The Constituents of Life

John Dupré

Spinoza Lectures

Will Kymlicka - *States, Natures and Cultures*
ISBN 90 232 3224 0

Manfred Frank - *Selbstbewußtsein und Argumentation*
ISBN 90 232 3278 X

Richard Rorty - *Truth, politics and 'post-modernism'*
ISBN 90 232 3279 8

Albrecht Wellmer - *Revolution und Interpretation*
ISBN 90 232 3426 X

Axel Honneth - *Suffering from Indeterminacy*
ISBN 90 232 3564 9

Seyla Benhabib - *Transformations of Citizenship*
ISBN 90 232 3724 2

Hilary Putnam - *Enlightenment and Pragmatism*
ISBN 90 232 3739 0

Judith Butler - *Giving an Account of Oneself*
ISBN 90 232 3940 7

Nancy Fraser - *Reframing Justice*
ISBN 90 232 4155 X

Hubert Dreyfus - *Skilled Coping as Higher Intelligibility in Heidegger's Being and Time*
ISBN 978 90 232 4378 6

John Dupré - *The Constituents of Life*
ISBN 978 90 232 4380 9

The Constituents of Life

John Dupré

2007 Van Gorcum

ISBN 978 90 232 4380 9

The Department of Philosophy of the University of Amsterdam presented the Spinoza Lectures by John Dupré in May and June 2006.

Design: Anneke de Bruin, Amsterdam
Cover: Crasborn, Valkenburg a.d. Geul
Printing: Royal Van Gorcum, Postbus 43, 9400 AA Assen

TABLE OF CONTENTS

Acknowledgement

This book is a slightly revised version of the Spinoza Lectures delivered at the University of Amsterdam in May and June 2006. I am very grateful to the Philosophy Department of the University of Amsterdam for giving me the opportunity to deliver the lectures and to spend a most enjoyable and productive term in that beautiful city. I would like especially to thank Professor Frans Jacobs, Head of the Department, for his excellent hospitality during the visit, and Ria Beentjes and Willy van Wier for taking care of all the practical needs connected with my visit with exemplary efficiency. Conversations with Michiel van Lambalgen, Martin Stokhof, Beate Roessler, Gerard de Vries, Veit Bader, and Tine Wilde contributed greatly to both my enjoyment and my intellectual well-being. Teaching a postgraduate seminar on Philosophy of Biology with Wolfram Hinzen provided an invaluable opportunity to discuss some of the issues with an appropriately sceptical audience.

The work from which these lectures derives is greatly indebted to many colleagues at Egenis, The Economic and Social Research Council (ESRC) Centre for Genomics in Society, with whom I have been discussing many of the issues addressed for several years, most especially Barry Barnes, Steve Hughes, Christine Hauskeller, Staffan Mueller-Wille, Lenny Moss, Paula Saukko and Jane Calvert. Much of the work presented here derives from research undertaken there, and I am very grateful to the ESRC for its continuing support. My greatest individual debts are to Maureen O'Malley, with whom I have been collaborating for some time on philosophical topics in microbiology and systems biology, and without whose expertise and insight on those topics I would not have been able to write these lectures; and to Regenia Gagnier, who as always read and provided uniquely insightful comments on drafts of the work, and whose stimulation and encouragement is an essential background to all my philosophical work.

Spinoza Lecture I

The Constituents of Life

The title of these talks, The Constituents of Life, refers to the things that are the subject matter of biology: organisms, the systems, organs, cells and molecules to be found within them, and the larger systems, such as species or ecosystems which they, in turn, compose. It might not be obvious that there is much for a philosopher to say on this subject. We are all familiar enough with these things at a common sense level, but it is surely for biologists to provide us with more sophisticated insight into what these things do and how they do it. Yet attempting to provide philosophically adequate accounts of these various categories has proved extremely difficult, and such difficulties have been a major topic for my won academic specialty, the philosophy of biology. In these lectures I shall consider some of these kinds of things and the philosophical difficulties they present. A wider aim will be to try to locate some fundamental problems in our conception of life and its constituents, problems that more generally explain these difficulties in understanding central biological categories.

It is natural and traditional to think of life in terms of a structural hierarchy. We analyse an organism into a set of interacting organs and systems – livers, hearts, brains, circulatory systems, immune systems, and so on – and these in turn into smaller structural components, most notably cells. Cells, in turn are understood as enormously complex ensembles of interacting molecules. And this picture extends in both directions. Molecules are complex structures of atoms; organisms are components of species, ecological systems or social groups. And so on.

This vision has undeniably been fundamental to the extraordinary success the sciences have achieved in advancing our understanding of the natural world. This success has often been taken to lend support to a more general reductionist scientific methodology. Reductionism, in its classical form, is the explanation of the behaviour of complex entities in terms of the properties of their parts, and some philosophers have taken this position to its logical conclusion and suggested that ultimately the world is, in principle at least, fully describable and intelligible in terms of the smallest microphysical particles it contains.

Reductionism has, however, been much criticised, including in the past by myself.[1] I shall not explicitly pursue this critical project today. In opposition to reductionism I have, over a number of years, defended a quite different, pluralistic perspective.[2] According to this perspective there are many different kinds of things in the world, from physically simple things like electrons or quarks, to very complex things such as planets, elephants, or armies. Many or all these things, in my view, have equal claims to reality. As the basis of this position is the idea that many or all such entities have causal powers that are not simply consequences of the way their physical components are fitted together. This perspective gives biology, in particular, autonomy from the physical sciences. One objective of these lectures will be to explain and defend this point of view.

Let me begin by pointing out what is perhaps the deepest difficulty with the reductionist hierarchy. Contrasting with the idea that life consists of a hierarchy of things, we may observe that it is more realistic to consider it as a hierarchy of processes. In a typical cell in a human body many thousands of chemical reactions are taking place every second. Molecules are constructed, reshaped, or dissolved. The cells in which they reside, divide, develop, and die. All of these countless events take place within a much longer process, the life cycle of the organism: conception, birth, death, and an exquisitely complex sequence of stages in between. And as these life cycles give rise to new life cycles through reproduction[3] we begin to glimpse a much longer process still, evolution. This reminds us that these life cycles are not a sequence of replicas but rather a sequence of similar but subtly different processes, Just as the process that is the life cycle of an organism changes constantly, partly in reaction to the demands put on it by its environment, so the sequence of life cycles changes in response

1 See especially The Disorder of Things: *Metaphysical Foundations of the Disunity of Science,* Cambridge: Harvard University Press, 1993.

2 Op. cit. *See also Humans and Other Animals,* Oxford: Oxford University Press, 2002.

3 I should emphasise that by 'reproduction' I include, for the case of organisms such as ourselves, much more than the biological process which is the primary referent of this term. Following so called 'Developmental Systems' theorists, I take the concept of reproduction appropriate for evolutionary thinking to include everything that is required for the replication of the lie cycle. In the human case this might include, for instance, schools and hospitals. See S. Oyama, P.E. Griffiths, and R.D. Gray, *Cycles of Contingency: Developmental Systems and Evolution,* Cambridge, Mass.: MIT Press, 2001.

to the longer term and greater changes to the environment – changes constituted most significantly by the changing patterns of life surrounding it.

Reductionism has, from its beginnings, been greatly inspired by our success in building machines, and even philosophers who have abandoned the epistemological dream of reductionism, the explanation of everything in terms of physics, still often adhere to versions of mechanism, the view that the functioning of complex systems, including biological systems, should be understood by analogy with machines.[4] So it is worth reflecting for a moment on how different the workings of a machine are from the hierarchy of processes that I have just sketched. The parts of a machine are not unchanging, of course, but their changes constitute a relentless and one directional trend towards failure. A good machine starts with all its parts precisely constructed to interact together in the way that will generate its intended functions. The technical manual for my car specifies exactly the ideal state of every single component. As friction, corrosion, and so on gradually transform these components from their ideal forms, the functioning of the car deteriorates. For a while these failing components can be replaced with replicas, close to the ideal types specified in the manual, but eventually too many parts will have deviated too far from this ideal, and the car will be abandoned, crushed, and recycled.

Reductionism is almost precisely true of a car. We know exactly what its constituents are – they are listed in the manual – and we know how they interact: we designed them to interact that way.

4 Since delivering these lectures I have had occasion to look more closely at an influential version of mechanism that has been promoted recently by a number of philosophers, especially in a series of recent papers by Carl Craver and collaborators. (See P.K. Machamer, L. Darden and C.F. Craver, "Thinking about Mechanisms", *Philosophy of Science,* 67: 1-25, 2000; C.F. Craver, "Beyond Reduction: Mechanisms, Multifield Integration, and the Unity of Science", Studies *in the History and Philosophy of the Biological and Biomedical Sciences,* 36: 373-396, 2005.) This explicitly anti-reductionist mechanism is generally very congenial to the perspective developed in these lectures. The term 'mechanism' is used to stress the importance of distinguishing a set of interacting constituents that must be understood at several different structural levels. It seems to me that the disanalogies with machines that I stress in the text are sufficiently important to make the choice of term unfortunate, though this is, of course, a matter of no more than terminological taste.

Reflection on the dynamic and interacting hierarchy of processes that constitute life should make us suspect that a very different picture is required.

An extreme reaction to this disanalogy might be that we should question the very idea of dissecting life processes into static things. I shall not take such an extreme position. One reason I shall not is that, most strikingly in the last few decades, mechanistic and even reductionistic explanations have provided extraordinary insights into living processes. Indeed, our understanding of the molecular mechanisms underlying living processes has been growing at a rate that perhaps exceeds any explosion of knowledge in the history of science. This growing understanding of the mechanical or quasi-mechanical interactions of molecules promises ever growing abilities to intervene in life processes, for example in combating disease. Certainly the processes of life are highly dependent on these mechanisms. It is even arguable that science is inescapably mechanistic; certainly its most impressive and uncontested achievements have been based on mechanical models. But even if this is all true, the great differences between living things and machines should tell us something very important about such scientific insights. Mechanical models, assuming fixed machine-like ontologies, are at best an abstraction from the constantly dynamic nature of biological processes. And it is this pervasive fact about biological science that is central to explaining the philosophical difficulties in characterising the constituents of life that biologists hypothesise. If, indeed, science is essentially an examination of mechanisms, this points to ultimate limits in the ability of science to understand life. In the next lecture, however, I shall briefly consider some scientific ventures which promise a more realistic approach to biological processes.

Let me summarise the problem that I now want to address. The reductionist believes that in the end there is nothing in the world but the stuff of which things are made – let me call this basic physical stuff. Of course, the reductionist does not say, bluntly and absurdly, that houses, for example, don't exist. The claim is rather that a house is, ultimately, nothing but an aggregate of physical stuff, and all the properties of any house can, in principle, be fully explained by appeal to the properties and relations of basic physical stuff. So there is a possible, microphysically grounded, account of the world which would have no need to mention houses. I am insisting, on the contrary, that there is a whole hierarchy of increasingly complex things that really exist, and that have causal powers that are not reducible to the mechanical combination of the powers of their constituents. Yet I have also claimed that the things we distinguish in our descriptions of life, at least, are always to some extent abstractions from the dynamic processes that ultimately constitute life. This second claim may seem to undermine the reality of the members of the biological hierarchy to which the first claim attributes causal powers. I must now try to show how these theses can be reconciled.

Let me start with a very brief and abstract answer, and then illustrate what I mean with a homely example. The processes of life are of course massively heterogeneous. This heterogeneity is expressed, for example, when we inventory the thousands of chemical species to be found at any instant in a cell. Although such an inventory is a static snapshot of a dynamic entity – at best an idealised description of the cell, therefore – the molecules we distinguish are more or less transient foci of causal power, real nodes in the astonishingly complex causal nexus that drives the cellular processes. Crucially, they are not merely nodes in an upward flowing casual cascade from the microphysical, but equally in a downward flow of causal influence from complex things to simpler things. Now the homely example, from a very high level in the causal hierarchy.

Readers familiar with South Central Amsterdam will be familiar with Albert Cuypstraat. This street has an unusual capacity to attract people, a capacity which, I suggest, has significant similarities to the ability of a flower to attract bees, or the ability of a magnet to attract iron filings: all are causal powers of individual things. The particular causal power of Albert Cuypstraat will be obvious to anyone wandering around the streets in the immediate vicinity: while there will be a light scattering of people in these surrounding areas, immediately one reaches Albert Cuypstraat one will encounter a dense throng. The reason is no mystery, of course: this is a busy street market. The market could not exist without the people (and stalls, and products) that make it up, but equally there are properties of the market itself that attract the people to it.

The powers of this market are exactly matched to the powers of the people it attracts. They must know it is market, for instance, and how to get there. These are not difficult accomplishments: I myself

managed to acquire them within a few days of arriving in Amsterdam. But of course I had acquired many of the necessary skills years ago: knowing what a market is, how to buy things, and so on. My return to the market to forage after my first accidental encounter with it is, however, a more complicated achievement than, say, returning to a place where I had previously discovered edible berries. I would not be similarly drawn to return to a place where I had seen delicious looking food through the window of a private house, for instance, and I would not return to the market at four o'clock on Sunday morning. The market is a social institution of a kind that I have learned to negotiate reliably. By learning this I have also become – willingly, I should add – susceptible to the attractive casual powers of this institution. The market depends for its existence on the people who go there to buy and sell; but it is simultaneously the power of the market that attracts the people that constitute its continued existence. And, insignificant though these may seem, the market effects changes in the people it attracts – it may determine, for example, what they eat for dinner. This is the sort of thing that I mean by a node in the causal nexus. I shall suggest that this model, incorporating the development of two-way causal interaction between a complex thing and its constituents, is the right model for interactions at many different levels of structural organisation.

There is no better example of the consequences of the shift from a static to a dynamic view of life than the influence of Charles Darwin's revolutionary ideas on the subject of his most famous work, biological species. It may perhaps be thought that sorting organisms into species is more like constructing the automobile company's model catalogue than a parts list for one model. But, first, a majority of philosophers concerned with biology now hold that species should be seen as individual things, components of the evolutionary process.[5] And, second, sorting organisms into kinds raises many of the same issues as sorting, say, molecules or parts of molecules into kinds: classification is an essential part of scientific activity at any level of organisation. The classification of organisms is both the most widely discussed and the most ancient such project – indeed a project that some believe was delegated to Adam when God invited him to name

5 Classic statements of this thesis are M. Ghiselin, "A Radical Solution to the Species Problem", *Systematic Zoology* 23: 536-544, 1974. D.L. Hull, "Are Species Really Individuals?" *Systematic Zoology* 25: 174-191, 1976.

the animals. One crucial point that will emerge from consideration of this topic, and which should be less surprising viewed in the light of the general problem of abstracting objects from processes, is that there is no uniquely correct way of classifying organisms: different investigative interests dictate different and often cross-cutting modes of classification.[6]

There is an ancient philosophical tradition that understands classification as involving the identification of the essence of things of a kind: the essence is a necessary and sufficient condition of being a thing of that kind and also the feature that most fundamentally explains the properties characteristic of that kind. So, for instance, a certain atomic structure might be both necessary and sufficient for a piece of stuff to be iron and, at the same time, provides an explanation of why that stuff has properties – being magnetic, being easily oxidisable, and so on – characteristic of iron. Whether or not such an idea works for chemistry, one thing that almost everyone now agrees on is that nothing similar works for biology.[7] A sufficient explanation of this failure is the agreement that one biological kind can evolve gradually into another. The identification of a kind of organism existing at this moment is an abstraction from a continuous process linking these current organisms through time to a long series of very different organisms and, indeed, if we trace evolutionary history back to a common ancestor and thence forward to the present, connecting any two currently existing kinds of organism. There is no way of understanding this link as consisting of a definite number of distinct types, each defined by its unique essence.

Just as evolutionary theory has put an end to certain traditional ideas about biological classification, so it underlies more contemporary views. What most contemporary theorists agree is that biological classification should reflect the evolutionary relationships

6 For further explanation of this view, see Dupré, *Humans and Other Animals,* chs. 1 and 2.

7 A classic argument of this kind is D.L. Hull, "The Effect of Essentialism on Taxonomy: 2000 Years of Stasis", British Journal for the Philosophy of Science 15: 314-326; 16: 1-18, 1965. A number of recent commentators have suggested that the essentialism attributed to pre-Darwinian thinkers by recent anti-essentialists is something of a caricature, but this of course only strengthens the anti-essentialist position. See, e.g., M.P. Winsor, "Non-essentialist methods in pre-Darwinian taxonomy,". *Biology and Philosophy* 18: 387-400, 2003.

between different kinds of organisms. Evolutionary history has traditionally been represented as a tree, with branches representing evolutionary divergences and the smallest twigs representing the most finely distinguished kinds, species. For a while the dominant view, the so-called Biological Species Concept especially associated with one of the twentieth century's most influential evolutionists, Ernst Mayr, reflected a theory about the mechanism of evolutionary divergence.[8] The separation of branches of the tree, it was supposed, required that organisms on different branches be reproductively, and hence genetically, separated from those on other branches. Thus species were thought of as a reproductively connected group of organisms, reproductively isolated from all other groups. Unfortunately this idea often fits poorly with biologists' sense of what constitutes a species. Many groups of what seem to be well defined species in fact show continuous reproductive links and, on the other hand, what seem like homogeneous species often divide into separate populations with little or no reproductive connection between them. In addition there is a major problem with asexual species, the members of which appear to be reproductively isolated from everything except their direct descendants and ancestors.

Since the 1960s an alternative programme has advocated a more direct relationship between the evolutionary tree and biological classification. So-called cladistic classification, or cladism, a version of this idea and increasingly the dominant school among taxonomists, aims directly to identify the branching points in the evolutionary tree.[9] Ideally, a distinct name would be given to any set of organisms lying between two branching points on the tree. The terminal branches will be the species. Because the patterns of branching in different parts of the tree can be very diverse, this often fails to reflect prior notions about how many species there are and how different

8 See, e.g. E. Mayr, *Animal Species and Evolution,* Cambridge, Mass.: Harvard University Press, 1963.

9 The *locus classicus* for this idea is Willi Hennig, (1966). *Phylogenetic systematics.* Urbana: University of Illinois Press. Cladism is generally understood as a form of phylogenetic classification that insists that all groups by monophyletic, which is to say that they mustinclude all and only the descendants of an ancestral species. Less rigorous versions of phylogenetic classification, sometimes referred to as 'evolutionary taxonomy', relax this requirement so that it is possible to deny such apparently paradoxical claims as, for instance, that birds are a kind of dinosaur. The arguments below apply to both versions of phylogenetic classification.

they are from one another. But, cladists have tended to conclude, so much the worse for our existing notions about species.[10]

Before continuing with the discussion of classification, I must now introduce a topic that will be important throughout these lectures. There is an English expression, 'the elephant in the room'. The elephant refers to a problem which, as is the way with elephants, is extremely obvious, but which, for whatever reason, all participants in a discussion decide to ignore. There is an elephant in the room of biological classification – indeed it is an elephant that can be found in many areas of biology and which I shall rudely point out at several points in these lectures. So let me now describe this elephant.

This elephant is not one large object, but a huge number of very small ones, the microbes. Microbes have been the only kinds of organisms on this planet for the majority, perhaps 80%, of the history of life. And they continue to be the dominant life-form. It is calculated that even by sheer biomass microbes continue to constitute over half of contemporary terrestrial life. And the most extreme terrestrial environments remain too hot, cold, dark, or chemically hostile for other life-forms.

I should explain what I mean by a microbe. For now I shall think of microbes as including all single-celled organisms though I shall suggest later that this concept is not unproblematic. Two of the three branches of what is generally considered to be the most fundamental division among organisms consist of microbes. These are the Superkingdoms, or domains, Bacteria and Archaea. The third domain, the Eukarya, is also mostly composed of microbes, so-called protists, but also includes multi-cellular organisms, animals, plants and some fungi. To emphasise their almost cameo role against the backdrop of microbial life, I and my collaborator on this topic Maureen O'Malley are attempting to popularise the word 'macrobe' to refer to those organisms, such as ourselves, that are not microbes. It seems absurd that we should have a word for the great majority of life forms, but none for the small minority that this word excludes.[11]

10 A variety of philosophical discussion of the main positions on the nature of species can be found in M. Ereshefsky, *The Units of Evolution: Essays on the Nature of Species,* Cambridge, Mass.: MIT Press, 1991; and R.A. Wilson, *Species: New Interdisciplinary Essays,* Cambridge, Mass.: MIT Press, 1999.

11 For this proposal and more detailed elaboration of most of the points about

I should now explain the relevance of this elephant to classification. Both the Biological Species Concept and cladistics have difficulties with asexual reproduction. The problem has already been noted for the Biological Species Concept. Cladistics is threatened in a somewhat different way. To see this we need to look more carefully at what is meant by asexuality. Sexuality is normally thought of, biologically, as a device through which two parents contribute genetic material in the production of a new individual. Asexuality, by contrast to this, is often thought of as parthenogenesis, the production of offspring by a single parent. Sexual organisms sometimes abandon sexuality in favour of the latter method of reproduction, sometimes use it as an optional alternative. But even more than a device for facilitating genetic collaboration, sexual reproduction is part of a system for restricting the flow of genetic material. As the Biological Species Concept, with its emphasis on reproductive isolation makes clear, sexual macrobes go to great trouble to make sure that their gene exchange takes place with very similar organisms. Indeed one influential descendant of the biological species concept is called the mate-recognition concept, recognising the diversity of mechanisms by which organisms, macrobes anyhow, ensure that they find the right partners for genetic collaboration.[12] The asexuality typical of microbes[13] should be seen by contrast to this aspect of sexuality. As has become increasingly clear over the last several decades, from the perspective of genetic exchange, microbes are not so much asexual, as massively promiscuous. Microbes have a number of different mechanisms for exchanging genetic material, and they use them fully. They have mechanisms for so-called conjugation, exchanging genetic materials in a way analogous to macrobe sexuality; DNA is transferred from one organism to another by phages, viruses specific to microbes; and

microbes made here and later in these lectures, see M. O'Malley and J. Dupré, 'Size Doesn't Matter: Towards a Philosophy of Microbiology', *Biology and Philosophy,* forthcoming 2006; and J. Dupré and M. O'Malley, "Metagenomics and Biological Ontology", *Studies in the History and Philosophy of Biological and Biomedical Sciences,* forthcoming 2007.

12 See H.E.H. Paterson, "The Recognition Concept of Species", in *Species and Speciation,* ed E. Vrba. Transvaal Museum Monograph No. 4. Pretoria: Transvaal Museum.

13 In this and the following paragraph, my references to microbes apply mainly to the simpler organisms, the Bacteria and Archaea, lacking nuclear membranes, which are generally referred to as Prokaryotes. Matters are somewhat more complex and diverse for microbial Eukaryotes (protists). I use the term 'microbe' since it is much more familiar, and no serious confusion is likely to be engendered.

many microbes can pick up free, or 'naked', DNA from the environment. These mechanisms can facilitate DNA exchange between distantly related forms, even across the three domains at the base of biological classification. Because of the prevalence of these processes, typical microbes will include genetic material from numerous distinct lineages.

The problem with the phylogeny of microbes, then, and one reason that few if any microbial taxonomists endorse cladism, is that there is no unambiguous evolutionary tree on which to superimpose a taxonomic system: microbes have too many diverse ancestors.[14] Or, at any rate, they do if any past organism from which they derived genetic material is counted as an ancestor. Microbes for a long time seemed practically almost impossible to classify simply because of their dimensions. The development of tools capable of providing detailed inspection of genomes offered a solution to this problem. Comparison of microbial genomes would allow biologists to track the phlyogenetic histories of particular bits of microbial genome sequence, and infer the phylogeny, the evolutionary history, of microbes. In the early days of genomic classification of microbes a set of ribosomal genes was identified as particularly suitable for this purpose, and these continue to this day to provide an important resource for classificatory work. However, it is also becoming clear that the phylogenetic history produced using these genes is to an important extent an artefact of that choice. Using different genomic criteria the same organisms can appear in very different parts of the phylogenetic tree. This should be no surprise. What it indicates is merely that the genetic relations between microbes do not really form a unique tree at all, but rather a web. It may be useful for particular purposes to represent the evolutionary relations between microbes in the form of a tree, but we must remember that this is an abstraction from a much more complex reality.

14 This remains a controversial matter among microbiologists. A strong advocate of the impossibility of defining a microbial phylogeny is Ford Doolittle (see e.g., W.F. Doolittle, "Phylogenetic Classification and the Universal Tree", *Science* 284, 2124-28, 1999.). An influential resister is Carl Woese, the scientist responsible for distinguishing between the microbial superkingdoms Archaea and Bacteria mentioned above. As will be clear, I find the former argument compelling.

I do not, in fact, believe that there is a uniquely correct way of classifying even macrobes,[15] but the case is even clearer for microbes.[16] The failure of evolution to provide us with a unique and unequivocal method of biological classification enables us to see that there are many real discontinuities across the vast spectrum of different organic forms. And different discontinuities can ground different ways of classifying these, suited to different purposes, again both scientific and mundane. Certainly we can imagine that God, had he created the plants and animals, would have known how many distinct kinds he had come up with. Phylogenetic classification can be seen as a device that might have reconciled this ancient doctrine, to some degree, with post-Darwinian biology. But it cannot do that job. It may be an irreplaceable approach to biological classification, but it is not the only one possible, and it is an abstraction from the real complexity of biological relations. Once it is clear that only under quite special circumstances does evolution determine a unique way of classifying organisms, we should reject the cladist's indifference to the convergence of evolutionary theory on existing categories. Classifications serving different biological interests – ecology rather than evolution, for instance – and even more practical interests such as those of the forester, the herbalist, or the chef may equally be grounded in distinct natural discontinuities.

To mention one practical issue that is easily misunderstood by failing to understand this point, we might consider the problems of biological conservation. One might imagine that the aim of conservation is to save as many species as possible. Though I don't claim to know what the goal should be – I'd guess that it would be a mixture of aesthetic, utilitarian, ethical, and probably other aspects – the simple idea just mentioned surely won't do. Most fundamentally this is because it is incoherent: there is no unique way of counting the species. But even ignoring this, from any sensible conservation perspective not all species are equal. Apart from quite legitimate aesthetic arguments that the loss of tigers or gorillas would be more serious than the loss of one member of a large group of beetles, the former are plausibly far more biologically distinctive than the latter. There is, at any rate, no absolute conception of the species that contradicts

15 See my *Humans and Other Animals,* op. cit., chs. 3 and 4.

16 This argument is spelled out in greater detail in O'Malley and Dupré, "Size Doesn't Matter" op. cit.

this idea. Conservation of microbial diversity is an issue, and potentially a very important one, that has hardly been considered – perhaps because the actual objectives of conservationists typically are predominantly aesthetic.

Does the denial that species represent a unique division of biological reality mean that they are unreal, or play no part in biological explanation? I have mentioned the widely held view among philosophers of biology that species are not kinds at all, but individuals. This view is linked to the idea that species are branches of the evolutionary tree and therefore inherits the limitations of that idea. However to the extent that the evolutionary tree has branches at all it is sometimes useful to think of species as spatio-temporally extended individuals that can be identified with these branches. It is useful for theorising much of macrobial evolution, and macrobial species should often be treated as individual things with significant causal powers. But macrobial species can be treated as individuals, because they do things: for example, they speciate, divide into two distinct species. Processes of macrobial speciation, the emergence of new biological forms, are often very real, and important for understanding biological diversity. Contrary to one popular idea, speciation is not always a slow or gradual process. About half of the species of flowering plants, for instance, appear to have arisen by a process of polyploidy, the doubling in size of the genome.[17] Such a process creates instantaneous infertility with the ancestral species, and may produce immediate changes in the phenotype. Sometimes this is the doubling of the genome of a single parental organism, sometimes it happens through the hybridisation of two related plants. Because many plants are self-fertile, the appearance of a single polyploid individual may, if circumstances are propitious, found an entire new species.

The preceding case provides a nice reductive explanation of organismic diversity in terms of molecular processes. But species also play a part in explanations at their own level, and can be affected by their involvement in processes that could be thought of as at a higher level. Species interact with one another as, for instance when the members of one prey on or parasite the members of another. This complex interaction will help to determine the dynamics of the size of the

17 See K.L. Adams and J.F. Wendel, "Polyploidy and Genome Evolution in Plants", *Current Opinion in Plant Biology* 8: 135-141, 2005.

participant species. In the longer term, interactions between predator and prey species will direct the evolution of each – as is well documented in the phenomenon referred to, in an unfortunately common militaristic vein – as an evolutionary arms race. The lineages of cheetahs and gazelles, for instance, exhibit ever greater speeds as their lives depend on capturing or escaping one another. These may be interpreted as examples of large complex things – species – interacting with one another, but their significance does not depend on this interpretation. The more general point is that classifying a thing as a cheetah identifies a set of processes in which it can be involved. Classifying it in other ways might identify different processes. Such possibilities of multiple, perhaps cross-cutting, classification become more salient as classification becomes less determinate. This will be most clearly the case among the microbes.

Particular characteristics of human societies have also affected biodiversity in ways that are best described by identifying species, indicating a very different kind of interaction in which species (or merely organisms by virtue of being members of a species) may be involved. It may be that if tigers go extinct it is in part due to the belief, among significant proportions of the human species, that consuming tiger penises has great medical benefits. However decisive this factor may or may not be, it is certainly entirely possible for quite specific human beliefs to affect the trajectory of a non-human species, and there are surely many real instances of this happening. Beliefs about the relative desirability of rainforests and marketable timber are exterminating species as I write. The practice of selective breeding, now often involving targeted intervention at the molecular level, provides another obvious set of examples.

I want now to turn to a quite different biological concept, the concept of a gene, and again want to demonstrate the lack of any unique motivation underlying this concept, and the consequently distinct kinds of object that may serve these diverse concerns. Historically, the concept of a gene was introduced in the context of the experiments on breeding in the early twentieth century, deriving from the rediscovery in 1900 of Mendel's work. The gene was a hypothetical object that explained the distinctive patterns of inheritance of features of organisms discovered by Mendel. It was thus conceived as the transmittable cause of a specific phenotypic difference. For some time it remained a matter of debate whether genes should be thought

of as material things at all, or rather conceived instrumentally as mere calculating devices. But by the time of the unravelling of the structure of DNA in 1953, it had become widely agreed that genes were material things and that they were located on chromosomes. This classical or Mendelian concept – the underlying cause of a difference – remains in use today, particularly in medical genetics, but as knowledge in molecular genetics has expanded exponentially in the last half century it has actually become more difficult to relate the classical gene to any particular molecular entity.[18]

Consider, for instance, the cystic fibrosis gene. This is a recessive gene, meaning that to suffer its effect, the severe congenital disease cystic fibrosis, one must receive the gene from both parents. The gene pretty accurately obeys Mendelian patterns of inheritance. But what is it? That is a harder question. Cystic fibrosis results from the failure of the body to make a particular protein, cystic fibrosis transmembrane conductance regulator, involved in the production of channels that conduct salt through certain membranes. The cause of this failure is a defect (in both copies) of a bit of DNA sequence called the CFTR gene. However, there are more than a thousand known defects in this sequence that produce cystic fibrosis, though producing variably severe symptoms. So the Cystic Fibrosis gene is actually a large set of variations in a bit of DNA sequence. A set of variations is at least an unusual kind of object.

The gene CFTR, on the other hand, is a rather different kind of thing. It is generally defined as a sequence of 188,698 base pairs on the long arm of human chromosome 7. This sounds a much more material kind of thing. However, it should be noted that there is no exact sequence of base pairs necessary to constitute a functioning CFTR gene. The genetic code, as is well known, is redundant, so that many changes will have no effect at all on the functioning of the gene, and there are very likely to be changes that do make a difference to the transcription of the gene, but do not prevent its proper functioning. In short, then, the CFTR is a definite stretch of DNA sequence, though one that allows a good deal of variation; the cystic fibrosis

18 Problems with the concept of a gene are discussed by Lenny Moss, *What Genes Can't Do,* Cambridge, Mass.: MIT Press, 2003, and in the essays in P. Beurton, R. Falk, and H.-J. Rheinberger (eds.), *The Concept of the Gene in Develoment and Evolution,* Cambridge: Cambridge University Press, 2000.

gene is any of a large set of dysfunctional variations in the same part of the genome – or, perhaps, as we shall see, in quite different parts of the genome.

The CFTR gene, at any rate, looks a good deal more like the sort of thing people expect a gene to be in the age of genomic sequencing: a specific part of the genome with a specific molecular function. However, as so often in biology, things are not as simple as they may seem when we try to generalise this concept of the gene. A hint of the trouble can be seen in the fact that the number of genes in the fully sequenced human genome is currently estimated as being somewhere in the range of 20,000-25,000. It is often noted that this is a much smaller number than had been assumed necessary for the estimated number of gene-related human traits. But a more important puzzle for the moment is the vagueness of this estimate. Why, one may wonder, can they not just count them? The complexities that stand in the way of this task can only be sketched in the briefest way, but even such a sketch will be sufficient to make my overall philosophical point.

Properly molecular genes are often thought of as a part of the genome that codes for a particular protein. However, as a definition this raises numerous problems. Typically genes (in macrobes, at any rate) are composed of alternating sequences called exons and introns. After the gene is transcribed, into RNA, the introns are edited out, and the exons are then translated into protein molecules.[19] However, in many, perhaps most, cases there are alternative ways of splicing the exons into finished RNA sequence, and some bits may be left out. Further changes may be made either to the RNA sequence or to the subsequently produced proteins. In some cases elements of other genes may be incorporated. Thus the relation between molecular genes and proteins is not one to one, but many to many. Some genes are involved in making hundreds of distinct proteins.

It still might seem that the genes could be counted, even if they were then found to have much more diverse functions than might once have been supposed. But things get worse. First, genes can overlap. So a certain sequence can be part of two quite distinct primary

19 There is much to be said, and a good deal that has been said, about all these semantic metaphors – editing, transcribing, translating, coding – but that is not a topic I shall address here.

RNA transcripts with quite different subsequent histories. Worse still, DNA is not always read in the same direction. So a sequence may be part of one gene read in the normal so-called 'sense' direction, but part of another when read in the opposite anti-sense direction. Philosophers Paul Griffiths and Karola Stotz have investigated empirically how many genes biologists claim to see in problematic bits of sequence and the answer, perhaps unsurprisingly, is that different biologists see different numbers of genes.[20] It would perhaps be possible to regiment the concept sufficiently so that the answer to such questions could be decided mechanically, but this would only conceal the real philosophical problem: Nature has declined to divide the genome into a unique set of constituent entities. Different, overlapping, and non-contiguous elements of the genome are involved in different biological functions. A realistic conclusion is that a molecular gene is any part of a genome that a biologist has some reason to talk about. (Just as, indeed, it is sometimes said that a species is any group of organisms a competent taxonomist decides to put a name to.)

In fact this discussion has only scratched the surface of the diversity of entities that may legitimately be referred to as genes. The protein coding genes that I have been discussing make up only a few percent of the DNA in many macrobial genomes, including our own. Until quite recently it used to be said that the remaining large majority of the genome was junk, a testament to the pernicious activity of genetic parasites.[21] It is becoming increasingly clear that much of this so-called junk serves important biological functions. At the very least it is essential for structural features of the genome. But it also appears that many parts of it are transcribed into RNA and that these RNA molecules play important roles in the functioning of the cell. It

20 K. Stotz, P.E. Griffiths, et al., "How scientists conceptualise genes: An empirical study." *Studies in History & Philosophy of Biological and Biomedical Sciences,* 35: 647-673, 2004.

21 These were, and still sometimes are, thought of as the truly selfish genes – merely competing with one another to occupy space in the genome. The transposable elements, the discovery of which eventually won Barbara McClintock the Nobel Prize, come closest to realising this image, apparently concerned only with making more copies of themselves in the genome and often constituting a large proportion of the genome. Even these, however, are increasingly suspected of serving some more 'altruistic' purpose – of contributing something to the wider organism.

has also been know for a long time that non-coding sequences in the genome serve to regulate the expression of protein coding sequences, and a growing number of different kinds of such regulatory sequence are now distinguished. So, in short, there are many very different kinds of sequence that molecular biologists have reason to distinguish, and hence many different kinds of genes.

Nature, then, no more determines how to divide the genome into genes, than she does organisms into species. Particular parts of the genome can, however, as in the other examples I have considered, provide nodes on the causal nexus that are appropriate points of focus for particular investigative purposes. Reductionist explanation of the power of genes is familiar enough. Indeed, it was in large part the chemical explanation of the stability, complexity, and replicability of the DNA molecule that made the description of its structure such an extraordinary scientific achievement. A moment ago I pointed to the much more specific way in which various genetic anomalies help to explain a disease such as cystic fibrosis. Less familiar is the extent to which the DNA in a cell is in constant two-way interaction with other constituents of the cell. What was for a long term known (approvingly) as the Central Dogma of molecular biology was the view that information flowed in one direction only, from DNA to RNA to protein. In accordance with this dogma it was supposed that the function of RNA was primarily to carry information from DNA to proteins. But today the study of the vast number of different RNA molecules and their influence on gene expression is one of the most rapidly developing fields in molecular biology. One class of these molecules, the so-called RNAi's, which can block the expression of a coding gene, are currently considered one of the most exciting prospects for molecular medicine. Many proteins, too, interact with nuclear DNA and affect the transcription of particular sequences.

It is still often supposed that the genome is the ultimate director of the process by which an organism develops, a supposition expressed in metaphors such as blueprints, recipes or programmes. This is, in fact, an expression of the reductionist philosophy that I reject. For reductionism, a complex process such as organismal development can only be explained by causal influences from smaller constituents. DNA is seen as the largely unchanging structure that mediates this transfer of casual power from below. The vision of DNA as a node through which causal influence passes both upwards and downwards

of course contradicts this picture, but does so on the basis of increasingly undeniable scientific evidence. Genes, in the end, are the diverse, nested, and overlapping sites in the genome where these casual influences are focused, at different times, in different ways, and often in different ways at the same place.

A good way to get a sense of the implications of this picture is to contrast it with the reductionist picture of genetics that has grounded an extremely influential view of evolution but one that must now be seen as highly simplistic. If the genome were indeed an unchanging repository of information, then from the perspective of evolutionary theory one could see evolution as simply a temporal sequence of genomes. The organisms for which they were described as the blueprints, or the recipes, would develop as the genome dictated, and take their chances in the lottery of life, and the best ones would be selected. But all they would pass on to the next generation were the most successful blueprints. Those familiar with the writings of Richard Dawkins should recognise this picture. Only the selfish genes in their immortal coils live on through evolutionary time.[22]

This picture should already look suspect when one sees that DNA is only one, admittedly very important, component of an interacting system of molecules, and that the whole system is passed on to the offspring in the cytoplasm of the maternal egg. But it now turns out that changes to the DNA itself that occur during the life of an organism can be transmitted to offspring. The best studied of such processes is methylation, a chemical modification of the DNA that prevents the expression of particular gene sequences. This is most familiar in stories about imprinting, the differential methylation of paternal and maternal DNA, claimed to reflect competing male and female evolutionary interests. But it is by no means restricted to this.

22 See Richard Dawkins, *The Selfish Gene,* Oxford: Oxford University Press, 1976, and many subsequent books. There is a good deal of much more sophisticated theoretical work on evolution currently under way, though none unfortunately, that threatens to compete with Dawkins's sales volumes. Particularly important recent contributions include M.J. West-Eberhard, *Developmental Plasticity and Evolution,* New York: Oxford University Press, 2003; and Eva Jablonka and M.J. Lamb, *Evolution in Four Dimensions: Genetic, Epigenetic, Behavioral, and Symbolic Variation in the History of Life,* Cambridge, Mass.: MIT Press, 2005. Recent insights that have undermined the simplistic, gene-centred picture of evolution are not, of course, limited to those briefly mentioned in the present text.

One famous case is a study of the effects of maternal care on rats. Absence of such care, especially of licking by the mother, produced nervous, fearful offspring and, unexpectedly, these characteristics appear to be passed on to the offspring of the neglected animals.[23] It has been verified that maternal care produces methylation of genes in the hippocampus, though the mechanism by which this change is passed on to subsequent generations remains obscure. There is also a famous, though also still controversial case in recent human history, the Dutch famine of 1944-5. Unsurprisingly, mothers who experienced this famine tended to have small babies. Much more surprisingly, their generally well-fed children also tended to have small babies. Many have concluded that transmitted methylation patterns induced by the shock of malnutrition explain this phenomenon.[24] Students of such phenomena are even beginning to call themselves neo-Lamarckians, transgressing perhaps the most inviolable taboo of twentieth century biology.[25] The so-called epigenome, the set of inherited mechanisms that determine how genes are expressed, is another booming area of research. One of the successor projects to the Human Genome Project is the Human Epigenome Project, that aims to map the methylation sites on the human genome. Epigenomics more generally, the study of the interactions between the cellular environment and the genome, is poised to become an even more significant significant field of research than genomics itself.

The picture I have tried to sketch will not please those who are wedded to the crystalline clarity that the mechanistic vision of life offers. Shifting levels of organisation with shifting, metamorphosing and even indeterminate constituents may seem like unlikely materials for understanding the exquisitely ordered and robust phenomena of life. And causal processes running upwards to exploit the diverse and specific capacities of countless chemicals and structures and downwards to provide externally enforced constraints on the actions

23 F.A. Champagne, I.C. Weaver, J. Diorio, S. Dymov, M. Szyf, M.J. Meaney, "Maternal Care Associated with Methylation of the Estrogen Receptor-Alpha1b Promoter and Estrogen Receptor-Alpha Expression in the Medial Preoptic Area of Female Offspring", *Endocrinology,* 147: 2909-15, 2006.

24 See, e.g., G. Vines, "Hidden Inheritance", *New Scientist,* 2162: 27-30, 1998.

25 E.g., E. Jablonka and M. Lamb, *Epigenetic Inheritance and Evolution: The Lamarckian Dimension,* New York: Oxford University Press, 1995.

of those structures and chemicals, may seem to be hopelessly intractable objects of real insight and understanding.[26]

Yet it is also worth considering how inadequate the mechanistic paradigm is for understanding these phenomena. As I explained with the example of a car, deterioration and failure are the inevitable history of a machine. Organisms, while perhaps all die in the end, show no such inevitable tendency. Some live for millennia with no obvious deterioration of vital functions, and it is now a matter of lively debate whether human aging is an inevitable process of deterioration, or rather a biological function that we might, if we chose, find ways of subverting. There are powerful reasons for thinking that emancipation from the mechanistic paradigm is a precondition for true insight into the nature of biological processes.

I am not, of course, the first person who has offered more complex and dynamic visions of life than are possible within the constraints of mechanism, and I shall end this chapter with one such vision that expresses, with a poetic elegance to which I can only aspire, a view remarkably congruent with much of what I have said today. The words are from Walter Pater, the British aesthete and philosopher, written in his *Conclusion to the Renaissance,* in 1893:

> "What is the whole physical life … but a combination of natural elements to which science gives their names? But those elements … are present not in the human body alone: we detect them in places most remote from it. Our physical life is a perpetual motion of them – the passage of the blood, the waste and repairing of the lenses of the eye, the modification of the tissues of the brain under every ray of light and sound – processes which science reduces to simpler and more elementary forces.

26 Carl Craver and William Bechtel, leading proponents of the contemporary mechanism mentioned in note 3, above, reject top down causation but equally, and for similar reasons, deny bottom-up causation. They consider causation, strictu sensu, to be a concept only applicable within one structural level. Both top-down and bottom up 'causation' they prefer to describe as 'mechanistically mediated effects'. Again, I suspect our substantive views are quite similar, though I am inclined to a much more catholic conception of causality, and am somewhat sceptical of any sharp divide between same and different levels. (Carl F. Craver and William Bechtel, 'Top-Down Causation without Top-Down Causes', *Biology and Philosophy,* forthcoming 2006.)

Like the elements of which we are composed, the action of these forces extends beyond us: it rusts iron and ripens corn. Far out on every side of us those elements are broadcast, driven in many currents; and birth and (gestation) and death and the springing of violets from the grave are but a few out of ten thousand resultant combinations. That clear, perpetual outline of face and limb is but an image of ours, under which we group them – a design in a web, the actual threads of which pass out beyond it. This at least of flame-like our life has, that it is but the concurrence, renewed from moment to moment, of forces parting sooner or later on their ways… It is with this movement, with the passage and dissolution of impressions, images, sensation, that analysis leaves off – that continual vanishing away, that strange, perpetual weaving and unweaving of ourselves."

Spinoza Lecture II

The Constituents of Life

In the previous chapter I tried to explain some of the difficulties in defining central concepts in biology, and also offered a general hypothesis as to why these difficulties arise. The general hypothesis is that many of these difficulties stem from the conflict between on the one hand, life itself as a hierarchy of dynamic and constantly changing processes and, on the other hand, our scientific understanding as grounded on a picture of mechanistic interactions between fixed and statically defined components. While not wishing to deny the extraordinary insight that mechanistic models have provided into life processes, I tried to explain the deep differences between living systems and the machines that have been such a central source of inspiration for science generally. Mechanistic models have given us extensive knowledge of many of the elements of which living systems are composed, but they are inadequate to provide a full picture of life as a dynamic system.

Key concepts in biology, I suggested, are static abstractions from life processes, and different abstractions provide different perspectives on these processes. This is a fundamental reason why these concepts stubbornly resist unitary definitions. They specify, more or less, the level at which we are abstracting, but nature does not determine for us a unique mode of abstraction. This problem is central to explaining the philosophical difficulties that have been found in attempts to provide unique definitions of two central categories in biology, the species and the gene, difficulties which I summarised in the last chapter. While arguing that there is no unique and privileged way of dividing biological reality with these terms, I claimed nevertheless that there were many and diverse real biological entities falling under these concepts. In the most important cases, this reality consists in the more or less transitory focus that such entities, for example the particular parts of genomes sometimes identified as genes, provide for causal processes. But these entities must be understood not only as inheriting causal powers from their structural components, but also as recipients of causal influence from the larger entities of which they are part. This two-way flow of causal influence through a shifting and diverse array of entities presents a very different picture of life from the pristine mechanism which still influences so much scientific thinking. In the second part of this chapter I shall say something about how we might conceive the prospects for scientific progress when confronted with such a picture.

In the first part of the chapter, however, I shall enrich the general view being developed by looking at some crucial levels between the extremes of species and gene so far discussed. I shall begin with probably the most generally familiar kind of biological entity, the individual organism.

At first sight it will seem quite obvious that I, or my cat, or George W. Bush, are discrete biological entities whatever else is. To adapt US Supreme Court Justice Potter Stewart's famous remark about pornography, "I know one when I see one". But when we consider a little more closely what is to be included in these entities, matters become less clear. A natural way of describing the limits of the individual, John Dupré, would be to imagine the surface that includes all the parts that move together when John Dupré moves, and treat all the material included within that surface as part of John Dupré. This is a good legal definition: if someone violates that space, for example with a sharp instrument, they are considered grossly to have violated my rights.

In the previous chapter I introduced an 'elephant in the room' – the microbes, the overwhelming majority of living things. The elephant is still very much in the part of the room I am now describing. Within the surface I just mentioned, my own, 90% of the cells are actually microbes. Most of these inhabit the gastro-intestinal tract, though within that, and elsewhere in the body, are a wide variety of niches colonised by microbial communities. Because of the diversity of these microbial fellow travellers, as many as 99% of the genes within my external surface are actually bacterial.[27]

We are sometimes told that the human body should actually be considered as a tube, so that the inside of my mouth or gut should rather be considered as part of the outside surface of my body. This will certainly reduce the microbial load in our own self-images, but it is somewhat counterintuitive. Considering the legal perspective just mentioned, it would be a very implausible defence against charges of various forms of serious sexual assault, for example. And it is biologically questionable. If we examine the inside surface of the gut we will discover complex and ordered communities of bacteria without

27 J. Xu, J.I. Gordon, "Honor Thy Symbionts", *Proceedings of the National Academy of Sciences* 100: 10452-10459, 2003.

which the interface between ourselves and the things we eat would be seriously dysfunctional. Our symbiotic microbes are essential to our well-being. Particularly interesting is the growing understanding that symbiotic bacteria are required for our proper development. It was recently reported, for example, that environmentally acquired digestive tract bacteria in zebrafish regulate the expression of 212 genes.[28] In fact, for the majority of mammalian organism systems that interact with the external world – the integumentary (roughly speaking, the skin), respiratory, excretory, reproductive, immune, endocrine, and circulatory systems, there is strong evidence for the coevolution of microbial consortia in varying levels of functional association.[29] For these reasons, some biologists are now proposing a second human genome project – the human biome project – that will catalogue all the genetic material associated with the human, including that of their microbial partners. At any rate, as a functional whole, there is much to be said for thinking of the whole community that travels around with me as a single composite entity.

I have mentioned that microbes remain by far the most versatile and effective chemists in the biosphere. The ability of multicellular organisms like ourselves to process food is entirely dependent on their cooperation. Being high on the food chain, we humans tend to consume highly processed foods that require less help from our microbial symbionts to metabolise than would more challenging inputs. It is worth recalling, though, that if we eat, say, cows, the amino acids we absorb were synthesised by microbes in one of the four stomachs of the animal, and if perhaps for moral reasons we prefer to get our amino acids from beans, this is only possible due to the nitrogen-fixing activities of bacteria associated with the roots of these plants. At any rate, as is familiar to users of powerful antibiotics, deficiencies in our gut bacteria are a serious problem even for human digestion.

Let me add a brief word about plants in this context. The best-known metabolic capacity of plants is the ability of their long captive microbial symbionts, chloroplasts, to capture the energy of sunlight. But plants also feed through their roots of course. And the roots of

28 J.F. Rawls, B.S. Samuel, J.I. Gordon, "Gnotobiotic Zebrafish Reveal Evolutionarily Conserved Responses to the Gut Microbiota", *Proceedings of the National Academy of Sciences,* 101: 4596-4601, 2004.

29 M.J. McFall-Ngai, M.J., "Unseen Forces: the Influence of Bacteria on Animal Development", *Developmental Biology* 242: 1-14, 2002.

a plant lie in the midst of some of the most complex multispecies communities on the planet. Bacteria and fungi not only form dense communities in the soil surrounding plant roots, but are also found within the roots themselves. These fungi put out nutrient seeking tendrils, or hyphae, through the roots and into the outside soil. This is perhaps the most striking illustration of the idea that the interface between large multicellular organisms and their environments is typically mediated – and essentially so – by microbial communities. Given the permeation of the boundaries between ourselves and the external environment by comparably complex multi-species communities, and their essential role in managing the chemical and cellular traffic across these boundaries, we should at least question our intuitive sense that they are not part of us.

There is a theoretical ground for the assumption that an individual should be taken to exclude its obligatory symbionts. This might be stated as the thesis, one individual, one genome. It is sometimes said, for example, that a group of trees, originating from the same root system, is *really* only one individual. The motivation for this stipulation is an evolutionary perspective, according to which evolution involves selection between genomes, but the stipulation is a problematic one. First, there are clearly important alternative perspectives that must sometimes be accommodated. From an ecological perspective, for example, we should surely prefer one trunk, one tree. We might also note that even complex animals are to some degree genomic mosaics. In extreme cases this may be the result of abnormal reproductive events, as is known to the great cost of a few human mothers who have failed genetic tests for parenthood of their children.[30] Transplant surgery, including blood transfusion, produces genomic mosaicism in some humans. More mundanely, mutation during development produces some genomic diversity, and in plants that reproduce vegetatively – for example by root suckers or rooting branches, this may provide material for natural selection.

30 There are several cases of mothers having lost custody of children on the grounds that they were 'proved' not to be the biological mothers as a result of different parts of their mosaic genomes appearing in the child and the genetic test for parenthood. See H. Pearson, "Human genetics: Dual identities", *Nature* 417: 10-11, 2 May, 2002.

It is also a familiar fact that the same genome may pertain to different individuals. Close to home, we do not consider monozygotic, or so-called identical, twins to be a single individual. This draws attention to the very important point, but not one I shall dwell on today, that there is much more to development than the unfolding of the genome.[31] But much more generally, we should again remind ourselves of the elephant. The vast majority of organisms do not produce an entirely novel genome in the process of reproduction. Although microbial genomes are extremely fluid over time, the basic process of reproduction is one of genome duplication. In short, the relationship between genomes and organisms is not one to one, but at least one to many. I want to suggest that it is a further but well motivated step to admit this relationship as many to many: not only can one genome be common to many organisms, but one organism can accommodate many genomes.

I hope I have anyhow said enough to dispose of a simple criterion that might give a simple answer to the question, What is an organism? The correct answer, I suggest, requires seeing that there are is a great variety of ways in which cells, sometimes genomically homogeneous, sometimes not, combine to form integrated biological wholes. The concept of multicellular organism is a complex and diverse one which, incidentally, provides no conceptual obstacle to the broader conception of the human individual sketched above.

The broader ramifications of this suggestion are once again best discerned by looking more closely at the elephant. I have described microbes as single-celled organisms, and this is how the organisms to which I wished to refer are generally conceived. However there are compelling grounds for revising this view. Microbes are most commonly found as parts of communities, containing either one or many

31 This point has been fundamental to recent critiques of classical models of evolution, from the mid-twentieth century synthesis to Dawkins's gene selectionism. See the 'Developmental Systems Theory' developed in S. Oyama, *The Ontogeny of Information: Developmental Systems and Evolution,* Cambridge: Cambridge University Press, 1986; and S. Oyama, P.E. Griffiths, and R.D. Gray (eds.), *Cycles of Contingency: Developmental Systems and Evolution,* Cambridge, Mass.: MIT Press, 2001. For application of this perspective to criticism of evolutionary psychological theories of human nature see my Human Nature and the Limits of Science, Oxford: Oxford University Press, 2001.

distinct types of microbe, communities which approximate many of the familiar features of multicellular organisms.[32]

There are a number of converging types of evidence that support this perhaps surprising proposal. Probably the most important has been mentioned, the frequency of genetic exchange between microbes, particularly within associated communities of microbes. An increasing number of biologists are beginning to suggest that the genetic resources of a microbial community should not be thought of as partitioned into individual genomes in individual cells, but are rather a community resource, a genetic commons.

The project of metagenomics, the attempt to collect all the microbial genetic material in entire environments has become quite widely discussed, not least due to its association with Craig Venter, the leader of the free enterprise wing of the Human Genome Project, and a larger than life figure in contemporary biology. Venter embarked on a well-publicised expedition around the world's oceans collecting microbial genetic material from the water.[33] This is sometimes seen as no more than a collecting or gene-prospecting exercise, and indeed Venter discovered an astounding quantity of unfamiliar genetic material and many previously unfamiliar types of microbes. But as suggested by my reference to a genetic commons, and by the great mobility of genetic material between microbial cells, it is but a small step from the metagenome, this totality of local genetic material, to the metaorganism, a multicellular organism composed of the community of microbes that shares this resource.

The clearest context in which to present the idea of microbial metaorganisms is with the phenomena of biofilms. Biofilms are closely integrated communities of microbes, usually involving a number of distinct species, which adhere to almost any wet surface. Biofilms are ubiquitous, from the slimy rocks and stones found under water and the chemically hostile acid drainage of mines, to the internal surfaces of drinking fountains and catheters; indeed biofilms are

32 J.A. Shapiro, (1998), "Thinking about Bacterial Populations as Multicellular Organisms", *Annual Review of Microbiology,* 52: 81-104, 1998. For discussion see O'Malley and Dupré, op. cit.

33 J.C. Venter, et al., "Environmental Genome Shotgun Sequencing of the Sargasso Sea", *Science* 304: 66-74, 2004.

where most microbes generally like to be.[34] In addition to the genetic exchange already discussed, the constituents of biofilms exhibit cooperation and communication. These are most clearly exemplified by the phenomena of quorum sensing, in which microbes are able to determine the numbers of cells in their communities and adjust their behaviour – including reproduction – in appropriate ways. The general idea can perhaps be best appreciated by quoting a scientific paper on a very familiar kind of biofilm:

> Communication is a key element in successful organizations. The bacteria on human teeth and oral mucosa have developed the means by which to communicate and thereby form successful organizations. These bacteria have coevolved with their host to establish a highly sophisticated relationship in which both pathogenic and mutualistic bacteria coexist in homeostasis. The fact that human oral bacteria are not found outside the mouth except as pathogens elsewhere in the body points to the importance of this relationship. Communication among microorganisms is essential for initial colonization and subsequent biofilm formation on the enamel surfaces of teeth and requires physical contact between colonizing bacteria and between the bacteria and their host. Without retention on the tooth surface, the bacteria are swallowed with the saliva. Through retention, these bacteria can form organized, intimate, multispecies communities referred to as dental plaque.[35]

For readers unconvinced that these should be counted as multicellular organisms, I invite reflection on the diversity of forms of multicellularity. Most familiar are the plants and animals, and it would certainly be possible to enumerate a series of major differences in the way these two prominent groups of organisms are organised. Or consider the third taxonomic group generally acknowledged to include multicellular organisms, the fungi. Fungi are generally divided into single celled organisms, yeasts, and a variety of multi-celled forms, such as mushrooms. But the multicellularity of fungi is a rather

34 This is more clearly true of terrestrial than pelagic microbes, very large numbers of the latter being found as individual 'planktonic' cells. However, it also appears that a high proportion of these are in an inert state, and only become active in the context of microbial communities.

35 P.E. Kolenbrander et al., "Communication among Oral Bacteria", *Microbiology and Molecular Biology Reviews,* 66:486-505, 2002.

simple matter. Fungi form threadlike chains of cells, called hyphae, which generally exist in tangled mats, called mycelium. Some varieties occasionally organise their hyphae into much more ordered structures such as the familiar mushrooms that function to disperse fungal spores. This is a far less complex form of multicellularity than that exhibited by the many differentiated cell types of plants or mammals. A mushroom is actually much more similar to the fruiting structures of such social bacteria as the myxobacteria, which also form colonial structures not unlike mushrooms for purposes of spore dispersal – though the cooperative hunting of other bacteria also reported in some species of myxobacteria perhaps suggests a more complex sociality than that of fungus cells. It is worth mentioning that some complex multispecies organisms have been familiar for a long time, most notably the lichens, symbiotic associations of photosynthetic algae or bacteria with a fungus. Anomalous from the perspective of a traditional dichotomy between unicellular organisms and monogenomic multicellular organisms, these seem quite unproblematic from the point of view of a more comprehensive understanding of multicellularity.

Multicellularity, even in the traditional sense just mentioned, is an enormously diverse phenomenon. Cells of different kinds organise themselves into a vast diversity of cooperative arrangements, with variously rigid structure, developmental trajectories and so on. What is particular surprising for traditional biological thought in the case of microbes, however, is that these cooperative ventures typically involve cells from quite diverse parts of the traditional phylogenetic tree – though as I have suggested this is not really unusual even for the more familiar multicellular organisms. The diversity of multicellular organisation should be no great surprise. Leo Buss, perhaps the pre-eminent theorist of biological individuality, claimed some 20 years ago that multicellularity had evolved independently around 17 times.[36] In fact it is fair to say that forming cooperative associations is very fundamentally what cells do. It is worth considering, as some biologists indeed have, that while no doubt evolution depends on competition among cells, it may be that what they primarily compete over is their ability to cooperate with other cells.[37] Altruism, in

36 L.W. Buss, *The Evolution of Individuality,* Princeton: Princeton University Press, 1987.

37 L. Margulis, *Symbiotic Planet: a New Look at Evolution.* New York: Basic Books, 1998

its technical biological sense of assisting another organism at some cost to oneself, far from being a fundamental problem for evolutionary biology, may turn out to be ubiquitous in the living world.[38] From this perspective, the diverse communities that make up microbial biofilms and even the diverse communities that constitute a properly functional plant or animal, including its mutualistic microbial communities, can quite properly be considered multicellular organisms.

Throughout these lectures, the one level of biological organisation I have constantly referred to without qualification is the cell. And in fact it does seem clear that this is the most unproblematic such level. If there were any unique answer to the question, what are the constituents of life? that answer would have to be cells. Cells are enormously diverse things, of course, but everything on the standard representation of the tree of life is a cell or is composed of cells. The problems I have indicated with a naïve conception of the organism derive from the complexity and diversity of the relations among very diverse sets of cells, but they do not problematise the idea of the cell as the basic constituent of these various associations.

I have to note, however, that hiding behind my now familiar elephant is yet another elephant. I have hardly mentioned the living forms that are not cellular and that are even more numerous than the cellular microbes, namely the viruses and related objects. Viruses have been found associated with every organism studied and they outnumber any other class of biological entities by at least an order of magnitude. Estimates of the numbers of viruses on Earth are in the range of 10 to the power of 31 – a 1 followed by 31 zeros.[39] This number is probably incomprehensible to non-mathematicians: it has been described by one virologist as amounting to 250 million light years of viral genes placed end to end.[40] The number of viruses on Earth probably exceeds the number of cellular microbes by at least an order of magnitude.

38 A less radical but still controversial claim for the prevalence of altruism, grounded in broadly conventional evolutionary thinking, is E. Sober and D.S. Wilson, *Unto Others: The Evolution and Psychology of Unselfish Behavior,* Cambridge, Mass.: Harvard University Press, 1998. The present suggestion departs considerably further from the widespread scepticism about cooperation.

39 F. Rohwer, and R. Edwards, "The Phage Proteomic Tree: a Genome-Based Taxonomy for Phages", *Journal of Bacteriology,* 184: 4529-4535, 2002.

40 G. Hamilton, "Virology: The Gene Weavers", *Nature* 441: 683-685, 8 June 2006.

It is sometimes said that viruses are not living things at all. And it is true that they often exist in an entirely static state in an inert crystalline form that can hardly be said to be living. On the other hand they are the most efficient replicators of their genetic material on Earth. It has been estimated that anything up to 50% of marine bacteria are killed every day by pathogenic viruses, phages, and in this process the hostile virus produces thousands of replicas of itself for each bacterium it destroys.[41] It has been suggested that perhaps ten to the power of 24 viruses are produced on Earth every second. And these massive replication rates are tied to high mutation rates and almost unlimited mutation mechanisms. Viruses can thus evolve at rates that are far beyond even what is possible for cellular microbes. It is a familiar observation that the HIV virus evolves significantly in the body of a single host, a fact that provides enormous obstacles to the development of effective therapies. When compared with our own 20-30 year generation spans, populations perhaps 22 orders of magnitude smaller, and handfuls of offspring, it is clear that viruses have abilities to explore the space of chemical possibility that organisms such as ourselves could hardly dream of. It is thus no surprise that viruses are the greatest producers and reservoir of genetic diversity on Earth.

Even more interesting than the viruses that kill their hosts are the ones that don't. Most viruses live in stable relations with their hosts but all viruses reproduce themselves by exploiting the chemical resources of their hosts, and many also insert their genetic material into the host genome. Since they may also incorporate DNA from their host's genome into their own, they can readily transfer DNA from one organism to another. Although viruses are generally quite specific in their hosts they can, as is well-known, transfer to new hosts. When they do this they can transfer DNA from one species of organism to another. I mentioned earlier the so-called junk DNA that makes up most of the DNA of eukaryotes such as animals or plants. As well as being increasingly clearly not junk, this is in fact material mainly or entirely of viral origin.[42] It has been noted, for example, that the main differences between human and chimpanzee genomes

41 M. Breitbart, and F. Rohwer, "Here a Virus, There a Virus, Everywhere the Same Virus?", *Trends in Microbiology,* 13: 278-284, 2005.

42 L.P. Villarreal, "Can Viruses Make Us Human?", *Proceedings of the American Philosophical Society,* 148: 296-323, 2004.

are not, as might have been supposed, in coding sequence, but in non-coding regions derived originally from viruses.[43] The enormous powers of viruses to evolve and their ability to insert genetic material into the genomes of cellular organisms has led some biologists to speculate that it is viruses that are the prime movers of major evolutionary change or, at any rate, the main providers of novel biochemical resources. It is beginning to seem possible that, just as microbes are the expert metabolists in nature, so viruses are the leading evolvers. And as microbes provide us with indispensable chemical services, it may be that viruses provide us with comparably significant evolutionary services. At any rate, without disputing the fundamental importance of cells as foci of causal power and organisation that make possible the complex biological structures and communities which we most naturally think of as biological objects, it is important not to forget the much larger number of non-cellular biological objects that spend their time moving genetic material into and between cells. The extraordinary biological capacities of microbes and viruses pose a very interesting question as to what it is that familiar multicellular organisms do well enough to exist at all unless, indeed, to adapt a well-known idea of Richard Dawkins, they are vehicles for carrying their microbial masters around, or niches constructed as residences for microbial communities.

~

My message so far may seem discouraging with regard to the prospects of real biological understanding. It is true that the very ability to discern the complexities I have sketched in these lectures displays the remarkable power of the instruments scientists have devised to explore the workings of life. Scientists have revealed in exquisite detail the structures of biological molecules and their modes of interaction with other molecules. And they are compiling comprehensive inventories of these molecules. Yet the sheer number of constituents thus discovered combined with two other problems that I have tried to emphasise throughout these lectures, presents a problem of almost

43 Villarreal, op. cit.

inconceivable complexity. The first problem is that even as we discern these multitudinous constituents of living things, their biological significance cannot be fully discerned without a view both of the causal powers derived from their own structure and the causal powers of the larger systems in which they participate. We cannot properly appreciate the biological properties of a virus or a bacterium, say, without understanding both the chemical processes found within it and the much larger systems of which it is a vital constituent. The second problem is that even the inventory of causally significant objects at a particular level is not something fully determined by nature, but may vary according to the kind of question we want to ask. Nature is not divided by God into genes, organisms or species: how we choose to perform these divisions is theory relative and question relative.

It is, then, possible to achieve remarkable insights into life processes, but is there any way we can ever hope to fit them together into an integrated understanding of how a living thing, even a single living cell, functions? There is an exciting project that is currently receiving a lot of attention and investment and which does have this aspiration, integrative systems biology. I shall now say a few words on this topic.

First, what is systems biology?[44] It is often said that it is nothing new. General systems theory is generally traced to Karl Ludwig von Bertalanffy in the mid twentieth century, and a number of biologists, perhaps most notably the American theoretical biologist Robert Rosen, have developed ideas that were at least important precursors of contemporary systems biology.[45] At another extreme I have heard biologists say that systems biology is no more than a new name for physiology. The difference, I think, and perhaps this is a case where a difference of degree becomes a difference of kind, is the vast quantity of data, especially molecular data, that is available to the current theorist. Indeed it is not too cynical to say that a major motivation for this entire project is the question, now that we have all these terabytes of molecular data, what do we do with them?

44 For a more detailed discussion of systems biology, see M.A. O'Malley and J. Dupré, "Fundamental Issues in Systems Biology", *BioEssays,* 27: 1270-1276, 2005.

45 L. van Bertalanffy, "An Outline of General System Theory", *British Journal for the Philosophy of Science* 1: 134-165, 1950. R. Rosen Dynamical Systems Theory in Biology, New York: Wiley *Interscience,* 1970.

The proposed answer, very broadly put, is that we employ teams of biologists, mathematicians, and computer scientists to work out how we can create adequate mathematical representations of the multitude of diverse objects discerned in a biological system, and then use these representations to explore and better understand the real processes. The envisaged subjects of such models range from chemical subsystems within cells, to whole cells, complex organisms, and even organismal communities. Microbial communities of the kind described earlier are one attractive target, and are indeed currently one of the major areas to which the very large contemporary investments in systems biology are being directed. Optimists about this project envisage that we may eventually have good enough models of human cells and systems that we can use them for example to test drugs in silico. Apart from likely economic advantages, this prospect looks an attractive way, especially in the UK, to avert the unwanted attentions of animal rights activists. In light of the problems I have been describing, one very clear advantage of this project is that in principle, at least, it seems possible for such models not only to represent systems as dynamic, but even, in principle again, to represent the interactions of different scales of dynamic systems.

There are two extreme perspectives on how systems biology should proceed – though no doubt most actual attempts at its implementation will lie somewhere in between the extremes. At one extreme is a broadly reductionist approach that simply aims to find ways of representing as much molecular data as possible, including of course interactions between molecules, and calculate what happens when all of this is put together. At the other extreme, many biologists hold that the only chance of making such a project work is to appeal to some much more general principles as to how systems work in order to have some way of deciding, among the vast mass of biological data, what is important and what is not.[46]

My own sympathies lie some way towards the latter end of the spectrum. Scientific modelling is not like building a scale model of a ship, where the ideal outcome is to produce an exact miniaturisation of the original object. Rather scientific models are successful to the extent that they identify the factors, or the variables, that really matter. I have emphasised throughout this work that the objects we dis-

46 See again O'Malley and Dupré, *Fundamental Issues,* op. cit.

tinguish in biological investigations are generally abstractions from the complexities of dynamic biological process. The models we are now considering, then, are abstractions of abstractions – selections among the first level of abstractions that we hope may provide us with approximations of the full functioning of biological objects. While the complexity of biological phenomena is forcing us to develop new kinds of models capable of including very large numbers of factors, it is unimaginable that that this project can be extended to the extreme of in silico miniaturisation. Some theoretical principles or assumptions will surely be needed to guide this second level process of abstraction. Moreover, in opposition to the extreme reductionist understanding of systems biology, and as I have emphasised throughout these lectures, one cannot understand the biological significance of a molecule without appeal to larger structures with which it interacts.

On the other hand, I am extremely sceptical of the idea that there are general laws of systems that can be applied equally to economic systems, weather systems, biological systems, or even, for that matter all biological systems. I don't claim to have an argument that there couldn't be such laws, I just see no reason why there should be. Apart from their various levels of complexity, different systems work in very different ways – I'll say a bit about this in a moment. I suggest that what we need as a theoretical infrastructure for modelling biological systems – and here is an argument that the topics I have discussed in these lectures really matter – is ontology. We cannot expect to understand the behaviour of molecules in a cell unless we have a clear idea of what sorts of molecular objects there are and, even more importantly, what kinds of larger structures they are particularly suited to interact with. Let me try to make this clearer by returning to an example that, though it may well be as complex as, or even more so than, the molecular economy of the cell, is in certain respects much more familiar, the behaviour of people.

I shall limit myself to a very modest aspect of human behaviour, one I touched on in my first lecture, the movement of people round a city. I mentioned then the Albert Kuyp and its curious capacity to attract large throngs of people. I noted that this is a capacity that it only has at specific times – a matter I myself confirmed with a pointless trip there on Queen's Day, a major public holiday. It also differentially attracts certain people. Some households still exhibit a division

of labour between wage earners, domestic workers, and free-riders (children). A market has a much stronger capacity to attract people belonging to the second of these categories. For reasons that could in principle be explored, even among those who regularly buy food, some are much more drawn to street markets, others to supermarkets. And so on. The point I want to make is only that we would have no chance of modelling the human movements around a city merely by a detailed inventory of the dispositions even of every individual person. These movements are constrained and promoted by a huge number of physical and institutional structures: roads, shops, schools, playgrounds, city councils, and so on. And importantly, my visit to the Albert Kuyp will make a difference to me and, however minimally, the market. Depending on my experience there I will be more or less likely to return and may eat different foods. And my activities will make marginal differences to the experience of the traders with whom I interact. Aggregated with the activities of many other visitors, this will ultimately affect the continued participation of traders in the market, and so on.

It is not entirely fanciful to compare this scenario with the insides of a cell. Millions of molecules move about this microscopic space, and their movements are also constrained by a complex cellular infrastructure of organelles, membranes, and so on. Ribosomes, for example, structures that host the production of protein molecules, are sites where amino acids and messenger RNA molecules congregate – not, of course, because they have intentions or plans but either because they are transported there by molecules with that particular function, or simply because they tend to stick when they bump into these structures. Just as my movement around a city has to be understood as relation between my internal dispositions and the infrastructure that surrounds me, so the traffic of molecules around a cell is jointly determined by the capacities that derive from their molecular structure and features of the cellular environment, including most notably the density of other molecular species and the membrane topography of cellular infrastructure.

I do also want to keep in view that disanalogies are important too. We should not assume that the same principles will emerge from a system connected by social conventions and language as one based on chemical interactions. And we might usefully remember that there are much simpler systems, for example the weather, based on

purely physical interactions. Though we are becoming quite good at modelling such simpler systems, they are hardly simple: the best models used to predict global warning take several months to run on our fastest computers.[47] And useful information from such models, information, that is, that is understandable to the many and diverse non-expert consumers of meteorological information, must also be conveyed in terms of abstractions from the flux of process, for example hurricanes, cold fronts, droughts, or showers.

In earlier critiques of reductionism I have suggested that what reductive explanations do is explain the causal capacities of things.[48] Where they go wrong is when it is supposed that this is sufficient to explain what they actually do – which capacities are exercised when – something that will typically require detailed knowledge of the context in which they are placed. But I now want to say that this is too simple. My capacity to deliver the Spinoza lectures required both internal capacities of mine (I hope!) and institutional facts about the University of Amsterdam without which no one could have such a capacity. Similarly, the important capacity of messenger RNA molecules to adhere to ribosomes requires both the chemistry of RNA and the presence, and salient features, of ribosomes. Even the capacities of things are produced jointly by internal structures and features of the context in which those capacities are to be exercised.

So an adequate model of, say, a cell, must at least be rich enough to include the mutual determination of properties of objects at different structural levels. If this is true, it may seem to imply that there can be no stopping place short of the entire biosphere. If cells have properties partially determined by, at least, the organisms of which they are part, and organisms by the larger associations of which they are part, then everything mutually determines everything else. But in practice things are not necessarily this bad. What these points do indicate is the importance of deciding what is a sufficiently isolated system to be a plausible target for modelling – this, of course, is part of the process of abstraction that I just mentioned was inevitably central to such modelling projects. And as many theorists have pointed out, partial isolation of systems, or modularity, is very probably a

47 Richard Betts, talk to the Migrations and Identity Symposium, Exeter University, November 2005.

48 *The Disorder of Things,* op. cit.

necessary feature of any partially stable system above a fairly low level of complexity, so we can expect some help from nature in providing appropriate boundaries for our abstractions. At the lower end there is little doubt that intrinsic chemical properties can be taken as brute facts from the point of view of biology – little benefit is likely to accrue from trying to explain variations in the behaviour of biological molecules by appeal to quantum mechanics, though no doubt such explanations are also possible. If we are trying to model a complex organism there are probably many features of the behaviour of cells that can be treated as given, though which these are will be a difficult part of the work of creating such models.

The upper end is harder. The clearest example of a plausible target for a system model is the individual cell. As mentioned earlier, this seems the most unproblematic layer of organisation into individuals in the biological hierarchy, though even here the difficulties consequent on the sensitivity of cells to biological context should not be underestimated. A strictly reductionist approach to cellular differentiation, for instance, would not easily have appreciated the importance of geometrical distortion by neighbouring cells as clusters of cells divide. Defining the outer limits of an organism, for reasons already discussed, are much harder. For many aspects of organismic function – digestion, immune response, or development, for example – it may prove that an adequate model requires treatment of the whole biome, and in the case of ourselves this would include not only traditionally human cells but those of our vast array of microbial, including viral, commensals. This is a good point at which to introduce the final question I want to discuss.

The question is, granting that ontology matters, why is this a matter for philosophers? Isn't biological ontology a matter for biologists? With some important qualifications, I would answer Yes to this last question. The substantive claims I have made in these lectures have mostly depended for their plausibility on my, no doubt limited, understanding of biological issues. Why not leave it to those with less limited understanding?

Part of the answer that I want to stress is that ontology is much less simple a matter than might appear, even to those with a deep understanding of the facts. Indeed my earliest work in the philosophy of biology, concerned with showing the deep differences between the

taxonomies of organisms required for biological investigations, and those required for everyday life, was part of a project of showing that biological ontology is seldom simple or obvious. One way in which it is less simple than generally assumed is that it is equivocal. My example of the human biome, can illustrate the point. To understand human development, human susceptibility and resistance to disease, or human digestion, this may be the narrowest sufficient system to consider. But for examining human behaviour it is likely that a more traditional conception of the human, ignoring microbial symbionts, will be appropriate (this will perhaps be a relief to psychologists, philosophers and others who study human behaviour). This is not only because – though this is a fascinating fact – the nervous system is the only part of the human biological system that is not currently believed to have coevolved with a commensal microbial community. It is more simply because from the point of view of behaviour, most of biology can be taken as given. The question why my arm goes up when I decide to raise it is an enormously difficult one for physiology, and a perpetually intriguing one for philosophy, but it can be taken for granted by most scientific students of behaviour. Of course my microbial associates will sometimes directly affect my behaviour: when my gut flora are unhappy, my behaviour will be much restricted. But it will usually be sufficient to note that I have an upset stomach without going into cellular details.

A central concept for addressing the ontological issues I have been considering is that of a boundary. The boundary to which I have just alluded is one that screens behaviour from the details of cellular chemistry and thereby intercommunity cooperation. But this is a boundary that screens behaviour but not, for instance, disease. So ontological boundaries are relative to the issues with which we are concerned, which is a central part of the reason why there is no unique ontology. To return to another of my favourite examples, how we divide organisms into kinds or species (which, in some instances, coincides with dividing them into individual things), depends on why we are doing it.[49] Note also that divisions of organisms into species amounts to discrimination of things – evolutionary lineages – just to the extent that there are real biological boundaries in place, that

49 And, in this case, what kind of biological entities we are doing it for: there is no reason to assume that the best principles for classifying birds will also be well suited to bees or bacteria, and many reasons for doubting it.

is to say, the boundaries that block the flow of genetic information. Microbial lineages are less plausibly treated as things than are some macrobial lineages, exactly because the boundaries of the former are so permeable, especially to genetic material. When I suggested earlier that the cell was the most unequivocal constituent of life, I might also have said that the cell membrane is, for a very wide range of theoretical questions, an effective boundary.

The clearest example of the importance of ontology to biology is in the theory of evolution. The issue that has been most extensively discussed by philosophers of biology for the last 30 years or so is the so-called units of selection problem: given that evolution is driven by natural selection, what are the things that selection selects. An idea that has been enormously popular in this regard has been that promoted so effectively by Richard Dawkins, that the units of selection are bits of DNA.[50] For a variety of reasons philosophers have almost uniformly rejected this idea, most importantly on the ground that it assumes a simplistic view of the relation between genes and organisms. Prior to the popularity of gene selectionism, it was assumed that the targets of selection were organisms. Nowadays it is widely held that the answer must be that there are a variety of levels – genes, organisms, and probably groups of organisms.[51] I think it may be possible to reinstate something like the idea that organisms are the primary target of selection, but with three very important qualifications. First, as I have argued today, we should not take it as obvious what the organisms are. It may be that the typical 'organism' is really a community of coevolved cell types. Second, as a number of biologists have argued in recent years, organisms do not evolve in passive response to their environment. The evolutionary 'niche' to which an organism is adapted is as much a product of the organism as a cause of the organism's adaptation.[52] This should seem a natural idea in the context of these lectures. Central to the niche of a bird, for instance, is its nest; but it didn't evolve to occupy the nests that happily turned out to be lying around, but rather modified its environment to provide the resources to which it is adapted. This is most obvious of all

50 Richard Dawkins, The Selfish Gene, Oxford: Oxford University Press, 1976.

51 For a good survey of recent thinking on the so-called Units of Selection problem, see K. Sterelny and P.E. Griffiths, *Sex and Death: An Introduction to the Philosophy of Biology,* Chicago: Chicago University Press, 1999, esp. chs. 3-5.

52 See F.J. Odling-Smee, K.N. Laland, and M.W. Feldman, *Niche Construction: The Neglected Process in Evolution,* Princeton: Princeton University Press, 2003.

for humans for whom the environmental niche includes schools, hospitals, and suchlike, all of which play an essential role in the life cycles of individual humans. And, third, as is indicated by my reference to life cycles, we must avoid seeing the organism as a static thing with a fixed set of properties. Organisms are generated, develop, reproduce, age and die. All of these stages are adapted in different ways to the niche with which the organism is coevolved.[53] Thus if I say that the organism is the normal unit of selection, it should be understood that the concept of organism involved is far removed from a naïve and static conception of a living individual. This organism is a process – a life cycle – rather than a thing; it may be a community of distinct kinds of organisms rather than a monogenomic individual; and it must be understood as conceptually and of course causally linked to its particular environment, or niche, which both contributes to the construction of the organism in development, and is constructed by the organism through its behaviour.

Simplistic understandings of evolution, often based on naïve views of deterministically understood genes as units of selection, can underlie bad, and even dangerous science. I have argued in some detail to this effect about the distressingly influential project of so-called Evolutionary Psychology.[54] This is the view that our basic behavioural dispositions are best understood by reflecting on evolutionary forces that acted on ancestral humans in the Stone Age. Although I can't rehearse these arguments here, the most fundamental failure of this programme is its grounding in an antiquated view of evolution based on a crude ontology of genes with deterministic developmental capacities, and isolated, self-interested individuals.

This suggests another reason why this may be fit work for philosophy. Most biologists, and for good reasons, are strongly focused on very specific problems. The best biologists often do concern themselves with ontological issues, and it is their work in these moods that is often most valuable for the kibitzing philosopher. But even these biologists may be constrained by their disciplinary expertise. It is rare to find someone professionally expert in microbiology, vertebrate evolution, and immunology, say. One might recall Ernst Mayr,

53 This point has been made most clearly by so-called Developmental Systems Theorists (see note 3 above for references).

54 See my *Human Nature and the Limits of Science,* op. cit.

one of the most distinguished evolutionists of the twentieth century, dismissing Carl Woese's revolutionary reformulation of taxonomy into the three domains I described in the last lecture, with the remark that 'Woese was not trained as a biologist and quite naturally does not have an extensive familiarity with the principles of classification'.[55]

And, finally, the discussion of ontology, what there is, can benefit from the availability of a set of conceptual resources that have been (and are still being) hammered out for centuries by philosophers – natural kinds, individuals, causation, and so on. I would argue that one of the best tests of the value of such tools and the value of the very abstract modes of argument philosophers have used to discuss them, is their ability to throw useful light on the much more concrete and specific issues that concern different societies at different times. And no set of issues, I suggest, should concern our own society, at this time, more than the remarkable insights into nature being offered by contemporary biology.

In these pages I have tried to sketch a synoptic view of biology – a survey of the kinds of things that constitute the biological world and the kinds of relations they have with one another. My only consolation for my undoubted inadequacy for this task is the thought that perhaps no one is properly adequate. Our scientific understanding of life processes is growing at a breathtaking rate, and our ability to synthesise and assimilate this understanding will have fundamental effects not only on everyone's understanding of life, but even on the future trajectory of human life. I, at least, am convinced that the task is of sufficient importance to outweigh the undoubted risks of getting things wrong.

55 E. Mayr, "Two empires or three?" *Proceedings of the National Academy of Sciences* 95: 9720-9723, 1998.