

Thesis Proposal: From a microbiological point of view  
Laura R. Franklin-Hall

Modern biological science has been extraordinarily successful in figuring out how the natural world works. But that science has also been extraordinarily complicated, both as a social activity and in the workings of nature uncovered. The nine essays outlined here embrace both the methodological and metaphysical tangles of contemporary microbiology, finding in them philosophical lessons about natural kinds, realism, explanation, hypothesis/experiment relations, and other topics in the philosophy of science.

Table of Contents:

PART I: KINDS AND INDIVIDUALS

*Chapter 1: Plato's Joints*

*Chapter 2: Bacterial, Sex and Systematics*

*Chapter 3: Motley Molecular Kinds*

PART II: EXPLANATIONS

*Chapter 4: Malleable Mechanisms*

*Chapter 5: Developmental Reductions*

*Chapter 6: Modules and Explanations*

PART III: EXPERIMENTS

*Chapter 7: Exploratory Experiments*

*Chapter 8: Ersatz Animals*

Part IV: Conclusion

*Chapter 9: Interests and Kinds; Explanations and Experiments*

## PART I: INDIVIDUALS AND KINDS

### *Chapter 1: Plato's Joints*

Plato's Socrates says in the *Phaedrus* that it would be wonderful to “be able to cut up each kind according to its species along its natural joints, and to try not to splinter any part, as a bad butcher might do”(265e). In the *Statesman* Plato's interlocutors make the similar suggestion that kinds should be divided from one another “limb by limb, like a sacrificial animal”(287c).

Although important to Plato for slightly different purposes, this metaphor has been famously used to illustrate the divisibility of the natural world into objective kinds—such as electrons, carbon atoms, and homo sapiens—containing objects with like essential properties. It has been thought that by dividing the world at its joints, we can lay bare these natural kinds; when we fail to divide at objective boundaries, we splinter the world's kinds like an incompetent butcher. In accordance with this metaphor, each bone in the animal body is likened to a category of things in the natural world. The claim that there is an objective, unique set of joints at which to physically separate the parts of the animal parallels the claim that there is an objective, unique set of natural categories—of natural kinds.

Some say that this metaphor is apt, and others that it is “unsavory rubbish” (Hacking 1991, 111). Less attention has been given to the source of the metaphor itself. It is this that I will explore in my opening chapter. Animal bodies are complex systems that have less straightforward osteological parts than the idealizing Plato might have imagined. I will argue that there is no unique set of bone joints in the adult vertebrate.<sup>1</sup> Things become even more confusing when considering the joints in a juvenile or in an arthropod. By looking at what joints are, what they divide, and the different ways that these divisions can be made, we can learn about the presence or absence of bona fide anatomical boundaries and of natural anatomical parts.

Plato's comment that we should try to cut at *the* joints presumably suggests that there is just one set of joints that should be used to partition the body into parts. Following Plato's lead, let us consider whether a *butcher* would actually say that animals have such a unique set of joints. In fact, there does not appear to be one good way to “joint” an

---

<sup>1</sup> I will restrict the discussion to the joints of bones for reasons of brevity. Note that the case for establishing natural joints in bones appears to be much easier than that of finding “joints” or bona fide boundaries between other anatomical parts. Bones joints are considered to be the *most* objective and stable of the natural divisions in the body (Smith and coworkers 2005). Thus, difficulty finding natural joints even there does not bode well for the objectivity of other anatomical divisions that we will not be able to discuss in depth here, such those between different parts of the digestive tract, such as between the *stomach* and the *small intestines*.

animal in the butcher shop. The following illustration shows how butchers in different countries, trained in different butchering traditions, differently carve the same cow:




Figure 1: Comparative beef butchering (adapted from CSC Research 1997)

These differences lead to substantially different meat parts at output, with no T-bone steak to be found in the British cow and no roast to be found in the American one. Although not the whole story, one reason it is so easy to multiply-carve a cow is that many butchers don't hesitate to cut *through* bone. In the case of the T-bone steak, the "T" in the steak results from a bilateral cut directly through a vertebrae (*not* a cut between vertebrae).<sup>2</sup>

Yet ancient butchers, perhaps the ones Plato had in mind, were more anatomically-minded; they did not cut through bone. Meat was removed from the bones without touching knife to bone at all (Ikram 1995, 117). Some "meat scholars" believe that ancient carving practices were different from modern ones because of the limited tools that were available:

In antiquity apparently most muscles were left intact. Even in the context of the modern, per-industrial village, it is more common for butchers without power saws to cut around bones than through them. The preparation of steaks and chops [, a modern practice] is, by contrast, at variance with anatomy, and a far simpler means of dividing a carcass is to follow the natural paths of muscles and cut them only where they need cutting, at their skeletal attachments. (Gilbert 1988).

While ancient butchers did appear to carve animals following "natural paths," it isn't clear that there was only one set of such paths to follow. As Salima Ikram (1995) observes in her account of jointing in ancient Egypt, "each country has different methods of dividing up an animal for consumption"(113). Unfortunately, most practices of the

---

<sup>2</sup>All information on modern butchering technique is from Mettler (2003).

ancient butcher remain veiled in obscurity. The only clear sense we have of these practices comes from ancient Egypt, where, to the delight of meat scholars, individual cuts of meat were separately embalmed.

Yet the prospect that ancient butchers followed natural joints and still came up with different ways of carving the animal suggests that we should look more closely at the biology of animals to make sense of what is going on and to better understand whether animals have natural parts. The second half of this chapter will consider whether there are natural joints between bones from a biological perspective, discussing whether we ought ultimately to be conventionalists, pluralists, or realists about osteological joints. Here I can only describe my general approach.

The term ‘bone’ is used by biologists both to refer to a *substance* called bone and to *individual* bones. Individual bones, as they are commonly described, are not composed entirely of the substance *bone*; they often contain cartilage and blood-cell producing bone marrow. Given this double usage, two methods of determining osteological joints might be envisioned. One method would determine what the substance ‘bone’ is, and then call a joint any bona fide boundary between one continuous bone substance and another. Alternatively, we could determine what an individual bone is, and then call a joint any location in which two bones make contact with one another. Both of these approaches will be pursued here.

The first approach will require understanding the identity of the substance biologists call bone. A potential problem is that there is no set of properties bone has that non-bone does not possess (Hall 2005, Chapter 5). Some have hoped that further investigations will uncover a molecular characterization of bone, but this has yet to be provided. After pointing out this problem, I will assume in this chapter that there is some way of characterizing bone substance, and then consider what bones and joints would be like, to a first approximation, if bones were just blobs of bone substance. Would these bones bear any resemblance to conventionally characterized bones? Would the boundaries of these bones be bona fide or fiat boundaries?

The second approach, inspired by Aristotle’s *Parts of Animals*, will consider bones to be parts of a special sort, that is, parts that are unified, functional units. This unification could lead them to be prioritized over arbitrary undetached parts. In order to explore this possibility, I will use Peter van Inwagen’s account of Aristotelian parts. I will then apply a modified version of his account to bones in order to determine whether bones should be considered to be genuine parts of animals.

Preliminary research indicates that both of these approaches fail to uniquely specify individual bones and their joints. This result should be interesting in itself. The philosophical consequences of the inquiry for a discussion of *natural kinds* (rather than natural individuals) should not be overstated; if animal joints are not objective, unique, or otherwise fail to have the properties that natural boundaries are thought to have, this might only show that Plato’s metaphor isn’t worth holding on to. My aim in this chapter is not to unseat a grand monism or thesis of strong natural kinds. It is to better understand the

anatomical world, and perhaps to suggest that believers in strong natural kinds should consider looking beyond animal joints to more appropriate metaphors.

### *Chapter 2: Bacteria, Sex and Systematics*

Biological natural kinds are of long-standing interest in philosophy (Kripke 1972; Aristotle: HA and PA) and the philosophy of biology (Mayr 1963; Hull 1992). Early on, species and higher taxa were thought to be defined by timeless essences which could be discovered through observation. The recognition by Charles Darwin and others of intra-species diversity as well as the evolution of species over time made the search for essences seem futile. Biological systematists met the challenge of evolution with an assortment of proposals for categorizing a changing natural world. As the objectivity of higher taxa began to seem implausible, research focused on *species* definitions alone. Systematists aspired to find some relation between organisms, or between an organism and its environment, that could both *demarcate* species and *explain* their cohesion. Keeping with an evolutionary perspective, these relations would not appeal to essential properties.

Finding such criteria has been difficult, but it could have been much harder. Save in a few footnotes, philosophers have exclusively concerned themselves with organisms like themselves – metazoans. However, much of terrestrial life, by far the majority measured either by mass or by census, is unicellular. Prokaryotes and protists have different lifestyles and forms of reproduction which are of relevance to a species definition. The two most popular kinds of species concepts, the biological species concept (BSC) and phylogenetic species concepts (PSC), are not applicable to many prokaryotes.

Why is this? Sex, to put things bluntly. Although not required for reproduction, bacteria exchange genetic and other hereditary materials. The extent of this transfer has only become apparent recently, and is now considered widespread, occurring not merely within species, but between genera and even over the major prokaryotic divide – archae and bacteria. The genes transferred are not always of the same families (e.g. metabolic, housekeeping, surface protein, etc.), coming from many of the thousands of genes each organism has. Exchange has led to the acquisition of fitness-enhancing traits, such as antibiotic resistance. It can also lead to inviable organisms and death.

The horizontal gene transfer (HGT) described above poses difficulties for BSC and PSC. BSC, originally proposed in Mayr (1963), defines species as sets of potentially interbreeding individuals. The thorny modal term ‘potentially’ has proven difficult to cash out. And some metazoan groups, namely ring species, have been presented as counter-examples to the concept. But those concerns aside, how does BSC fare for defining bacterial species? Not well. It cannot describe bacterial species unless we are willing to bite one of two bullets; 1) make all bacteria part of one super-species, 2) stipulate that some genes are essential, abandoning a central tenant of post-evolutionary systematics. BSC leads to one big species because any given bacterium, as an empirical matter, is potentially a mate and parent of any other bacterium, donating some hereditary material to them via conjugation, transformation, or transduction. By the standards of

BSC, they must all be considered conspecifics. To avoid this unseemly consequence, we can stipulate that some genes are essential, and require that *those* genes in particular be exchangeable for bacteria to be part of the same species. In this case, bacteria would likely be grouped in smaller kinds. This is because some genes, particularly those used in metabolism, need to be compatible with the recipient's physiology, substantially restricting the number of viable pairings.

PSC, which defines species as groups of organisms on the same (suitably cut) branch of the tree of life, also cannot be applied to bacteria without a commitment to essential genes. To apply PSC, each organism needs to be situated on the tree of life in some unique location. Ancestors will be nearer to the root; descendants will be on branches. Organisms not of the same ancestry are located on different branches. Given HGT, a bacterium cannot be uniquely situated on such a tree, because different parts of its genome have different histories. If we define these trees as graphs, they are often not isomorphic when drawn for different genes. In general, differences between the trees for gene A and gene B occur whenever there has been an event in which *either* gene A *or* gene B, but not both, was transferred. Evidence indicates there have been many such events. To obtain a *unique* tree and a unique classification, genes will have to be chosen. Some biologists have been willing to do that, claiming that the 'house keeping' genes should be used for organism taxonomies. Such a commitment is perfectly reasonable, but it comes with a price: the abandoning of an interest-free systematics.

The inapplicability of PSC and BSC to non-metazoans has been acknowledged before not *because* of sex, but for its *absence*. Papers on biological systematics distinguish between *sexual* and *asexual* populations<sup>3</sup>, proceeding to sidestep putatively asexual bacteria in favor of the sexual metazoans. As we've seen, bacteria cannot any longer be considered asexual. But nor do they have the sexual regularity required for the reasonable application of accepted species concepts. So until an alternative is provided, this leaves the field of microbiological classification open to—gasp—pragmatic considerations.

### *Chapter 3: Motley Molecular Kinds*

Just as naturalists have partitioned the world of animals into species, and chemists have partitioned the atoms into elements, molecular biologists have recently produced classification systems for biologically important molecules. Focusing particularly on proteins, this chapter examines molecular classifications based on structural, functional, and genealogical standards. It considers how these classifications are different from systems in other domains, explains how they are related to one another, and asks whether they track natural kinds.

First I will provide some examples of the classification systems under consideration. The Gene Ontology Consortium has used three methods of protein/gene classification, in

---

<sup>3</sup> This dichotomy is assumed in (Ehrlich 1992; Valen 1992; Wiley 1992) and many others. In (Mayr 1992) the dichotomy is between *biparental* and *uniparental* reproduction.

terms of molecular function, biological process, and cellular location. On the other hand, the Protein Families database (Pfam) and Protein Fingerprints (PRINTS) provide classifications based on protein primary sequence. The Structural Classification of Proteins (SCOP) project has established classifications based on secondary and tertiary structures.

On the traditional picture, natural classifications that partition the same domain are hierarchically related, such that higher level classes can be defined in terms of disjunctions of lower level kinds, as is the case with metals and elements (Elder 1994). Interestingly, it seems that these protein classification systems are not hierarchical, but instead cross-cut one another. This even holds within the family of “structural” classifications. While the incommensurability of functional and structural classifications is no surprise, one might expect broadly structural classification to be relatively orderly. However, there doesn’t seem to be a basic structural classification, on which other structural classifications can be defined. As an example, consider two classifications, one that partitions proteins based on similar geometries, and the other based on similar amino acid sequences. Neither of these can be used to extensionally capture the species-level classifications of the other. Proteins with identical primary sequences can take a variety of secondary and tertiary folding patterns (Wadsworth and coworkers 2003). Conversely, geometrical classifications, which attempt to trace out the electron-boundary of a folding protein, can sometimes classify different primary sequences in the same geometrical kind because of similar higher-level folding. Thus, at least among the classification systems now available, there is no bottom-level on which other classifications can be defined.

A look within each individual classification system reveals additional problems for the hierarchical conception. Gene Ontology classifications, for example, do not require that lower level categories have only one parent class. Consider the GO classification based on biological process. Cytochrome C is classified under the biological process of *oxidative phosphorylation*, as well as *cell-death induction*. Multi-parent classification is allowed in secondary structural classifications as well; proteins are often members of multiple classes depending on the different folding patterns found along the length of the molecule. In this way, proteins are seen as charm-bracelets – each has a different set of charms, and is classed according to each. This sort of system is useful when each “charm” is linked to properties that are important in understanding behavior of the object. In such cases, there is no single set of Kripke-like essential properties from which all other important properties flow.

Once we have clarified the structure of each classification system, and the non-hierarchical relations between them, we can consider whether there are any metaphysical implications of the discussion. While the systems under consideration were constructed with different purposes in mind (e.g. drug discovery, cell modeling, genetic networks), and with different breadths (e.g. enzymes, mouse proteins, all known proteins, etc.), many customary criteria of natural kinds—such as their use in explanatory generalizations and the projectibility of their positive instances—still seem to apply. For example, secondary and tertiary structural classifications have been used to infer the therapeutic activity of potential drugs on known chemical targets. Objects partitioned in

functional classifications based on domain-binding abilities are used in explaining the behavior of gene regulatory networks. Because these classifications satisfy many (although not all) traditional desiderata for natural kinds, they should not be considered merely nominal kinds. If these kinds are not nominal, but ‘natural’ in some sense, we ought to be metaphysical pluralists about this domain.

John Dupré (1993) has tried to show that biological categorizations are cross-cutting and non-hierarchical by examining species-level classifications. The argument here, based on protein classification, is not subject to a common criticism of the Dupré argument. Dupré’s examples of cross-cutting kinds come from juxtaposing natural language categorization and scientific categorizations (e.g. *Opuntia* (p. 27-28), *Lepidoptera* (p. 28), cedar (p.35)). Some think that this has no bearing on the pluralism of purely *scientific* classification. Obviously scientists and cooks have different interests, and thus focus on different sets of properties when establishing classifications. But science itself, the critics say, will reveal the scientifically important properties and objective kinds. Biomolecular classification are thus a better example of pluralism, since it reveals the disparate classifications and correlated interests within even a small scientific community.

## PART II: EXPLANATIONS

### *Chapter 4: Malleable Mechanisms*

Biologists often claim to be investigating the mechanisms that they hope will explain biological phenomena. Not surprisingly, they rarely state precisely what a mechanism is. In this chapter I will first explain philosophical views of biological mechanisms. Then, I will consider the applicability of these different conceptions of mechanisms to a putative mechanism, the membrane “pump”,  $F_0F_1$  ATP synthase. Despite the use of mechanistic language when describing this system, most formal conception of mechanism are inapplicable. Finally, I will examine an interesting debate among cell biologists about whether we ought to have a broadly mechanistic approach to the cell in the first place. This I think will reveal that most accounts of mechanisms leave out an important part of what biologists mean when they use the term.

Somewhat un-illuminating descriptions of mechanisms abound: “a mechanism is simply a causal process”(Fehr 2002), or “mechanisms are entities and activities organized such that they are productive of regular changes”(Machamer and coworkers 2000). Others are inapplicable to biological systems, such as Salmon’s view that a mechanism must involve the transfer of a conserved quantity. The first step in this project will be to review the various conceptions of mechanisms articulated by philosophers (e.g., Bechtel and Richardson 1993, Craver 2002, Darden 2002). I will consider what kinds of causal and explanatory claims go along with these conceptions. Do they assume that causation is only happening at the lowest-level? Do they insist that mechanisms be robust to environmental variation? Do they claim that mechanisms function independently of one another?

I will proceed to examine a particular mechanism from membrane biology, the working of the  $F_0F_1$  ATP synthase. To explain the behavior of many membrane proteins, including this one, biologists explicitly call upon metaphors of motors and pumps, classic artifactual mechanisms. Interestingly, the actual functioning of the molecular ‘motor’ that drives the ‘pump’ is understood at the intra-molecular level, where traditional attributes of mechanisms (e.g. such as being made of separate, impenetrable parts) completely break down. Do standard concepts of mechanism even apply to this case? If not, does this suggest that mechanistic language is heterogeneous, with no set of properties applying to all putative biological mechanisms? If so, it could be interesting to examine what role, if any, the concept of a mechanism plays in molecular motor research. It is possible that mechanistic language, even if vague or ambiguous, helps biologists conceptualize their objects of research, and design experiments that interfere with them in interesting ways, as has been claimed is the case for the putatively vague gene concept (Rheinberger 2000).

In spite of the sometimes broad use of mechanistic language revealed above, biologists do at times distinguish between mechanistic and non-mechanistic explanations. Examining such a situation might be instructive for those trying to figure out what concept lies behind this ever-popular word. To that end, I will look at a long-running debate in the cell biology community about the best way to conceptualize the cell. There are two general approaches to cell physiology, one that we can call the ‘mechanistic orthodox view’, and the other the ‘phase-transition heterodox view’. On the orthodox view, the cell is seen as a container of aqueous solution encased in a semi-permeable plasma membrane laden with pumps and channels. These channels act for the most part independently of one another, and they specifically enable the cell to carry out certain tasks – signaling, solute transport, depolarization, etc. Whenever scientists need to explain the presence of some molecule in the cell, they can posit a pump or channel (a “special widget” according to critics (Pollack 2001, p. 97)) able to transport it as needed. These supposedly molecule-specific passage-ways have even been used to account for the concentrations of molecules that cells don’t encounter in their natural environments.

The heterodox view<sup>4</sup>, on the other hand, which has been promoted by Gilbert Ling (1962, 1984, 1992, 2001) and most recently by Gerald Pollack (2001), imagines the cell encased in a highly permeable membrane, and filled with a gel-like substance, not with bulk-phase water familiar to us from the bathtub. Instead, polar water molecules are intricately

---

<sup>4</sup> I cannot overemphasize the extent to which this is a minority view; it is not taught in universities, and even if its mentioned by leaders in the field, it is treated with disgust. On the other hand, it has survived for 40 years, and continues to have proponents, albeit often in bioengineering or fluid-flow departments rather than in biology departments. So I am not examining it because I think that it is right (although I do find some of the arguments for it to be intriguing), but because I think that the dispute can show us more of what scientists mean when they talk about mechanisms, and also illuminate the background commitments of present day biologists, presuppositions that shape how they conduct their research. If one is raised in a traditional cell theory environment, it can actually become difficult to see what a non-mechanistic, yet non-vital force explanation would even look like.

organized in layers by charged proteins within the cell. Distributed effects of this pervasive gel-substance are said to explain many of the same phenomena that pumps have been posited to explain, such as the low intra-cellular sodium concentrations in resting muscle cells, and the high sodium concentrations after excitation. These explanations don't call upon particular localizable protein objects, responsible for each physiological effect. Instead, vast webs of protein and fluid, which can undergo rapid phase-changes when their environment (pH, temperature, solvent concentration) is changed only slightly, explain cellular behavior. Proponents find this exciting because they feel their explanations are more unified or parsimonious than those of the orthodox. They also don't provide particular mechanisms to explain distinguishable aspects of cellular behavior, which at times seems unsatisfying. Reviews of this theory by traditional biologists indicate that they not only disagree that phase-transitions provide the best explanation for cellular behavior, but also deny that intelligible explanations are being given at all, claiming that the view is full of "disinformation" and "nonsense"(Miller 2001).

It is admittedly possible that biologists think the theory is non-sense because it is non-sense. However, my study of the theory indicates that it makes sense (albeit it is possibly wrong), but it uses different kinds of explanations than those most cytologists are used to. The heterodox view explains cellular events via a bulk process that accounts for many effects at once. As is sometimes complained, these explanations do not provide explicit mechanisms. What is meant by a mechanism here? As this project is still in its initial stages, I am not yet prepared to give a full characterization. My intuition is that when biologists point out mechanisms, they want the chain of causes to be somewhat independent of the chains that explain other phenomena in the system. Chain independence is not usually required by mechanistic conceptions, such as those of Craver and Darden. When biologists look for mechanisms, it isn't just that they are looking for explanations in terms of the sequential contact of molecules with one another, but they are looking for effects that are partitionable and independent (assuming certain background conditions are held fixed). There are of course evolutionary reasons to have this preference. We can more easily imagine how a complex system could have gotten put together (either by an intentional agent or by natural selection) if the parts of the system responsible for individual functions could have been independently constructed.

### *Chapter 5: Developmental Reductions*

The contemporary discussion of developmental reduction began after Alex Rosenberg took to heart an intriguing suggestion by the developmental biologist Lewis Wolpert. In his [1994], a two page opinion piece in *Science*, Wolpert considers whether we might one day be able to 'compute the embryo', to predict all changes in embryonic morphology using only the properties of individual molecules and their interactions. Rosenberg embraces this possibility, and hints with some well-chosen examples that we are on our way to such an understanding. Wolpert himself is less optimistic. He worries that tracking the movements and interactions of all cellular constituents will prove impossible. On his view, computing the embryo will be feasible only "if a level of complexity of

description of cell behavior can be chosen that is adequate to account for development but that does not require each cell's detailed behavior to be taken into account"(572)<sup>5</sup>.

Laubichler and Wagner (2001) responded to Rosenberg's claims by pointing out some developmental phenomena that they did not think had been explained by molecules and their interactions. For example, they claim that it is the functional context of the engrailed pathway that explains how it can establish both the anterior-posterior compartment boundary in *Drosophila*, and also the eye-spot organizer in the butterfly. Frost-Arnold (2004) then argued that once we are allowed to spell out the cellular context in molecular terms, these examples do nothing to impugn the reductionist position. He suggests that the only foreseeable scenario in which developmental generalizations will be irreducible to molecular interactions would be if the angle of the gravitational field to the early embryo is causally relevant to the development of that embryo. Since the orientation of the embryo to gravity cannot be spelled out in purely molecular terms, the positioning of the embryo's major axis would then be non-reducible.

Following Sarkar, all parties agree that, because biological theories aren't the axiomatizable systems originally envisioned, questions about the success of reductionism in biology are about the success of *explanatory* reductionism, a practice that I will follow here. With this in mind, we can identify three problems with the recent discussion of developmental reductions to molecular terms; taking them seriously leads us to be less optimistic about the ability of a structural molecular biology, now or in the future, to fully explain developmental regularities.

Problem 1: The reducing theory has been insufficiently or misleadingly characterized in discussions of developmental reduction. The following quote from the most recent in a series of papers on the reducibility of developmental biology admits that the reducing theory has not been a topic of discussion:

A reductionist about development would have to maintain that the mechanisms or rules that govern interactions between molecules are sufficient to derive or 'explain' in some sense the regularities at the developmental level of description (in combination with the F-terms). *What sorts of 'mechanisms' are these in the case of molecular biology? I will not go into detail about this, since neither Rosenberg nor Laubichler and Wagner do.* [my italics](Frost-Arnold 2004)

The author admits that neither his account, nor that of his predecessors, deal with the actual mechanisms in molecular biology used to explain developmental facts. Instead, he appeals to a general account of molecular reduction from Sarkar (1998) in order to understand what the reduction base consists in. Sarkar takes the reducing realm to be what he calls 'macromolecular physics' (p. 136), described by the following four rules: i) Weak interactions rule: "the interactions that are critical in molecular explanations are very weak"(p. 149). ii) Structure determines function: "the behavior of biological

---

<sup>5</sup> Thus it is curious that many have taken Wolpert's program to support the reduction of developmental biology to *molecular* terms.

macromolecules can be explained from their structures”(p. 148). iii) the importance of molecular shape: “these structures, in turn, can be characterized entirely by molecular size and, especially, shape, and some general properties (such as hydrophobicity) of the different regions”(p. 149). iv) Lock-and-key fit for molecular interactions: “molecules such as protein molecules forming larger structures or enzymes interacting with substrates interact when there is a lock-and-key fit between the two surfaces”(p. 150).

It is no doubt true that some parts of developmental explanations refer to molecular shapes and their binding (such the Helix-Loop-Helix structure of the homeobox binding domain), but mention of occasional molecular interactions does not a molecular explanation make. Although they refer to molecules at certain junctures, explanations of many of the examples given are not in molecular terms, but rather in terms of functionally characterized genetic switches. Developmental biologists use one kind of objects ad nauseum to explain the progress of development: transcription factors. “Master genes” (e.g. *Pax-6*, *Dlx-1*) are supposed to function by setting off a cascade of gene transcription events in which the protein product of one gene is the transcription factor for another. Some protein products of master genes are transcription inhibitors, and a ‘logic’ is being developed to take into account the combined actions of a series of transcription factors at a given promoter, to determine the transcriptional consequences of prior protein expression. These explanations and descriptions in terms of a logic, exactly the same as a circuit diagram, pick out their objects of interest by the roles they play in expression networks and not by their molecular identities. Of course, in any given case a transcription factor does have a molecular constitution. But there isn’t anything all transcription factors have in common except that they cause gene transcription. Inasmuch as genes themselves can only be functionally defined, a molecular reduction is even less promising (Kitcher 1982; contra Waters 1994; Kincaid 1997).

Problem 2. A second problem with the reductionist account concerns the developmental explanandum. Rosenberg, followed by other authors, rephrases the question of *explaining organismal development* into the question of *being able to predict the molecular changes in a particular embryo*. That is, given we have information about the location and constitution of every molecule in a particular zygote, can we predict all subsequent molecular locations? This involves abandoning any explanation of regularities in the development of all members of a species, let alone all animals or plants. This is a rather ironic narrowing of the explanation of interest, given recent excitement about producing explanations for the development of all vertebrates, or even all multi-cellular organisms, based on evidence that there are universal patterns in development, both those that are, and aren’t, accounted for by common history (Carroll 2005). But once the explanandum is changed in this way, it seems that the reductionist wins from the beginning – any physicalist would agree that a token explanation of molecular location can be given in terms of past molecular location and (still undiscovered) rules which govern interactions between molecules.

Problem 3. Finally, even if we accept Rosenberg’s presentation of the explanandum, we can still question the claim that biologists will ever model embryonic change at the

molecular level<sup>6</sup>. Wolpert himself thinks the super-cellular level might be the best level for such modeling attempts. There has not yet been any success in creating purely molecular computational models of developmental events. The best developmental computational models around are couched at the cellular level or higher (Sumar and Bentley 2003). These, at times remarkably predictive models, actually suggest that there is something important about these higher levels, such that they are sufficient to predict morphological change during development. Why this might be the case is addressed in more detail in the following chapter.

### *Chapter 6: Modules and Explanations*

The practice of biology (and science in general) depends on the incorrectness of both Hume and Leibniz. A naïve (although frightfully common) reading of Hume has him believing that the observable world is “loose and disconnected”. If there weren’t real causal dependencies between observables, all would be lost for science. Leibniz, on the other hand, imagined that all monads “express” the states of all other monads in the world. Hence, to understand even one object, all others must be understood. Neither this hyper-connected world, nor the unfortunately disconnect world of Hume, does the experimentalist much good. If each object effects all others, to understand the dynamics of the universe, the number of interactions to keep track of would go up exponentially with the number of objects, even assuming that all relevant interactions are pair-wise.

Of course, Leibniz was right that everything in the universe does effect everything else. As Herbert Simon has pointed out, early physicists were able to get a grip on the solar system only because they were lucky enough to live in a planetary system seemingly arranged to be easy to figure out, given the gravitational laws that govern it: the sun is \*much\* more massive than its satellites, there aren’t any other really massive objects near by, and the satellites of the planets are quite close to them relative to their distances to the sun. The fact that the basic forces all fall off quickly with distance is also rather useful for many explanatory tasks.

Biological systems are not generally understood in terms of the basic physical laws. Yet pseudo-Leibnizian problems still lurk. It’s very hard to discover a cure for asthma when you not only have to figure out how the bronchial tubes contract because of pollen-related immune reactions, but you also must worry about how hearing the word “wheeze” effects the patient’s hormone levels, which then modulate her bronchial constrictions. But there

---

<sup>6</sup> The molecular level might be unsuitable to such calculations due to variability between molecularly indistinguishable molecules (indistinguishable when defined in terms of their component atoms and their topological connections). Doinchi and coworkers (1996) found that when they measured the enzymatic activity of individual molecules, different tokens of the same molecule as defined by molecular formula varied in activity over a ten-fold range. This indicates that a specification at the chemical level might still be inadequate to the task of fine-grained prediction, since it leaves underdetermined relevant properties about each individual molecule. These differences might completely wash out at the cellular level.

is increasing evidence that biological systems are not generally Leibnizian – there are relatively independent causal systems, called modules, whose activities buffer or isolate the system from certain other systems within the organism. This applies not only to physiological processes, but also to those of development. The parts of a module don't have to be located in the same section of an organism – some immune processes are modular, and the tight causal connections between the objects that participate in these processes is based on chemical specificity. It is also important to remember that not all organism parts are modules: the left side of Andrew is not a module of his body.

Biologists have made various suggestions for how to define these modules: “modules are building blocks of interacting elements that operate in an integrated and relatively autonomous manner”(Schlosser & Wagner, 2004, 520). Others have thought that biological modules are emergent entities that should be defined epistemically: “Modules are composed of many types of molecule [sic]. They have discrete functions that arise from interactions among their components (proteins, DNA, RNA and small molecules), but these function cannot easily be predicted by studying the properties of the isolated components”(Hartwell and coworkers 1999). I will begin this chapter by considering these different accounts.

I would like to find a way to characterize modules which will account for the obvious proclivity of biologists to explain biological processes at particular levels of granularity, and not at others. Higher-level descriptions of systems might not be preferable merely because they are simple enough for our little minds to grasp, but because they capture objectively self-contained causal systems. Bob Woodward, Daniel Dennett and Art Garfinkel, and Philip Kitcher have also attempted to explain why some higher-level generalizations or patterns can be explanation-worthy, and I intend to consider their views in the process of developing my own.

The existence of biological modules was not inevitable, and their ubiquity has come as a surprise<sup>7</sup>. In addition to helping us understand whether there are preferred levels of explanation, an investigation into biological modules should also be interesting because of other roles they are beginning to play in biological theory. The modularity of a system can sometimes be used in explanations of evolutionary change or stasis. For example, the variable modularity of development processes is thought to explain the stability of the phylotypic stage within an organism lineage. The processes of early development, it turns out (contra Haeckel) in many phyla has substantially diverged, as have processes of late development, but a certain stage in the middle of embryonic development involves a

---

<sup>7</sup> If there hadn't been modules, and thus if there hadn't been reasons to decompose biological systems in one way rather than in another, molecular reductionists might have been right that the best explanations would have been at the molecular level. If the causal web were not unified at levels below that of the whole organism, and all parts of the organism regularly interacted with all other parts, then the only way of understanding what was happening would be to provide a model that included all molecules in the system.

particularly large amount of interaction between all parts of the embryo. Evolutionary tinkering at this stage could have disastrous consequences on the entire organism. Thus, this stage has maintained its structure much more than have other stages in which interactions are more local (Raff 1996, chapter 6 and 10).

### PART III: EXPERIMENTS

#### *Chapter 7: Exploratory Experiments*

Philosophers of experiment have acknowledged that experiments are often more than mere hypothesis-tests, once thought to be an experiment's exclusive calling. However, we still lack adequate characterizations of the alternative roles that experimentation can play in scientific inquiry. Drawing on examples from contemporary biology, I here characterize circumstances in which *exploratory experimentation* can be fruitful. Conceptual arguments, examples from genomics research, and a computer modeling project [under development], all indicate that there are circumstances in which experimentalists ought to leave precise theorizing behind and conduct experiments about which no predictions can be made because no hypotheses have been constructed.

Philosophers have paid close attention in recent years to the details of scientific experiment (Hacking 1983; Franklin 1986; Galison 1987; Franklin 1990; Steinle 1997). They have noticed various roles that experiments play in scientific investigation besides the traditional Popperian one of putting scientific theories to the test. For example, Peter Galison (1987) emphasizes that much of a scientist's time is spent carrying out experiments whose main purpose is to establish that some piece of experimental apparatus is functioning as desired. Friedrich Steinle (1997) notes that during periods of theoretical turmoil, experimentation can move into an exploratory mode in which scientists do not test theories. Instead, they conduct undirected experiments whose results might be used to inspire an entirely new theory or research program.

Despite these helpful emendations to (and at times rejections of) the Popperian picture, important contemporary experimental activity has been neglected by philosophers interested in experiment. Modern biological techniques, like DNA microarrays and proteomics, allow scientists to experiment in an exploratory mode even in times of normal science. Exploratory experimentation is feasible, and desirable, as a result of "high throughput"<sup>8</sup> data collection and analysis technologies, available only in the past ten years. It is interesting to examine how scientific methodology changes as the technologies used in data collection and analysis themselves change. New methods don't merely allow scientists to measure new *kinds* of things, but they also change the ways they ask questions and approach experimentation more generally. As a consequence, the conditions are now ripe to reconsider the attractiveness of a neo-Baconian scientific picture. Many of the circumstances that made a data-driven approach descriptively incorrect and normatively unattractive have altered.

---

<sup>8</sup> However ugly, this is the established term from the biological literature.

DNA microarrays, an example of high-throughput technology, allow simultaneous quantitative measurement of the RNA produced by every ORF in a cell. Pioneered by Patrick Brown's lab at Stanford in the late 1990's, the technique exploits inexpensive inkjet printer technology and cell purification techniques that are widely available, and has consequently become enormously popular in the microbiological community. The microarrays are two-dimensional grids of oligonucleotides, usually 20mers, which have been spotted on small glass slides. Each 20mer is a unique substring of a particular coding sequence. mRNA produced by cells in two conditions (experimental and control) is translated into differently colored cDNA and added to a microarray. If a given cDNA is complementary the oligo on a particular spot, it will preferentially hybridize to it. Pictures are taken of the slides when illuminated by light of the appropriate frequency, revealing the relative amount of RNA of each type expressed by the cells. The amount of light reflected from each spot corresponds to the expression level of one gene. Scientists compare results in the control and experimental conditions to determine the effects of some intervention on expression.

Microarrays can contain up to 25,000 spots, far more than the number of genes in many organisms, and approaching the number of genes in mammals. Consequently, the entire 'transcriptome' can be surveyed at one go. Prior to DNA microarrays, scientists could measure the expression of a given gene using the Southern blot. This could be done with accuracy for one, or at a most a few, genes at a time. Similar transformations in data production have occurred for other important molecular measurements, such as protein composition, glycosylation, post-transcriptional modifications of mRNA and others.

How has experimental practice changed since microarrays and like technologies have become widely available? Exploratory experimentation has become more common. I consider experiments exploratory when the experiment is not intended to test a hypothesis. This is most obvious when scientists claim that they don't have any hypotheses about the values they will measure. Scientists talk of 'going fishing' in the laboratory, looking for *some effect* of their favorite intervention. Case studies of classic microarray experiments, such as DeRisi et al. (1997), show how scientists can design experiments to test the effect of an intervention without predicting what this effect might be. Such a strategy was implausible in the past because any given intervention has only a limited range of effects. It is important for scientists that their interventions produce some measurable effect. When only a small amount of data can be collected, scientists carry out experiments in which they had reason to think an effect *will* be obtained. That requires sticking close to an established theory in a way not needed when all conceivable dependent variables can be measured.

Of course, hypothesis and theory still play some role in exploratory experimentation. We can find three roles of hypothesis in experiment, only one of which is modulated by high-throughput technology. First, hypotheses and theories are required to interpret data acquired with the help of complex instrumentation. Second, background theories and research programs select the objects and processes that are considered important and worthy of inquiry (Kuhn 1970). They also provide general knowledge about how those important objects and processes work. For example, the central dogma of molecular

biology, even in some modified form, is a presupposition of microarray experiments, and the general framework of contemporary biology holds that measuring protein and mRNA levels would be a useful experimental approach. Third, hypothesis directs research when scientists design experiments to directly test hypotheses. This is the hypothesis/experiment relation HT technology has shown to be contingent. Of course, experiments *are* often intended to test hypothesis, and that's good. However, *sometimes* this is a matter of convenience, not a matter of logic. I suspect that wide-spread use of hypothesis to direct experiment has been misinterpreted as evidence for a falsificationist logic<sup>9</sup>. The close coupling between hypothesis and experiments can be better explained as a pragmatic requirement of low-throughput experimentation in a scientific community that doesn't value negative results.

### Chapter 8: Ersatz Animals

The use of animals<sup>10</sup> as experimental subjects, although thought to be unavoidable, is not often welcomed. Utilitarians such as Jeremy Bentham have held that the suffering of animals should have some moral weight even if it is frequently out-weighed by benefits that follow from it. St. Augustine believed that causing an animal to suffer unnecessarily was wrong because of the effect that causing this suffering had on the *agent* causing the pain. In short, although there is no consensus on exactly what suffering *is* unnecessary, it is virtually universal to believe that unnecessary suffering should be avoided.

In the light of this moral cost and the need to conduct animal experiments for scientific reasons, many university ethics guidelines, and organizations such as PETA,<sup>11</sup> have urged scientists to instead “experiment on”<sup>12</sup> computer models of animals or animal organs which simulate animal physiology and behavior (Biever 2006). According to the Howard Hughes Medical Institute, “researchers should reduce the animals needed to the minimum number to prove the research valid, that they *should replace animals with cell cultures or computer models when possible.*”<sup>13</sup>

In this chapter I will explore the ethical implications of such substitutions. Is experimenting on a computational model any more ethically acceptable than experimenting on fluffy? The answer to this question obvious seems to be a hearty *yes*, and I do not aim here to convince you otherwise. However, I will probe whether there are

---

<sup>9</sup> “The theoretician puts certain definite questions to the experimenter, and the latter, by his experiments, tries to elicit a decisive answer to these questions, and to no others. All other questions he tries hard to exclude”(Popper 1963).

<sup>10</sup> For reasons of brevity, I will be using *animals* to mean *non-human animals*. This should not be taken to suggest that I am denying our animal nature.

<sup>11</sup> <http://www.petaindia.com/cexp.html>

<sup>12</sup> There is an interesting literature on whether it is even possible to “experiment on” computer models which I will only be able to address in passing here. We will have to assume that it is possible to “learn something” from in silico experimentation for such substitution projects to get off of the ground.

<sup>13</sup> <http://www.hhmi.org/research/bioethics/animal/>

any moral costs to such experiments, even if they do not compare that those that result from live-animal experiments. As in other circumstances in which the use of emerging technologies could alter the ethical calculus, it is worthwhile to ponder the *potential* ethical consequences of a proposed course of action even before it is commonly pursued or even possible to pursue. I intend this essay to broach the issue of the ethics of in silico experimentation and clarify a number of potential responses to it.

The first step in evaluating the ethical implications of in silico experiments is to determine what they are. I will characterize the different kinds of computational models under development and consider how they are used to carry out “experiments”. I will then use the literature in strong A-life and strong AI to consider whether any of these models has the potential to be considered either *alive* or *conscious*. Is there any reason to think that these “simulations” of living things—or thinking things—are living or thinking themselves?

Based on this analysis, I will survey ethical approaches to the value of animals, and determine which, if any of these theories would give ersatz animals moral weight. Utilitarian theories of animal welfare, such as that articulated by Peter Singer, might well consider these experiments to be ethically relevant. On the other hand, theories which value animals based only on their membership in endangered species would not.

#### PART IV: CONCLUSION

##### *Chapter 9: Interests and Kinds; Explanations and Experiments*

Discussions in this thesis address interest- or context- sensitivities in the individuals, natural kinds, explanations and experimental techniques found in the work of contemporary biologists. This chapter will concentrate on the theoretical underpinnings of this pluralism. First, I would like to examine whether natural-kindness is itself a natural kind, a topic broached by Witmer and Sarnecki (1998). Can we define what it is to be a natural kind a priori, or must we somehow figure out what natural-kindness is from our scientific practice? Could the a posteriori nature of, and difference of opinion about, natural kindness trickle down, leading to pluralism about the natural kinds themselves? A related question concerns when multiple classification systems should be taken as evidence for promiscuous realism about these domains, as Dupré or Kitcher would have it, or merely an instrumentalism about biology in general, as Rosenberg might conclude. Finally, I need to address a possible tension between the morals I draw about modular explanation and those about natural kinds. I think there is reason to believe that development is somehow “best” explained at the level of certain developmental modules. Is this sort of preference consistent with my pluralism overall? I don’t think so, although this will surely depend on how I fill out the account modular explanation.

## References

Aristotle (1984). The Complete Works of Aristotle: the revised Oxford translation. Princeton, NJ, Princeton University Press.

Glennan, S. (2002). Rethinking mechanistic explanation, Philosophy of Science **69** (2002) (Suppl.), pp. S342–S353.

Carroll, S. (2005) Endless forms most beautiful: The New Science of Evo-Devo. New York: Norton & Co.

CSC Research (1997) "Successful business componentisation implies finding the natural joints," <http://lef.csc.com/foundation/library/comp/RP06.asp>.

Culp, S. (1989). "Theory Structure and Theory Change in Contemporary Molecular Biology." British Journal of the Philosophy of Science **40**: 459-483.

Darden, L. (2002). Strategies for discovering mechanisms: Schema instantiation, modular subassembly, forward/backward chaining, *Philosophy of Science* **69** (2002) (Suppl.), pp. S354–S365.

DeRisi, J. e. a. (1997). "'Exploring the Metabolic and Genetic Control of Gene Expression at a Genomic Scale.'" Science **278**: 680-686.

Dupre, J. (1993). The Disorder of Things: Metaphysical Foundations of the Disunity of Science. Cambridge, MA, Harvard University Press.

Ehrlich, P. (1992). Differentiation of Populations. The Units of Evolution: Essays on the Nature of Species. M. Ereshefsky. Cambridge, MA, MIT Press.

Elder, Crawford L. (1994) "Higher and Lower Essential Natures", American Philosophical Quarterly 31:255-265.

Fodor, J. (1983). The Modularity of Mind: An Essay on Faculty Psychology. Cambridge, MA: MIT Press

Franklin, A. (1986). The neglect of experiment. Cambridge, Cambridge University Press.

Franklin, A. (1990). Experiment, right or wrong. Cambridge: Cambridge University Press.

Frost-Arnold, G. (2004). "How to be an Anti-Reductionist about Developmental Biology: Response to Laubichler and Wagner." Biology and Philosophy **19**(1): 75-91.

Galison, P. (1987). How Experiments End. Chicago, University of Chicago Press.

Gehring, W. J. (1998). Master Control Genes in Development and Evolution. New Haven, Yale University Press.

Ghiselin, M. (1992). A Radical Solution to the Species Problem. The Units of Evolution: Essays on the Nature of Species. M. Ereshefsky. Cambridge, MA, MIT Press.

Giere, R. N. (1999). Science without Laws. Chicago, University of Chicago Press.

Giere, R. N. (1999). Using Models to Represent Reality. Model-Based Reasoning in Scientific Discovery. L. Magnani. New York City, Kluwer Academic.

Gilbert, A. S. (1988) "Zooarchaeological observations on the slaughterhouse of Meketre" *Journal of Egyptian Archaeology*, 74: 69-90.

Hacking, I. (1983). Representing and Intervening. Cambridge, Cambridge University Press.

Hacking, Ian (1991) "A tradition of natural kinds," Philosophical Studies 61:109-126.

Hall, Brian K. (2005) *Bones and Cartilage: Developmental and Evolutionary Skeletal Biology*, Elsevier Press: San Diego, CA.

Heyer, L. J. (2003). Discovering Genomics, Proteomics, and Bioinformatics. San Francisco, CDHL Press.

Hartwell, L. J. Hopfield, S. Leibler, and A. Murray (1999). "From molecular to modular cell biology". Nature 402: C47-52.

Hull, D. (1992). A Matter of Individuality. The Units of Evolution: Essays on the Nature of Species. M. Ereshefsky. Cambridge, MA, MIT Press.

Ingber, D. (2001). "Putting the cell biology establishment on the stand". Cell 109: 688-689.

Kincaid, H. (1997). Individualism and the unity of science: essays on reduction, explanation, and the special sciences. Lanham, MD, Rowman & Littlefield.

Kitcher, P. (1982). "Genes." British Journal of the Philosophy of Science **82**: 337-359.

Kripke, S. A. (1972). Naming and Necessity. Cambridge, MA, Harvard University Press.

Kuhn, T. S. (1970). The structure of scientific revolutions. Chicago, IL, University of Chicago Press.

Kumar S. and P. Bentley. (2003) On Growth, Form and Computers. Amsterdam: Elsevier Academic Press.

Laubichler, M. D. (2001). "How Molecular is Molecular Developmental Biology? A Reply to Alex Rosenberg's Reductionism Redux: Computing the Embryo." Biology and Philosophy **16**(1): 53-68.

Ling, G. (1962). A Physical Theory of the Living state: The Association-Induction Hypothesis. New York: Blaisdell.

" (1984). In Search of the Physical Basis of Life. New York: Plenum Press.

" (1992). A Revolution in the Physiology of the Living Cell. Malabar, FL; Krieger Publishing Company.

" (2001). Life at the Cell and Below-Cell Level: The Hidden History of a Fundamental Revolution in Biology. New York: Pacific Press.

Longino, H. E. (2002). The Fate of Knowledge. Princeton, NJ, Princeton University Press.

Mayr, E. (1963). Animal Species and Evolution. Cambridge, MA, Harvard University Press.

Mayr, E. (1992). Species Concepts and Their Application. The Units of Evolution: Essays on the Nature of Species. M. Ereshefsky. Cambridge, MA, MIT Press.

Mettler, John J. (2003) *Basic Butchering of Livestock & Game*. Storey Publishing: North Adams, MA.

Miller, C. (2002). Book Review: Cells, Gels and the Engines of Life. The Physiologist 45: 511-514.

Pollack, G. (2001). Cells, Gels and the Engines of Life. Seattle, WA: Ebner & Sons.

Popper, K. (1963). Conjectures and Refutations: The Growth of Scientific Knowledge. London, Routledge.

Quine, W. v. O. (1969). Natural Kinds. Ontological Relativity and Other Essays. New York, NY, Columbia University Press.

Raff, R. (1996). The Shape of Life: Genes, Development, and the Evolution of Animal Form. Chicago, IL, Chicago University Press.

Rheinberger, H. J. (2000). "Gene Concepts: Fragments from the Perspective of Molecular Biology". in The Concept of the Gene in Development and Evolution. eds., Beurton, Falk and Rheinberger. Cambridge: Cambridge University Press.

Rosenberg, A. (1997). "Reductionism Redux: Computing the Embryo." Biology and Philosophy **12**: 445-470.

Ikram, Salima (1995), *Choice cuts: meat production in ancient Egypt*, Peeters Press: Leuven, Belgium.

Sarkar, S. (1998). Genetics and Reductionism. Cambridge, Cambridge University Press.

Steinle, F. (1997). "Entering New Fields: Exploratory Uses of Experimentation." Philosophy of Science **64** (Proceedings): S65-S74.

Stossel, T. (2001) "Cell Biology: Manifesto for a Cytoplasmic Revolution". Science **293**: 611.

Valen, L. V. (1992). Ecological Species, Multispecies, and Oaks. The Units of Evolution: Essays on the Nature of Species. M. Ereshefsky. Cambridge, MA, MIT Press.

Wadsworth, Jonathan D.F., Andrew F. Hill, Jonathan A. Beck, and John Collinge (2003). "Molecular and clinical classification of human prion disease", British Medical Bulletin, **66**: 241-254.

Witmer, D. Gene and John Sarnecki (1998). "Is Natural Kindsness a Natural Kind?" Philosophical Studies **90**: 245-264.

Waters, K. C. (1994). "Genes Made Molecular." Philosophy of Science **61**(2): 163-185.

Wiley, E. (1992). The Evolutionary Species Concept Reconsidered. The Units of Evolution: Essays on the Nature of Species. M. Ereshefsky. Cambridge, MA, MIT Press.

Wimsatt, W. C. (1998). "Simple Systems and Phylogenetic Diversity." Philosophy of Science **65**: 267-275.

Woodward, J. (2003). Making Things Happen: A Theory of Causal Explanation. Oxford University Press.