This is a draft manuscript. Comments are most welcome,
but please do not quote without permission.

Lecture 5
Reduction and Higher Levels of Organization: A Rapprochement

As I have stressed throughout these lectures, in seeking to understand how a mechanism performs its activity, one needs to decompose the mechanism into component operations and component parts. These components constitute a lower level of organization, and research which appeals to lower levels of organization to explain phenomena at higher levels is commonly characterized as reductionistic. But, as I will try to demonstrate in this lecture, mechanistic research is Janus-faced. As Wimsatt, following Simon?, proposes, it is possible to be both a reductionist and an emergentist. By an *emergentist* I mean recognizing that there are processes occurring at higher levels of organization that are different from those found at lower levels which must be studied at their own level and serve to constrain behavior of components at lower levels. Research at higher levels of organization has a kind of autonomy—it studies, using its own modes of investigation, phenomena different from that studied at lower levels and not completely explicable in terms of the phenomena studied at lower levels. Recognizing such autonomy, however, does not require any spooky metaphysical posits such as vital forces.

Levels of Organization in a Mechanism

Most discussions of reduction begin with a conception of levels in nature as demarcated by the various sciences. On such a view, physics addresses the most basic level, chemistry a level above that, and biology, psychology, and sociology each address successively higher levels. But a moment's reflection reveals that typically there are, in fact, a host of different levels of phenomena studied in each of these broad disciplines. In chemistry there are atoms, molecules, and macromolecules. And in biology there are cell organelles, cells, tissues, organs, organ systems, organisms, ecosystems, and the biosphere.

This approach, starting with the disciplines of science and then focusing on the phenomena addressed by the disciplines, assumes a stratification of nature as a whole. A theorist who has made perhaps the greatest effort to demarcate levels that stratify nature as a whole is William Wimsatt. At the foundation of Wimsatt's approach is a mereological perspective according to which things at higher levels are made up of things at lower levels.

> By level of organization, I will mean here compositional levels—hierarchical divisions of stuff (paradigmatically but not necessarily material stuff) organized by part-whole relations, in which wholes at one level function as parts at the next (and at all higher) levels.

This will also be the key to the view to be developed here. But Wimsatt is concerned to develop general criteria by which one can relate different phenomena as being at the same level. He therefore characterizes levels in terms of interactions between denizens of the level, with a level relating entities that maximally interact with each other. As a result of being points of maximal interactions, he also maintains that levels are "*local maxima of regularity and predictability in the phase space of alternative modes of organization of matter*" and hence points in the hierarchy of levels at which we would expect sciences to develop. Thus, he says levels "*are constituted by families of entities usually of comparable size and dynamical properties, which characteristically interact primarily with one another, and which, taken together, give an apparent rough closure over a range of phenomena and regularities.* (7)

Wimsatt also proposes that, at least amongst organisms, evolutionary forces will lead entities to congregate into levels. In order to operate efficiently, an organism needs to be able to predict reliably how things in its environment will behave, and accordingly it will evolve so as to interact with entities at a given level. For example, organisms typically interact mostly with each other, cooperating or competing, acting as a predator or serving as prey, etc. Organs, on the other hand, interact mostly with other organs within the body. Organisms may interact with the organs of another organism's body (e.g., a surgeon operating on a person's heart, or a predator selective consuming specific organs of the prey), but these interactions, according to Wimsatt's, are less frequent and we can differentiate the two sets of more frequent interactions as demarcating different levels.

Wimsatt acknowledges a number of ways in which this framework is complicated. Especially with higher levels of organization, cross-level interactions increase in frequency and at these stages, he advocates alternative conceptions such as what he terms perspectives and causal thickets (Wimsatt, 1974). Perspectives, as Wimsatt characterizes them, cross levels, relating for example the parts to a whole system. As I have been talking about mechanisms, looking at something as a mechanism involves such a cross-level perspective, relating parts of a system to the activity to the whole system. Focusing on something as a mechanism is thus, on Wimsatt's analysis, to adopt a perspective.

Although the attractiveness of being able to characterize levels across nature is clear (we would, for example, have a well-structured ontology), there seem to be enormous problems with the endeavor. Presumably molecules would comprise a level of organization that Wimsatt is trying to capture. But many of the reactions molecules engage in are with individual atoms which are, for example, being added to or removed from the molecule. As well, they frequently interact with presumably higher-level structures such as membranes. In fact, insofar as mechanisms are ubiquitous in nature, so are interlevel interactions since what mechanisms do is coordinate the activity of the parts of the mechanism to accomplish an activity performed by the mechanism. And, as I will emphasize in this lecture, in carrying out their activities, mechanisms frequently impose constraints which alter the behavior of their parts.

The notion of levels is important for understanding mechanisms, but it may not be possible or important to characterize levels across nature (Craver). Rather, insofar as we are trying to understand mechanisms, what is critical is being able to identify levels in the context of a given mechanism. And this is much more readily accomplished. As I noted in the first lecture, we identify mechanisms in terms of the activities they are performing. As a result of these activities, the mechanism is engaged with other entities, many of which will themselves be understood as mechanisms. For example, the mitochondrion receives glucose, fats, and amino acids as well as ADP from other parts of the cell and generates ATP as well as expelling water and carbon dioxide. The ATP is in turn used by other components of the cell as a source of energy. The interactions between the mitochondrion and other cell organelles constitute a level, and the operations performed by the constituents inside the mitochondria that we appeal to in explaining how it performs this activity—the Krebs cycle, the electron transport chain, the cristae membrane, and the ATPase, characterize a lower level. Many of these constituents are, as I have emphasized, themselves mechanisms, and their constituents and the operations they carry out characterize a yet lower level. And the activities of the mitochondrion take places in the cell and make it possible for the cell as a whole to perform activities and thereby engage other cells. In this manner, one can identify levels locally in terms of what the components do and what other entities they interact with.

One of the features that Simon (1996) and Wimsatt emphasize about levels, that entities at a level are of roughly a common size, falls out of this perspective in which levels are arrived at by decomposing a mechanism. Parts are typically smaller than the whole of which they are a part, so as we take a mechanism apart, we encounter a set of entities at a smaller size-scale. And as we take these components themselves apart, we again drop to a smaller size-scale. But we can also understand why it is only approximately true that entities at a level are of a common size. There can be a range of sizes of the components that interact in a mechanism—the chemiosmotic process involves protons forming a gradient over a membrane, and both of these entities figure in the account of that process. The atoms of the membrane are themselves formed of protons, but those figure at a much lower level while the protons comprising the gradient are at the same level as the membrane itself.

The reason it is important to start the identification of levels from the perspective of a mechanism is that the levels that are identified are local in character. Within a mechanism we can identify the components and the operations they perform, but there is no well-defined way to relate those components as being at the same level as entities outside the mechanism—specifying, for example, whether an enzyme inside the oxidative phosphorylation system is at a higher level, a lower level, or the same level as a chip in a computer. Such local identification of levels is all that is required for understanding mechanisms. It may simply not be possible to characterize levels more globally. But even this characterization needs to be fleshed out further so as to make clear how, even on a local level, we are differentiating the components of a mechanism from the mechanism as a whole. This will be critical for our later discussion of how different levels affect each other.

The activity of a mechanism versus the operation of its components

One of the well-recognized challenges in relating theories at different levels is that such theories employ different vocabulary both to designate the entities at different levels and to depict what the entities do. Chemical substances, for example, bond with each other, cell organelles carry out activities such as protein synthesis or formation of ATP, cells respire and divide, and organisms navigate environments and mate. Similarly, turning to the brain, synaptic terminals release and take up neurotransmitters, neurons spike, populations of neurons process specific features of stimuli, and brains form beliefs and desires.

Such shifts of vocabulary are not accidental. They reflect the fact that in order for a mechanism to perform an activity, the parts of the mechanism must perform other operations. It is helpful to look at the process for the point of view of a designer—either an engineer or mother nature. The challenge is to make something that carries out an activity not already performed by other things. To do this, entities which perform other activities are coordinated together in the mechanism to carry out the new activity. Frequently the operations of the parts and the activities of the whole mechanism are radically different from those of any of its components and a different vocabulary is required to describe what they do. Even in cases in which we use a common term to describe an activity of a component and of the whole (both cells, and sets of chromosomes divide), the term is used differently. Chromosomes pair up and are pulled to separate spindles, whereas cell division involves a process of segregating the new cells, including their chromosomes, within membranes.

Discovering the kinds of operations which parts of a mechanism can perform and developing a vocabulary to describe them is not a simple task. First explorations at a given level of organization often mischaracterize the kinds of operations that can be performed there. In a particular graphic example, the first attempts to identify intermediates in the process of fermentation referred to the chemical intermediates as themselves being fermented. Fortunately, research in organic chemistry had reached a point by the end of the 19th century so as to characterize the types of chemical operations that were generally performed on organic compounds—oxidations, reductions, decarboxylations, deaminations, etc.[1] These are operations involving whole subunits of organic molecules. This provided the necessary vocabulary biochemistry required to develop accounts of metabolic processes in organisms. The challenge for biochemistry was to determine how such operations were linked together to accomplish the various overall tasks that occur in living systems—generating useable energy reserves, synthesizing proteins, producing secretions, etc. (In this process yet other basic biochemical operations, such as phoshorylations and transphosphorylations, had to be discovered.)

It is still a challenge to develop the right vocabulary to describe the component activities in the brain that enable it to perform cognitive activities. The core idea that the brain is an information processing system has proven very fruitful. But what are the basic operations in information processing? For the most part, theorists have tried to employ analogies to activities at higher or lower levels. Many of the ideas stem from the way in which whole organisms process information. Turing and Post developed their ideas of

---

[1] Much earlier in the 19th century techniques for determining the elementary chemical composition of the major organic compounds had developed, but the analysis at the level of changes in the elementary composition of organic compounds is too low a level of analysis to be useful in characterizing the activities of biochemistry.

computation on the model of how they took humans to perform computations and many of the ideas in information processing psychology, such as symbol processing (Newell, 1980) or operations on symbols in a language of thought (Fodor, 1975), have likewise drawn upon the activities whole agents perform. This is much like characterizing the chemical operations involved in fermentation as fermentations. Neural network theorists, on the other hand, have focused on activities of individual neurons (McCulloch & Pitts, 1943; Rumelhart, McClelland, & Group, 1986). Both of these are probably inappropriate for describing the information processing activities in the brain. They probably involve operations at a level above that of individual neural activity but below the activities performed by cognitive agents (Bechtel, 1994). Given the prominence of the columnar structure in much of cortex, it seems plausible that what we need is a way of characterizing information processing operations carried out by columns. But at present we are lacking such a vocabulary.

So far I have focused on the differences between the vocabulary used to characterize the components of a mechanism and what they do and that used to designate the mechanism as a whole and what it does. But the vocabulary differences are just a symptom of the fact that the components perform different operations than does the whole mechanism. Broadly speaking, both components and a whole mechanism engage in causal activities, and there has been a great deal of interest in philosophy in trying to generate an analysis of causation. But beyond the minimal fact that what causes do is bring about change, which is not a particularly informative statement which threatens to be a tautology, there may be little that can be said to characterize causation per se. Rather, there are different sorts of causal operations, and it is these we can flesh out more specifically (Cartwright). For example, what enzymes do is catalyze reactions. What the release of neural transmitters at a pre-synaptic cell does is increase the probability of depolarizing the membrane of the post-synaptic cell. What firing of MT neurons does is carry information about motion in a particular part of the visual field to other brain areas. Thus, underlying the differences in the vocabulary used to characterize entities at different levels is the difference between the operations these entities perform.

From Component Operations to the Activities of a Mechanism

How is it that entities engaged in one kind of operation inside a mechanism enable the mechanism as a whole to perform a different type of activity? The key to this is the fact that the components of a mechanism are organized so that the operations of the different components are coordinated with each other. If there were no organization, then there would be mere aggregation and the activities of the entities taken as a whole would be merely a sum of the activities of the components. There are some things that just aggregate in this manner. For example, the mass of a pile of sand is just the aggregate of the mass of the individual grains.

Wimsatt (1986) proposes thinking in terms of aggregative systems as the null case in terms of which we can recognize what organization involves. His four conditions for a system being aggregative are:

1. Intersubstitutability of parts;

2. Qualitative similarity with a change in the number of parts;
3. Stability under reaggregation of parts; and
4. Minimal interaction among parts (pp. 260-268).

Two of these are particularly important in understanding mechanisms as organized systems. First, a mechanism departs from aggregativity on the first criterion since a critical feature of being a mechanism is that different components perform different operations. But the key component on which a mechanism departs from aggregativity is in terms of interaction. Since the parts carry out different operations, they must relate what they do to each other in order for the whole mechanism to perform its activity.

Once we start conceptualizing a mechanism as involving parts interacting with other parts, then the particular order in which they interact can be critical. One component depends on the operation of another before the conditions for its own operation arise. This is where organization becomes critical. If one component depends upon another, then the components must be organized so that the different components receive what they need from each other. But often the requirements on organization are more complex. A given operation may depend on several other operations, and these might in turn depend on the first. In such cases, these multiple components must all be properly related to each other for the whole mechanism to perform its activity.

In the next part I will elaborate on the critical role organization plays in providing regulation to the operations within a mechanism. But at this point the key point that mechanisms perform activities that are different form the operation of their components is clear. With the simplest mode of organization, in which components operate sequentially upon the products of the previous components, we have systems that as a whole carry out an operation which none of the components alone could do. Assembly lines, for example, make whole cars, whereas an individual worker may only mount the engine. In that case, one might counter that the worker could indeed make a whole car. But not without a great deal of additional training through which he or she would master the skills of all the workers on the assembly line. It is the coordination of those operations that enable either the assembly line as a whole, or the individual craftsperson, to produce a product.

Higher-level control in a mechanism

There is a strong temptation for humans to conceptualize the operation of a number of components in a serial fashion, with the first component carrying out an operation on a substrate, preparing it for the next component to perform its operation, etc. But in most biological systems, the organization is more complex than that. Products that we might think of as being prepared downstream in the mechanism may be critical for operations earlier in the mechanism. Historically such complex means of interaction have typically been discovered by recognizing the impossibility of accomplishing a task in a simpler way. For example, attempts to decompose metabolic processes into linear strings of reactions often resulted in production of products which could not be further broken down by similar reactions. One way to overcome this obstacle was to propose that these terminal products might interact with earlier steps to produce a substance that could again be operated on. The Krebs cycle is such an example. Whether or not these difficulties also provide the correct explanation of the initial evolution of such integrated systems, they enabled researchers to discover them. And a key feature of such systems is that causal loops such as Krebs identified give rise to tight coupling between various components in a mechanism.

Often these couplings play a critical role in regulating the behavior of the components of the mechanism. Without regulation, the components would simply continue their operation. For example, a furnace would continue generating heat as long as it had a source of fuel and heat. The introduction of a thermostat, which sends negative feedback cutting off the supply of fuel when heat is not required, allows the operation of the furnace to be regulated. Negative feedback systems, which excited the interest of cyberneticists (Wiener, 1948), are just the simplest examples of such coupled systems in which different components can be regulated by others in the system, or demands on the system as a whole.

With coupled systems the activities of the mechanism as a whole become even more distinct from the operations of the parts. Fermentation and oxidative phosphorylation are systems not just for producing ATP, but ATP as needed by the energy consuming parts of the cell. The individual enzymes in the mechanism each catalyze their reaction whenever the necessary components are available. They do not regulate themselves.
As a result of the relations between the components, one operation is able to affect another, thereby allowing the mechanism as a whole to be regulated so that ATP is produced as, and only as, needed.

As I noted above, we tend to conceptualize the activity of a mechanism linearly. In part this is due to the fact that our conscious thought is serial (Bechtel & Richardson, 1993). But it is also due to the fact that it is very difficult to figure out what role interactive processes are playing in a mechanism. For example, it has long been recognized that there are as many recurrent or backwards projections in the brain as forward projections. (The backwards, forwards, and lateral connections generally project from and to different layers of cells in a given column, so that the system is not simply interactive but can also be construed as a hierarchy. See (Felleman & van Essen, 1991) Yet it has been difficult to determine in general what the significance of these backwards projects is. Accordingly, most accounts of neural processing treat the system as a forward processing system. The account of the ventral stream of the visual processing system I developed in the second lecture construes the system as beginning with the identification of features in the visual stimulus and in successive stages of processing recognizing higher level features such as color and shape, eventuating in the recognition of objects. Ideas as to the significance of backwards projections are usually developed as modulations of the forward processing account. Computational modeling with artificial neural networks, for example, suggests that backwards projections may bias activity in lower processing systems so as to provide input to the higher level processor that enables it to generate a more sharply delineated response (McClelland & Rumelhart, 1981; Rumelhart & McClelland, 1982). Analyses of various different attention systems in the brain also rely on the backwards projections to emphasize certain lower level processing, again biasing the input to later processing areas (Corbetta, 1998; Posner, Grossenbacher, & Compton, 1994). In these accounts, the backwards projections are construed as modulating what is conceptualized primarily as a forward processing system.

Some critics maintain that the ubiquity of backwards and lateral connections in systems such as the brain mean that we need to adopt a very different conceptual stance toward such systems. The alternative they propose is that we construe the system as a dynamical system, using the modeling tools offered by dynamical systems theory (Port & van Gelder, 1995; Van Orden, Pennington, & Stone, 2001). It is certainly true that analytical tools of dynamical systems theory are a great bone in trying to understanding complex, interactive systems. But as Fred Keijzer has argued, applying dynamics to biochemistry

> only works because we have a theory of chemical reactions and enzymes which specifies the parameters to be included in the mathematical model. The same goes for the interactions between an organism and its environment which amount to behavioral regularities. . . . [F]irst the relevant parameters which govern those regularities will have to be discovered. DST needs to be supplemented by a theory of the implementing substrate which specifies the relevant parameters and variables for the generation of behavior. (Keijzer, 2001, p. 187)

Dynamical systems theory is not a magic bullet for understanding complex, integrated systems, but a tool that can be combined with attempts to differentiate components and how they relate to one another so as to understand what such a system can accomplish. We can hope, though, that a combination of such tools will, in time, give us a better understanding of the significance of the contribution of backwards as well as forward and lateral projections in the nervous system.

Even if we retain the perspective of construing a system as primarily a forward processing system with feedback modulating the activity of the components, we are led to focus on the overall mechanism as one in which components are regulated by the demands on the mechanism as a whole. This is an instance of what some theorists, following Donald Campbell (1974), characterize as *top-down causation*. The reason to characterize it as top-down is clear—the mechanism as a whole is at a higher level, and its engagement with its environment is what results in the regulation of the components of the mechanism at a lower level. Calling such relations *causation*, however, is more suspect and leads quickly to problems (Craver and Bechtel, forthcoming).

Since the idea that higher level processes and constrain lower-level one is regarded by many as problematic, let me provide a couple quick examples that will show how natural this perspective is within the mechanist framework. Today there is a combination of DNA molecules in this room that has never been together before. But to explain the presence of the DNA in this room we do not focus on the sorts of activity in which DNA typically engages—replicating, creating RNA, etc. Rather, we would focus on the activities of academics and the institutions in which they participate. At that level we can provide an account of each of us coming to this room. And when we came, so did our DNA. It is a part of us; specifically it is a part of many of the mechanisms that comprise us. Consider another example. After this lecture you might engage in some strenuous physical activity. At that point the mechanisms of anaerobic and aerobic metabolism in you will start to operate at far higher levels than at the moment. The physical activity in which you engage places constraints on the component mechanisms that comprise you, and somewhere down this hierarchy we will reach the level at which individual cells are confronted with excess ADP, providing a trigger to initiate the processes of glycolysis and oxidative phosphorylation.

Looking up, looking down, and looking sideways

We are now at a stage where we can recognize how the mechanistic perspective provides for a rapprochement between reductionism and the semi-autonomy of higher levels of organization. The mechanistic perspective is reductionistic in that when one seeks to explain how a mechanism performs an activity in its particular context, one decomposes the mechanism into its parts, focusing on the operations they perform and how they are coordinated so as to yield the activity of the mechanism. And this is an iterative process—when one wants to explain how a component of a mechanism performs its operation in the context of the mechanism, one again decomposes it into components and their operations.

Yet, as I have been at pains to insist, the components do not perform the same activities as the whole mechanism. The components, as they are organized, account for the ability of the mechanism to perform its activity. When the mechanism as a whole performs its activity, it is interacting with entities outside the mechanism itself. Moreover, just when and how it performs this activity is partly determined by conditions impinging upon the mechanism from outside. These are not part of the lower-level analysis but need to be determined if we are to understand the behavior of the mechanism. Thus, in addition to looking down to the components, we need to look sideways to other entities to understand just what the mechanism does. An account at the level of these lateral interactions between the mechanism and things outside it is needed as well as the account about operations within the mechanism.

These two accounts, one at the level of the mechanism itself, one at the level of the mechanism's components, are semi-autonomous from each other. The autonomy involves the fact that each provides information additional to that which the other provides, and generally does so using different tools of inquiry. For example, when studying the cognitive system as a whole, and detecting the regularities in its behavior, researchers use behavioral tasks, but to study the operations in it researchers resort to such techniques as lesioning it, stimulating components, or registering the activity of components (through cell-recording or neuroimaging). But the autonomy is only partial, in each direction. What activity the mechanism performs is in large part determined by the operations of its parts. But their behavior, particularly in a mechanism with regulation in it, is due to the factors impinging on the mechanism as a whole.

Finally, if the mechanism itself is incorporated into a mechanism as a component, then one must also look up to the analysis of that mechanism to understand the behavior of the component. The operation of the higher-level mechanism can place constraints on the operation of the lower-level mechanism that result in it behaving differently than if the higher-level mechanism were not operating or if the component had been removed from the mechanism.

Theory-reduction versus mechanistic perspectives on reduction

So far I have focused on how reduction is conceived within the mechanistic perspective and how it is compatible with not just recognition of higher levels of organization as semiautonomous but as regulating behavior of components at lower levels. Although this seems entirely natural within the mechanistic perspective, it is radically at odds with how reduction is often understood. Reduction is often construed as undercutting any autonomy, except pragmatic, for higher levels and as allowing only upward propagation of effects from lower to higher levels. What lies behind this usual construal?

Within philosophy of science, most accounts of reduction focus on laws or theories, construed as statements or axiomatic systems (Nagel, 1961).[2] A science focused on a given level of organization, on this view, advances a set of laws, perhaps axiomatized. Reduction consists of developing a set of deductive relations between the laws of one science and those of another. Since the vocabulary in which the laws of different sciences are phrased is different, a necessary pre-condition for a derivation is to establish bridge laws that relate the vocabulary of the two theories. (Construing these as bridge *laws* conceals the fact that they characterize mereological relations and that often a critical aspect of the mereological relation is the manner in which the lower-level parts are organized within the higher-level structure.) Furthermore, insofar as the laws of the higher level sciences only apply under restricted conditions, one must specify boundary conditions. With the bridge laws and boundary conditions, reduction involves deriving the higher level laws from the lower level ones.

Within this framework, it is relatively easy to understand the primacy offered to the lower level. Since all higher-level laws, on this construal, are derived from lower-level ones, they are not autonomous. Moreover, they do not provide any additional information than is provided in the lower-level laws. And since the lower-level laws are more general, there is a clear sense in which they are to be preferred. The utility of the higher-level sciences is purely heuristic and pragmatic. Heuristically, one might need to rely on empirically discovered higher-level laws in the process of discovering the full lower-level account. But once the lower level account is available, it does all the explanatory work. Pragmatically, one might still employ the vocabulary and laws of the higher-level science for convenience, see they may be less complex and more easily stated, but they have no true autonomy on this account. There is certainly no place in this framework for higher levels to constrain lower levels. Any higher-level processes that might be viewed as constraints are just derived consequences of lower-level processes which already impose the constraint.

---

[2] There have been a number of attempts to develop more sophisticated versions of the theory reduction model. See, for example, (Causey, 1977; Churchland, 1986; Hooker, 1981; Schaffner, 1967; Schaffner, 1993). For the purposes of the contrasts to be drawn here, the differences between these different construals of theory reduction are not critical.

There is a fundamental difference between the way lower-levels are viewed in the theory-reduction and mechanistic perspective. First, the theory-reduction account treats the lower level as extending across all domains whereas within the mechanistic framework a level is confined to the set of entities and operations occurring within a mechanism. Second, the theory-reduction account assumes that it is possible to provide a complete account of the phenomena at the lower levels in terms of laws. These laws must specify how the entities referred to in the laws will behave under all possible circumstances. If such laws existed, then one might be able to derive from them how all higher-level entities comprised of these components would behave.[3] The account of lower-level entities envisaged in the mechanist account is much more restricted. The lower-level account is always context-bound—it characterizes how an entity behaves in a specified context. From a mechanistic account of the lower level activity of a component under the range of conditions found in a mechanism, there is no prospect of deriving a general account of upper-level phenomena. Only when the account of the operation of the component is combined with the account of the operation of other components and accounts of the way these are organized in the mechanism and the conditions under which the mechanism is operating can one draw conclusions about phenomena at the higher level.

Critical to the theory-reduction account is its characterization of lower-level laws. Is there reason to think that there are such lower-level laws for science to discover? There are reasons to be dubious. First, when one thinks of prototypical lower-level laws, such as Newton's force laws, it looks like they have universal scope and determine the behavior of everything. But that picture is extremely misleading. At best, they characterize particular relations between lower-level entities. Even if one assumes that these relations hold outside of the experimental set-ups in which they are demonstrated (which, for example, Cartwright doubts—see (Cartwright, 1999), they do not enable us to predict or explain the behavior of lower-level entities in general. In most real world circumstances, multiple factors are interacting to produce behavior. For example, multiple gravitational forces plus potentially electromagnetic forces will impinge on the same entity and we lack laws to characterize behavior under these conditions (Cartwright, 1983).

A second reason for being dubious focuses on what a complete lower-level theory (that is, one that fully specified how everything would be behave) would have to be. It would have to specify how an entity would behave under all possible circumstances, including those found in all different kinds of mechanisms that might be produced. That is not, however, the way in which a theoretical account at a given level is typically developed. It typically seeks to find regularities in the behavior of a particular entity under a restricted range of circumstances. Most theories in physics do not consider the conditions arising in biological organisms but in relatively spartan circumstances in which the entities are in

---

[3] There would still be the question of where the boundary conditions came from. If they aren't themselves derived from the lower-level laws and initial conditions, but are discovered empirically, then the sufficiency of the lower-level is undermined and higher levels enjoy a kind of autonomy even within the theory-reduction account.

relative isolation from each other. When practitioners of lower-level sciences engage collaboratively with practitioners of higher-level sciences, they have to consider these more specific conditions and develop accounts that work for them. This involves more than *mere* application of general principles. It often involves utilizing instruments and research techniques developed to study lower-level entities under the new conditions to develop accounts of what happens under these conditions.

What is discovered through such collaboration is not already part of the lower-level knowledge-base and is not generally considered to add to it, but rather to comprise part of an interfield theory. One might consider trying to expand the body of theory at the lower level to incorporate all forthcoming discoveries of how lower-level entities behave in all possible circumstances, but we should note that the result may not be a very cohesive body of knowledge and would not resemble knowledge currently produced in lower-level disciplines. It would be a new construction, built to support the kind of derivation of laws of higher-level sciences that the theory-reduction model requires. It is far from clear what form this product would look like, but it certainly would not resemble the kinds of knowledge generally produced in the lower-level sciences. The theory reduction-model is not reduction to physics, as currently conceived, but to something whose form and content we do not know.

Conclusion

The last comments make it clear why the reduction embraced in the mechanistic perspective is very different from that advocated in the theory reduction model. The turn to lower-levels in the context of developing an understanding of a mechanism is not to general theories from which all the claims at the higher level can be derived. It seeks rather to understand how a component of a mechanism carries out the operation specific to it by considering how its components operate and cooperate under specified circumstances to produce the behavior of found in the mechanism. This is a restricted inquiry, and requires as complements the development of knowledge about how the components are organized together and how the mechanism of which they are part engages in higher level activities which may impose constraints on the operation of the components. A mechanistic perspective requires looking down, up, and sideways.

References

Bechtel, W. (1994). Levels of description and explanaiton in cognitive science. *Minds and Machines, 4*, 1-25.

Bechtel, W., & Richardson, R. C. (1993). *Discovering complexity: Decomposition and localization as scientific research strategies*. Princeton, NJ: Princeton University Press.

Campbell, D. T. (1974). 'Downward causation' in hierarchically organised biological systems. In F. J. Ayala & Dobzhansky (Eds.), *Studies in the philosophy of biology*: Macmillan Press Ltd.

Cartwright, N. (1983). *How the laws of physics lie*. Oxford: Oxford University Press.

Cartwright, N. (1999). *The dapled world: A study of the boundaries of science*. Cambridge: Cambridge University Press.

Causey, R. L. (1977). *Unity of science.* Dordrecht: Reidel.

Churchland, P. S. (1986). *Neurophilosophy: Toward a Unified Science of the Mind-Brain.* Paper presented at the Cambridge: MIT Press/Bradford Books.

Corbetta, M. (1998). Frontoparietal cortical networks for directing attention and the eye to visual locations: identical, independent, or overlapping neural systems? *Proceedings of the National Academy of Sciences, USA, 95*, 831-838.

Felleman, D. J., & van Essen, D. C. (1991). Distributed hierarchical processing in the primate cerebral cortex. *Cerebral Cortex, 1*, 1-47.

Fodor, J. A. (1975). *The language of thought*. New York: Crowell.

Hooker, C. A. (1981). Towards a general theory of reduction. *Dialogue, 20*, 38-59; 201-236; 496-529.

McClelland, J. L., & Rumelhart, D. E. (1981). An Interactive Activation Model of Context Effects in Letter Perception: Part I. An Account of Basic Findings. *Psychological Review, 88*, 375-407.

McCulloch, W. S., & Pitts, W. H. (1943). A logical calculus of the ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics, 5*, 115-133.

Nagel, E. (1961). *The Structure of Science.* Paper presented at the New York: Harcourt, Brace.

Newell, A. (1980). Physical symbol systems. *Cognitive Science, 4*, 135-183.

Port, R., & van Gelder, T. (1995). *It's about time*. Cambridge, MA: MIT Press.

Posner, M. I., Grossenbacher, P. G., & Compton, P. E. (1994). Visual attention. In M. Farah & G. Ratcliff (Eds.), *The neuropsychology of high-levle vision: Collected tutorial essays* (pp. 217-239).

Rumelhart, D. E., & McClelland, J. L. (1982). An Interactive Activation Model of Context Effects in Letter Perception: Part 2. The Contextual Enhancement Effect and Some Tests and Extensions of the Model. *Psychological Review, 89*(1), 60-94.

Rumelhart, D. E., McClelland, J. L., & Group, t. P. R. (Eds.). (1986). *Parallel distributed processing: Explorations in the microstructure of cognition. Vol. Foundations*. Cambridge, MA: MIT Press.

Schaffner, K. (1967). *Approaches to reduction.* Paper presented at the Philosophy of Science, 34, 137-147.

Schaffner, K. (1993). *Discovery and explanation in biology and medicine*. Chicago: University of Chicago Press.

Simon, H. A. (1996). *The sciences of the artificial* (Third ed.). Cambridge, MA: MIT Press.

Van Orden, G. C., Pennington, B. F., & Stone, G. O. (2001). What do double dissociations prove? Inductive methods and isolable systems. *Cognitive Science, 25*, 111-172.

Wiener, N. (1948). *Cybernetics: Or, control and communication in the animal machine*. New York: Wiley.

Wimsatt, W. C. (1974). Complexity and organization. In K. F. Schaffner & R. S. Cohen (Eds.), *PSA 1972* (pp. 67-86). Dordrecth: Reidel.

Wimsatt, W. C. (1986). Forms of aggregativity. In A. Donagan & N. Perovich & M. Wedin (Eds.), *Human nature and natural knowledge* (pp. 259-293). Dordrecht: Reidel.