

Realism, Pluralism and Naturalism in Biology

John Dupré,
Egenis, University of Exeter.

In: *Rethinking Order: After the Laws of Nature*, edited by Nancy Cartwright and Keith Ward. London: Bloomsbury, 2016.

Abstract

Questions of realism take on a rather different form in the philosophy of biology from their more traditional context in the philosophy of physics. Problems arise not so much in relation to the existence of the entities that biologists discuss, but in terms of the partial, idealized or abstracted, models that they use to describe them. Given this partiality of representation, are we entitled to take such models as true descriptions of the world? There is a burgeoning literature on models in the life sciences, some of which will be assessed in this chapter. I shall argue for a qualified realism about the main classes of such models. However it is important to recognise that biological models aim at the truth, but not the whole truth. This is why, which is my second main thesis, we should always expect multiple models in biology, providing partial and sometimes complementary perspectives on the reality they aim to illuminate. This pluralistic realism provides a framework for understanding a wide range of issues in the philosophy of biology. Two that I shall briefly discuss are the ineliminable role of context in understanding the nature of a biological entity, and the fundamentally processual nature of living systems. This chapter concludes with a brief remark on the relation of realism to the wider issue of naturalism.

Realism and Models

Realism in philosophy of biology poses a quite different question from that traditionally addressed in the context of physics. The question whether belief in quarks, neutrinos, and suchlike is warranted by the successful application of theories that employ them remains a contentious one. Bas van Fraassen (1980) is one prominent philosopher of science who has continued to argue that belief in unobservable entities is unwarranted. This does, however raise a very tricky problem of deciding exactly what constitutes the observation of an

entity, and probably more philosophers now would take a demarcation between the unproblematically real and the philosophically debatable as suggested by Ian Hacking's famous remark, 'If you spray electrons, then they are real' (Hacking 1983, p. 24). If we can do things with, or to, an entity, then we are entitled to believe it exists.

From this point of view there should be no serious debate about the reality of nucleic acids, amino acids, and so on. It is not that they are 'directly' observable: though we can observe these large molecules in one sense through high-powered microscopes, we cannot observe that they are nucleic acids or polypeptides any more than we can 'directly' observe that a track in a cloud chamber is an electron. But we can do lots of things with DNA molecules. Just as one example, transplantation of alien genes into organisms is often ineffective because the host organism preferentially uses different versions of synonymous codons. To address this, scientists can rewrite the transplanted DNA using the preferred codons, and this often enables the transplanted gene to be expressed (Gustaffson et al. 2004). If you can rewrite it, it is real.

Interesting questions of realism arise in biology not about the existence of entities but about the extent to which representations of biological entities, or biological models, correspond to the world. There is a parallel here with a somewhat different traditional question, whether the laws or theories postulated by science are true, a negative answer to which question is often motivated by the so-called pessimistic induction on the history of science, the claim that all past theories have turned out to be false, so ours probably will too. The problem does not arise in quite this form for biology, because most philosophers of biology are sceptical of whether there are any laws or theories in biology, and have described biological ideas in terms of models that are held to represent their target phenomena. Whereas laws have traditionally been taken to be quite literally true (or false), the relation between models and their intended targets is at least more complex.

In what follows I shall refer freely to biological models. However I need to acknowledge that my use of the term is somewhat cavalier in relation to the extensive and often subtle discussion of this topic in recent philosophy. Michael Weisberg (2007) distinguishes scientific modelling from 'abstract direct representation' (ADR). The former involves exploration of a structure

that is in some way analogous to the system it is supposed to illuminate, but the elements of which do not claim to represent anything in reality. Many models in evolutionary theory or population ecology are of this kind.

Somewhat similarly, Tarja Knuuttila (2005), focusing especially on economic models, holds that scientific models are best seen as cognitive tools, designed to explore a structure taken to be similar in some respects to a target system in the world, but by analogy rather than any direct correspondence. Abstract direct representation, on the other hand, involves representation of specific entities that are supposed to exist, but which are described in ways that abstract from some of their features that are taken to be unimportant given the purposes for which the representations are to be used. In fact, according to Weisberg, “theorists who practice ADR typically aim to give complete representations” (2007, 229), though they never fully achieve this goal. I am slightly sceptical of this distinction, at least in that I suspect that many cases may be difficult to allocate unequivocally to one category or another. But given that my present topic is realism, my interest is in ADR, or in strategies well to that end of the spectrum, rather than what Weisberg refers to as models *strictu sensu*. As will become, I am also sceptical about the goal of complete representation, which may contribute to this partial disagreement with Weisberg.

Biological models in my broad sense are, at any rate, a motley crew. Some biological models are concrete entities, like the highly standardised model organisms that are taken to represent, more or less, much wider classes of organisms (Leonelli and Ankeny 2013), or the sticks and balls arranged to form representations of molecular structures. Others are sets of equations, used to represent phenomena as diverse as the relations of population numbers in ecological systems or the concentrations of chemical species in metabolic processes. Systems biologists try, sometimes, to integrate such mathematical representations into models of much wider biological systems. Structural models may be used to explore, perhaps using computer technologies, interactions between biological molecules. And so on. The models that will be my concern here will be biological models that it is at least strongly tempting to interpret realistically, for example structural models of molecules, metabolic pathways, or cells.

What is common to all these kinds of models is that, in contrast to the dominant traditional interpretations of scientific laws¹, they are not expected or intended to be strictly true. Just as a map will emphasise particular features of the environment and ignore others, on pain of the fate immortalised in Borges's story of the perfect, and thus completely useless, map (Borges 1973), models focus on particular features of the phenomenon they aim to represent, and ignore others. This suggests a simple answer to the question whether a model is true: if the features it aims to represent correspond exactly with the features in the system or entity represented, then the model should count as true. Unfortunately, models do not typically come with a definitive statement of their intended application, so this criterion is not readily applicable. Indeed, models are often used to explore phenomena by determining empirically the limits of their application, a project that would be rendered unintelligible if the intended scope of application were given in advance.

If science, or at any rate biology, does indeed provide us only with partial and abstracted models of phenomena, are we entitled to claim that they are realistic? As noted above, some models are not intended to be realistic. A paradigm area of science in which this is plausible is neoclassical economics (Knuuttila 2005). But there is an obvious and revealing contrast between economics and models in a field such as molecular biology. No one doubts that there are real economic agents, and no one supposes that there are perfectly rational agents. The object of models involving the latter is not to identify what kinds of entities or processes there are in the world, but to propose a concept in terms of which the aggregation of the behaviours of real agents may be better understood. Whether the model is successful depends on whether typically, or in aggregate, the behaviour of real agents approximates to that of the ideal agents in the model. If it does, then the model may be a useful device for predicting or explaining economic outcomes in the real world. But it does not seem that there is any point at which such a model is intended to correspond to the world.

Compare a model of, say, the citric acid or Krebs cycle, a set of chemical reactions central to metabolism for all aerobic organisms. We do not start here with an uncontested entity that we want to understand better, but rather we postulate a process involving a set of (real) entities whose interactions may help to explain a wide range of empirical phenomena. Either

before or after formulating the model we will hope to confirm that these entities do in fact occur in the cells to which the model is intended to apply. We then may confirm, most likely in vitro, that they react together in the required ways. Gradually we gain increasing confidence that we have identified something that really happens in living systems. The purpose of the model is to represent. To believe that it succeeds is to interpret it realistically. It is necessary to bear in mind, however, that a realistic interpretation of the Krebs cycle model doesn't require that something identical to the reactions occurring in a test tube happens in the living cell, nor that the descriptions of entities in the model are complete descriptions of the entities in the world. In vivo, this metabolic process is interconnected with multiple other processes that provide the metabolites and employ the products of these particular reactions. The rate of the reactions isolated in this particular cycle will depend on these contextual processes as well, thereby, its impact on the larger systems within which it is embedded. The model aims to represent a part or an aspect of what happens, but not to correspond exactly with anything real.

Pluralism

This incompleteness of models leads me to my first main thesis: we should welcome and expect a plurality of models of biological phenomena. If we imagine that science gives us the whole truth, we are likely to suppose that the ultimate objective is a single model that represents everything about a phenomenon of interest, as the classical picture described in Chapter 1.2 suggests. But once we see that models abstract particular features of interest or of relevance to particular questions from the complexity of natural (especially living) systems, this ideal of completeness can be seen to be quite misguided. Different interests will call for different models; and for phenomena of high complexity multiple models will be desirable regardless of any specific questions or interests².

Let me illustrate the pluralism I have in mind with the example of the genome. Nobody now doubts that genomes exist, though at an earlier stage of enquiry successful models of the genome were instrumental in their movement from hypothetical entities to unquestioned existents. Indeed, the possibility of providing multiple consistent though independent models of the genome are now a central part of what puts their existence beyond question (Barwich 2013). Consider, then some of the models of genomes commonly employed. Most familiar are the sequences of four letters, representing the nucleotides

Adenine, Cytosine, Guanine, and Thymine. These sequences, we suppose with good empirical reason, often correctly represent the order in which the nucleotides appear in a real molecule. If they do this, then what they say is true. The representation serves to predict or explain, among other things, the subsequent formation of strands of RNA and of polypeptides, with specific sequences of molecular constituents related to that of the DNA.

A quite different representation of the genome presents the general structure of the DNA polymer and its double helical structure, displaying the various chemical bonds that hold the structure together, including for instance the weak bonds between base pairs that are broken when the two strands separate. Such a model, often in the form of a concrete three-dimensional model, is typically used to instruct students about the structure and function of the DNA molecule. It would be possible simultaneously to represent this molecular structure and a specific sequence though, given the different functions that these models serve, it is hard to imagine any occasion for doing so.

The above three-dimensional structure of the DNA even if, for whatever strange reason, extended to billions of base pairs length with a sequence of nucleotides mapping some actual genome (or better now, chromosome, since most genomes contain a number of these subunits), is still far from a complete representation of the genome (or even single chromosome). Actual genomes contain much more than DNA. First we should include the histones, proteins that form the structural core of the chromosome, and around which they are densely spooled, as is required to enable 2 metres of DNA (in the human case) to fit into a few tens of micrometres of cell diameter. Second, many molecules are, at any time, attached to the genome in ways that crucially determine its behaviour. Various of these bond to particular parts of the genome inhibiting or enhancing the expression of particular sequences. Of special interest are the small molecules that modify either the nucleotides themselves, or the histone core, in the various processes generally referred to as epigenetics. These modifications change the chemical composition and shape of the chromosome and, by thereby exposing or restricting access of particular sequences to the transcription machinery, help to determine the functioning of the genome.

This last point draws attention to a very fundamental issue: the genome is not a static object, but highly dynamic. Its shape and its behaviour are constantly changing. All of the representations I have just sketched, as indeed any model that represents a structural feature of the genome, are in fact snapshots, frozen in time, of an instant of this genomic process. To imagine a 'complete' model of the genome it would not be sufficient to somehow integrate all of these various perspectives in three dimensions, but one would need also to extend the representation in time and present the dynamics of the system. Some of these dynamic elements form parts of the typical development of the wider cell or organism, others reflect adaptive responses to contingent features of the wider environment. Both these dynamic aspects are vital parts of the function of the whole. I shall return to this question of living dynamics below.

One reason it is important to be clear on these issues is that it is very common to take a partial representation as the whole truth, or at least the sufficient truth, about some object of enquiry. A striking example of the problems this can cause was the genetic determinism widely propagated in the course of the human genome project. It is possible that some of the notorious statements by prominent scientists³, describing the human genome as the book of life, or the blueprint for the human, were best understood as propaganda for funding of the project. However, these ideas have certainly disseminated into public understanding so that sequences of As, Cs, Gs and Ts have achieved an almost sacred status, one visible, for instance, in certain arguments to block the development of genetically modified foods⁴. Even if by 'genome' in these contexts were meant the full, four-dimensional, concrete entity, including all the aspects discussed above, these statements would be hyperbolic. The genome is a fundamentally important element of all known biological systems, but it is still only one of many without which these systems could not be properly understood. To reduce the organism merely to the sequence of nucleotides in the genome, however, is wholly absurd.

It is, incidentally, easy to see the attraction of the fetishization of genome sequence, as it provides an acceptable way to short circuit the problem of development, and the reproduction of biological form. The question how organisms develop to become as similar as they do to their parents is as fundamental and difficult as any in biology, and the idea that the explanation was there from the start, whether, as in ancient visions of an animalcule in the

head of a sperm, or as information 'encoded' in the genome, is perennially appealing. This attraction is obvious in Richard Dawkins's (1976) appealing but misguided reduction of evolution to genetics: if the development of the organism could be fully encapsulated in genome sequence, it could be ignored from the perspective of genetic models of evolution. Unfortunately life is not that simple.

The moral of all this for present purposes is the necessity of cautious realism. It is always important to remember that scientific models present only partial truths and abstract from a great deal that is important. But it would be as absurd to deny that genome sequence corresponds to some aspect of reality as it would be to suppose that it provided the whole truth about any real entity. Once this simple point is recognised it becomes almost inevitable that we will want multiple models of a real entity, models that address some of the aspects ignored by a single model. Just the same argument could be directed at the stereochemical and epigenetic models of the genome mentioned above. These models reflect an aspect of reality, but only an aspect. Thus it is not only possible to be both a pluralist and a realist about scientific models, but it is clear that we should be both.

It is worth mentioning in passing that I believe that just the same applies to so-called fundamental physics. Theoretical articulations of reductionism generally start with the premise that physics provides models that apply to absolutely everything, so that ultimately we should be able to understand absolutely everything by appeal only to these models (or perhaps rather by some set of laws). Physicalism, the widely held view that there is nothing wholly non-physical, does imply that there are aspects of everything that are physical. Organisms obey the law of gravity, for instance. But that the physiological, ecological, etc. properties of organisms also fall within the purview of physical laws or models seems to me a view that reveals a deep misunderstanding of the nature of scientific representation, and certainly should not be implied by a reasonable physicalism. The extrapolation from models developed to understand the behaviour of the simple microphysical systems constructed in laboratories, or of physical aspects of gross mechanical or chemical systems as found, for example, in astronomy, to everything whatever, is as lacking in justification as the extrapolation I have just considered in genomics. This, unfortunately, is an issue detailed discussion of which is beyond the scope of our book (but see Dupré 1993; Butterfield 2011).

I have explained why actual models used in molecular biology are partial and abstract from the complexities of the real world. On the other hand such models should presumably not be inconsistent with one another, and if not there seems no reason why we should not aspire to produce some summation of partial models that ultimately acquires the status of a complete model. I do not claim that there is a logical impossibility to some four dimensional representation of every structural feature of some genome, extended through time to represent the dynamic functionality of the whole system. Such things have sometimes been supposed to exist in the mind of God. But even should there turn out to be such a thing, it is not the goal of science to explore the mind of God. More seriously, even such a four dimensional model would only represent the history of one particular genome subject to a very likely unique set of causal influences. There is no reason to think it would have any general utility. At any rate, such supposedly complete representations play no part in the actual business of doing real science. Yet this is no reason to deny that actual, partial scientific models may correspond to reality in the specific ways that they are intended to do so.

Context and Process

In fact, as just noted, even the complete model imagined in the previous paragraph would not provide the kind of universal predictive and explanatory import sometimes imagined by advocates of a complete and final science. I want now to discuss two reasons for this. First is the question of context. Even a complete model of the genome in the sense just discussed would not be a sufficient description for all scientific purposes without some specification of context; context partially determines the properties of a biological entity. Second, and closely connected to the first point, I want to re-emphasize the importance of process. I have already noted that a full description of an entity such as a genome must be four-dimensional, extended through time. But a living process is not a self-contained thing, with its trajectory determined by its internal structure. Its persistence is determined in part by its interaction with its environment. The fact that models of biological structure abstract from the dynamic features of the entity represented reinforces the extent to which the character of the entity is dependent on its wider context.

Even if a model were able to represent everything within its intended domain, it must stop somewhere. So, for example, even a ‘complete’ model of a cell would not include a full description of the surrounding biological or physical conditions. But a cell, or for that matter an organism, is an open system maintained in thermodynamic disequilibrium by its interactions with the environment. So a representation of such an entity is inevitably inadequate to a full understanding of its behaviour. This simple point is already sufficient to show that biological representation is only partial. However, there is a deeper point: the external context in which a biological entity is positioned does not only partially determine its behaviour through interaction between entity and context, it may contribute to determining what the entity is. Or so I argue.

Consider protein-coding genes. One quite legitimate representation of such a gene is in terms of sequence. But genes are also identified in terms of function. In their origins in Mendelian genetics, genes were identified in terms of the difference they made to a phenotype—the genes for eye colour or bristle number, for instance, studied in classic research on *Drosophila* by Thomas Hunt Morgan and his associates. A particular sequence is capable of determining whether a fly has a particular eye colour, but only given a wide range of background conditions in the cells of the fly. In fact the same sequence could occur in a quite different organism, with quite different effects on the organism phenotype.

Identifying genes as being “for” phenotypic features, is in many ways problematic, as has been extensively pointed out by philosophers of biology (see, e.g., Dupré 2012, pp. 105ff; Griffiths and Stotz). But the same problem arises for more proximate functions. It is natural to think of a “protein-coding” gene as coding for a particular protein. But it is now well known that a gene may code for many, sometimes even thousands of, different proteins as RNA transcripts are rearranged in the process of splicing, or subject to various kinds of post-transcriptional editing. So again, a description of a gene in terms of the protein it codes for is underdetermined by genetic sequence, and is made true, in part, by further conditions of the cellular environment.

Even the protein does not provide a stopping place for this contextual determination of functional descriptions. It is increasingly clear that proteins can display a wide range of different functions depending on their location within the cellular milieu (Jeffery 1999). It is fascinating that this

phenomenon is widely labelled as protein “moonlighting”, suggesting that the protein has a proper function, and is doing something different in its spare time, an implication that surely reflects a residual assumption that a biological entity has at least a primary function determined by its structure. But of course, though one function may be primary in the sense that it occurred earlier in evolutionary history, this has no relevance to a current understanding of the way a system works. So the question whether, for example, some entity is a structural protein or an enzyme, will often depend on the broader context in which it is located.

The upshot of these examples is that no spatially restricted description of a biological entity is sufficient for a full understanding of what it is, still less of its behaviour. There are no closed biological systems, and their openness involves ways in which their behaviour is partially determined by their context. This issue has an important bearing on discussions of reductionism. The suggestion that the behaviour of a biological system could be explained or predicted in terms of its molecular parts is clearly vulnerable to the kind of problem just presented. Reductionists are liable to respond that we have just assumed too narrow a scope for the reductive base, and a successful reduction will require inclusion of the relevant aspects of the context. But of course the same problem is sure to arise for the elements of the context that are then included, and there is no reason why this process should terminate.

Given the increasing implausibility of practical reduction as the posited reducing base expands in this way, at this point the reductionist will usually retreat to supervenience, the idea that the behaviour of the whole is determined by a sufficient extent of the underlying molecular or microphysical reality. At this point I am inclined to say that as we move to something approaching a thesis of supervenience on the microphysical state of the entire universe, we move towards a vacuity that verges on the meaningless. Certainly such global supervenience is not an idea with empirical implications. But for present purposes it is sufficient to conclude that actual biological models will always and inescapably be partial, and their applicability will be sensitive to features of their context.

As noted earlier, another perspective on the protean character of biological entities can be gained by suggesting that they are generally better understood as processes than as things. A cell, for example, is a process that originates

with the bifurcation of a mother cell, and ends in bifurcation or apoptosis. During the period of its existence it will adopt various structural configurations and contribute to various functions. Its existence is determined not by some property or properties that it exhibits as long as it exists, as we normally suppose for a substantial thing, but rather by a variety of processes that sustain its integrity. In fact, everything in biology is dynamic, and an appropriate metaphysics sees the living world as composed not of things, but of processes (Dupré 2012, chs. 4-5). Rather than see processes as involving things undergoing changes, it is better, at least in biology, to understand things as processes stabilised to some degree over relevant time scales. The 'things' in terms of which we often describe processes are processes stabilised on time scales relevant to the process we are considering (Baptiste and Dupré 2012). Given the current state of physics it is plausible to see the world as composed of processes 'all the way down'; but this is an issue that can be left moot for present purposes.

To different degrees, all models abstract from the four-dimensional character of biological processes. Many or most biological models are of course dynamic to the extent that they represent the way a system develops over time, perhaps depending on some external variables. But this is already too linear an approach to capture the full dynamic complexity of the system. As J.S. Haldane (1931, p. 22; cited in Nicholson and Gawne 2013) put the point: "Structure and functional relation to environment cannot be separated in the serious scientific study of life, since structure expresses the maintenance of function, and function expresses the maintenance of structure." I do not want to argue that it is impossible to provide a representation of all this dynamic complexity, but only to insist that there is very likely no reason to do so. Once again, we need only recall the specific purposes for which we represent parts or features of living systems, and the uselessness of the 'perfect' Borgesian map.

Realism and Truth

My general point so far has been that biological models are partial, but nonetheless the entities that they refer to are, often enough, real. This does leave an awkward question, however, whether the things they assert are true. If I present a model involving perfectly spherical cells, it may be acceptable to say that I am referring to the actual, not quite spherical, entities in the world. But surely if I say they are spherical, I say something false. Generally, it is

impossible to avoid the conclusion that if a model is intended to assert the existence of the states of affairs it portrays, then it will almost fail in this intention.

Considerations of this sort have led many philosophers to deny that statements based on scientific models are generally true. Several responses to this problem are possible. One would be to collapse Weisberg's (2007) distinction between models and abstract direct representation towards the former, and claim that in both cases we are dealing with some kind of cognitive instrument that illuminates the world by analogy rather than by direct description (Knuuttila 2005). Some go further and argue that scientific models are best seen as a kind of fiction and can be illuminated by exploiting ideas about literary fiction (e.g. Frigg 2010; for a general discussion of this strategy see Godfrey-Smith 2009); but this strategy, while it may be defensible for certain classes of models, is surely inappropriate for the kinds of partial representation I have been considering. My own preference is for a pragmatic approach that takes scientific statements as true, or perhaps just correctly assertible, if they correspond with the world in the intended respects. To illustrate with a very simple example it is true, in appropriate contexts that glucose is $C_6H_{12}O_6$, and also that fructose is $C_6H_{12}O_6$. No chemist would be tempted to conclude that glucose was identical to fructose, though there might be contexts in which the differences between these sugars was irrelevant. Partial representations license assertions, but only in a suitably limited class of contexts. The extent to which scientific descriptions are adequate to their intended uses is always a matter open to further investigation. A telling and tragic example is provided by the chemical thalidomide, which normally exists as an equal mixture of two optical isomers, mirror image structures, an apparently very subtle difference. One of these isomers is an effective treatment for morning sickness; the other causes drastic developmental abnormalities. For most purposes the three-dimensional structure, not specifying either of the optical isomers, would be a fully adequate chemical description; for pharmacological purposes it was clearly catastrophically insufficient. Whatever view is taken on truth, the crucial point is that the limitations, known or unknown, of particular scientific perspectives must always be borne in mind.

Naturalism

A wider context in which discussions of realism figure significantly is the general philosophical thesis of naturalism. Unfortunately naturalism is a highly contested concept. My own view is that a more useful thesis is anti-supernaturalism (Dupré 2012, Ch. 1). That is to say, naturalism is the view that a proper ontology should be restricted to entities that have some empirically accessible connection to the natural order; it is a negative thesis rather than a positive one. The objective should be to leave the characterisation of the natural order as vague as possible, while allowing enough substance to exclude such entities as deities or immaterial souls. Of course it is not an a priori truth that there are no deities or souls. If Jupiter were to appear in public view and begin casting thunderbolts and mating with swans, our evaluation might rapidly change. Theology might become a respectable branch of empirical science. But while Gods continue to play no detectable role in public life, a naturalist will exclude them from his or her ontology. The point of the vagueness is to prevent substantive positive theses about the world to be smuggled in under cover of this purely negative thesis.

Materialism, for example, is sometimes taken to imply no more than the exclusion of the immaterial, where this has very much the same import as the supernatural. But there is also a tendency to equate naturalism with materialism and thence often to reductive physicalism. Concepts such as emergence or so-called downward causation, the determination of the behaviour of parts by wholes of which they are parts, are said to be 'spooky', indicative of the supernatural. On the contrary, these seem to me clearly contingent and empirical hypotheses about how the natural world works. The motivation for naturalism is often, and appropriately, the view that the broadly empirical methods we think of as scientific have provided an epistemology for exploring the world vastly more successful than the appeal to supernatural agents lying outside the natural order. Our ontology, it is then argued, should be based on successful epistemology. This motivation makes clear why some kind of realism is necessary for such a motivation of naturalism. Without a realistic interpretation of science no link can be made from successful science-based epistemology to ontology. As I hope to have shown, while it is not trivial to provide a realistic account of biological methodology, there are no insuperable obstacles to doing so, and indeed the successes of modern biology require some kind of realistic interpretation.

Bibliography

- Baptiste, Eric and John Dupré. "Towards a Processual Microbial Ontology." *Biology and Philosophy*, 2013, 28:379-404.
- Barnes, Barry and John Dupré. *Genomes and What to Make of Them*. Chicago: University of Chicago Press, 2008
- Barwich, Ann-Sophie. "Fiction in Science? Exploring the Reality of Theoretical Entities." In: *Defending Realism: Ontological and Epistemological Investigations*, edited by Greg Jesson, Guido Bonino, and Javier Cumpa. De Gruyter Publishers, 2013, pp. 291-309.
- Beatty, John. "The Evolutionary Contingency Thesis." *Concepts, Theories, and Rationality in the Biological Sciences: The Second Pittsburgh- Konstanz Colloquium in the Philosophy of Science*, edited by Gereon Wolters and James G. Lennox. Pittsburgh: University of Pittsburgh Press, 1995, pp. 45-81.
- Borges, Jorge Luis. *A Universal History of Infamy*. Translated by Norman Thomas di Giovanni. London: Allen Lane, 1973.
- Butterfield, Jeremy. "Laws, Causation and Dynamics at different Levels." *Interface Focus (Royal Society London)* 1, 2011, pp. 1-14.
- Dawkins, Richard. *The Selfish Gene*. Oxford: Oxford University Press, 1976.
- Dupré, John. *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. Cambridge, MA: Harvard University Press, 1993.
- Dupré, John. *Processes of Life: Essays in the Philosophy of Biology*. Oxford: Oxford University Press, 2012.
- Frigg, Roman. "Models and Fiction". *Synthese*, 2010, 172:251–268.
- Godfrey-Smith, Peter. "Models and Fictions in Science." *Philosophical Studies*, 2009, 143:101-116.
- Griffiths, Paul and Karola Stotz. *Genetics and Philosophy: An Introduction*. Cambridge: Cambridge University Press, 2013,
- Gustaffson, Claes, Sridhar Govindarajan and Jeremy Minshull. "Codon bias and heterologous protein expression." *Trends in Biotechnology*, 2004, 22:346-53.
- Hacking, Ian. *Representing and Intervening*. Cambridge: Cambridge University Press, 1983.
- Haldane, John S. *The Philosophical Basis of Biology*. London: Hodder and Stoughton, 1931.
- Knuuttila, Tarja. "Models, Representation, and Mediation". *Philosophy of Science*, 2005, 72:1260–1271.
- Leonelli, Sabina and Rachel A. Ankeny. "What Makes a Model Organism?" *Endeavour*, 2013, 37:209-212

Jefferey, Constance J. "Moonlighting Proteins." *Trends in Biochemical Sciences*, 1999, 24:8-11.

Mitchell, Sandra. *Biological Complexity and Integrative Pluralism*. Cambridge: Cambridge University Press, 2003.

Nicholson, Daniel and Richard Gawne, "Rethinking Woodger's Legacy in the Philosophy of Biology", *Journal of the History of Biology*, 2013.

Sober, Elliott. "Two Outbreaks of Lawlessness in Recent Philosophy of Biology." *Philosophy of Science*, 1997, 64 (Supplement): S458-S467

van Fraassen, Bas. *The Scientific Image*. Oxford: Oxford University Press, 1980.

Weisberg, Michael. "Who is a Modeler?" *British Journal for the Philosophy of Science*, 2007, 58:207–233