

**CHAPTER
T W O****It Is Not Possible to Reduce
Biological Explanations to
Explanations in Chemistry
and/or Physics***John Dupré*

In this paper, I argue that the traditional notion that complex systems, such as those found in biology, can be fully understood from a sufficiently detailed knowledge of their constituents is mistaken. My central claim is that the properties of constituents cannot themselves be fully understood without a characterization of the larger system of which they are part. This claim is elaborated through a defence of the concepts of emergence and of downward causation, causation acting from a system on its constituent parts. Although much of this argument can be read as having only epistemological or methodological force, the final section of the paper defends a more robust metaphysical reading: even purely metaphysical understandings of reductionism such as are commonly represented by supervenience theses are misguided.

1 Introduction: No Need for Special Biological Laws?

Kenneth Waters (1990) has referred to the “antireductionist consensus” in the philosophy of biology, so it is perhaps not too surprising that I find myself in agreement with most of what Evelyn Fox Keller says. At the beginning of her paper, she says that she “could as easily have gone the other way,” and I would say that, perhaps, she might better have chosen the other side. However, there are some passing statements in her paper with which I am inclined to differ, and I shall explore a few of these to see whether there may be some significant disagreement after all.

Like Keller, I am a materialist. That is to say, I do not believe there is any kind of stuff in the world other than the stuff described by physics and chemistry. There are no immaterial minds, vital forces, or extra-temporal deities. Keller writes, however, that as a materialist she is “committed to the position that all biological phenomena, including evolution, require nothing more than the workings of physics and chemistry.” Even as a materialist, I’m not sure I feel committed to this; but, of course, that depends upon exactly what the title question means. A little unpacking of this question may help to reveal where (if anywhere) there is a serious difference between Keller’s position and my own.

One could start with a trivial interpretation of this sentence. If “the workings of physics and chemistry” meant no more than the workings of things that were made of physical or chemical stuff, then a materialist, such as myself, could hardly deny it. But an ambiguity immediately appears in the phrase “physics and chemistry.” It could be read—as I just have read it—ontologically, as referring to the things of which physics and chemistry speaks. And for a materialist, therefore, it refers to the entire material world.¹

But on a quite different reading, one might more naturally think of physics and chemistry as scientific domains, traditions of inquiry, or suchlike. Then, it is far from trivial to claim that biology requires no more than the workings of physics and chemistry. Why should there *not* be biological workings that are quite different from those of physics and chemistry? And here, I appear to be in agreement with Keller when she denies that biology could be derived from the *theories or laws* of physics and chemistry.

As a matter of fact, much of the recent history of reductionism, as a philosophical doctrine, has addressed the relation between theories or, perhaps, laws of nature. Classical versions of the doctrine held that the relation in question was logical: laws of biology should follow deductively from the laws of physics or chemistry. Within the philosophy of biology, something that has surely received the status of a consensus is that no such derivations are plausible. One holdout against this consensus might be Alex Rosenberg (2006), who seems to suggest that the laws of physics, supplemented by the principle of natural selection, would suffice to derive the whole of biology. I discuss this position critically elsewhere (Dupré, 2007). Even ignoring the rather serious problem that, as Keller notes, there do not seem to be many, or any, laws in biology, there is a further problem that the concepts employed in different sciences seem to be incommensurable. As David Hull observed over 30 years ago (Hull, 1974), the relations between Mendelian genes and molecular genes are many/many. So Mendelian genetics, a scientific project still very much alive in medicine and agronomy, employs concepts that are incommensurable with those that pertain to the molecular entities that, in some sense, underlie the Mendelian phenomena. And subsequent developments in molecular genetics have suggested that the problems are more severe even than Hull could have known (Barnes & Dupré, 2008).

If the question of reductionism were merely a question of whether all of biology could be derived from the laws of physics, then we could confidently assert that the issue had been resolved. As I noted in the preceding paragraph, it cannot. However, this is clearly not what Keller has in mind and, despite her ambivalence on the topic, she has some willingness to be counted as a reductionist. So, again, what are we to make of “the workings of physics and chemistry”?

2 The Reductionist Principle

Suppose we are interested in the ecological system that includes lynxes and hares. Everyone can agree, I assume, that, among the constituents of this system are lynxes; and everyone can agree (or everyone I am concerned to argue with) that the lynxes are made entirely of physical stuff. There is nothing else to be made of. The first of these propositions is the one that the anti-reductionist will tend to emphasize, while the reductionist will be more inclined to stress the second. But where should we look for a definite disagreement between these opposing camps?

We have agreed that the lynx is made only of physical stuff. Sometimes this is expressed as the claim that the lynx (or anything else) is “nothing but” an array of physical parts. While this claim might be endorsed by most reductionists, even the moderate reductionists who have abandoned deductive relations between successive theories, it is liable to be treated with more suspicion by anti-reductionists. So we might try to separate the claim about constitution from the “nothing but” claim. What more is there to what a thing *is* than what it is made of?

Of course, one answer to the last question that, again, everyone can agree on is that it matters how the physical (or chemical) constituents are put together. In fact, to make a lynx they have to be put together in a stunningly intricate way; and a pile of chemicals that happened to be the very same molecules that could, properly assembled, constitute a lynx would be no more than an inert heap of stuff. So the reductionist’s claim should be that the lynx is nothing but a collection of physical parts *assembled in a certain way*. So here, finally, is a proposition that we might expect the reductionist and anti-reductionist to disagree about: if we knew everything about the chemicals that make up a lynx, and the way they are assembled into cells, organs, and so on, we would, in principle, know everything about the lynx. Reductionists will generally endorse something like this, whereas anti-reductionists will deny it. Let me call this, with the specific biological system lynx replaceable by any system we care to investigate, the *reductionist principle* (RP).

An important feature of RP is, of course, the phrase “in principle.” Certainly, no one knows how to explain all the properties of a complex organism in terms of the properties and arrangements of its parts; the question is whether this is simply a reflection of the underdeveloped state of our current biology, or whether there are deeper obstacles, obstacles in principle, that will continue to prevent us from doing this. The kinds of principles involved will distinguish a variety of different versions of anti-reductionism or, as it is often called, *emergentism*, the belief that there are emergent properties, properties that could not have been predicted (even “in principle”) from a complete knowledge of the constituents of a thing and their internal relations. (I should note, however, that this probably still does not capture any disagreement between Keller and me, as she acknowledges explicitly the existence of emergent properties. It may be that we differ as to what they are.)

3 Strong Emergence

One conception of emergence, championed recently by Mark Bedau (2003), proposes the obstacle to explanation as the lack of a general principle connecting features of

the constituents to features of the whole, but holds on to the reductionist intuitions with the idea that a fully detailed *simulation* of the interactions among the constituents would generate the behavior of the whole. The behavior of the whole is fully determined by the behavior of, and interactions between, the parts; but the only way to get from the latter to the former is by a complete simulation. Bedau calls this kind of emergence—which shares with reductionism the insistence on the dependence of the whole on the properties of the parts—*weak emergence*. Strong emergence, in contrast, denies such dependency.

I propose, here, to defend strong emergence. That is to say, I want to deny that the behavior of the whole is fully determined by the behavior of, and interactions between, the parts. And hence, the elements of behavior that are not so determined are what we don't know when we know everything about the parts and the way they are assembled; and thus, finally, what violates RP.

At this point, we need to be rather more careful with the relations between *dispositions to behavior* and *behavior*. No one believes that the behavior of a whole is, in general, determined solely by the properties of its parts, even for the most paradigmatically mechanical systems. A properly functioning grandfather clock, say, the action of which is powered by a slowly falling weight, will not function if the weight is supported so as to disconnect its gravitational force from the action of the clock. With few, if any, exceptions the properties of parts translate into the behavior of wholes under specific circumstances. So, the most any reductionist should claim is that dispositions of the whole are determined by properties of constituents, together with appropriate surrounding conditions. With the important qualification that these dispositions may be probabilistic rather than deterministic, let me, for the sake of argument, concede this much. Have I, then, conceded what is important about reductionism? In earlier work, I have distinguished further between the case in which probabilistic dispositions involve determinate probabilities of specific behavior, and the case in which the behavior is possible, but in which there is no reason to think that it occurs with any determinate probability (Dupré, 1993, Ch. 9). As a matter of fact, I see no reason why the second case should not be characteristic of much that happens, an idea to which I shall return briefly and tangentially in the conclusion to this paper.

One way of glossing the previous remarks about context is to note that many, at least, of the dispositional properties that appear to fall within the range of the preceding discussion are relational. An elephant gun has the capacity to kill elephants. This is a property that depends on many features of the gun and many features of elephants. One could deprive the gun of this capacity by fitting all elephants with suitable armor plating or, indeed, by killing them all so that nothing any more has the capacity in question. But it is natural to think that there is an intrinsic capacity of the gun, the capacity perhaps to project a lead pellet of a particular mass at a particular velocity in a determined direction, that underlies the relational capacity. Relational capacities of a thing are quite obviously not reducible to any information about the parts of that thing alone, since they depend also on features of the other party to the relation. But the reductionist might reasonably propose that all the intrinsic properties of the thing are reducible to properties of the parts, and that the relational properties were deducible from a knowledge of the intrinsic properties of the things related. The hardness of an elephant's skin, the distance from the skin to organs

necessary for the maintenance of life, and so on, combined with the intrinsic capacity of the gun to project a leaden missile, together entail the capacity of the gun to inflict fatal damage on elephants.

So, here is a possible ontological picture for the reductionist. Imagine arranging entities in a traditional ontological hierarchy: elementary particles, atoms, molecules, cells, organisms, etc.² At each level of the hierarchy, we can determine a set of intrinsic properties of the relevant entities. From the intrinsic properties of entities at one level, say atoms, and the relations between the atoms, we can infer the intrinsic properties of molecules. Thus, ultimately, the intrinsic properties of everything are consequences of the intrinsic properties of their constituents and, ultimately of their smallest (physical) constituents. Or anyhow, this will follow as long as we can take care of the relational properties smuggled into the story. (Keller probably would not endorse this reductionist picture. For she writes that biological explanations assume “the dependence of the identity of parts, and the interactions among them, on higher-order effects.” As will emerge very shortly, I take this to capture a fundamental deficiency in the reductionist picture.)

4 Complex Relations in Biology

Unfortunately, of course, taking care of the relational properties is not an easy matter. Sometimes these are simply a matter of location, as with the relative positions of the elephant and the gun, and spatial relations presumably belong comfortably in the realm of the physical. But many relational properties in biology are not that simple. Consider, for example, the characteristic properties of enzymes. An enzyme is a catalyst, generally a protein but in some cases an RNA molecule, which facilitates a generally highly specific biochemical reaction in a cell. The mechanisms by which many enzymes work are well-understood and involve a variety of spatial and electrochemical interactions between the enzyme and its substrate. Enzymes typically have an active site, a small part of an often very large and complex protein, which binds to the substrate and changes its spatial configuration or electrical charge pattern in a way that reduces the activation energy of the reaction the enzyme catalyzes.

We may, no doubt, assume that intrinsic properties of the substrate and the enzyme are sufficient to explain the capacity of the latter to act as a catalyst on the former. However, the intrinsic properties of a large and complex molecule such as a protein will very likely allow it to catalyze many different reactions. And as a matter of fact it has become clear that many proteins do, in fact, have multiple functions. Ramasarma (1999) lists over 50 proteins (or groups of closely related proteins) with more than one known function (also see Jeffery, 1999). These alternate functions include a range of activities in addition to serving as enzymes, such as binding or transporting various molecules, inhibiting various cellular processes, or forming subunits of larger proteins. Glyceraldehyde-3-phosphate dehydrogenase (GAPDH), for example, a common “housekeeping” enzyme, is believed to act as an acyl phosphatase, an esterase, a protein kinase, and a Uracil-DNA glycosylase, in ADP-ribosylation, microtubule-binding, t-RNA binding, amyloid protein binding, and membrane binding (Ramasarma, 1999). The number of *possible* functions of a protein molecule seems, in principle, quite indefinite.

Hence, finally, a complete knowledge of the physical and chemical properties of a protein will certainly not tell us what a protein does. When we know what the protein does, chemistry may certainly be able to explain how it does it; but that is a different matter. The distinction between explaining *how something does what it does* and explaining *what it does* was central to my earlier critique of reductionism (Dupré, 1993). The idea is developed further by Pigliucci and Kaplan (2006).

Note, also, that it is common practice to say that a protein *is* an esterase, a protein kinase, etc. In fact, the primary name of the protein just mentioned, GAPDH, indicates its role in catalysing the transformation of glyceraldehyde-3-phosphate; most proteins are named by reference to (one of) their functions. But being a GAPDH, a molecule with that particular catalytic function, requires not merely a particular chemical structure, but an environment in which there is glyceraldehyde-3-phosphate to transform.

I take it that the point I have just been trying to make is part of what Keller means when she writes of “the dependence of the identity of parts, and the interactions among them, on higher-order effects.” But this dependence points, I believe, to a fundamental *objection* to reductionism. Chemistry, alone, cannot tell us that a particular protein is a GAPDH rather than one of the countless other functionally defined things it might have been. To be a GAPDH requires, in addition, an environment that includes the other elements that make the performance of its specific function possible.

The point can perhaps be more intuitively illustrated by thinking of the quite different case of human capacities. As is most definitively argued in Wittgenstein’s *private language argument* (1953/2001), the ability most humans have to speak a language is one that would be impossible, in principle, without the existence of a linguistic community of which they were part. Countless human capacities—to write a check, make a promise, play chess, and so on—depend for their possibility on the existence of a social context in which conventions or rules create the conditions for such activities. No amount of knowledge of my physics, chemistry, neurophysiology, or the like, could determine whether I was able to write a check. But the point also applies, in important ways, to less socially embedded physical activities. My physical properties do not determine, for example, whether I am able to move vertically through buildings. If my legs function in the standard way, my ability to do this will depend on the availability of staircases. If I use a wheelchair, my vertical mobility will depend on the provision of ramps or elevators. Again, the capacities of a thing, as opposed to the countless merely possible capacities, can be seen to depend on the relationship between the thing and the environment in which it exists.

It may seem that by conceding that the actual capacities—the capacities that become actual rather than merely possible in an appropriate environment—can be explained in physical or chemical terms, I have conceded everything the reductionist really cares about. Certainly, I hope I have conceded what is necessary to account for the extraordinarily successful practices of scientists engaged with molecular aspects of biology. Here, however, it is useful to recall a banal point, though one that occasionally gets overlooked in such discussions, that we are talking about science, not Nature. Biological explanations are part of biology, not part of the world, and biology, like any other science, is an articulated conceptual structure and not a repository of things-in-themselves. I shall move into metaphysics and attempt to say a little about Nature at the end of this paper, but, for now, I shall continue to address science. And the

fact that biology—a science—works with *concepts* that depend on the larger systems of which they are part, as well as on their constituents, is a fatal objection to the claim that “it is possible to reduce biological explanations to explanations in chemistry and/or physics.”

5 A Misinformed Slogan and Its Contributions

The preceding point can be developed, as well as some of the main dangers of reductionist thinking illustrated, by considering the view that the genome contains *all* the information required to build an organism. In fact, this view is still too-often promulgated by scientists who should know better, and widely asserted in popular science writing. Probably not a lot of people, biologists or otherwise, who have thought seriously about such things still believe this. However, a close look at why it is so deeply mistaken will be a useful way of elaborating the difficulties with reductionism. I won't dwell too long on the tricky concept of information. As developmental systems theorists, in particular, have pointed out for some time, the concept faces an impossible dilemma (an excellent introduction to development systems theory is Griffiths & Gray, 1994; Oyama, 1985, is the *locus classicus*). If it is interpreted in the everyday semantic sense, then it is obviously false. Genomes contain meaning only in a highly metaphorical sense. But if it is interpreted in the technical sense of information theory, according to which, very roughly speaking, a source conveys information about a target when it reduces the uncertainty about the state of the latter, then there is no sense in which the genome carries information that does not apply equally well to everything else that is necessary for features of the genome to have their normal effects on the cell. Without a complete transcription mechanism, for instance, the genome carries no information; and taking the genome as part of the channel through which the information flows, the transcription machinery carries information about the same targets. In what follows, I shall occasionally use the concept loosely to refer simply to casual determination.

But this rather technical issue points directly to a more fundamental difficulty even with less pedantic interpretations of the idea under discussion. Strings of DNA, or even real genomes replete with histones and other structural elements that make up real physical chromosomes, do nothing on their own. Their involvement in the production of proteins is as part of a system that includes a very large number of additional molecules and cellular structures, and although there are very special and biologically important features peculiar to nucleic acids, singling them out in the way suggested by the reductionist slogan makes no sense. There are several simplistic ideas that contribute to the continued popularity of the slogan, however, and I shall briefly discuss three of these.

The first is what was named by Francis Crick (1958), presumably with a trace of irony, the *central dogma of molecular biology*. According to the dogma, information flows from DNA, to RNA, to protein sequence, but never in the opposite direction. Although one should question whether merely specifying a list of amino acid sequences that constitute proteins would be sufficient to specify the entire organism, the doctrine that DNA is only a source of information, entirely immune to the influence of signals from its cellular environment, provides a powerfully reductive

perspective on the economy of the cell. But, at any rate, interpreted this way the dogma is entirely misguided. Apart from its dependence, already referred to, on the transcription mechanism that produces mRNA sequence from DNA sequence—not to mention the mechanisms including complex cellular structures, ribosomes, that facilitate the production of proteins in accord with mRNA sequence—the behavior of the genome is affected by countless other molecules in the cell. It does certainly remain the case that the important activity of the genome is providing a template for transcription of RNA molecules, but which RNA molecules are transcribed and in what quantities is dependent on interactions with many other constituents of the cell, and the actual structure of the DNA in the genome is constantly being modified by these other constituents. Information, then, flows constantly to the genome from RNA and protein molecules.

The importance of RNA molecules in modifying genome behavior is a field still in its infancy, but one that is thought by many to be likely to revolutionize cell biology. A second simplistic view that is relevant here is the idea, once widely cited, that most of the DNA in the genome, perhaps up to 98% of it, is “junk.” It has been realized for some time that only a very small proportion of the DNA in the genome provided sequence that ended up translated into protein structure. In accordance with the central dogma, it was concluded that most of the rest lacked any function at all and that, perhaps, its presence reflected no more than competition for space in the genome among genetic elements that played no part in the functioning of the organism (Dawkins, 1982). However, it now appears that the large majority of the genome, at least 70%, is transcribed, and as more is understood about the variety of RNA molecules in cells and their diverse functions, it becomes increasingly imprudent to assume that these RNAs—and, therefore, the genomic DNA from which they derive—may not play essential roles in the economy of the cell.

At this point it might be said that, even acknowledging that the function of the genome is affected by numerous RNA and protein molecules, the latter are derived from genomic DNA sequence, so that all that has been added to the idea of the DNA as controlling molecule is a few feedback loops. Everything still begins with the genome. This leads to the third simplistic view I want to mention. Sometimes, it is imagined that all that passes from one generation to the next, through reproduction, is the genome. Clearly, this presupposes the idea that the genome contains all the information necessary to build an organism, since the organism is built, and nothing is there to build it from but the genome. But this view is quite wrong, as well. The smallest “bottleneck” in the developmental cycle is a single cell, the zygote. And this cell contains all the machinery necessary for the functioning of the DNA, and the rest of a normal cellular complement of molecular constituents and sub-cellular structures. The cell is an evolved structure which, far from being assembled through instructions contained in the DNA, is a product of several billion years of evolution. Naturally occurring DNA is always and everywhere found as part of such exquisitely complex evolved systems.

This brings me back to the central point that occupied the earlier parts of this paper. The capacities of DNA are not merely consequences of its molecular constitution, but are simultaneously determined by the systems of which DNA molecules are part. The best way to illustrate this point is by considering the ways that we divide the genome into functional parts, genes.

6 Genes

The concept of a gene has two rather different traditions of use. Its origin is in the breeding experiments, especially on fruit flies, of the first half of the twentieth century, though the classic experiments of Mendel in the 1860s are generally considered canonical precursors of this work, and this research tradition is often referred to as *Mendelian genetics* (Sturtevant, 2001). In this context, the gene was a theoretical term used to track inherited features of the organism under study. The “gene for red eyes” was the hypothetical cause of a pattern of inheritance of red eyes. Flies with red eyes were assumed to possess this gene, and this, through Mendelian patterns of inheritance and theories of dominance and recessiveness, explained the quantitative characteristics of the pattern. Such a gene made the difference between having and not having a particular trait, in this case a particular eye color, and it is an important point about this research tradition that genes were always, and only, difference makers. Where a kind of organism showed no variation in a trait, there were no genes for that trait to investigate.

It is uncontroversial that, in this tradition, genes were what they were only in a very complex context. Even a quite deterministic view of the action of genes would need to allow that only deep within the body of a fly would any molecular entity actually make a difference in eye color. The same entity (if indeed it is an entity) might appear in another organism with a quite different effect, or indeed might appear as a production stage in a chemical factory making parts for organic computers. Being a gene for red eyes is very far from being an intrinsic feature of a bit of chemical.

Of course, the reason for this is clear enough: we have identified the gene in terms of its effect in a much larger system. Nowadays, we are inclined to think of genes rather as sequences of nucleotides, and conceived that way surely they are simply chemical objects. And indeed, we might suppose that the genes for red eyes and suchlike could now be identified with sequences of nucleotides describable quite independently of the biological context in which they appeared. As a matter of fact, however, this turns out not generally to be possible. One of the major areas in which Mendelian genetics remains a thriving tradition is in the medical genetics that addresses single gene disorders. A standard example is cystic fibrosis, which is caused by a recessive gene with very serious health consequences. However, there is no particular sequence that corresponds to the gene for cystic fibrosis. Rather, cystic fibrosis results from a large range of mutations in a certain transcribed DNA sequence (known as the *cystic fibrosis transmembrane conductance regulator gene*) any of which render the gene incapable of producing a functional product. If both copies of the gene are dysfunctional, the disease state ensues. Currently, over 1,000 such mutations have been catalogued. But not every possible mutation of this stretch of the genome will render the gene dysfunctional, and in fact the severity of the disease will vary according to the precise mutation. So, whether a piece of DNA sequence is a cystic fibrosis gene is determined only by how it functions in the entire organism.

It might be supposed that things would be much more straightforward from the perspective of the second tradition, molecular genetics. In this research tradition, which developed out of the identification of DNA as the material out of which genes were made and the resolution of the chemical structure of that molecule, genes are generally identified as specific sequences of nucleotides. Surely being such a

sequence is something that occurs quite independently of any context? I have already explained part of the problem with this thought. Insofar as molecular genes are identified as templates for particular proteins, and proteins are distinguished by (one of) the specific function(s) they carry out in the cell, then the identity of the gene, as it is conceived as a gene with the function of providing a template for a particular protein, is, again, tied to a molecular context. And it is under some such description that we are motivated to distinguish a particular sequence as a discrete entity, a gene. But the difficulty goes deeper than this, arising from fundamental limitations to the gene concept itself.

In the early days of molecular genetics, and nicely encapsulated in the central dogma, it was supposed that genes were discrete DNA sequences that specified precisely the amino acid sequence for a particular gene product. Subsequently, however, things have proved far more complicated. First of all, it was discovered early on that genes (in eukaryotes, anyhow, the division of life that includes, among other things, animals such as ourselves and plants) were not typically unbroken coding sequences, but parts of the coding sequence (known as *exons*) were interspersed with non-coding parts (known as *introns*). The whole sequence, exons and introns, is transcribed into RNA; but, subsequently the introns are excised by further processing machinery. Subsequently it was discovered that in many (probably most) genes this excision process could be carried out in different ways, resulting in different RNA products, so-called *alternative splicing*. These RNA products were also liable to undergo further “editing,” alteration of details of their sequence, prior to being translated into amino acid sequences. And these amino acid sequences might subsequently be assembled into a variety of different functional proteins. So, different and discontinuous parts of the sequence, initially thought to be a gene, might end up in a range of different functional proteins.

And this is not the end of the relevant complications. The early picture had discrete genes, separated by non-coding (either regulatory or “junk”) sequences in a reasonably orderly sequence. But it now seems that coding genes may overlap one another; they may be embedded within the intron of another gene; and they may sometimes be read in both directions, as so-called sense and anti-sense genes, and a particular sequence might be part of both a sense gene and an anti-sense gene. A particular part of the genome might be part of several different genes. In fact, it now seems that something like 23,000 “genes”—in the sense originally assumed in molecular genetics—are involved in the production of perhaps as many as a million proteins. It is hardly surprising that the process from the former to the latter is not straightforwardly linear.

I have presented these complications in terms of genes, but the fact is that they raise serious worries as to whether there is really a coherent concept of the molecular gene at all. The philosophers Paul Griffiths and Karola Stotz carried out a research project in which they presented scientists with DNA sequences involved in various complexities of the kinds just described, and asked them how many genes there were in these sequences (Stotz & Griffiths, 2004; Stotz, Griffiths, & Knight, 2004). There was little consensus as to the right answer. The best way to understand this finding is certainly open to debate, but my own view would be not that it showed that there were no such things as genes, but that distinguishing part of the genome as a gene only makes sense in relation to some function that particular bit of sequence serves

in the general functioning of the cell. There is, therefore, no objectively unique division of the genome into genes. Again I conclude, the conceptualization of the genome, as an object of study and as divisible into discrete functional constituents, requires that it be placed in the wider *context* with which it interacts.

7 Causation

We can also approach the question of the reducibility of the biological by looking at intuitions about causation. One of the intuitions underlying reductionist thinking is that, whereas it is natural to think of parts of an entity as causally explaining the behavior of the whole, the reverse, causal explanation of the behavior of the parts in terms of features of the whole, so-called *downward causation*, is somehow considered mysterious. (Downward causation has been the subject of philosophical debate for some time, generally dated from a proposal of Donald Campbell [1974]. Interest has been greatly increased recently in the context of systems biology, of which more below.) So it seems natural to explain, for example, the movement of my arm in terms of a series of biochemical processes leading to the contraction of bundles of fibers attached to parts of the skeleton. This constitutes a classic causal/mechanical explanation in terms of pushes, pulls, hinges, and suchlike. It seems strange to many, on the other hand, to suppose that the whole organism of which the muscle tissue is part could somehow cause the necessary molecular activities. Of course, philosophically untutored intuition may find the second possibility quite natural. The naïve explanation of my arm's going up is that I intended to reach for a book, say, which explains the bodily movement in terms of a feature of the whole, its intention. If the whole person is capable of raising the arm, and raising the arm is caused by (among much else) calcium being pumped into the sarcoplasmic reticulum, then it appears that the person is capable of causing calcium to be pumped into the sarcoplasmic reticulum.

Downward causation seems a very natural way to think of much of what I have been saying about molecular biology. What causes the human genome to behave in the particular ways it does—for example, various sequences being transcribed or not at varying rate, changes in conformation and spatial relation of chromosomes, and so on—is a variety of features dispersed over the surrounding parts of the cell. The behavior of the part is to be explained by appeal to features of the whole.

Another example that fits naturally into this picture is the phenomenon of protein folding. A major problem in molecular biology is to explain the transition from an amino acid sequence to the baroquely complex structure that results as this sequence folds into a three-dimensional shape. The topology of this structure is essential to the proper functioning of the protein, yet in many cases it appears to be strongly underdetermined by the chemical properties of the links between successive amino acids. It is known that many proteins require specific collaboration from other proteins, known as chaperones, to accomplish this complex feat. One might argue that this was simply another interaction, between the folding protein and the chaperone, fully compatible with a traditional reductionist perspective.

However, even if, as is probably a great oversimplification, interaction with the correct chaperones were all that was required for correct folding, the kind of argument considered with regard to protein function applies equally here. There is a very specific environment, in this case one replete with appropriate chaperones, which endows the amino acid sequence with the capacity, or disposition, to fold in a particular, functionally desirable, way. Still more is it a specific environment that disposes the various relevant parts of the genome to produce, in the end, an appropriately folded protein. And again, this environment is not something that could possibly be generated *de novo* by the genome but, on the contrary, it is one that took a few billion years to evolve. The cell, I think we must say, with all its intricate structure and diverse contents, is what causes these contents to behave in these life-sustaining ways (Powell & Dupré, 2009).

8 Systems Biology

A scientific development that has brought these issues of downward causation, emergent properties, and reductionism to the fore is the rapid growth of *systems biology*. Systems biology can be seen as a response to the growing realization that the accelerating avalanche of molecular data from ever-faster gene sequencing and comparable technologies for assaying RNAs and proteins had not been matched by similar growth in the ability to assemble these data into adequate models or explanations of larger-scale phenomena. Systems biology was conceived as a collaborative effort among molecular biologists, mathematicians, and computer scientists to attempt to provide such integrative understanding. In earlier work (O'Malley & Dupré, 2005), my collaborator and I distinguished between top-down and bottom-up tendencies in current systems biology. The former, generally reductive tendency hopes to build up more global understanding by gradual integration of information from molecular censuses and knowledge of molecular interactions. Top-down systems biologists doubt whether this can be done, and insist on the need for more general principles that emerge at higher levels of organization, and constrain the behavior of constituents.

This is a relatively crude dichotomy, of course, and the consensus among biologists involved in these projects is that some combination of the two will be needed for systems biology to succeed (Krohs and Callebaut [2007] offer criticism of the dichotomy just mentioned). This is exactly what should be expected in the light of the preceding discussion: the capacities of parts will require explanation through reductive, bottom-up approaches; but a top-down approach is required to understand their actual behavior, and to identify the capacities that need to be explained. However, whereas some top-down systems biologists hold that this is a matter of identifying laws that govern complex systems, my own prejudice is that the top-down part is more a matter of higher-level description of particular systems. One cannot infer the behavior of a cell by treating it as a bag of chemicals; one might begin to make progress by describing the intermediate structures – viz., ribosomes, Golgi apparatuses, and so on – and the heterogeneous distribution of various molecules in relation to such things.

9 Metaphysical Coda

Reductionism is inspired both by observations of the methods of science and by more purely metaphysical reflections or intuitions. Though, as I have argued, I don't doubt that reductionist methods have a vital role in science, I think this role can be overstated, especially in biology. A science such as molecular biology tells us not only how particular entities come to have the complex capacities they do, but also how complex systems enroll some of these capacities to create stability, order, and function. In doing so, I have suggested, those systems constrain and causally influence the behavior of their molecular constituents.

An influential movement in recent philosophy of science has attempted to describe biological systems in terms of mechanism (e.g., Bechtel, 2006; Machamer, Darden, & Craver, 2000). I am generally sympathetic to this movement, and these accounts have strong parallels, for example, with the view of top-down systems biology mentioned at the end of the preceding section. Although these recent accounts of mechanism do not rest a great deal on the implicit parallel with machines, this parallel does have serious disadvantages, as well as some virtues.

On the positive side, machines, like organisms, exploit capacities of their constituents to create order and predictable behavior. But there are important differences. The machines we construct typically have a fixed set of parts, and those parts are invariably subject to decay and failure over time. Organisms, in contrast, constantly renew and replace their parts, often with different ones. Organisms have life cycles; machines have only a linear progression toward decay.

However, both machines and organism illustrate one very important point. Order is difficult to achieve. Machines achieve it with all kinds of ingenuity, and auxiliary devices that anticipate, and sometimes prevent, the common causes of failure. Organisms maintain order with stunningly complex arrays of interacting parts, the "Bernard machines" eloquently described in Keller's paper, and much else besides. Also, I maintain, these order-preserving systems work by creating synergies of mutual determination between different levels of organization. Although this last point may indicate a fundamental difference between machines and organisms, there is a crucial point in emphasizing their similarity, namely, to indicate the dubiousness of an intuition that underlies much reductionist thinking. This is the idea that order is everywhere, i.e., that everything is determined by the unvarying capacities of microscopic constituents; and it is not at all borne out by a close study of the systems that do manage to maintain order and predictability. Biological order, I argue, is the extraordinary achievement of systems honed by billions of years of evolution. It is not something that comes for free with the determinism of the physical and chemical worlds.

Throughout this paper, I have been concerned to engage with reductionism as a serious aspect of scientific methodology. I have tried to produce an account that does justice to the undoubted importance that working scientists attach to reductionist methods, while avoiding philosophical conclusions that go beyond what such a concern requires. However, a great deal of philosophical discussion involves much weaker notions that have no such connections with scientific methodology, actual or even imaginable. Perhaps the most widely discussed, and certainly of no threatened relevance to science, are various theses of "supervenience," a form of reductionism

sometimes considered so weak that any sane person must accept it. Indeed, supervenience is often thought to be a paradigmatic form of *anti-reductionism*. A thesis of the supervenience of the biological on the physical asserts that, however inaccessible are principles connecting lower levels to higher levels, nevertheless, the biological depends on the physical in the sense that for any biological system there is a physical state that constitutes it, and wherever we were to find an identical physical state we would find an identical biological state.

It should be obvious what my worry with such a position will be. Perhaps this would be true for any closed biological system, but then there are no closed biological systems. This is one way of understanding the dependence of the identity of biological entities on context that I have emphasized in this paper. Bounded biological systems do not supervene on their physical parts because aspects of what they are depend on the context with which they interact, a context always extending beyond any predetermined boundaries. Perhaps I should concede that everything in the universe supervenes on the total physical state of the universe? Perhaps. But, here, we are so deeply into the domain of purely speculative metaphysics that I more than happy to remain agnostic.

Postscript: Counterpoint

I still find the differences between myself and Evelyn Fox Keller to be slight. I shall make some very brief responses to the points that Keller raises in the postscript to her paper. While I agree that “it is important to distinguish *interactions*, *behavior*, and *properties*,” the properties with which I am most concerned are capacities, and capacities, I take it, are defined in terms of the behavior that is their characteristic exercise. Moreover, capacities, I argue, are jointly determined by intrinsic features of an entity and by features of its environment. So though interactions, behavior, and properties should be distinguished, they are, I would argue, intimately connected.

On downward causation, I cannot see anything in Keller’s postscript with which I disagree. It may be that something in the main body of my text gives a contrary impression, but I am unable to see the disagreement. I am, of course, happy if there is none.

Finally on the question of context and interaction, I agree that there is no absolute or objective distinction here. However, it seems to me that there is a useful pragmatic distinction. I have used the word “context” to refer to features of an object’s environment that are necessary to confer on the object a particular capacity, as just explained. Interactions are simply the exercise of such capacities with relation to some other entity that will presumably constitute all or part of that context. So clearly there is no way of saying of a particular feature of the environment that it is objectively part of the context rather than something with which the entity interacts. However, I would say that the project of characterizing the entity, which I have said requires reference to the context, and the project of describing what, on a particular occasion, it does, namely, interact, are distinct activities. And it is only as a means to discriminate these intellectual activities that I would defend the distinction in my text, such as it is, between context and interaction.

Acknowledgments

I am grateful to Maureen O'Malley, Alex Powell, and Evelyn Fox Keller for comments that led to substantial improvements in the final version of this paper. I gratefully acknowledge support from the Economic and Social Research Council (ESRC) and the Arts and Humanities Research Council (AHRC). The research in this paper was part of the programme of the ESRC Centre for Genomics in Society (Egenis).

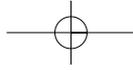
Notes

- 1 Incidentally, the disjunction here and the inclusive disjunction in the title are somewhat problematic. If strong doctrines of reductionism are true, then chemistry is reducible to physics. Hereafter, I shall generally refer just to physics, though, of course, any practical scientific project will surely be happy with a reduction to chemistry.
- 2 I won't worry here exactly what are the appropriate steps in the hierarchy. As a matter of fact, I am increasingly skeptical whether there really is such a hierarchy rather than a number of points at which we have found it useful to abstract objects of particular degrees of complexity. For example, we typically think of multicellular organisms as distinct from the multitudes of microbes with which they are in an obligate symbiotic relationship, but this is not appropriate for all purposes. The argument below concerning the contextual dependence of the identity of biological entities points more generally in this direction, but I won't develop the point any further here.

References

- Barnes, S., & Dupré, J. (2008). *Genomes and what to make of them*. Chicago: University of Chicago Press.
- Bechtel, W. (2006). *Discovering cell mechanisms*. New York: Cambridge University Press.
- Bedau, M. (2003). Downward causation and autonomy in weak emergence. *Principia*, 6, 5–50.
- Campbell, D. (1974). "Downward causation" in hierarchically organized biological systems. In F. Ayala & T. Dobzhansky (Eds.), *Studies in the philosophy of biology* (pp. 179–186). London: Macmillan.
- Crick, F. (1958). On protein synthesis. *Symposia of the Society for Experimental Biology*, 12, 139–163.
- Dawkins, R. (1982). *The extended phenotype*. Oxford, UK: Oxford University Press.
- Dupré, J. (1993). *The disorder of things: Metaphysical foundations of the disunity of science*. Cambridge, MA: Harvard University Press.
- Dupré, J. (2007). Is biology reducible to the laws of physics? Review of Alexander Rosenberg, *Darwinian reductionism: Or, how to stop worrying and love molecular biology*. *American Scientist*, 95, 274–276.
- Griffiths, P., & Gray, R. (1994). Developmental systems and evolutionary explanation. *Journal of Philosophy*, 91, 277–304.
- Hull, D. (1974). *Philosophy of biological science*. Englewood Cliffs, NJ: Prentice-Hall.
- Jeffery, C. (1999). Moonlighting proteins. *Trends in Biochemical Sciences*, 24, 8–11.
- Krohs, U., & Callebaut, W. (2007). Data without models merging with models without data. In F. Boogerd, F. Bruggeman, J. Hofmeyr, & H. Westerhoff (Eds.), *Systems biology: Philosophical foundations* (pp. 181–213). Amsterdam: Elsevier.

- Machamer, P., Darden, L., & Craver, C. (2000). Thinking about mechanisms. *Philosophy of Science*, 67, 1–25.
- O'Malley, M., & Dupré, J. (2005). Fundamental issues in systems biology. *BioEssays*, 27, 1270–1276.
- Oyama, S. (1985). *The ontogeny of information: Developmental systems and evolution*. Cambridge, UK: Cambridge University Press.
- Pigliucci, M., & Kaplan, J. (2006). *Making sense of evolution: The conceptual foundations of evolutionary biology*. Chicago: University of Chicago Press.
- Powell, A., & Dupré, J. (2009). From molecules to systems: The importance of looking both ways. *Studies in History and Philosophy of Science Part C (Biological and Biomedical Sciences)*, 40, 54–64.
- Ramasarma, T. (1999). Is it fair to describe a protein recruited for many cellular chores as “moonlighting” and “promiscuous”? *Current Science*, 77, 1401–1405.
- Rosenberg, A. (2006). *Darwinian reductionism: Or, how to stop worrying and love molecular biology*. Chicago: University of Chicago Press.
- Stotz, K., & Griffiths, P. (2004). Genes: Philosophical analyses put to the test. *History and Philosophy of the Life Sciences*, 26, 5–28.
- Stotz, K., Griffiths, P., & Knight, R. (2004). How scientists conceptualize genes: An empirical study. *Studies in History & Philosophy of Biological and Biomedical Sciences*, 35, 647–673.
- Sturtevant, A. (2001). *A history of genetics*. Cold Spring Harbor, NY: Cold Spring Harbor Press.
- Waters, K. (1990). Why the anti-reductionist consensus won't survive: The case of classical Mendelian genetics. *Proceedings of the Biennial Meeting of the Philosophy of Science Association*, 125–139.
- Wittgenstein, L. (1953/2001). *Philosophical investigations*. Malden, MA: Blackwell Publishing.



UNCORRECTED PROOF

