



Seattle Pacific University  
**LIBRARY**  
*Discover, Create, Share*

Seattle Pacific University  
**Digital Commons @ SPU**

---

Honors Projects

University Scholars

---

January 1st, 2014

# Emergence, Scale, and the Layered Model of Biological Systems

Katie Morrow  
*Seattle Pacific University*

Follow this and additional works at: <http://digitalcommons.spu.edu/honorsprojects>

 Part of the [Biology Commons](#), and the [Philosophy of Science Commons](#)

---

## Recommended Citation

Morrow, Katie, "Emergence, Scale, and the Layered Model of Biological Systems" (2014). *Honors Projects*. 17.  
<http://digitalcommons.spu.edu/honorsprojects/17>

This Honors Project is brought to you for free and open access by the University Scholars at Digital Commons @ SPU. It has been accepted for inclusion in Honors Projects by an authorized administrator of Digital Commons @ SPU.

EMERGENCE, SCALE, AND THE LAYERED MODEL  
OF BIOLOGICAL SYSTEMS

by

KATIE MORROW

FACULTY ADVISOR, PATRICK MCDONALD  
SECOND READER, REBEKAH RICE

A project submitted in partial fulfillment  
of the requirements of the University Scholars Program

Seattle Pacific University

2014

Approved \_\_\_\_\_

Date \_\_\_\_\_

## ABSTRACT

The layered model of the world—the view that the physical universe is ontologically stratified into a hierarchy of levels—has commonly been applied in a variety of philosophical contexts, particularly in discussions of reductionism about causation, properties, or theories in science. In this paper I question whether this model, as traditionally understood, adequately reflects a contemporary scientific understanding of the world. Utilizing the layered model, philosophers have tended to focus on composition as the salient interlevel relationship, and to describe systems at temporal instants; while biologists stress the importance of the spatiotemporal scale of description and the environment in governing complex processes. These distinct frameworks result in widespread conceptual disconnects regarding issues like how questions about reduction are articulated, how we assess different scientific methods, and how we understand concepts like emergence and downward causation. I illustrate these differences with various examples from the philosophical and biological literature, including emergence as a case study. I conclude that the traditional layered model is inadequate for discussing the structure and behavior of complex biological systems.

## Acknowledgements

I am very grateful to my first reader, Dr. McDonald, for continuous support and discussion of ideas throughout the process of thinking about this paper; and for providing helpful feedback on several prior drafts. Thanks also to my second reader, Dr. Rice, for helpful comments on a previous draft. I gathered many of the scientific works, and most of the philosophical ideas I discuss, from the various classes in biology and philosophy I have had the opportunity to take at SPU.

## 1. INTRODUCTION

There are at least four different (though interrelated) questions about reductionism in the philosophy of science literature. (1) The first is methodological: is it predictively or explanatorily better for scientists to model system behavior in terms of interactions of parts, or in terms of system-level features? For example, there have been long-running disputes in numerous applications whether system-level models can be as epistemically secure as methodologically “reductive” models which describe system behavior in terms of the smallest feasible components. This question has direct implications for scientific methodology and is therefore of interest both to scientists who model complex systems (e.g. population ecologists, evolutionary biologists) and to philosophers interested in the epistemology of science. (2) Another issue is what relationship the theories and concepts of the special sciences have to fundamental physics (i.e., to base-level theory). That is, can theory from the domain of biology in some sense be derived from or closely mapped to theory utilizing only physical terms? If biological theory cannot be so reduced, we might further wonder what this means about biology: is it a mere approximation of physical theory, which cannot be applied in complex scenarios due to our epistemic limits; or does plurality of theory in the sciences reflect something about the physical world itself? (3) A third question is whether properties or states of higher-level entities (objects, processes, or events) can be reduced to, or perhaps identified with, those of their smaller constituent entities. This question is of interest in the metaphysics of mind, as well as more generally the metaphysics of complex, organized physical objects, like biological organisms. (4) A final issue is whether higher-level objects/properties have real causal powers, or whether they are merely epiphenomenal on fundamental physics. Many philosophers are interested in this question in application to human agency (i.e., the causal efficacy of our mental states), but we might also

worry about this question in application to biological systems more generally. One reason for concern about this latter issue is that humans *are* biological systems, so it might be a problem for agency if it turns out that nothing *qua* biological has causal powers. A second reason, which will be more important to philosophers of biology in particular, is that biologists freely make causal claims about the systems they investigate, a fact which would need to be explained away should philosophers claim that systems *qua* biological do not have causal efficacy.

It is clear that these questions about reduction have important consequences for the methodology of science, the epistemology of science, and the metaphysics of physical systems. I am interested in many of these consequences, but I will not attempt to directly resolve them in this paper.

My interest here is instead in the way in which these questions have been framed in the philosophical literature, particularly the literature pertaining to biological systems or theory. My reason for taking these issues from a step back is that I think the way we pose the questions can influence what positions come out looking most plausible. When a certain framework is used very commonly to articulate all sorts of important questions about reduction, it seems like a good idea to stop and ask whether that framework is a good one.

The framework I mean is this: all of the above questions about reduction have routinely been described as questions about *interlevel* relationships. I have followed this convention in stating the four questions in the first paragraph. To further illustrate, reductionists (about the various questions) typically make claims such as: *higher-level* models are predictively worse than *lower-level* models; *higher-level* theories, or properties, reduce to *lower-level* items; or the causal powers of *higher-level* objects or properties can be reduced to or identified with those of items on a *lower level*, perhaps the *base level* of physics. So, the notion of levels is used both to

frame these questions and to state the views that might be taken on them. It is this widespread assumption that objects, properties, and/or theories occupy a series of lower to higher levels which I will dispute in this paper.

Although the notion that theories, objects, and properties occur in ordered levels (or layers) is widespread, and although questions about reduction which employ the model have generated much discussion, the layered model itself is infrequently discussed. This may be because (first) it *appears* to be neutral with respect to viewpoints about reductionism. The mere statement that there are levels does not seem to entail that (an item occupying) a given level does or does not reduce in any manner. Second, talk about levels is common in science, so a layered model might be thought to be the general view utilized by empirical science. And, third, the general idea that there is stratification within science (think of the physics-chemistry-biology hierarchy) might seem to be so obviously correct that it is not itself an interesting object of discussion. So, whatever the reason, it has not frequently been discussed what assumptions the layered model imports into these conversations, and whether these assumptions are really in line with the explanatory structure of contemporary science.

I suspect that the first assumption—the neutrality of the model with respect to reduction—may not be entirely true, but I will not discuss that issue here. My target in this paper is the second item (and, by implication, the third). I will argue, contrary to what I take to be widespread assumption, that the layered model does not adequately capture the structure of physical systems, as seen especially from the perspective of biology. Utilizing a variety of case studies, including the highly contested issue of emergence, I will argue that the general framework biologists use to model complex system behavior is incompatible with features of the layered model philosophers have commonly employed in their discussions of reduction. My

project is closely related to that of Angela Potochnik and Brian McGill (2012), who have similarly argued against the empirical utility and accuracy of a strict hierarchy of levels framework, especially as applied to the explanatory context of ecology. Their paper provides a starting point for my project, although I will treat a slightly different set of issues and cases than they do.

Before I get started, I should explain why I think the conceptual apparatus of working biologists should have any bearing on the framing of very abstract philosophical issues. I have two things to say here. First, insofar as some philosophical positions make predictions or recommendations about scientific methodology, actual scientific practice can serve as a sort of test of the philosophical view. In section 3.2, for example, I argue that the layered model, at least as philosophers have commonly understood it, makes a prediction about the best way to scientifically predict the behavior of a complex system. Looking at the scientific work can fact-check this prediction, which I argue turns out to be false.

Of course, empirical science will underdetermine the possible positions we might take on many philosophical issues (e.g., regarding causal reduction). Still, where scientists and philosophers are both interested in the same objects—like the structure and behavior of complex biological systems—we might go for more than mere factual consistency between our philosophical and scientific views. We might further ask how well the philosophy coheres with the scientific conceptual apparatus. I think it constitutes reason to search for a better metaphysical framework if that framework fails to sit together well with the assumptions scientists have to make to conduct their empirical work—even if the metaphysical thesis is abstract enough to guarantee it will not directly contradict any factual claim made by a scientist. I will try to show in this paper, by considering a variety of examples, that the way philosophers

have often conceptualized systems *goes together very awkwardly* with the way biologists do so. I think this constitutes reason to consider revising how we approach some of the philosophical questions, especially as they pertain to biological systems.

Before I begin to discuss the layered model, I also want to state, for the sake of clarity, what I will not be arguing for in this paper. I have said I will not be resolving any questions about reductionism. Of the four questions about reduction I mentioned above, in the first paragraph, the first two are epistemological and the last two metaphysical. In fact, what I argue in this paper will suggest more of an antireductionism about the epistemological issues, or what otherwise might be characterized as a soft or methodological antireductionism in the context of science. On this view, scientific theory is expected to remain “patchy” rather than ultimately being reducible or collapsible down into a single basal physical theory. I will not directly argue for this thesis in this paper, but I believe it may be a consequence of my argument. Since I will reject the notion that explanatorily salient physical properties occur in neatly stratified layers, this looks like it might disrupt our ability to conduct neat downward “mapping” among scientific theories.

As far as the metaphysical questions about reductionism go, proposing any answers is beyond the scope of this paper. I will only discuss how these questions have been framed in philosophy, i.e., in terms of interlevel relationships. My argument about the levels framework might have implications for these issues, but for the purposes of this paper I remain officially neutral: my argument does not depend on any reductive or antireductive metaphysical views.

I want to make this clear at the beginning of the paper because below I will assume, based on the testimony of biologists, that there are emergent properties. Often, especially in the philosophy of mind literature, emergence is framed as a (very) strong metaphysical

antireductionist stance: the view that there are certain novel properties which arise in complex systems which are causally irreducible to properties of the sum of the system's microphysical components. However, I will argue that this strong emergentist view is distinct from, and in fact incompatible with, a "weak" scientific view of emergent features of systems.<sup>1</sup> So, it should be understood that if I refer to something as an emergent property, I am referring—as I will explain below—to an empirically discoverable feature of the system, and that calling it "emergent" does not commit me to any particular metaphysical views beyond what I explicitly argue for in this paper.

One might wonder, if I intend to discuss emergence only in a metaphysically inert sense, why it is interesting to discuss in a philosophical context. The reason is that I think how we understand system phenomena like emergence (and its partner, downward causation, the view that "higher-level" properties can causally influence events on a lower "level") depends on whether we are utilizing a layered model or an alternative, more biologically sophisticated scale-oriented model. Emergence, that is, happens to be convenient for illustrating how the two models differ. I do think that emergence in the weak sense is interesting for its own sake in some philosophical contexts, especially on the borders between the philosophy of science and theoretical science, e.g. when system modelling is under investigation. However, since this is not my interest in this paper, I will utilize emergence just in its capacity as a case study.

This paper will have two main parts. First I will characterize the layered model by looking at how it is being used within the philosophical literature, especially in the case of emergence. I will then provide a brief argument that the layered model cannot appeal for support

---

<sup>1</sup> "Weak" and "strong," as qualifiers of emergence and related concepts, are used inconsistently throughout the literature. In this paper I will use "strong" to distinguish views which explicitly contain a controversial metaphysical thesis from "weak" views whose immediate intent is only to describe empirical findings.

to scientific usage of the term “level.” In the third section, I will argue that the layered model is at odds with the explanatory structure of biology. To do this I will discuss a variety of examples from the biological literature, and argue that features of these cases conflict with the assumptions of the layered model.

## 2. THE LAYERED MODEL

The layered model portrays the physical universe as stratified such that entities can be ranked as higher or lower than each other, where lower-level items are those which are more basic, universal, or fundamental, and higher-level items are those described by the special sciences. Furthermore, physical objects are arranged such that smaller (lower-level) ones compose the next-larger kinds of objects, a level up. Higher-level objects are thus dependent both for their existence and for their features on lower-level composition. Moreover, this hierarchical arrangement is supposed to be relevant to the discussion of causal and other metaphysical relationships within and among physical objects.

This view of the structure of the physical world is, I think, supposed to reflect the organization and methodology of the sciences, which are thought to exemplify a similar hierarchical arrangement. Scientific theories applying to the most basic entities (i.e., the laws of fundamental physics) are considered most basic, and have variously been thought to ground, underlie, or potentially reduce a stratified series of higher-level theories pertaining to complex objects like molecules, cells, and brains. Although it is important not to conflate issues pertaining to theory reduction with issues pertaining to metaphysical reduction, the overall picture of a neatly stratified hierarchy has been shared by discussions of both issues. I think that the layered model of physical objects gains a certain amount of its credibility from the received

notion that the theoretical apparatus of science is layered. In this paper, I focus on reasons to think the world is not ontologically stratified, rather than directly discussing relationships among theories. Still, because of the tight connection between how we see the structure of science and how we see the structure of the physical, I think it is helpful to discuss the layered model in conversation with the findings and methodology of actual scientific practice.

As I said above, little recently has been written about the layered model, although it has been widely utilized by philosophers. An exception to this is a helpful paper by Jaegwon Kim entitled “The Layered Model: Metaphysical Considerations” (2002). We should take note of Kim’s warning that “it isn’t easy to come up with a neat and satisfying general model of levels that will serve useful philosophical purposes” (2002: 3). There have not been any detailed proposals since some highly idealized models, now clearly too neat to be true, were developed by philosophers during the early to middle part of the last century (including the British Emergentists, e.g. C. Lloyd Morgan (1923); also Oppenheim and Putnam (1958); see Kim (2002) for discussion of these models).

Because of the general unclarity surrounding levels, my aim here is not to present what I think is the best way to formulate a layered model, or to argue against any particular formulation. Rather, my goal is to see what background assumptions are in fact being made by philosophers utilizing this framework in recent applications.

In his (2002), Kim criticizes the mentioned historically important views and ultimately proposes we take a somewhat deflated, localized view of the world’s stratification. Historical views involved discrete, universal layers, such that every object (or property) in the physical universe could be ranked as higher or lower with respect to every other object (property). The most obvious issue for such a model of universal levels is, given that only select entities have

similar special-science descriptions, it may not be the case that all entities are directly comparable to each other with respect to their levels. For example, both a rabbit-sized rock and a rabbit are made of molecules. Does the rock occupy a higher level than its molecules? If so, does it occupy the same level as the rabbit's cells, or the rabbit? Or what about a computer, which is highly complex but organizationally dissimilar from organic life? These questions do not seem to have good answers.

Furthermore, entities we would intuitively put on the same level may have different constitutive levels. Take, for example, the organism. Some organisms are single cells, like any prokaryote, or a yeast or amoeba. Others are multicellular individuals like the rabbit. For some plant species, the organism is a distinct individual, like a single tree. However, other species are able to reproduce clonally via their roots, such that, for example, an entire stand of aspen can constitute a genetically identical and physically continuous entity. (Is that one organism or many?<sup>2</sup>) Another unusual organism, the plasmodial slime mold, is macroscopic and multinucleate but not composed of distinct cells—that is, it lacks internal cellular membranes. So, as we can see, not only is it sometimes unclear what counts as an organism, but different things that all clearly count as organisms—humans, bacteria, and slime molds—will have radically different decompositions. How could we decide whether, say, a bacterium is on the same level as my whole body or as one of my cells? I think, again, the answer is simply that

---

<sup>2</sup> Having raised the question, I will note that the answer I have heard given by biologists—who must sometimes give a stipulative answer to these sorts of questions in order to conduct their research—is that it depends on one's research question. For studying canopy cover in an ecological context, each trunk is considered a distinct individual. For studying population genetics, the whole stand is an individual. This sort of question-specific categorization of entities raises a somewhat distinct problem for a levels worldview, the simplest reason being that if there is no objective (context-independent) fact about what constitutes an individual tree, there is also no fact about what lower-level entities compose and/or determine the features of a tree. Alternatively, the lower-level base for the tree changes for each different phenomenon to be explained, which seems unparsimonious as a metaphysical framework. See Dupré (1993) for extended discussion of this issue and (a strong view of) its implications for scientific explanation and reduction.

many different kinds of entities, on account of different internal structure, are not directly comparable with respect to levels.

Having reflected on some problems similar to the ones given above as well as some in addition, Kim concludes that preserving a meaningful view of levels requires a localized and top-down approach: “we first pick a nomic kind of interest to us and go [i.e., look downward] from there, rather than start with a comprehensive levels ontology and then try to locate each object, or kind...” (2002: 20). That is, it looks hopeless to universally stratify the world and then see where each item fits into an overall hierarchy. Nevertheless, if we select a given entity, we can still see how that entity relies on the lower-level entities featured in its structure. In this way we can retain hierarchical stratification on a case-by-case basis.

A second issue for a traditional levels model has to do with the messiness of stratification within objects. It has been pointed out by very many authors (including Kim (2002) and Potochnik and McGill (2012); Beckner (1974) also gives a helpful treatment of hierarchies) that most macroscopic objects are not, in Beckner’s terminology, perfect hierarchies. A perfect hierarchy is one in which objects at level  $L$  are composed exhaustively of objects at level  $L-1$ , which fails to hold for biological objects. For example, a human body has no full decomposition into organ systems, but rather is composed of organ systems plus various fluids. A cell is similarly best seen as being composed of a mixture of objects we would intuitively place on different levels, from free ions to large organelles.

However, it is not clear that an idealized perfect hierarchy is a necessary feature of a levels model. It only seems to be necessary (assuming physicalism) that every object have a complete decomposition into objects *some* level below (unless the item in question is already basal). Furthermore, it remains that most biological objects are roughly hierarchical, so that

much of an animal's behavior can be explained in terms of the coordinated behavior of organ systems, and organ systems are largely conglomerates of cells, and cellular behavior is driven by organelles, and so on. Based on these observations, we might think that a sufficiently qualified model of levels will retain empirical adequacy and usefulness.

The above two fixes—a move towards localization, and the allowance that objects are not perfect hierarchies—do bring the model better in line with empirical science, and should be uncontroversial. But so far we have only said that objects are composed of successively smaller parts; this is hardly an interesting layered model. There are some stronger further assumptions which are made when a model of levels is applied to specific problems. I will illustrate these further features of the model by considering the case of emergent properties.

Emergence is a good case study because it is a classic position on certain interlevel relationships. The view in philosophy is a form of strong causal antireductionism. It is sometimes thought to be a mere philosophical working out of a concept borrowed from the biological sciences. In fact, however, I will later argue that the way emergence is typically defined in philosophy makes it inconsistent with how biologists understand emergent features of biological systems. Metaphysical definitions of emergent properties have construed them as intrinsic and static features, whereas emergence in biological applications pertains to dynamic, extrinsic features of systems. I will later attempt to explain how these distinct conceptions have their origins in the way philosophers have utilized the layered model.

Emergent properties are typically defined in philosophy along the following lines. For my purposes here, I will give only a rough and partial characterization. Emergent properties are: *mereologically supervenient* features of systems which are *qualitatively novel* with respect to lower-level features and are *not causally reducible* to the features of the system's lower-level

parts alone or in conjunction (see Francescotti 2007; Kim 1999; McLaughlin 1997; O'Connor 1994). I will define each italicized part of the definition sketch in what follows.

First, the phrase “qualitatively novel” is intentionally vague. It is merely intended to capture the intuitive idea behind any emergence view, which points to the fact that large and complex enough arrangements of matter will begin to exhibit interestingly new kinds of behavior.

Second, the causal irreducibility clause is typically thought by philosophers to require that emergent phenomena cannot be given a mechanism or functional explanation on any lower level. Biologists, by contrast, tend to assume that the processes they study can be given a mechanism, that is, processes can be described in terms of interactions among smaller-scale objects which compose relevant systems. Nevertheless, they also think select features count as emergent: so, they think emergence describes features of certain systems which have nothing to do with our ability to mechanize. For example, a certain process in the heart may be considered emergent because the behavior is novel with respect to the heart's components taken *individually*, yet the process is explained in terms of the mechanism involving the way those components interact in a given organization. This interpretation differs from the metaphysical view, on which the causal powers of the heart would have to be inexplicable in terms of the heart's molecular components, even in a given organization, in order to count as emergent. I will further discuss this biological case below, in section 3.

Sandra Mitchell, in a paper on emergence in biology, argues that differences of interpretation regarding emergence, such the one I just described, point to some problems for philosophers. She poses the following question: “If the philosophical analyses that dismiss the reality of emergent properties are correct, then why have descriptions of emergent properties in science become so widespread?” (2012: 173). Mitchell has in mind arguments like those of

Jaegwon Kim to the effect that an interesting emergence view requires a sort of reflexive downward causation which seems implausible—perhaps even incoherent (e.g. Kim 2000, 2006). But matter-of-fact references to emergence seem to be common enough in the biological literature to indicate that biologists do not take the phenomenon to be anything extraordinary. The reason for this puzzling discrepancy, Mitchell argues, is that scientists and philosophers mean very different things by “emergence.” This can result in two problems. First, philosophers might be in danger of misinterpreting scientific findings, and philosophers with different amounts of scientific background might end up talking past each other. But, furthermore, the metaphysical conception of emergence (Mitchell argues) does not apply to the systems which biologists claim exhibit emergence. While the metaphysical view may be internally consistent, she finds it troubling that it appears to fail to apply to these real physical systems, i.e., that the view is out of touch with empirical fact. I am similarly concerned about this issue, which I am worried may generalize from the single case of emergence to the whole levels framework which has been used to develop various reductive and antireductive positions in philosophy.

To return to the last part of the definition of emergence, emergent properties are thought to mereologically supervene, which is to say, they supervene on properties of the system’s (“lower level” or microphysical) parts, taken in conjunction. Roughly, supervenience holds between two classes of properties when any object intrinsically alike in terms of one class (here, features of the object’s lower-level composition) must also be alike in terms of the other class of properties (here, higher-level or emergent features). That is, anything with the same physical composition will have the same emergent properties. This statement seems innocuous, insofar as the supervenience of all features of the world on the microphysical is an assumption of physicalism.

However, the explicit appeal to mereology imports the following assumption: that emergent features of systems supervene on just the features of their own parts, that is, without considering any features of the environment. This means emergent properties, on the standard metaphysical definition, can only be intrinsic features of objects. I will argue later that this is a bad restriction since biologists understand emergent features to be extrinsic, i.e., partially dependent on the environment.

Furthermore, the supervenience clause is typically (at least implicitly) taken to put the emergent feature and its lower-level subvening base in temporal lockstep. That is, systems compositionally identical *at a time* must be the same with respect to their emergent properties (if they have any of those). For example, here is a statement by Kim from a discussion about emergence; notice the temporal claims:

*M\**, as an emergent, must have a basal (physical) property *P\** from which it emerges; *M\** cannot be instantiated unless some appropriate basal condition, say *P\**, is present; moreover, the presence of *P\** by itself guarantees that *M\** will be instantiated at that time, *no matter what has preceded this occurrence of M\**. That is, as long as *P\** is there at the time, *M\** will be there at the same time... (2006: 557; emphasis in original).

Kim is here in the process of making a specific argument about mental causation (*M\** stands for a mental property, though it could be read as any other emergent property), but mind is not my immediate concern here. Instead, notice that his assumption is that emergent features are fixed within temporal slices, absent information about the preceding or following states of the system.

To summarize, what this emergence case illustrates is that when it comes to particular objects or systems, some philosophers have tended to model system behavior in terms of causal relationships among temporal sections through the (local) compositional hierarchy of levels. It is usually reported (or clearly assumed) that mereological supervenience holds at those temporal instants, such that any system compositionally alike at an instant of time will also be alike in

terms of salient systemic features. This further requires that the explanatorily salient features of the system under investigation be intrinsic features, since the subvening base is typically construed as including only features of the object's internal structure.

Although I have only discussed one example for the sake of space, I do not think I have cherry-picked an unusual case. Emergence is just one of many related discussions about reduction and the "levels" of causation (including downward causation), and the discussed assumptions—that causal structure is adequately captured within a framework of composition and temporal slices—appears to me to be made throughout this literature.

## 2.1. SCIENTIFIC USE OF "LEVELS"

Above I have attempted to show how philosophers have been utilizing a layered model. On this model, the salient relationship for investigating complex systems is interlevel composition, and the parts making up a higher-level object at a given time fix the salient higher-level features of the object at that time. I want to now step back and make some general comments about the idea that the world is layered. In this section I will give a brief negative argument regarding layers: the philosophical layered model does not fall out of any scientific usage, so it cannot derive its plausibility from science.

I think the ordering of levels is usually expected to reflect the purported ordering (from general to special) of scientific theory. But, it is assumed in many philosophical applications that layers are objective ontological features of the world, that is, classes of physical objects and/or their properties are stratified. This in itself seems a little strained, since scientific theory often divides up the world orthogonally to the way we would intuitively class entities into ontological layers. To illustrate, one can apply the scientific study of mechanics, electrochemistry, and

ecology all to organisms (a single intuitive level). Ecologists might appeal to ions and plant communities (different intuitive levels) in a single explanation of some ecosystem process. The same laws of physics, similarly, can apply to a tiny ball or a star. So, at a glance, it is not clear we should expect the organization of science to ground or even roughly mirror the sort of ontological layering philosophers are after.

It must therefore be something about compositional relationships among entities which grounds their layering: either the part-whole relationships among objects, or asymmetric determinative relationships among properties. As far as I know, no one has given a thorough, nonarbitrary, widely accepted account of how this could work, i.e., what facts about the world could make it the case that certain objects/properties count as objectively higher-level than others (see Kim 2002). This is a serious explanatory gap.

Although I think this gap in the model is worth pointing out, I have no argument that it cannot be filled. What I will argue is that levels (in any clear and interesting sense) are not typically important parts of the explanatory structure of biology, and further (in the next part of the paper) that the way the layered model has been utilized in philosophy clashes with the way scientists typically frame the structure of complex systems. So, although it may be that some notion of levels which is satisfactorily internally worked out could be produced, and that such a model could be relevant to the discussion of select abstract issues, I question whether a levels view can be made adequate for discussing the structure of complex physical systems.

With that background, I will now briefly argue that the layered model does not derive from scientific use of “levels.” In colloquial or scientific usage, “level” can refer to all sorts of things which are only vaguely related to each other and to philosophical usage. The term often pertains to sequences of inclusiveness, abstractedness, or theoretical broadness, in application to

descriptions or explanations of system behavior. Alternatively, it can apply to series of larger-scale objects or processes, with or without invoking composition. I will give some examples of various scientific uses.

Here are examples of the use of “levels” in science. (A) Geneticists often speak of explanation on the level of genes versus phenotype. The phenotype consists of gene products, so moving between these two modes of description is not really moving up-down in the sense of composition, but rather going back and forth between two causally interconnected networks. (B) We might discuss the description of an entity’s behavior at the level of physics (i.e., in terms of basic physical magnitudes) versus at the level of biology (e.g., in terms of behaviorally salient facts). Here we are considering one and the same entity under the conceptual schemes of different sciences. (C) We might compare a model that predicts system behavior at the level of a whole population versus at the level of individual interactions. In this case our different models vary in granularity. This case approaches the layered-model sense of “level,” except that differences in the temporal scale of interactions may be built into the two models. More examples of scale-related issues, and why they problematize the layered model, will be considered in what follows. (D) Ecologists often talk about the biological hierarchy moving from spatially smaller to larger objects: atoms all the way up to the biosphere. It bears noting, though, that it is an oversimplification to consider each level that of a larger kind of object, especially at larger scales. Populations, for instance, are arguably not objects. So, these levels are of explanatorily salient kinds, not of metaphysically comparable sorts of physical entities. Alternatively, (E) ecologists might use “level” almost synonymously with “scale” in application not to different sorts of entities but to relevant processes within a single entity. For example, an

ecosystem process that occurs slowly and over a greater area (e.g. succession) might be called higher-level with respect to a faster, more localized process (e.g. nutrient flow).

I think the above should make clear how inconsistent levels talk tends to be. This is alright when scientists use the notion of levels only as a metaphor for the way we can describe or resolve entities and processes on different scales, from different perspectives, or with differing abstractedness. It should go without saying that this sort of unclarity is problematic when using a layered model to frame philosophical discussions. I want to stress that, contrary to what is apparently sometimes assumed, the philosophical levels model is not contained within scientific usage of the term. Scientists do not use “levels” with any sort of consistency, let alone with any particular ontological structure of the world in mind.

### 3. SYSTEM STRUCTURE IN BIOLOGY

As I argued previously, many philosophers seem to assume that properties are related across levels via mereological supervenience, and that carving up systems based on composition and time slices adequately preserves their causal structures. Complications for the scheme arise, I will now argue, when we consider complex systems, perhaps especially biological systems. One reason for this complication is that features which are determined in part by object-environment relationships seem to be very important to the causal structure of biology. Relational features, of course, will not be subject to mereological supervenience. Furthermore, many biological processes are dependent on the spatiotemporal scale at which we describe them. Restricting description to fixed time intervals, as is typically done under the layered model, will cause us to lose sight of system processes that occur over longer timescales (processes which might be considered “higher-level” on the traditional scheme). As a result, the levels framework,

which models objects based on their composition and in temporal slices, excludes from consideration many of the features of systems which are most important in many biological contexts. I will attempt to clarify and provide support for these claims in what follows. In the four subsections below, I will discuss as examples (1) the extrinsic nature of biological function; (2) the relevance of recent work on so-called computational emergence to the question of methodological reductionism; (3) the case of emergence as understood by physiologists; and (4) the importance of spatiotemporal scale and the environment to the explanation of ecological systems. For each of these examples, I will explain why the traditional layered framework poorly fits the structure of the biological explanation.

### 3.1. STRUCTURE VERSUS FUNCTION

One of the first things drilled into one's head in a typical introductory biology curriculum—before “always use a control,” or “never infer causation from statistical correlation”—is that structure determines function. Actually, unqualified, this statement is false. It is important to draw out how it is false, because functions are highly important in all sorts of biological explanation: they (apparently) do much of the causal work in biological systems across scales. How biologists understand system behavior will to some extent always be tied up with how they understand function.

The reason structure fails to determine function is that a function is the *actual* causal role which is played by a part of a system, whereas structural arrangement fixes only the *possible* causal roles a part might have. What determines the actual function of some structure is its external environment, e.g., other parts of an organism of which a molecule is part.

John Dupré expresses this thought nicely by saying that “whereas the structure of an entity looks downward, as it were, to its constituents and their relations, the function of an entity looks upward to its role in a system or entity of which it is part” (1993: 125). He gives the following example. Hemoglobin is a protein which has been selected in a variety of vertebrates for its function, namely, transporting oxygen. But there is no determinate genetic or molecular sequence corresponding to hemoglobin, because its sequence has changed throughout the course of evolutionary history, and mutations can result in sequence differences even among (e.g.) contemporary humans—as in those with sickle cell anemia. Because there is an indeterminate spectrum of similar molecules which might carry oxygen more or less efficiently when placed in an appropriate environment like a human circulatory system, there is also no determinate disjunction of molecular sequences which perform the function of oxygen transfer.

The outline of the point above will be familiar as an illustration of multiple realizability, which has commonly been used as an argument against the possibility of theory reduction. But I want to stress a somewhat different point about the case. Various molecular arrangements perform the job of transporting oxygen only because of a contingent evolutionary history. At some point in evolutionary history, the sequence or the human body might have changed slightly such that the hemoglobin molecule began to perform some other function. It is furthermore possible that exactly that same molecular sequence might have ended up performing two different functions at two different places or times in one and the same organism, as a result of changes in the organism’s developmental processes. In fact, it is very common for the products of a gene to perform several unrelated roles during development, a phenomenon known as pleiotropy. So, the structure of a molecule does not determine or even necessarily strongly

constrain what functional role it might end up performing on an organism, a role which is determined by the precise way in which it interacts with its complex surroundings.

Consider a related example: the way we classify genes. We tend to speak of the allele (i.e., the specific genetic sequence) for X, where X is either a type of molecular product or a more apparent trait like height or color. But, as is well known to anyone who has worked in any capacity with genetics, there is rarely a nice one-to-one mapping between genes and gene products. There can be multiple gene sequences which produce the same trait; and the same gene sequence can produce different protein products in different organisms, or in the same organism in different places or at different times. This latter fact is because the sequence gets heavily edited (parts get chopped out, molecular bits stuck on) in the process of turning the original DNA sequence into a functional protein.

To further illustrate, consider the allele that makes a hypothetical flower purple. The genetic sequence produces a purple pigment only because of the vast and complex protein-making machinery that exists in each flower cell. So, the same segment of DNA floating in a vacuum is not intrinsically a gene for anything; it is just a molecule. And the same DNA sequence in another organism might be a gene for something, but there is no in-principle reason or even (without further information) any very strong empirical likelihood it would be involved in producing a purple pigment or anything functionally related to one.

On a layered framework, one would intuitively consider the molecular structure to be lower-level and the biological function to be a higher-level feature. However, as we have seen, the functions of proteins or DNA sequences fail to mereologically supervene. The surrounding environment in addition to mere composition is needed to explain why some molecule performs

some function, i.e., some causal role in a biological system. This is the first way in which a layered model fails to be consistent with the explanatory structure of biology.

One utilizing a levels framework might object to these sorts of cases as follows: it may be that properties like *being an oxygen transporter* or *being a purple-pigment allele* fail to supervene on molecular structure, but in any causal interaction involving some token molecules, their causal capacities—and so the course of the interaction—are fully determined by their structures. That is, the causal capacities of any molecule will supervene on structure; what is underdetermined is merely which of those capacities is exercised, which (trivially) depends on the surrounding environment.

I agree with the above; however, more needs to be said, because it is not the case that one can adequately explain all biological phenomena by appeal only to structurally fixed capacities. For example, in the first case I gave, a causally important feature of a hemoglobin molecule—its *capacity* to bind and release oxygen under specific circumstances—is fixed entirely by its intrinsic molecular structure. However, this fact alone does not allow us to give evolutionary explanations. Biologists argue a trait was selected when it gives some reproductive advantage to organisms in a given environment. And to give a reproductive advantage, a structure must perform some function in the organism. Natural selection is notoriously blind to *potential* causal capacities. Humans as a population would not have hemoglobin in our circulatory systems if it did not actually transport oxygen in ancestral organisms. And one cannot explain the fact that hemoglobin *actually* transports oxygen—that is, that its function is oxygen transport—by appeal only to structural facts. So, if we are to take seriously evolutionary explanations of biological phenomena, we must allow for the appeal to features of systems' parts which are not fixed by the parts' structural arrangements.

A further objection that might be raised is that evolutionary explanations are not *really* causal explanations, or at least they are merely higher-level integrations over the detailed causal processes really going on in fundamental physics. And those basic physical processes require appeal only to structural facts, not to functions. It is merely a result of our computing limits that we humans have to appeal to function to understand system behavior. So, an objector might claim, my argument relies on an unwarrantedly strong assumption about how literally we can interpret biological explanations, such as evolutionary theory.<sup>3</sup>

I do not think this sort of objection undermines my argument in this paper. I am not taking a metaphysical stance on what entities are *really* doing the causal work. I am merely taking a stance on how biologists explain biological processes. Insofar as some view about micro-macro relationships must be presupposed by any discussion about reduction, it seems as if the issue of levels must be resolved prior to discussing whether and how biological entities might reduce, causally or otherwise. That is, one cannot appeal to a certain reductive view in order to prop up the levels model, since reduction (as standardly formulated) just is a view about *interlevel* relationships. The question I am asking here is how we should *formulate* these questions about reductionism. I am assuming that we ought to take biologists seriously about the nature of biological phenomena—including how causal processes within biological systems are best understood—before we discuss whether they are in fact subject to reduction. My view on the layered model does not rule out the possibility that all causation is really going on among microphysical entities; however, it might suggest a different way of understanding *how* microphysical processes account for biological processes, should we take a causally reductive view.

---

<sup>3</sup> The potential for this sort of objection to my project was pointed out to me by Dr. McDonald.

To continue, above we saw that functions are not fully fixed by structure, though causal capacities are. In fact, there seems to be a spectrum from biological properties which are largely constrained by intrinsic structure, to those which are largely determined by the environment. Genetic and molecular functions probably sit somewhere in the middle. Examples on the more extrinsic end of the spectrum might include the fitness of a specific trait, which is determined for the individual by average probabilistic effects on reproduction over the whole population (although the nature of fitness is controversial; but see Dupré 1993, ch. 6, for one perspective). A less contentious example is camouflage, which is not determined by the color of an organism but rather by its relationships to the color of the environment and perceptual and behavioral features of predators (Potochnik and McGill 2012).

Moreover, the entire theoretical structure of ecology is full of explanatory appeal to extrinsic features. For example, *being an apex predator* or *being a keystone species* (i.e., a species whose effect on ecosystem processes is disproportionate with respect to its biomass) are features of populations which are determined in part by their surrounding communities and/or geographical environment. Many processes are also strongly context-, density-, or rate-dependent, i.e., behaviors or effects of individuals will change based on contextual factors. For example, algae growing at a “normal” rate will tend to cause a net increase in dissolved oxygen within a body of water (because of photosynthesis), while a rapid, large algal bloom will cause a net decrease in oxygen (because a large mass of algae will die off and sink to the bottom of the water column, where its decomposition will use a lot of oxygen). Landscape ecology, the study of metapopulation dynamics, and island biogeography are whole subfields whose explanatory structure depends importantly on the way in which spatial context affects processes internal to systems (I will discuss these in more detail later on).

With all of these examples, I hope I have begun to illustrate why a layered model, which focuses on composition and intrinsic features of systems, cannot accommodate the explanatory structure of biology. I clearly have not raised any problems for physicalism, for levels in any loose or colloquial sense, or (at least not immediately) for various metaphysical reductive projects. However, the importance of extrinsic features in biology *is* a problem for the popular version of the layered model. One way to describe the problem is that, from a top-down perspective, we cannot have both the composition base and the supervenience base for system features at once, as the layered model wants. This is because the subvening base of important system features in biology often includes large chunks of the environment, so *mereological* supervenience fails to hold. This makes it impossible to adequately model a system's behavior just by decomposing it into its parts. Focus on the parts alone causes the model to neglect extrinsic features which, it turns out, do a lot of explanatory work in biology.

### 3.2. COMPUTATIONAL EMERGENCE

To strengthen my above point and clarify its implications, I want to transition from specific examples to a more general characterization of the explanatory situation I am describing. Biologists assume that any process explained by a system feature can be given a finer-grained mechanism, but this does not mean that classically “lower-level” descriptions of systems are *predictively* better in principle. More formally, given the claim that “system S has (extrinsic) property P which explains behavior B,” biologists would infer that the following is also true: “salient micro-constituents of S + salient parts of its environment E are organized in such a way

as to jointly produce B.”<sup>4</sup> Importantly, we cannot merely consider the micro-constituents of S if we want to know why S does B, i.e., we cannot assume that salient micro-constituents of S alone are organized in a way that is sufficient for (or even makes likely) B. Thus, explaining the process in terms of *finer-grained* interactions requires the consideration of (*spatially*) *more of the environment*.

This point can be put in both metaphysical and epistemic terms. First, a micro-description of S alone fails to fix all of the (relational) features which are involved in causally producing B. This is what I illustrated above with features like camouflage and *being a gene for trait X*, which are not fixed by the micro-construction of the object which has those features. Second, if we wish to predict the behavior of S, considering only interactions within S’s micro-configuration will severely underdetermine its future states, since doing so neglects important environmental factors. It is obvious that a structural micro-description of S alone will contain less predictive information than a description of the micro-state of all of S+E. But more subtly, it will also contain less information than a description of the macro-state of S, which integrates some information about the environment by incorporating some of the system’s extrinsic features (e.g., when we describe a frog as camouflaged). Here is another way to put the difference: the micro-description of S contains only the relationships among its micro-parts, while the macro-description of S contains extrinsic features of S *qua* whole system—that is, relationships between the whole of S and its environment. So, given a spatial scale identical with the size of the system (i.e., given the exclusion of the environment from direct consideration), a macro-description of the system could be better from a predictive standpoint, depending on how sensitive the system dynamics are to extrinsic features.

---

<sup>4</sup> Notice that this leaves various metaphysical issues—including issues about reduction—underdetermined. There need not be any particular relationship between parts of the mechanism and system property P; and the empirical biology on its own does not settle which entities are really doing the causal work.

The epistemic point I just made, I take it, is related to the thesis of some recent work on “computational emergence” (Bedau 1997, 2008; Huneman 2012; also see Mitchell 2010, section 3). These authors have pointed out that in systems exhibiting complex nonlinear or chaotic behavior (even if deterministic), predicting the macro-state (i.e., the overall structure) of a system at time  $t_n$  may require computing each successive micro-state of the whole system from  $t_0$  to  $t_n$ , which quickly becomes computationally infeasible. One cannot compress the calculation since each step is too sensitive to the immediately prior state of the system. More feasible predictions (and, given a ceiling on computational power, more accurate ones farther into the future) may instead be made by considering patterns of system-level dynamics for the kind of system under investigation.

As stated, this is just a thesis about epistemic limits, albeit ones determined by the nature of the system rather than by limitations of the observer. However, the thesis can be made stronger by incorporating the point I made above, namely, that micro-descriptions of real biological systems will neglect causally salient parts of the environment. The work on computational emergence tends to employ highly idealized, simplified computer simulations, in which environmental influence is either ignored or incorporated into the program. In a biological system, not only is the environment highly causally important, but the environment itself may behave in a nonlinear or chaotic manner.

This means that one *cannot* predict the future behavior of a cell, organism, or ecosystem given only its microstates. This amounts to, say, trying to predict my daily movements without considering my surroundings. You cannot predict where I might walk next if you do not know where the walls are in my house. This holds just as strongly even if my behavior is thought to be entirely causally produced by subatomic interactions.

To give a more realistic illustration, consider predicting the state of a lake several years into the future. Perhaps you are worried about the effects of fertilizer runoff. You might do a simulation of the lake's molecular structure to see how the incoming phosphate interacts with all of the other molecules constituting the lake. Alternatively, you might consider how phosphate tends to affect lakes from a whole-system perspective.

Stochastic environmental effects will be greater at finer resolutions. Recall that the environment causally influences the system and itself behaves in a hard-to-predict manner. Stochasticity—that is, apparent randomness, from the perspective of the model—generates an inherent amount of error in any prediction. So, the behavior of salient parts of the environment (perhaps, e.g., future weather patterns) is a stochastic or unpredictable factor in the lake model. Importantly, this effect *increases* with the granularity of the model, both because of the exponentially higher number of individual interactions one has to account for, and because (as I explained above) finer-grained descriptions will tend to incorporate less information about the system's extrinsic features. So, the second option—a whole-system model, rather than a micro-simulation—would give a better prediction even to an observer with perfect knowledge of the present state of the lake and unlimited computational power. Put differently, assuming relevant knowledge regarding the scientific “explanatory kinds” invoked in the descriptions, the macro-description contains better predictive information about the lake *in principle*.<sup>5</sup> This seems to be

---

<sup>5</sup> To clarify, this is from a local perspective (since above we saw that a layered framework is supposed to apply to local stratification). It might be thought that the *best* predictive scenario would be having (contrary to physical possibility) exhaustive knowledge of the universal microstate. Although this claim seems plausible, I think whether it is actually true is an open question, since it may depend on further unresolved issues such as whether and in what sense causation might be irreducibly probabilistic. As I will discuss later in this paper, macrostates tend to occur over a greater temporal as well as spatial interval. This *could* mean that error due to randomness or intrinsic uncertainty would accumulate more quickly for micro-based than for macro-based simulations even when there is no external environment, simply because the number of computational steps over a given time interval will be higher for a micro-simulation.

one reason whole-system models will always be useful in biology, even though any token process can presumably be given a finer-grained (and therefore more detailed) mechanism.

On the standard levels framework, macro-descriptions are limited to mereologically supervenient (i.e., intrinsic) features of a system. This being the case, micro-descriptions will look like they should always generate more accurate predictions, since they contain more structural detail. Historically it was a common assumption in the philosophy of science and among scientists that scientific predictions would continue to improve only as they pertained to smaller or more basic entities, but this has not been borne out by contemporary science. Both in principle and sometimes in actual practice, coarser-grained system descriptions can better predict the behavior of complex entities.

This clear failure of a prediction of the layered model constitutes one reason to reject it. To reiterate: if we assume, per the received layered model, that the microstructure of an object is the uniquely important factor fixing its properties and governing its causal interactions, then we would predict that finer-grained descriptions of systems will, in principle, always better predict their behavior. We would expect this to be especially true for systems which are highly organizationally complex, since the intricate micro-structure will be washed over by coarser-grained descriptions of such systems. This expectation, I have tried to show, is false *in principle*, even ignoring such practical issues as human computational limits. It is false because of the *very high sensitivity to the environment* of complex biological systems, whose behavior is inadequately modelled considering just intrinsic features.

### 3.3. BIOLOGICAL EMERGENCE AND DOWNWARD CAUSATION

What I said above begins to indicate that conceptualizing individual biological entities in terms of a downward-looking compositional stratification will not adequately capture their explanatorily important features. This is partly because such features may obtain only in certain environmental contexts. This is true for very many ordinary biological properties which feature in everyday biological explanations, like *being an allele for X*, *being camouflaged*, or *being an apex predator*.

So, I would emphasize that my case against the levels framework ultimately depends on what I take to be the typical understanding of system structure utilized at least implicitly by biologists, not on any unusual or controversial particular cases. However, I think emergence is a case that especially clarifies the problems with the layered model. In this section will try to see what biologists are best understood as meaning when they call a property emergent, and why this view conflicts with the layered-model version. To do this, I will look at some examples of emergence in physiology.

Denis Noble is a systems physiologist who, unlike most biologists, has written explicitly about causation and levels. He is a convenient source for my purposes here, since he is writing from the perspective of a working scientist and not a philosopher (that is, his insights seem to come from his experience in physiological modelling rather than from any extensive study of philosophy). I do not agree with everything he has written in his (2008) or elsewhere, but he makes some observations that will be helpful here.

Noble observes that molecules in a biological system often behave very differently from those in a nonbiological system, but only (we tend to assume) because they are parts of systems

that are organized differently to begin with, not because of any intrinsic difference—that is, the basic chemistry of the molecules remains the same.

Intricate feedback networks, which result from the initial physical organization of many biological systems, may be particularly important drivers of novel behaviors of system parts. In order to accurately model the functioning of such a system, we need to identify “the level at which such networks are integrated” (3012). What Noble means by this, I think, is that it is necessary to figure out which components of the system are minimally necessary to observe the behavior in question, which is produced by feedback among those components. When we know which parts participate in the behavior, it is then possible to model how those parts must be arranged and interact with each other to exhibit the function observed in nature. No such function will appear if we look at parts in isolation or consider too few parts at a time. Note too that the relevant organization for a given behavior may occupy various scales, i.e., the behavior may be exhibited at the scale of a cell or may be invisible until an entire organ is considered. This is a contingent, empirical matter, not a matter that can be settled by considering intrinsic features of isolated microphysical parts.

To illustrate the above claims, Noble discusses how electrochemical oscillation within the heart is produced. There is no single component of the heart that oscillates in isolation. (He notes that in mathematically modelling the beat of the heart, no individual component of the equations included an oscillator like a sine wave.) However, when a number of ions and protein channels are put together within a network of cells, rhythmic oscillation emerges within the system as a result of feedback between the electric potential of the cells and the channels that allow the movement of ions.

There are several points I want to make about biological emergence. I will give one of them here, and make further points after giving another example. The heart's oscillation is considered emergent because it is a feature that only occurs at the scale of the heart: it is "novel" with respect to the constituent parts of the heart. However, this claim does not entail any strong metaphysical antireductionism—that is, it is merely a "weak emergence" claim. It is not, for example, necessary to posit any causal interactions to explain the heart's rhythmicity other than those among individual system components (i.e., molecules) in a given organization. That is, the process can be given a complete "lower-level" mechanism as long as the components' arrangement is taken into account.

Metaphysically strong emergence, as I mentioned previously, is often thought to require reflexive downward causation. This is because emergence is framed on the layered model as a view about *static* (i.e., temporally sliced) causal directions; whereas biologists are not interested in "static mapping" of relationships, but rather in dynamic processes (Mitchell 2012).

Sometimes cases like that of the heartbeat will be said to involve downward causation, but what is meant by biologists in this context is only that one must take into account *organization* at the *scale* of the heart to understand why its component molecules interact in a certain *dynamic* fashion. I think this difference in the understanding of emergence reflects a difference with respect to the layered model.

A number of authors have in fact argued that metaphysically strong versions of downward causation do not look very applicable from the perspective of empirical science. So, if a version of emergence requires downward causation, then that sort emergence will also be inapplicable to real biological systems. Furthermore, downward causation seems to be directly

motivated by a layered model of macro-micro relationships, so seeing why it fails to clearly apply to biological systems will help to show why the layered model is inadequate.

William Robinson (2005) considers a variety of cases where there are purportedly causally competing micro- and macro-descriptions of the same process. He argues that what some philosophers have urged, namely that the organization of such systems should be thought of as exerting a downward causal effect on the system's micro-constituents, either turns out to be too weak or too strong. Unqualified, the view is too weak, since it makes downward causation too pervasive to be philosophically interesting: it occurs every time organization is important to system behavior, which would seem to be nearly always. Unfortunately, the only clear way to make the view stronger is to add a claim to the effect that the systemic organization violates or changes the laws pertaining to the constituents' behavior; put differently, the downward causal influence prevents us from giving causal mechanisms pertaining just to the micro-parts.<sup>6</sup> Neither option seems very appealing. The first version is trivial, but not many philosophers want to posit that systems can cause the *violation* of physical laws pertaining to their component micro-parts.

Robinson argues that in cases of competing descriptions, there are actually often *two* causal questions being answered. The first is simply how the laws of chemistry apply to the interactions of the micro-parts. The second is why the system is organized such that the parts interact in the way they do, rather than in some other way (which in biology might warrant an evolutionary explanation). That is, the fact that there are two *prima facie* competing causal stories about a system's behavior may simply reflect the fact that for a complete explanation, we

---

<sup>6</sup> I think the view I have stated here is weaker than—but something like it is entailed by—the synchronic and reflexive downward causal powers which Kim thinks must be had by strongly emergent properties (see, e.g., Kim 2000). That is, I take it the way a strong downward cause makes a causal difference is by interfering with lower-level mechanisms. I am doubtful whether the strong version of downward causation is even coherent, but you can hold that it is and still think, based on considerations I am discussing here, that science renders the thesis implausible in application to most real systems.

may want both a mechanism and an origin. And these two kinds of description apply to an organism on different spatiotemporal scales; they are not merely re-descriptions of the same process on different “levels.” Further, as I argued above, explaining the origins of a process in biology requires appeal to functional kinds, not merely to structural units which feature in the mechanism. So, as Robinson argues, to see the system-level feature as exerting strong downward influence on the parts is simply to misconstrue the explanatory task, at least as far as biologists would see it. Biologists are interested in dynamic processes on different scales, not in different-level explanations pertaining to static time slices. Moreover, seeing each explanation as occupying a different explanatory scope rather than merely a different “level” will prevent us from jumping to the conclusion that the two stories must be redundant.

Along related lines, Carl Craver and William Bechtel (2007) have argued that apparent cases of interlevel causation can be seen not to involve any sort of reflexive between-level causation, when we consider the way mechanisms are nested.<sup>7</sup> Often we might *think* a part has an effect on a whole or vice versa, when in reality we have conflated a state with its constitutive lower-level mechanism. So, for example, if a person has a heart attack which then (one might say) causes her whole body to die, what really has happened is that the cessation of blood flow to various other organs causes those to cease functioning, and the death of the person just consists in the cessation of function of her constitutive organs. Or (to give another of their examples), if someone starts playing a game of tennis, his bodily movements (we might be tempted to think) cause various biochemical processes to occur within his muscles. It is not, however, as if his playing tennis downwardly rearranges his constituent molecules. We can track all the

---

<sup>7</sup> This paper also proposes that there is layered stratification of mechanisms; see Eronen (2014) for a critique of the view that switching from nested objects to nested mechanisms can salvage layers. I think that Craver and Bechtel’s explanation of the examples I have borrowed are on the right track, and agree with Eronen that their attempt to frame their discussion in terms of actual levels just complicates an otherwise helpful treatment of mechanisms.

biochemical interactions through his nervous and muscular systems without explanatory remainder. Rather, the bodily movements involved in playing tennis are produced and constituted by the various interactions mediated at the chemical level. This, again, is at least what most scientists would tend to assume, the alternative being to posit frequent explanatory gaps within micro-descriptions of systems of complex arrangement.

That is all I will say here about downward causation. I think it is reasonable to conclude—though I have not given a complete discussion—that any version of downward causation which is strong enough to be interesting will also be, at very least, on tenuous ground with regard to its applicability to biological entities. If this is right, it is a major problem in turn for the applicability of strong views of emergence.

More importantly, I hope to have shown that strong downward causation fails to apply because it comes straight out of the levels model, which frames entities in a way that is largely unrelated to biological explanatory goals. Downward causation has been conceived as a way to save the causal relevance of macro-properties, or the explanatory relevance of macro-descriptions of systems, from reductive projects. This is because construing systems in terms of a static hierarchy of levels threatens to make the higher-level properties causally and explanatorily redundant. However, rather than positing strong downward causal powers—which is very controversial from both an empirical and a metaphysical standpoint—we might do better to reject the notion that macro-properties are merely higher-level integrations of the system's micro-state. Rather, describing a system in terms of its macro-properties incorporates information about extrinsic features of the whole system and its larger-scale context. So, again, salient macro-micro relationships in biology are not really interlevel relationships in the standard sense, but rather pertain to larger and smaller scales of description. Since downward causation

has been developed precisely as an *interlevel* relationship, it fails to apply easily to biological systems, whose dynamics are better understood as governed by the *scale* of description.

To return to emergence in biology: here is my second case from animal physiology. In our brains we have a region known as the SCN which is a large organized group of neurons responsible for our circadian rhythm. The cells on their own will oscillate, but they tend to deviate from the 24-hour cycle. When organized together in the SCN, they jointly produce an accurate, highly robust 24-hour rhythm. Furthermore, because of how the network of neurons is organized, some sort of oscillation can occur even if the protein components of the cells which are responsible for internal rhythmicity are removed. This is from a review paper on circadian rhythms: “When the individual [neural] cells are no longer rhythmic [due to gene knockout], the coupling pathways within the SCN network can propagate stochastic rhythms.... Thus, ... rhythmicity can arise as an emergent property of the network in the absence of the component pacemaker or oscillator cells” (Mohawk et al. 2012: 449). This is a purely scientific paper with no discernable metaphysical agenda. What they report is simply that a complex behavior of a system can arise due to how its parts are arranged, where those parts in isolation would not necessarily exhibit that sort of behavior, and where the mechanism of the behavior involves the system’s organizational features. This is exactly what we saw in the first example as well.

One lesson of this particular case is that it is a matter of empirical fact that the emergent feature occurs at the scale of the organ. Having an emergent property is not a matter of the organ’s being on a certain “level”; emergence pops up inconsistently. In similar systems, the mechanism for a behavior may simply involve intrinsic features of its components. But in the SCN, its rhythmic behavior is instead a result of relationships among components. This is why the authors quoted above explicitly report experimental results to the effect that (features of) the

rhythmic behavior of the SCN can be emergent, rather than being driven entirely by component cells which are intrinsically 24-hour rhythmic. Emergence in this sense thus has little to do with a particular *level* of organization and everything to do with the *actual* organization of a given system.

There is a further issue for the strong, levels-based emergence concept which arises with this example. The SCN, I take it, requires input from the visual system about when it is daytime in order to sync our circadian rhythm with the earth's daylight period. An SCN network sitting on the counter would therefore (as far as I understand) produce a rhythm, but not a properly aligned *circadian* rhythm. The same goes if we take a whole mammal and put it in an environmental chamber with no light; or if we put it in a freezer with a simulated 24h light cycle but with the air temperature kept at  $-30^{\circ}\text{C}$ . So, it bears noting, this emergent property does not supervene on features of the parts in conjunction, even if we add a "properly arranged" clause regarding the components. Rather, circadian rhythmicity occurs only if there is (1) a properly arranged network of neurons which is (2) imbedded within a reasonably typical mammalian brain and hooked up to various sensory inputs while (3) the brain exists within a reasonably well-functioning body which is living within a narrow range of further environmental conditions.

One might reasonably point out about the temperature case that a certain temperature of the SCN *itself*—not of the environment—is what is necessary for physiological function. However, physiological processes always take place over some temporal interval. While the *capacity* to oscillate under certain conditions may be a temporally instantaneous feature of a system, the oscillations themselves are not. And when we consider the diachronic causal story, it is clear that the system's having all of the intrinsic (structural) features necessary for oscillation at some time  $t_1$  does not guarantee it will have those same features at  $t_2$ . To put it crudely, one

could stick a mammal in the deep freeze at any time. And if something is not oscillating from  $t_1$  to  $t_2$ , then it is not oscillating at all, since oscillation requires a time interval. Since a mammal can maintain its body temperature only under a limited range of environmental temperatures, its external environment must remain within a favorable temperature range for a circadian rhythm to occur over the given time span.

Because of the causal relevance of factors like external temperature, it is always necessary to consider the environment when explaining biological processes. This is true even if all of the features which *generate* the process are features intrinsic to the system, like internal temperature and the arrangement of the components. Since the layered model considers only instantaneous determination (i.e., supervenience or similar), the layered model neglects the environment and assumes that the internal structure of the system (i.e., the micro-description) is causally sufficient for the state of the system at  $t_2$ . But in fact this is probably always false, since the environment can always causally influence the system when we take a diachronic perspective, that is, when we consider the system over appropriate time *intervals* rather than in a series of temporal slices. The instantaneous micro-configuration of a mammal does not guarantee that it is ever in a state of circadian oscillation, even if its state at the given time is appropriate for or consistent with the production of rhythmicity. On the other hand, saying that it is exhibiting circadian rhythmicity at some time guarantees that it is in that state at some other time as well. The issue is that the former description neglects the appropriate temporal scale for the process in question. I will return to this point in the section 3.4.

My cases have raised the following difficulty for the layered-model version of emergence. Looking straight downward from an emergent property (i.e., to the decontextualized structure of the system which has the property), we can retain a rough hierarchy of local compositional levels.

However, the system's structure alone may not exhaust the supervenience base for emergent properties, and the localized view furthermore neglects causally relevant features of the environment. On the other hand, if we look outward from an entity with an (extrinsic) emergent property we can cover the subvening base for the property as well as salient features of the environment. But, we do this at the expense of any consistent compositional hierarchy, since we will now be looking at a conglomerate of different kinds of objects whose levels of organization may not be commensurable (as I discussed in section 2). Moreover, in either case if we consider the micro-description in temporal slices, we may well lose the (dynamic) emergent processes entirely.

Forcing biological cases of emergence into the layered framework is thus both inconsistent with the claimed metaphysics of emergence and with the way biologists understand emergent properties. Emergence in biology can only be understood in terms of *diachronic processes* on an appropriate *scale* involving a *specified* structure (not merely: a structure on a given level). Such features are not nicely captured by the layered model, which forces systems into spatial and temporal slices rather than accommodating process-specific changes in the relevant spatiotemporal scale of description.

To be clear, I have not argued that any of the philosophers cited in section 2 who are working on (for or against) a strong emergence view are mistaken on their own terms. In the context of mentality, a localized (mind-brain) levels framework could be relevant, since this discussion is restricted just to the scope of the brain (i.e., the environment is not typically mentioned). However, I think much of the *general* plausibility of a layered model, which I think has been applied too hastily to other systems, comes from this single, really very exceptional case. We should be cautious about assuming that the conceptual work on emergence and

downward causation from the philosophy of mind can apply to the systems biologists investigate. Similarly, philosophers utilizing a framework of discrete, time-sliced levels should be cautious about co-opting biological examples in support of their views.

### 3.4. SPATIOTEMPORAL SCALE IN ECOLOGY

In this final section I will attempt to show that environmental context is not only important to the behavior of isolated biological objects, but that environmental features are explicitly part of the whole explanatory structure of several subfields of biology. The need for environmental context has been especially appreciated by ecologists, who work with systems in which the qualitative and quantitative nature of processes is notoriously scale-, resolution-, and context-dependent. Because of the strong dependence of certain ecological processes on scale and the environment, it is impossible to understand either the ecological theory or the objects involved in causal processes as occupying a hierarchy of stratified layers, which can be neatly linked up considering just interlevel composition. Instead, different processes require us to investigate systems at different spatiotemporal scales and within an open-ended environmental context.

The quotes in what follows come from a recently updated forest ecology text written for advanced undergraduates or beginning graduate students (Chapin et al. 2011), so they should present a fairly sophisticated and representative ecologist's take on ecosystem structure. Chapin et al. have this to say, for example, about *spatial context*: "Spatial patterns control ecological processes at all scales" (370). This is a sweeping claim; scientists are not ordinarily prone to sweeping claims. I will try to illustrate what they mean, and its implications for the layered model, with examples utilizing theory from island biogeography and metapopulation biology.

Chapin et al. allude to the fact that the shapes, sizes and arrangements of patches of a given ecosystem type (within the broader landscape) are very important influences on many ecological and evolutionary processes, including succession, migration, disturbance regime, and extinction rates. How isolated a patch is from other ecologically similar patches, for instance, influences how quickly individuals and species can arrive to colonize it. How small the patch is also limits the total number of individuals it can support. As a result, colonization rates are low and extinction rates are high on islands, which tend to be both isolated and small. These factors importantly change the genetics, evolutionary trajectories, and the complexity of the communities that can be supported by islands, as compared to *same-sized* patches of land which are part of a larger, spatially continuous region of the same ecosystem type (i.e., across parts of a large continent which have not been plowed). Additionally, for species which are distributed across relatively (but only partially) isolated patches, their likelihood of extinction may be far better investigated by considering metapopulation dynamics, i.e., the behavior of the partially connected group of populations. This is because individual patches may be periodically vacated and then recolonized even if there is a stable metapopulation. It may be worth noting that patches of land which are not literally islands can behave (ecologically and evolutionarily) like islands on account of size and isolation. This happens, for instance, when a tract of forest gets surrounded by cornfields. So, the example I picked is not limited to the single context of literal islands. For any patch of land, its *size* and its degree of *isolation* jointly govern numerous features of the basic ecology and evolutionary trajectory of the region.

Size may be a feature intrinsic to a patch, but distance from similar habitat clearly is not. “Higher-level,” external features of the landscape thus partially govern ecological processes within individual patches. It does not matter, furthermore, what compositional “level” the

relevant ecosystem processes are supposed to occupy: they could involve anything from molecules to communities of organisms. A layered model, which focuses on composition and intrinsic features, is inadequate for all of these ecological processes which can *only* be explained—within ordinary ecological theory, and not in “odd” cases, like the mammal in the freezer above—with reference to external context.

The authors of my forest ecology text also write this in the caption for the figure reproduced below (Fig. 1): “The study of any ecosystem process requires understanding at least one level below (to provide mechanistic understanding) and one level above (to provide context with respect to pattern of temporal and spatial variability)” (371). Their different levels correspond to the processes occurring within ecosystems at different scales. Notice that there is nothing hierarchical, in a strict mereological sense, about the “levels.” All of them involve ecosystem processes which are mediated by individual organisms (largely plants). They all in some sense involve entities flowing through the ecosystem, whether those entities are atoms like N and P or communities of trees. Furthermore, the important interactions in which ecosystem processes consist can involve entities across multiple scales, as when a plant takes up inorganic P from the soil—which originated in bedrock—and incorporates it in its leaf tissue, which is then dropped and decomposed by fungi and soil microbes. Their levels consist of *processes* within an ecosystem which are observed over successively larger spatiotemporal *scales*.

There is often a positive correlation between the temporal and spatial scales appropriate for observing a process (as indicated in the below figure), though this correspondence only holds very neatly for generalizations; for specific cases, appropriate spatial and temporal scales might come apart. This has been stressed by Potochnik and McGill (2012), who note that it is a matter of your research question and of empirical fact what scales are important as well as how finely

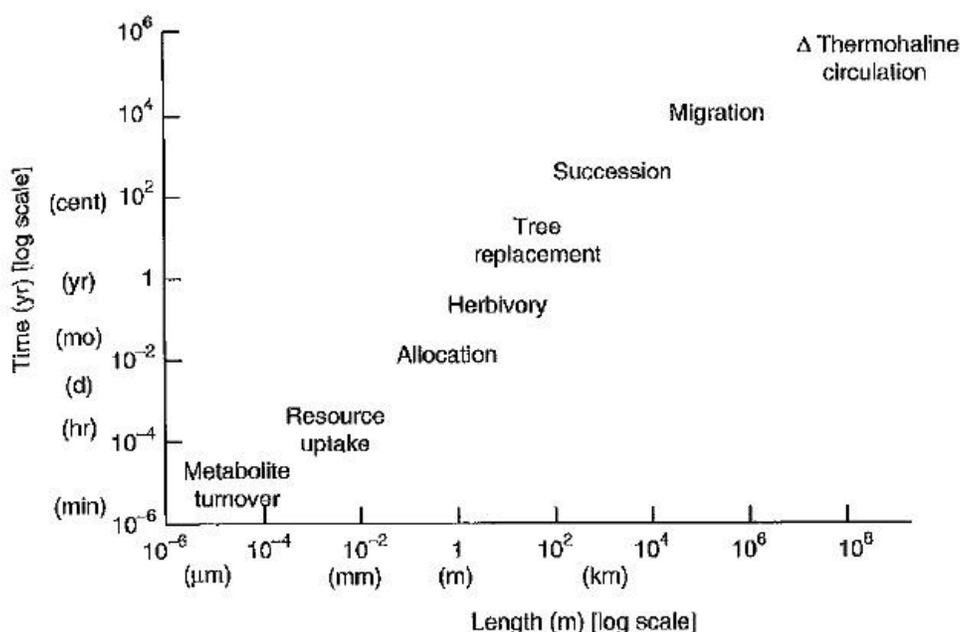


Figure 1. Spatial versus temporal scales for general kinds or “levels” of forest ecosystem process, showing a tight positive correlation. For example, succession—the whole-ecosystem process of changes over time in plant community and physical features of the soil following a disturbance—typically takes place over scales of thousands of kilometers and over centuries to millennia, while allocation—plants’ differential production of root, shoot, or reproductive mass in order to optimize nutrient uptake and reproductive success in a given environment—involves plant tissues at spatial scales on the order of just meters and can occur within single growing seasons. Some of these processes apply to animals, but the main focus here is on the behavior of individuals and populations of plants, which are most important to how ecologists classify and characterize the behavior of ecosystems. Source: Chapin et al. (2011, Fig. 13.2).

one must resolve differences in scale; it is not just a matter of how large component objects are. For example, they argue, the difference in size between a squirrel and a tree whose seeds it eats is important for their coevolutionary dynamics, but it makes no difference to the rate of their mutual northward movements in response to climate change. I think my physiology examples are relevant here as well: it is a matter of the specific process at hand whether the difference in scale between an individual cell and the organ is or is not important.

But, consider the fact that it is *often* the case that system processes occupy greater timescales than processes involving their microphysical constituents, an issue which I also

alluded to in the discussion of emergence. This leads to the following problem, as diagnosed by Sandra Mitchell, for the way the layered model is often utilized:

If we take a snapshot view of the higher and lower levels, then the dynamics of *how* the higher level is constituted and stabilized is lost. Contemporary sciences show us that there are processes, often involving negative and positive feedback or self-organization, that are responsible for generating higher-level stable properties, and these processes are not captured by a static mapping (2012: 177).

That is, if we take a time slice of a system, we will capture its microphysical organization, but we will be blind to the stable system processes which tend to occur over larger temporal intervals. Moreover, we will miss *how* these processes occur, which is the scientifically interesting and predictively useful issue. That is, interesting system processes may involve complex dynamics (like feedback) across longer time intervals, which are excluded from consideration by the layered model. Because of the complexity of biological systems, both the spatial and the temporal scales at which we will see stable, predictively useful processes are context-specific facts which must be determined empirically; one cannot make such discoveries by thinking about part-whole relationships from an armchair. Importantly, we cannot assume that we will be able to adequately model biological systems if we consider them only *at* instants in time, as has standardly been done on the layered model. Both spatial and temporal context are vital components of biological theory and are necessary for understanding the relationships among biological objects.

#### 4. CONCLUSION

The layered model portrays physical objects, and corresponding scientific theory, as stratified into a hierarchy of layers. On this model, the lower-level objects which compose higher-level systems are the most important determinants of the system's properties and behavior.

Furthermore, the system's micro-properties are supposed to fix its macro-properties on very small (instantaneous) temporal scales, which allows for clean mapping between micro- and macro-descriptions.

I have argued that these suppositions do not accord with the structure of complex biological systems and biological theory. Many properties of biological objects—including functions, which are necessary for evolutionary explanations—are extrinsic features, i.e., they are not fixed by the object's micro-structure. The phenomenon of emergence in particular is thought by biologists to involve extrinsic, scale-dependent, dynamic features of certain intricately organized systems. This contrasts with the metaphysical view of emergence, which, in accordance with the layered model, involves static upward and downward relationships. Finally, many ecological processes vary with spatial context and temporal scale, rather than with the intuitive "level" occupied by the objects which mediate the processes. Taken together, these sorts of observation make a cumulative case that the received layered model *very poorly fits* biological systems. Biological dynamics are driven by scale and environmental context, not just by an object's structure at a moment of time.

This is not to say that the layered model is internally implausible, lacks any good application, or can be conclusively demonstrated false via empirical investigation. Rather, taking a broad view of relevant issues, the model does not cohere well with how biologists conceptualize the systems they study. This suggests that the philosophical discussion of biological systems—including discussions about causal reduction and relationships among theories—might be improved in sophistication and in empirical adequacy by a departure from the layered model.

It would be a second project to see how discussions about reduction and causation within complex systems might be altered by a shift to a scale-based, contextual, dynamics-oriented framework. However, it is clear that if we wish to discuss whether some biological item reduces, framing the question as one about *interlevel* relationships might cause us to miss the relevant biological phenomena.

## REFERENCES

- Beckner, Morton. 1974. "Reduction, Hierarchies, and Organicism." In *Studies in the Philosophy of Biology*, eds. F. J. Ayala, and T. G. Dobzhansky, 163-176. Berkeley: University of California Press.
- Bedau, Mark A. 1997. "Weak Emergence." *Philosophical Perspectives* 11:375-399.
- , 2008. "Is Weak Emergence Just in the Mind?" *Minds and Machines* 18:443-459.
- Chapin, F. Stuart III, Pamela A. Matson, and Peter M. Vitousek. 2011. *Principles of Terrestrial Ecosystem Ecology*, 2<sup>nd</sup> ed. New York: Springer.
- Craver, Carl F., and William Bechtel. 2007. "Top-Down Causation without Top-Down Causes." *Biology and Philosophy* 22:547-563.
- Dupré, John. 1993. *The Disorder of Things*. Cambridge: Harvard University Press.
- Eronen, Markus I. 2014. "Levels of Organization: A Deflationary Account." *Biology and Philosophy*. DOI: 10.1007/s10539-014-9461-z [electronic preprint].
- Francescotti, Robert M. 2007. "Emergence." *Erkenntnis* 67(1):47-63.
- Huneman, Philippe. 2012. "Determinism, Predictability, and Open-Ended Evolution: Lessons from Computational Emergence." *Synthese* 185:195-214.
- Kim, Jaegwon. 1999. "Making Sense of Emergence." *Philosophical Studies* 91(1):3-36.
- , 2000. "Making Sense of Downward Causation." In *Downward Causation*, eds. P. B. Andersen, C. Emmeche, N. O. Finnemann, and P. V. Christiansen, 305-321. Aarhus: University of Aarhus Press.
- , 2002. "The Layered Model: Metaphysical Considerations." *Philosophical Explorations* 5(1):2-20.
- , 2006. "Emergence: Core Ideas and Issues." *Synthese* 151(3):547-559.

- McLaughlin, Brian P. 1997. "Emergence and Supervenience." *Intellectica* 25:25-43.
- Mitchell, Sandra D. 2012. "Emergence: Logical, Functional, and Dynamical." *Synthese* 185:171-186.
- Mohawk, Jennifer A., Carla B. Green, and Joseph S. Takahashi. 2012. "Central and Peripheral Circadian Clocks in Mammals." *Annual Review of Neuroscience* 35:445-462.
- Morgan, C. Lloyd. 1923. *Emergent Evolution*. London: Williams and Norgate.
- Noble, Denis. 2008. "Genes and Causation." *Philosophical Transactions of the Royal Society A* 366:3001-3015.
- O'Connor, Timothy. 1994. "Emergent Properties." *American Philosophical Quarterly* 31:91-104.
- Oppenheim, Paul, and Hilary Putnam. 1958. "Unity of Science as a Working Hypothesis." In *Minnesota Studies in the Philosophy of Science*, vol. 2, 3-16. Minneapolis: University of Minnesota Press.
- Potochnik, Angela, and Brian McGill. 2012. "The Limitations of Hierarchical Organization." *Philosophy of Science* 79:120-140.
- Robinson, William S. 2005. "Zooming in on Downward Causation." *Biology and Philosophy* 20:117-136.

### Appendix: Faith and Learning

Here I will briefly discuss how I think my honors thesis, or my work towards this paper, interacts with faith. I have not said anything directly related to religious beliefs as part of my project, so I will instead discuss my general viewpoint on faith-science dialogues.

I hope part of what I have achieved in my paper is a model of how I think interdisciplinary (philosophy/faith-science) conversation can take place. Some have taken a “segregation” approach to faith and science, proceeding as if their faith-based beliefs and their scientific beliefs have nothing to say to each other. I think this stems from the wish to insulate one’s religious beliefs from potential damage by apparently hostile scientific findings. This may not be a wholly badly-motivated reaction, but I think concerned and intellectually honest theists should take the opposite viewpoint. To subject one’s faith to rational consideration and the input of other disciplines seems to be a sign that one thinks those beliefs can actually hold up to scrutiny. For example, we should be willing to discuss whether belief in the existence of the soul, or in the special creation of species, holds up to scientific evidence and to good reasoning. If they do hold up, then so much the better for those beliefs. If not, then those who hold religious views should be willing to consider which beliefs their faith commits them to, and which beliefs are best revised in light of available evidence.

Furthermore, I think developing a robust Christian faith, with a full appreciation for the natural creation, involves a proper appreciation for science and for those parts of philosophy which discuss the natural world. So, I think it is important for Christian scholars, or scholars interested in issues in philosophy and theology, to be willing to entertain scientific viewpoints.

Throughout my coursework and my research for this project, I have been able to think about many of the ways in which faith, philosophy, and science have relevant things to say to

each other. I have spent a fair bit of time, as part of my study of biology, thinking about the way various Christian beliefs may or may not be compatible with evolutionary theory. As an ecology major, I am also particularly interested in the way religious mandates or religiously-motivated moral beliefs might influence our perspective on conservation. As part of philosophy classes, I have thought about the problem of evil, reasons for theistic belief, and whether we can make sense of some of the more puzzling Christian doctrines, like the Trinity and Incarnation. I have emerged from thinking about some of these issues feeling like I have a better grasp of the nature of the problem and what is at stake. About other issues I remain confused, but this only means I am motivated to continue thinking about them. Studying philosophy has caused me to appreciate the value of thoughtfully investigating very difficult questions rather than jumping to premature conclusions just in order to avoid uncertainty.

Part of developing thoughtful perspectives on such difficult issues involves appreciating input from multiple disciplines. In writing this paper, I have taken the view that philosophy can benefit from collaboration with scientific work. Considering scientific findings in a philosophical context—where discussions can get highly abstract—can help keep things grounded in empirical reality. I also think those coming from a scientific background should be interested to consider how some of their work might inform broader conceptual issues in religion and philosophy.

One's faith commitments might help to focus or motivate the issues in science and philosophy one thinks are most pressing. In my case, I am (at least provisionally) a theist, but I am otherwise (for philosophical reasons) a physicalist about everything in the universe, including about human persons. That is to say, I do not think we are immaterial souls; rather, we are exclusively composed of physical matter.

Being a physicalist seems to raise some novel theological and philosophical problems. In order to make sense of holding humans morally responsible for their actions, I think we should take a strong view of human agency—our ability to govern our own actions—as well as our ability to sometimes make choices freely. Christian theism furthermore involves the beliefs that humans (a) have some sort of special relationship to God and (b) might survive their deaths. Without being able to point to an immaterial soul as an easy explanation of all of these claims, I must think more carefully about the nature of the human person. Specifically, the physicalist has to consider those features which make us both continuous with and distinct from the other organisms with which we share our origins, and how our distinct features could possibly arise from our construction out of physical stuff. This is partly why I have been so interested in issues surrounding reductionism about biological systems. I think, for example, that if all of our causal powers are reducible to causal interactions among microphysical particles, this could pose a problem for human agency, which is then a problem for holding people morally accountable for their choices. This is one of the reasons why someone who is a theist and a physicalist might be concerned with some of the issues I discuss in my paper.