effect and that are represented as distinct from other causal systems in the world.

The pathway concept is often associated with a different causal investigative strategy. A first step in this pathway approach involves first identifying causal connections, or "pathways," across entities in some domain without specifying either an effect of interest or a causal starting point. In this approach there is an interest in creating a map of available causal connections in some context–a kind of roadmap, or what biologists might call a "network" or "landscape" of available causal routes. Unlike the mechanism strategy these maps do not represent a particular, discrete set of causal parts that all interact to produce a specific outcome. Instead they represent available (or potential) causal

[11] A third objection might claim that the "new mechanists" do view and account for pathways as a type of mechanism: they classified them as "etiological" mechanisms in contrast to "constitutive" mechanisms. Although this distinction is mentioned by some "new mechanists" the bulk of the literature on mechanistic explanation has focused on the constitutive type. One reason for this is that discussions of etiological processes in biology are not viewed by philosophers as paradigmatic mechanisms (and they are certain not often referred to as mechanisms by biologists, but more often as "causal pathways," as in the "causal pathways" of disease etiology.

connections that are relevant to a variety of explanatory outcomes and causal starting points. These are channels that can be navigated to get from any one point in the map to any other. So instead of identifying a particular explanatory target and "drilling down," these maps involve identifying a set of entities in some domain and "expanding out" by tracing their causal connections. These connections are "available" (or potential) in a way that differs from the actual causal components in mechanisms–they contain information about various causal possibilities as opposed to a single, circumscribed causal process that leads to an effect. In other words, these causal maps are not similar to car engines or watch mechanisms in the sense that they depict parts, which are all relative to a single main behavior of some system. They are more like a set of available freeways that some car or some entity of interest can travel along. Examples of these "pathway maps" are metabolic pathways found in pathway databases[12], stem cell pathway diagrams, anatomical illustrations of vasculature, lymphatic vessels, and neural tracts, and ecological food webs (as shown in Figures 2-5). In all of these cases scientists are concerned with representing widespread causal connections in some system without being tied to a single explanatory target. As these maps are intended to reveal these widespread causal connections they are often referred to as "connectance" diagrams, "wiring" maps, and connections that represent "global anatomic continuity" (Caswell 2005, vii) (Meyers, Charnsangavej, and Oliphant 2005, viii). On first impression these diagrams clearly seem to represent some type(s) of causal structure in the world. However, it might seem that this structure does not clearly meet the mechanistic criteria–at the very least these diagrams do not clearly identify detailed interacting parts of a system that produce single, particular behaviors. Alternatively, they do seem to depict sets of interconnected routes. Once some set of interconnected routes is captured, various explanatory questions can be asked and answered, which I turn to now.

4 Explanation: Pathways and mechanisms. Scientific explanations are often viewed as answers to particular types of why-questions, in particular, "explanation-seeking" why-questions (Hempel 1965). Within this framework, a why-question and its answer represent an explanandum and its explanans, respectively. I am going to suggest that there are some explananda (phenomena) for which pathway information is explanatory and mechanistic information is not. This identifies one type of "pathway explanation" that differs from mechanistic explanation.

4.1 Pathway explanation: Three examples. I will illustrate one type of pathway explanation with the pathway diagram shown in Figure 6. In this diagram, the letters and nodes represent entities in some domain, while the arrows represent the causal relationships they participate in. As I will soon clarify with some examples, this diagram contains pathway information in the sense of containing information about causal relationships in some area, where these relationships have pathway features (i)-(iv) discussed in section 3.1. In particular, the letters and nodes in this diagram represent entities, while the arrows reflect causal connections among them. Before I discuss this diagram further, it may already appear as though it contains information that answers why-questions that mechanistic information cannot answer. These include questions such as: How many different downstream products can substrate A produce? How many different upstream substrates can lead to the production of D? How many different ways are there to get from B to C? If the uppermost route from A to C in the diagram is blocked, what downstream products will excess sub-

[12]Examples of these maps include: Reactome, Kyoto Encyclopedia of Genes and Genomes (KEGG), WikiPathways, Nature Pathway Interaction Database (PID) and Pathway Commons.

strate A produce? Providing answers to these questions is very natural within the interconnected "roadmap" representation of "available" pathways. However, these answers require a representation of widespread, available causal interconnections and a kind of flexibility of causal starting-point or effect end-point, which conflicts with the mechanism concept. Mechanisms specify "actual" interacting components that make up a single, individual causal systems, which are relative to a fixed explanatory target (or causal starting point). These features of mechanisms prevent them from from capturing the widespread causal connections, which provide answers to the aforementioned questions. Of course, one might claim that these why-questions may not be of the "explanation-seeking" variety that are relevant to scientific explanation. In order to address this response, and further illustrate the nature of one type of pathway explanation, I use diagram 6 to consider three cases–two from science and one from ordinary life.

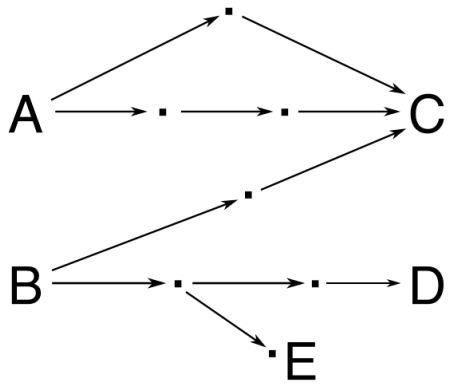

Figure 6: Pathway map: Ecological, anatomical, or "city-street" pathways

In a first example, the letters and nodes in diagram 6 represent species in an ecosystem, while the causal connections (i.e. arrows) between them represent prey-predator relationships. In particular, consider an actual scientific case where A is a bivalve species (clam), B is a species of crustacean, and C, D, and E are different species of fish (Stewart, Luoma, Schlekat, Doblin, and Hieb 2004). The diagram represents relationships between these species where downstream predators consume upstream prey and energy flows downstream, in the direction of the arrows. These species are located in the San Francisco Bay, which contains selenium–an element that is toxic to these organisms in high levels. It has been identified that fish species C contains high levels of selenium, while fish species D and E do not. Scientists want to know why this is the case–they want to know what explains these differences.[13] They explain this by citing the fact that fish species C is *causally connected* to species A, which contains high levels of this toxin. Alternatively, species D and E do not have high levels of selenium, because they are causally connected to primary producers with low levels of this compound, viz. species B (crustaceans), and not to species A (clams). In this manner, "exposures of top predators can be *explained* by food web relationships" (Stewart, Luoma, Schlekat, Doblin, and Hieb 2004, 4519). More specifically, they claim that the differences specified

[13]In particular, they explicitly ask: "Why did concentrations of Se differ so widely among predators in the Bay, and do those differences still occur? Does food web biomagnification of Se occur, and if so, why is it reflected differently in different predator species?" (Stewart, Luoma, Schlekat, Doblin, and Hieb 2004, 4519)

by this why-question are "explained by food-related variables" (such as those shown in Figure 6), background knowledge about how this toxin bioaccumulates along prey-predator connections, and the fact that of those species under consideration "predators feed differently" (Stewart, Luoma, Schlekat, Doblin, and Hieb 2004).[14] The differences in how predators feed is captured in the pathway information displayed in Figure 6. This pathway information is explanatorily relevant to the explanandum in this case, because it is information that "makes a difference" to it. If fish species C were no longer causally connected to A (or any primary producer with high selenium) it would no longer have high levels of this compound, and if D and E were connected to A, they would. These differences are not captured at the level of mechanistic information–the same pathways can be instantiated by similar or different mechanistic details and differences in these details do not track with changes in the explanatory target. In other words, so long as there are causal connections in the ecosystem represented by Figure 6, it does not matter "how" energy and selenium move along these pathways, how they are metabolized by organisms, or further details about how predators consume their prey–it just matters "that" these materials move through the ecosystem in the particular way captured by the causal pathways in this diagram. In other words, varying these lower-level mechanisms is not explanatorily relevant to the explanandum because changes in these details do not "make a difference to" or change the outcome, while varying the pathways in the diagram, would. I will explore these points further by analyzing two similar examples.

In a second case, the variables and nodes in Figure 6 represent spatial locations along lymphatic pathways in a human patient where these pathways trace the physical location and movement of lymphatic fluid through these vessels. It is discovered that cancer is present in location C, but not in the nearby locations D and E, and there is an interest in knowing why this is the case. Similar to the previous example, this is explained by the fact that A is the primary site of cancer and that this upstream site is causally connected to C, but not to D and E. The cancerous cells move along lymphatic vessels from A to C, seeding the growth of new cancerous off-shoots along the way. There is no cancer in D and E simply because there is no causal route that connects A to these sites. This method of explaining disease spread by appealing to anatomical pathways is found in other medical contexts with minor variations–both cancer and infectious material can spread along anatomical pathways such as lymphatic vessels, blood vessels, nerve tracts, and physically connected tissues (such as interconnecting spaces through the abdomen and thorax) (Meyers, Charnsangavej, and Oliphant 2005, 24). Researchers refer to these anatomical pathways as a "scaffold" and "interconnecting space" that disease processes use to navigate through the body. Tracking or following these pathways helps to identify (1) downstream locations of pathogenesis given an original upstream site, (2) upstream locations of original disease if only downstream pathogenesis is known, and (3) expected locations of recurrent disease (Meyers, Charnsangavej, and Oliphant 2005, 55). Answering these questions, and explaining the downstream location of disease requires a widespread view of the interconnections in this space, which reveals "potential" routes that a disease entity might travel.[15] Consider how scientists discuss the pathway concept in the context of this "cognitive framework":

[14]Tracer experiments suggest that these organisms acquire selenium through their diet and that it bioaccumulates (or builds up) along prey-predator connections–this compound is "propagated up the respective food webs" as predators acquire it from the prey they consume (Stewart, Luoma, Schlekat, Doblin, and Hieb 2004, 4519).

[15]Consider another example: Why does a clot in the carotid artery cause an infarct in brain vasculature (stroke) as opposed to a heart attack or pulmonary embolism? The reason is because the carotid directly irrigates the vasculature of the brain and not the heart and lungs. The clot gets logged in the small vessels

> "Understanding the pathways of extension of intraabdominal disease requires conceptualization of the interrelationship of this network..as one interconnecting space. This continuity provides avenues for the direct spread of disease...A disease process, regardless of its site of origin or cause (tumor, inflammation, etc.), upon gaining access to this interconnected space, is provided an anatomic avenue for direct spread...This unifying concept of direct spread underlies an understanding of the clinical appearance of abdominal disease at a distance from its site of origin solely by direct spread. Knowledge of the possible pathways of spread provides a rational system for a clearer understanding of disease process..." (Meyers, Oliphant, Berne, and Feldberg 1987, 601).

Here the emphasis is on depicting "one interconnecting space" that reveals various potential pathways for disease spread. In particular, the different ways that some locations are connected relative to others, figures in explanations of the variable spread and location of disease. This should clearly seem different from the discrete, isolable nature of mechanisms that capture individual causal structures, as opposed to a space of interconnected and potential causal routes. It is differences in the way that C is connected up in this space, relative to D and E, that explain why disease is present in the former but not the latter locations. In particular, the fact that C is causally connected to the primary site of disease, via identifiable anatomical pathways, and that D and E are not is what explains disease occurrence in the former, but not latter locations.

Third, consider an ordinary life case that captures a similar explanation to the those present in the above two examples. In this situation, the variables and nodes in Figure 6 represent cities and the arrows represent roads that connect them up. Cities A and B have corn factories and this corn is delivered to cities C, D, and E with vehicles that travel along the routes displayed in this diagram. It is discovered that city C has a supply of corn that is contaminated, while cities D and E do not. We want to explain why this is the case–why there is this difference in contaminated and uncontaminated corn across cities. This explanation is provided by the fact that a toxin has infiltrated the corn supply at factory A and that roadways connect the delivery of this supply to city C, but not cities D or E. An explanation of this difference is not provided by mechanistic, or lower-level causal information about this situation. In asking why city C has contaminated corn and why cities D and E do not, notice how unsatisfying it is for someone to tell you the intricate, fine-grained details of how the delivery vehicles were loaded, how they were off-loaded, and even how their engines work. These details do not matter for the difference in question–it does not matter whether this corn was delivered by vehicles that are electric, gas, diesel, front-wheel drive, rear-wheel drive, all-wheel drive, or the particular manner in which they were loaded and off-loaded with corn. The same roadways–with similar or different vehicles and loading and off-loading practices–would still give rise to the same difference in corn supply. Part of what this shows is that, changes in these mechanistic details do not "make a difference" to the explanandum of interest. What does make a difference is the higher-level structure of interconnecting roadways in this system and differences in how cities are connected up. If factory A was connected to C, D, and E, all of these cities would have contaminated corn. If it was not connected to these cities, none of them would have it. The pathway information that captures these higher-level causal features are what explain the differences in corn supply.

Although a more detailed treatment of these pathway explanations is best left for the subject of a separate paper, it will help to briefly mention a few features of these explanations. To be clear,

of the brain before it ever reaches the more downstream location of the heart and lungs.

these cases do *not* involve an interest in explaining "how" some entity travels along a pathway or "why" some pathways have the structure they have. Instead, there is an interest in explaining the difference in some feature across downstream locations. This difference is explained by the causal connections or pathways in some domain, because changes in these connections would create changes in this explanatory outcome.[16] The same cannot be said for the lower-level mechanistic information that instantiates these pathways. In fact, so long as the pathway relations remain fixed, lower level mechanistic information can vary without making a difference to the explanatory target. In each of these cases, information about the higher-level causal connections in some system explains particular outcomes that lower-level causal information cannot explain.

5 Conclusion: Mechanism and pathway as analogy. It is a fact about the world that it contains different causal systems with different features. I have suggested that these different features lead to distinct causal investigative strategies, explanatory why-questions, and possibilities for how such systems are represented, described, and discussed. Furthermore, it is a descriptive fact about biology that scientists in this domain use particular causal terms to refer to these distinct systems, as I have suggested they do with the "mechanism" and "pathway" concepts. This practice is coherent and it makes sense for various reasons–in particular, it makes sense in the context of analogizing these systems to structures in ordinary life that we are familiar with.

Consider that analogy is "a kind of similarity in which the same system of relations holds in two different examples" (Jee, Uttal, Gentner, Manduca, Shipley, Tikoff, Ormand, and Sageman 2010, 2). In this manner, analogy is often characterized as involving a mapping of structural features from a well-known base to a less well-known target (Gentner 1983, 157). This mapping can serve a number of purposes. A first main purpose is that it functions to transfer knowledge about some well-known domain to one that is unknown or poorly understood. In fact, analogy is considered one of the most effective strategies for "convey[ing] an entire system of relations in a new, unfamiliar example" and, because of this, it is often used and examined in educational settings or any situation where a new topic is described to a novice (Jee, Uttal, Gentner, Manduca, Shipley, Tikoff, Ormand, and Sageman 2010, 3) (Gentner and Smith 2012, 131). It is easy to see this being done in the mechanism and pathway cases–entry-level biology text books introduce these causal structures and their features with "machine" and "roadway" analogies, respectively. A second function of analogy is that it is used to highlight key features of a target system to more expert audiences who already have some familiarity with it. In this case, the analogy emphasizes features of the system that are relevant for the context of inquiry. This makes sense of the fact that we see explicit use of these analogies in high-level research publications, which focus on key features of some system for a particular purpose. For example, the "roadmap" analogy, and associated concepts of "routes," "transportation," and "flux," are used in metabolic research projects that focus on the rate of product formation and flow of metabolic material through a biochemical process. Alternatively, if there were an interest in capturing "how" an enzyme converted a particular substrate into some particular product, the "mechanism" concept would more likely be used. In this latter case, "flow" is not the primary feature of interest, but instead the local, interacting enzyme components that product a particular outcome–this leads to the characterization of enzymes as "molecular machines," with constitutive

[16]The explanatorily relevant information in these pathways cases is similar to the "connectionist" information discussed by (Bechtel and Levy 2013) and the topological information discussed by (Huneman 2018), although there are clear differences. Comparing these accounts to the present analysis is an interesting topic for future work.

interacting component parts. Finally, a third function of these similarities is that they figure in analogical reasoning and problem solving in new contexts. One example of this are cases of "inborn errors of metabolism," which are metabolic diseases that are understood with a "freeway" analogy that captures the pathological build-up of material (or "traffic" caused by a "roadblock") and those therapeutic measures that address it.[17] In this case, knowledge about solving problems in the context of freeway traffic is applicable to solving problems involving "traffic" on metabolic pathways. In all of these cases, the analogy is a sort of "psychological aid," which makes complex and (potentially) foreign features of biological systems more cognitively accessible and supports reasoning in a new domain by rendering the problem into a familiar context (Hesse 1966, 3).

It is not just that, as philosophers, we can choose to interpret these causal concepts as relying on analogy–scientists explicitly use these analogies in their work. These analogies pick out similar causal structures that arise in a variety of biological contexts. For example, I have demonstrated how the pathway concept is used to refer to a particular causal structure (with features i, ii, iii, and v) found in molecular biology, biochemistry, stem cell biology, developmental biology, biomedicine, and ecology. Scientists's use of similar analogies in different contexts is consistent with evidence that experts classify systems by shared causal structure even when they arise in different scientific situations (Rottman, Gentner, and Goldwater 2012). Being able to do this–and to have an "abstract understanding of causality" or ability to identify these "causal system categories," as empirical psychologists say–has advantages in the sense that these systems have similar implications for prediction, explanation, and control.

What consequences does this analysis have for philosophical accounts of biological explanation? My analysis indicates that biologists use a variety of causal terms to refer to unique causal structures in their field. Moreover, these unique structures motivate distinct causal investigative strategies, varying explanatory why-questions, and different explanatory patterns. We should want a philosophical account of explanation that accommodates this diversity–the diversity of causal structures in the world and our diverse techniques, methods, and strategies for managing them. We should expect a philosophical account of explanation to tell us why some causal details are explanatory and others are not, why some causal concepts are used in some situations, but not in others, and why scientists use a variety of causal concepts, as opposed to always using the notion of "mechanism." My analysis outlines an approach for doing this. As causal explanation is frequently understood as involving the explanation of some outcome by citing its causes, in some sense it should be unsurprising that different complex causal structures are likely to lead to different explanatory patterns. Standard accounts of mechanistic explanation can be compatible with this picture, but they should capture a biologically accurate conception of mechanism and leave room for the explanatory role of other causal concepts and structures. As many mechanistic philosophers suggest, it makes sense to understand mechanisms in biology as causal structures that involve constitutive relations, significant fine-grained detail, and causal-mechanical interactions expressed in terms of "force," "action," and "motion." This concept is well-representative of the majority of in-

[17] In these cases there is a blockage or "roadblock" along the metabolic pathways (induced by a gene mutation), which leads to upstream "traffic" and the "build-up" of some prior substance and the absence of the required downstream entity. This incessant build-up and the inability to produce downstream material can result in severe pathology, which is avoided by therapies that re-route this traffic around the blockage (via a bypass), shunt it into some non-pathologic downstream product, prevent flow into the blocked pathway, or find some alternative way to produce the final products downstream of the blockage. Examples of these diseases include phenylketouria, alcaptonuria, and glycogen storage disease.

stances where biologists refer to "mechanisms." However, if we accept this scientifically motivated picture of mechanism, we cannot also maintain–as some suggest–that "mechanism" also captures all other causal concepts in biology or that it is representative of all causal explanation in this domain, because clearly there are causal structures in this area with different features. This line would clearly undercut an honest depiction of the complexity of causal structures in biology and the nuanced and complex reasoning practices, investigative strategies, and explanatory patterns that they generate.

There may be a ready reply to these claims. Perhaps "mechanism" should be understood as a "catch-all" or "one-size-fits-all" concept, which is intended to distinguish any generic causal structure from those that are non-causal. Consider further, that there may be different "types" of mechanisms within this broad category–perhaps the "mechanism" examples I have discussed fall under some "mechanism$_a$" category, while the "pathway" cases fall under some "mechanism$_b$" category. This might be used to indicate how a mechanistic account could be flexible enough to capture all interesting causal structures and types of causal explanation in science. This approach faces serious problems. First, if "mechanism" is synonymous with any causal structure, how is mechanistic explanation different the generic claim that causes (or set of causes) explain their effects? Surely no account of causal explanation would deny this. Relatedly, if "mechanism" is short for "any causal structure" why not just say this? Mechanistic accounts have been motivated by the view that "mechanism" is some type of important or unique causal structure, but this approach distances the account from this motivation. Relatedly, these accounts are typically motivated by our view that mechanisms have unique features, which are not found in all causal structures. We see this in the fact that we refrain from calling single causes and causal chains "mechanisms." This leads to a second issue, which is that this "catch-all" approach simply does not accommodate how "mechanism" is used in biology. If the new mechanists want to base an account of biological explanation on a definition of "mechanism" that is divorced from biological use, they should have a good reason for doing so and they should stop referring to the biological use of this term to motivating interest in and suggest the credibility of their accounts. A third disadvantage of this approach, is that, by glossing over the use of the "mechanism" and "pathway" concepts in biology, we loose sight of a key strategy used in causal and explanatory reasoning–the use of analogy. By appreciating the role of this strategy biological reasoning we get a better understanding of the range and limits of mechanistic explanation and the character of other forms of causal explanation in this area–such as pathway explanation–that have gone unnoticed. In fact, appreciating the role of analogy in biology provides one fruitful way to understand the limits and range of a mechanistic explanation. Since the inception of these new mechanist accounts, there have been widespread worries about their over extension to fit various types of causal explanation (Woodward 2013; Dupré 2013; Skillings 2015; Halina 2018). These worries have been motivated by the view that, while mechanistic explanation is common and important, that is does not capture the whole of explanatory practice in the biological sciences. This paper gives one way to understand the nature and limits of mechanistic explanation, in a way that makes sense of biological reasoning, and accommodates the intended use and meaning of those causal concepts that commonly figure in this domain.

References

Andersen, H. (2014a). A field guide to mechanisms: Part I. *Philosophy Compass*.

Andersen, H. (2014b). A field guide to mechanisms: Part II. *Philosophy Compass*.

Ankley, G. T., R. S. Bennett, R. J. Erickson, D. J. Hoff, M. W. Hornung, R. D. Johnson, D. R. Mount, J. W. Nichols, C. L. Russom, P. K. Schmieder, J. A. Serrrano, J. E. Tietge, and D. L. Villeneuve (2010). Adverse outcome pathways: A conceptual framework to support ecotoxicology research and risk assessment. *Environmental Toxicology and Chemistry*.

Bechtel, W. (2011). Mechanism and Biological Explanation. *Philosophy of Science*.

Bechtel, W. and A. Abrahamsen (2005). Explanation: A mechanist alternative. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences 36*(2), 421–441.

Bechtel, W. and A. Levy (2013, October). Abstraction and the Organization of Mechanisms. Technical report.

Bechtel, W. and R. C. Richardson (2010). Discovering complexity. *The MIT Press*.

Bickle, J. (2006, August). Reducing mind to molecular pathways: explicating the reductionism implicit in current cellular and molecular neuroscience. *Synthese 151*(3), 411–434.

Bogen, J. and P. Machamer (2010). Mechanistic Information and Causal Continuity. In *Causality in the Sciences*. Oxford University Press.

Brigandt, I. (2013, December). Studies in History and Philosophy of Biological and Biomedical Sciences. *Studies in History and Philosophy of Biol & Biomed Sci 44*(PA), 477–492.

Caswell, H. (2005). *Advances in ecological research*. Elsevier.

Craver, C. and L. Darden (2013). *In search of mechanisms*. The University of Chicago Press.

Craver, C. F. (2007). Explaining the brain. pp. 1–329.

Darden, L. (2006). *Reasoning in biological discoveries*. Cambridge University Press.

Dupré, J. (2013). Living causes. *Proceedings of the Aristotelian Society Supplementary Volume LXXXVII*.

Fagan, M. B. (2013). *Philosophy of Stem Cell Biology*. Knowledge in Flesh and Blood. Palgrave Macmilan.

Gentner, D. (1983). Structure-mapping: A theoretical framework for analogy. *Cognitive Science*.

Gentner, D. and L. Smith (2012). Analogical Reasoning. In *Encyclopedia of Human Behavior*. Elsevier Inc.

Glennan, S. S. (1996). Mechanisms and the nature of causation. *Erkenntinis*.

Grotzer, T. A. and B. B. Basca (2003). How does grasping the underlying causal structures of ecosystems impact students' understanding? *Journal of Biological Education*.

Halina, M. (2018). Mechanistic explanation and its limits. In *Routledge Handbook of Mechanisms and Mechanical Philosophy*. Routledge Taylor & Francis Group.

Hempel, C. (1965). *Aspects of scientific explanation*. And Other Essays in The Philosophy of Science. The Free Press.

Hesse, M. B. (1966). *Models and analogies in science*. University of Notre Dame Press.

Huneman, P. (2018). Diversifying the picture of explanations in biological sciences: Ways of combining topology with mechanisms. *Synthese*.

Hutchinson, T. (2007). *Intelligent testing strategies in ecotoxicology: Mode of action approach for specifically acting chemicals*. European Centre for Ecotoxicology and Toxicology of Chemicals.

Jee, B. D., D. H. Uttal, D. Gentner, C. Manduca, T. F. Shipley, B. Tikoff, C. J. Ormand, and B. Sageman (2010). Commentary: Analogical thinking in geoscience education. *Journal of Geoscience Education*.

Kaplan, D. M. and C. F. Craver (2011). The explanatory force of dynamical and mathematical models in neuroscience: A mechanistic perspective. *Philosophy of Science 78*(4), 601–627.

Lehninger, N. and M. M. Cox (2008). *Principles of biochemistry*. Maxmillan.

Lewis, D. A. (1986). Causal explanation. In *Philosophical Papers Volume II*. Oxford University Press.

Machamer, P., L. Darden, and C. F. Craver (2000). Thinking about mechanisms. *Philosophy of Science*.

Mahadeo, D. C. and C. A. Parent (2006). Signal relay during the life cycle of Dictyostelium. *Current Topics in Developmental Biology*.

McClanahan, T. R. and G. M. Branch (2008). Food Webs and the Dynamics of Marine Reefs.

Meyers, M. A., C. Charnsangavej, and M. Oliphant (2005). *Meyer's dynamic radiology of the abdomen* (6 ed.). Springer.

Meyers, M. A., M. Oliphant, A. S. Berne, and M. A. M. Feldberg (1987). The periotoneal ligaments and mesenteries: Pathways of intradbdominal spread of disease. *Radiology*, 1–12.

Moore, J. C. and P. C. de Ruiter (2012). *Energetic food webs: An analysis of real and model ecosystems*. Oxford University Press.

Nersessian, N. (2002). Maxwell and "the method of physical analogy": Model-based reasoning, generic abstraction, and conceptual change. In *Reading Natural Philosophy*. Open Court.

OUP (2012). *Paperback Oxford English Dictionary*. Oxford University Press.

Pardee, A. B. (1994). Multiple molecular levels of cell cycle regulation. *Journal of Cellular Biochemistry*.

Richter, E. and T. Feyerabend (2004). *Normal lymph node topography: CT atlas*. Springer.

Robins, S. K. and C. F. Craver (2009). Biological clocks: Explaining with models of mechanisms. In *The Oxford Handbook of Philosophy and Neuroscience*. Oxford University Press.

Rottman, B. M., D. Gentner, and M. B. Goldwater (2012). Causal systems categories: Differences in novice and expert categorization of causal phenomena. *Cognitive Science*.

Schaffner, K. F. (2016). *Behaving: What's genetic, what's not, and why should we care?* Oxford University Press.

Skillings, D. J. (2015). Mechanistic explanation of biological processes. *Philosophy of Science*.

Smith, T. M. and R. L. Smith (2012). *Elements of ecology* (8 ed.). Pearson Education Inc.

Spirin, A. S. (2002). RNA polymerase as a molecular machine. *Molecular Biology*.

Stewart, A. R., S. N. Luoma, C. E. Schlekat, M. A. Doblin, and K. A. Hieb (2004). Food Web Pathway Determines How Selenium Affects Aquatic Ecosystems: A San Francisco Bay Case Study. *Environmental Science & Technology*.

Tabery, J. (2014). *Beyond Versus*. The struggle to understand the interaction of nature and nurture. MIT Press.

Thagard, P. (2003). Pathways to biomedical discovery. *Philosophy of Science*, 1–20.

Wentsel, R., N. Beyer, V. Forbes, S. Maund, and R. Pastorok (2008). A Framework for Population-Level Ecological Risk Assessment. In *Population-level ecological risk assessment*. CRC Press.

Wimsatt, W. C. (1974). Reductive Explanation: A Functional Account. *Philosophy of Science*, 1–41.

Wimsatt, W. C. (1976). Reductionism, levels of organization, and the mind-body problem. In *Brain and consciousness: Scientific and philosophical strategies*. Plenum Press.

Woodward, J. (2013). Mechanistic explanation: Its scope and limits. *Proceedings of the Aristotelian Society Supplementary Volume LXXXVII*.