Philosophy of Science for Biologists

Edited by

Kostas Kampourakis

and Tobias Uller

CAMBRIDGE
UNIVERSITY PRESS

University Printing House, Cambridge CB2 8BS, United Kingdom

One Liberty Plaza, 20th Floor, New York, NY 10006, USA

477 Williamstown Road, Port Melbourne, VIC 3207, Australia

314–321, 3rd Floor, Plot 3, Splendor Forum, Jasola District Centre, New Delhi – 110025, India

79 Anson Road, #06–04/06, Singapore 079906

Cambridge University Press is part of the University of Cambridge.

It furthers the University's mission by disseminating knowledge in the pursuit of education, learning, and research at the highest international levels of excellence.

www.cambridge.org
Information on this title: www.cambridge.org/9781108491839
DOI: 10.1017/9781108648981

First published 2020

Printed in the United Kingdom by TJ International Ltd, Padstow Cornwall

A catalogue record for this publication is available from the British Library.

ISBN 978-1-108-49183-9 Hardback
ISBN 978-1-108-74070-8 Paperback

Contents

Contributors

Ingo Brigandt
University of Alberta

Carol E. Cleland
University of Colorado Boulder

David J. Depew
University of Iowa

Michael R. Dietrich
University of Pittsburgh

Carrie Friese
London School of Economics and Political Science

Kostas Kampourakis
University of Geneva

Tim Lewens
University of Cambridge

Kevin McCain
University of Alabama at Birmingham

Emily C. Parke
University of Auckland

Erik L. Peterson
University of Alabama at Tuscaloosa

Anya Plutynski
Washington University in St. Louis

Angela Potochnik
University of Cincinnati

Barbara Prainsack
University of Vienna, and King's College London

Thomas A. C. Reydon
Leibniz University Hannover

Michael Ruse
Florida State University

Tobias Uller
Lund University

Preface

Is philosophy of science of any use to biologists? A well-known response is that philosophy of science is as helpful to science as ornithology is to birds. Whether or not it was Richard Feynman who actually said this does not affect the fact that many biologists that we have met, especially those older than us, would readily agree. Among these biologists one can find top researchers, with prestigious grants and publications, who think that any philosophical discussion is a waste of time. The experienced researcher, they would say, knows what has to be done; the inexperienced has to learn from the experienced ones in the lab or in the field. Whatever Kuhn or Popper said (they have not heard about Lakatos, or any philosopher after him) is irrelevant to the actual practice of science. Philosophy of science is, at best, a nice endeavor for retired scientists, if they decide to reflect upon their own career and work. Or so the story goes.

This response is a caricature, of course, and many biologists do not think like this. But even those who are not in principle opposed to philosophical reflection and discussion usually do little to promote it. They have data to analyze, papers to write, and grant proposals to submit. Science is a full-time job, and there is little time left for philosophizing, which thus becomes a luxury. However, our aim with the present book is to show that it is not a luxury but a necessity. Philosophical reflection is inherent in any scientific activity, and what is necessary is to guide the experienced researchers to make it explicit, and the inexperienced ones to understand it. We hope that the chapters in the present book show how important philosophy of science is for biology, and how much biologists will benefit from thinking and reflecting in a philosophical manner.

We must note that the chapters in this book cover some philosophical aspects only, focusing on those that we considered the most important for biologists, especially the younger ones, to understand. We begin with a chapter we wrote that sets the context by explaining in some detail why biologists should care about the philosophy of science. The next three chapters discuss some very fundamental issues: what constitutes an explanation in biology (by Angela Potochnik); what biological knowledge is (by Kevin McCain); and what the nature of theories and models in biology is (by Emily Parke and Anya Plutynski). Then we focus on concepts, devoting four chapters on the nature and role of concepts, discussing how biology concepts are used and transformed (by Ingo Brigandt); why it matters that many biology concepts are metaphors (by Kostas Kampourakis); how concepts contribute to scientific advancement (by David Depew); and how conceptual analysis can contribute to scientific practice (by Tim Lewens).

The subsequent chapters discuss the methods used in the life sciences (by Erik L. Peterson); how biologists study the past and why this kind of work can be as solid as experimental science (by Carol Cleland); what the basis of biological classification is (by Thomas Reydon); what the nature of scientific controversies in the biological sciences is (by Michael R. Dietrich); and what the relation is between facts and values in biological science (by Carrie Friese and Barbara Prainsack). Last, but not least, Michael Ruse, one of the founders of the field we call philosophy of biology, shares his fifty-year-long experience of doing philosophy of biology. We conclude with some practical suggestions of our own about how to teach philosophy of science to biologists.

We are of course indebted to the contributors to this volume for their high-quality chapters and their excellent collaboration. We are also indebted to Katrina Halliday, publisher for life sciences at Cambridge University Press, who supported this – rather unusual for the life sciences series – book project right from the start and toward publication. We are also very grateful to Olivia Boult and Sam Fearnley for their work and collaboration during production, as well as to Chris Bond for his meticulous copyediting of the book. Finally, we are indebted to each other for an excellent collaboration, the outcome of which you are now holding in your hands. We hope that you will enjoy reading it as much as we did.

Kostas Kampourakis and Tobias Uller

BOX 1.1 What Kinds of Questions Do Philosophers of Science Ask?

Philosophy of science is concerned with what science is, how it works, and what it can tell us. Some of the most fundamental questions concern very general features of science:

- What makes science different from non-science?
- How are scientific knowledge and understanding generated?
- How is science organized?
- What are the limits of science?

Major topics in philosophy of science are those that analyze and clarify the main components of scientific investigation, including

- What is a scientific explanation?
- How are scientific concepts used and transformed?
- What is the role of idealization in science?
- What is the relationship between theory and data?

Answers to these questions often require careful study of more narrowly defined questions, and most philosophers of science are therefore working on particular problem agendas that can range from quite general to very specific:

- What is the difference between reductionist and holistic approaches to the study of life?
- What is a biological mechanism?
- What do biologists mean when they refer to genes?
- What is the utility of Hamilton's rule in evolutionary theory?

Philosophers of biology are those philosophers of science who are particularly concerned with the biological sciences. Philosophy of biology did not begin in earnest until the 1970s, and earlier philosophy of science was largely concerned with physics or chemistry. However, philosophy of biology is now one of the main areas of inquiry and philosophers of biology have made important contributions to philosophy of science as well as biological theory. For examples, see the Further Reading section at the end of this book.

influence each other over time (see Chapter 13). Being scrutinized can feel uncomfortable for scientists, in particular if they believe that science is – or should be – free of such biases. As illustrated by several of the chapters in this book, this belief is not only mistaken but can actually be detrimental to science itself. The questions, methods, and models that a scientific community considers exemplar are shaped in part by shared attitudes and beliefs. Ignoring these social aspects of science can make biologists less well equipped to identify and solve scientific problems, and it can make them struggle to handle controversies between scientists and between science and other parts of society (see e.g., Chapters 7, 12, and 14)

For philosophy of science to become useful to biologists, scientists and philosophers need to find ways to communicate, share ideas and results, and perhaps occasionally work together. As biologists who have attempted to work with people from other disciplines will testify, collaboration is easier said than done. One main hurdle is simply ignorance about each other's work. Another is to become familiar with terminology and habits of mind that are often specific to particular disciplines. These hurdles can be overcome. Nevertheless, just as it will take time and effort for a cancer researcher to figure out if insights from evolutionary theory will be useful to her, it will take time and effort to figure out if philosophy of science will be useful to you. We hope that this collection of chapters will be helpful for those who are willing to dedicate their time. At the end of the book, we provide suggestions for further reading on both general topics in philosophy of science, as well as topics that are likely to be of particular interest to biologists. In the last chapter we also make concrete suggestions about how topics from philosophy of science could be taught to biology students.

1.3 Kuhn and Popper As Caricatures

Perhaps no philosopher of science is more familiar to scientists than Karl Popper. Popper was concerned with the big questions in philosophy of science, and his work has had a long and lasting intellectual impact, not the least on scientists (Lewens 2015). His idea that scientific hypotheses can never be proven, only falsified, is commonly introduced to beginners in the natural sciences as the fundamental feature of science. Falsification not only separates science from non-science but Popper also meant that the repeated failure to falsify hypotheses can account for the growth of scientific knowledge. Another philosopher of science who is likely to be mentioned in introductory science classes is Thomas Kuhn, famous for introducing the idea of a paradigm shift (Kuhn 1996). Scientists tend to be more ambivalent toward Kuhn since he emphasized that science is a collective, social endeavor where scientists sometimes appear irrational. But Kuhn's concept of paradigm shifts can be interpreted as a radical theory change introduced in the face of repeated falsification of established theories. For many biologists, this view of how science works – steadily securing knowledge through hypothesis testing and rarely interrupted by radical theory change when major hypotheses are disproven – may be the entire philosophy of science they are exposed to during their studies, perhaps even for their entire academic career.

No shame on Popper and Kuhn, and scientists are often taught a caricature of their work (like we just did!), but this is not really enough to understand how science works. Believe it or not, philosophy of science has progressed! While falsifiability remains an important litmus test for a scientific hypothesis, it is now widely recognized that the building of knowledge through falsification of *a priori* hypotheses is a poor characterization of many successful sciences, including biology. Scientific knowledge and understanding is generated through much more diverse standards and activities than envisaged by these early philosophers of science. There are good reasons for this diversity. The world is immensely complex, and humans are limited beings. Thus, it is reasonable that different scientific questions demand different approaches or methods. However, a diversity of scientific standards does not imply an absence of standards. It is important to understand what works and why.

In practice, biologists tend to pick up most ideas of what science is and how it works from fellow biologists, typically those who work on similar problems using similar methods. But if there is no universal standard of science, this can make it difficult to recognize or understand the importance of research that uses different standards or, for that matter, the limitations of one's own approach. Such failure can lead to inefficient science, missed opportunities for scientific breakthroughs, or even long and fruitless controversy. In what follows, we reflect on three features of science – its aims, methods, and concepts – to make a case for why biologists can benefit from insights gathered from the philosophy of science.

1.4 Scientific Aims

What are the aims of science? A short list would likely include description, classification, prediction, and explanation. Biologists *describe* and *classify* new species, molecules, and biological processes; they *predict* the effects of human activities on biodiversity or the spread of disease; and they *explain* how cells work and why populations evolve. A main reason for these activities is that many biologists ultimately strive to *understand* living systems, such as cells, organisms, and ecosystems. This understanding has practical consequences for technology, medicine, and many other features that make up societies, and it is therefore important far outside academic circles.

A phenomenon can be said to be understood when one can give it a satisfactory explanation (see Chapter 2).[1] Given that we explain phenomena all the time, it will perhaps come as a surprise to learn that it is neither obvious what it means to explain something, nor what, if anything, that makes scientific explanations different from everyday explanations. The traditional point of view on behalf of philosophers of science is that scientific explanations consist of statements that demonstrate that the phenomenon to be explained follows from natural law (Woodward 2017). This account of explanation is heavily influenced by physics, and biologists hardly find it very appealing since there is a widespread skepticism toward the existence of biological laws.

A more promising idea is that explanation is linked to causality, manipulability, and control (e.g., Woodward 2003).[2] It will feel natural to biologists to think of causes as difference-makers (Illari & Russo 2014). Rain causes seeds to germinate because if it had not rained the seeds would remain dormant. Loss of genetic variation causes population extinction because if it were not for the loss of genetic variation the population might have adapted to the environment. One view of scientific explanation is that it is achieved when the information provided by the explanation allows one to answer a range of such what-if-things-had-been-different questions (e.g., Woodward et al. 2003; Strevens 2008; Potochnik 2017). For example, an explanation for how ATP is generated

[1] Philosophers speak of the phenomenon to be explained as the *explanandum* and the sentences that do the explaining as the *explanans*.

[2] There are various versions of this theory of causal explanation, Woodward 2003 and Strevens 2008 are useful starting points.

may refer to biochemical features of glycolysis. This explanation reveals something about the causal tapestry of the world; the molecular detail makes it possible to grasp the consequences of a change in the concentration of pyruvate or the chemical structure of the reacting molecules. According to some philosophers, this is what it means to understand how ATP is generated, and the more what-if-things-had-been-different questions about ATP production we can answer the better we understand it.

Not all explanations in biology are mechanistic like this, however, and many explanations in biology look more like historical explanations (see Chapter 10). An explanation for the extinction of dinosaurs may refer to a meteorite that struck the earth and caused long-term changes in the earth's climate. Nevertheless, the reason why this explanation generates understanding is similar to the case of ATP; reference to the meteorite and its effect on climate makes it possible to grasp what would have happened to the dinosaurs if the meteor had not have struck the earth, or if it had been smaller, or if there had been no competing mammals around. There may be other kinds of scientific explanations, but being able to give answers to what-if-things-had-been-different questions appears to at least be one important feature of many scientific explanations.

A good thing about this notion of explanation is that one need not take truth too seriously. What really is "out there" may forever be out of reach, but representations of the world can be sufficiently good approximations that enable one to foresee what would have happened if things had been different. It is not always possible to support the explanation through active intervention, of course (this is difficult for the dinosaur extinction, for example). But scientists can nevertheless ensure that their theories are empirically justified – or true enough – by imagining and studying a range of different situations. This is why it is important that scientific theories are falsifiable; if a theory makes no falsifiable claims, it also appears impossible to predict the consequences of an intervention.

Another helpful feature of the causal theory of explanation is that it brings attention to the fact that scientists need to manage causal complexity (Potochnik 2017). Biological systems are enormously complex, and any representation of a living system will only capture some of its actual causes. This is in itself not a problem. In fact, too much detail makes it harder to grasp what would have happened if things had been natural selection in adaptive evolution, not the role of development, physiology, or behavior. The assumptions made in evolutionary theory tend to turn the latter into constraints; they can account for the absence of adaptive fit but not its presence.

This line of thought is so common to biologists that many take it for granted. However, a comparison to the explanations for the delayed train arrivals is a reason to treat this conclusion with caution. That is, that one particular idealization of evolution by natural selection privileges genes and natural selection does not imply that there is an inherent causal asymmetry in evolutionary processes (Laland et al. 2011). The role of proximate causes in adaptive evolution is in fact one of the most persistent controversies in biology (Amundsen 2005). Contemporary examples include the disagreement over the explanatory role of development, plasticity, extra-genetic inheritance, and niche construction in evolution (see Laland et al. 2014, 2015). One possible reason that these issues are difficult to resolve is that the genetic representation of evolution is commonly taken at face value, rather than being understood as an idealization designed to explain evolutionary phenomena in terms of natural selection. An increased awareness of the relationship between idealization and explanation may reduce the risk that causes that are idealized away become permanently neglected, facilitate capitalization of insights from other disciplines, and put a restraint on unproductive scientific controversy.

While there are good reasons why a biological phenomenon like adaptation can have several explanations, biologists may sometimes wish to determine which of a number of different explanations is the most satisfactory (see Chapter 3). Consider cichlid fish, famous for the ability to evolve very similar morphologies in different lakes (Seehausen 2006). Evolutionary biologists have demonstrated that this convergence happened because the local habitat and foods are often similar in different lakes, which favors a limited set of life styles such as bottom-dwelling grazers and open water predators (e.g., Muschick et al. 2012). Thus, natural selection explains the convergent evolution of cichlid fish. But biologists have also pointed out that some of the recurring features of these fish, such as the shapes of bodies and jaws, tend to be plastic (Schneider & Meyer 2017). That is, those characters respond to the habitat or diet that individual fish encounter during their lifetime. Some biologists believe that plasticity has contributed to the striking convergence between different species because plasticity can make some features more likely to become selected than other, perhaps equally fit, phenotypes (see West-Eberhard 2003). In terms of explanation, the question is whether or not an explanation for adaptive convergence in terms of plasticity *and* natural selection is better or more satisfactory than an explanation in terms of natural selection alone (Uller et al. 2019).

Picking the best of two explanations is easy if only one of them is backed up by empirical evidence (although negative evidence is not always enough to give up a hypothesis, as explained in the next section). Beyond this, there may be nothing inherent in these explanations for adaptive convergence that make one better than the other; perhaps picking the best explanation is simply a matter of one's interest (assuming both explanations are empirically justified). However, scientists sometimes prefer one explanation over others if it applies to many different phenomena rather than a few, if it is easy to understand, or based on its elegance or simplicity (Ylikoski & Kourikoski 2010). Explanatory standards such as these are important since the criteria for picking the best explanation influences what biologists consider knowledge (Chapter 3). Such explanatory standards can vary between biological disciplines (e.g., those concerned with mechanistic versus historical explanations), but these differences may not be well recognized by biologists. As a result, it can be difficult to assess the quality and scope of biological research that lies outside of one's immediate expertise.

These examples from evolutionary biology reveal exactly how difficult it is to do science. We have discussed scientific explanation at some length here because we believe that it is a good illustration of how philosophy of science can be helpful to biologists, but similar issues also come up for other scientific aims, such as classification (Chapter 11). In our opinion, the philosophy of science that is usually taught in introductory classes does not pay sufficient tribute to the challenges that scientists face when they attempt to produce knowledge and understanding of a world that is enormously complex. How biologists choose to represent the world influences the questions they consider worthwhile, how they organize their research, and the answers they look for. This means that scientific advancement requires the flexible use of methods and concepts.

1.5 Scientific Methods

It is probably obvious that biologists do not follow a single universal scientific recipe, but rather several more or less distinct approaches (see Chapter 9). Biological science certainly appears to frequently stray away from a strict method of falsification. A careful look behind the scenes of scientific papers that communicate the results of a test of a well-defined hypothesis will often reveal a process that looks very different to what we have been told science should look like. This is one reason why p-values are so problematic; scientists' formulation of hypotheses often develops in parallel with observation, practice, and data collection rather than in a strictly ordered fashion. It may be tempting to conclude that failure to adhere to the strict rules of hypothesis testing makes some biological research fundamentally flawed. A more optimistic view is that scientific practice simply reflects that there is a diversity of scientific methods, all of which may be appropriate. Regardless, it is important to examine how different scientific methods achieve scientific aims (e.g., understanding), what tools are available to meet these aims, and how well those tools actually work in practice. Biologists are highly engaged in these issues, in particular with respect to the appropriate use of statistics, lab- vs field-based methods, the use of model systems, and so on. These discussions may benefit from a greater attention to the literature on philosophy and history of science, which often has a fair bit to say on the matter (e.g., Chapter 10; Chapter 9).

Another peculiar break with falsification is that scientists only occasionally and reluctantly abandon a hypothesis when the data fail to support it. For example, a frequent failure to demonstrate that offspring of males with exaggerated sexual signals were fitter than other offspring did not generally make behavioral ecologists abandon the "good genes" hypothesis (Roughgarden 2009, pp. 213–224). This widespread practice is difficult to make sense of if biologists really believed in falsification as a method to scientific progress. Eventually a hypothesis may of course fail to be confirmed and become abandoned, but this will not result in a wholesale rejection of the theory. For example, the biologists that did conclude that there was little if any evidence for good genes rarely used this to argue against the theory of sexual selection (but see Roughgarden 2009). This is not irrational behavior on behalf of the scientists. Instead, examples such as these demonstrate that scientific theories are organized into more fundamental theoretical frameworks that are not revised on the same basis as a more specific hypothesis. Philosophers of science have introduced and use many different concepts to make sense of this feature of science, including thought styles (Fleck 1979), paradigms (Kuhn 1962), research programs (Lakatos 1978), and problem agendas (Love 2008; see Chapter 7).

An example from population biology can illustrate this point. During range expansions, populations on the front line are often small but rapidly growing. The combination of bottlenecks and rapid population growth can make particular genetic variants become very common even if they do not bring any fitness advantage (Excoffier et al. 2009). The formulation of this idea – sometimes referred to as allele surfing – relies on a set of population genetic principles (see Charlesworth & Charlesworth 2010). These principles derive from abstraction and idealization of complex biological processes, which are considered appropriate to solve a particular problem or kind of problem. Biologists may refer to both allele surfing and the set of population genetic principles as "theories," but it is only the former that can be falsified (theoretically by demonstrating that the conclusions do not follow from the premises, and empirically by demonstrating that allele surfing is not something that actually happens in real populations). Principles relevant to some discipline, methodology, or problem are typically not falsifiable. In fact, idealizations are commonly made despite full knowledge they are false (Potochnik 2017). Putting these "theories," such as population genetic theory, to the test means to assess how well those theories deliver satisfactory explanations, not to try to prove them wrong. In the next section, we argue that this is one reason why conceptual analysis can advance science.

Despite these reasons to doubt the classic view of the scientific method, many biologists adhere to its core features including, of course, the notion that experiments are the key to scientific knowledge and understanding. However, not even experiments are fundamental to all biologists. Many evolutionary biologists, for example, rely heavily on observation and "traces" of past events to reconstruct what happened and explain why it happened (Currie 2018; Chapter 10). Even molecular biology – perhaps the "ideal" reductionist science of biology – makes use of nonexperimental inference, for example, when species comparisons are used to substantiate claims that the activity of transposable elements ("jumping genes") causes genome evolution (Bourque et al. 2018). In her chapter, Cleland (Chapter 10) argues that experimental science is not as reliable as popularly thought, and historical science is not so unreliable. This may be welcome news to some biologists, and worrying to others. Regardless, it will be useful to become aware of the possible advantages and disadvantages of different scientific methods.

1.6 Scientific Concepts

Progress in the biological sciences is commonly driven by new discoveries and technological breakthroughs. The history of genetics provides many good examples, including PCR, high-throughput sequencing, and CRISPR-Cas. Nevertheless, data alone is often not enough, and conceptual analysis too can advance science (see Chapter 5, Chapter 8). This should perhaps not be surprising since concepts organize research agendas and feature in models, theories, and explanations. Since these are core features of science, understanding concepts may help to understand phenomena, or at least make science more effective at understanding phenomena. Concepts should not be confused with terminology, since a single concept can have several terms and the same term can be used for several concepts.

A good example of the latter is "gene," which in modern biology routinely refers to several different concepts (Griffiths & Stotz 2013; Kampourakis 2017). Biologists sometimes express the feeling that this multitude of meanings only creates problems and confusion. The solution, in their view, is a clear definition and a precise one-to-one map from concept to term. One illustrative example is many biologists' frustration over the term "epigenetics" (e.g., Deans & Maggert 2015; see Baedke 2018 for a conceptual analysis). For example, the journal *Cell Reports* requires authors to adhere to a strict definition of epigenetic (and epigenetics) and does not allow authors to refer to epigenetic as a stand-alone term.[3] This certainly removes ambiguity with respect to what biological feature that epigenetic refers to in papers published in this journal. This can be a good thing. Nevertheless, there are reasons to be skeptical toward this puritanism since concepts are more than

[3] This example is based on personal experience and email communication between authors and editors concerning a paper published in the journal (Tobi et al. 2018).

alternative representations that may be better suited to move beyond those limits. The literatures on developmental plasticity, extra-genetic inheritance, and niche construction are good illustrations of this endeavor (Laland et al. 2015).

This endeavor is complicated by the fact that it is possible to explore the role of plasticity, non-genetic inheritance, and niche construction in evolution without giving up on idealizations that will grant a privileged role of natural selection and genes (see Section 1.4). One particularly good illustration is the large body of research on "plasticity-led evolution" that has followed since the publication of Mary-Jane West-Eberhard's book *Developmental Plasticity and Evolution* (West-Eberhard 2003). Much of this research represents plasticity as a property of genotypes (biologists call these reaction norms; Levis & Pfennig 2016). As a result, any contribution of plasticity to adaptive evolution can be considered a *consequence* of natural selection on genes, rather than as a primary cause of adaptive change (see Uller et al. 2019). Examples such as these illustrate that simply "extending" a theory to include new phenomena need not resolve contention. It is also important to be aware that biologists' interpretative understanding of these phenomena will be dictated by their conceptual framework, that is, how they think about living beings.

1.7 Concluding Remarks

As biologists we feel that we have benefited from reading philosophy of science and engaging with philosophers. Not all biologists will feel the same, of course. But we do not believe that we are particularly unusual. Any field probably benefits from a diversity of perspectives, and this diversity tends to grow from encouraging reflection and critical assessment, not the least from scientists within the field. A little bit of philosophy of science is one way to make this happen.

The chapters in this book can be read in any order and where one would like to start depends on one's interest. Some chapters are easier to digest than others, and not everything will be to everyone's liking. Our advice to biologists is to look for issues that feel most relevant to their own work and begin there. If you are engaged in a field that is controversial, consider if the controversy could partly be dissolved through conceptual analysis or a more explicit formulation of the idealizations that are used for the phenomena you study. If you struggle to see the value of someone's research, or even an entire field, consider if your opinion is a result of different aims, problem agendas, or methods, or if it is shaped by preferences, values, and beliefs. If you look for new and exciting ways to tackle your problem, identify causes that are currently screened off, and alternative concepts and metaphors that may prove fruitful. If you are engaged in public outreach, consider if your research can be communicated more effectively by emphasizing the process by which knowledge is generated. Above all, stay curious, not just about biology, but also about the nature of the biological sciences.

References

Allis, C. D., Caparros, M-L., Jenuwein, T., Reinberg, D., & Lachlan, M. (2015). Epigenetics, 2nd ed. Cold Spring Harbor, NY: Cold Spring Harbor Laboratory Press.

Amundsen, R. (2005). *The Changing Role of the Embryo in Evolutionary Thought: Roots of Evo-Devo*. Cambridge: Cambridge University Press.

Baedke, J. (2018). *Above the Gene, beyond Biology: Toward a Philosophy of Epigenetics*. Pittsburgh, PA: University of Pittsburgh Press.

Bourque, G., Burns, K. H., Gehring, M., Gorbunova, V., Seluanov, A., Hammell, M., et al. (2018). Ten Things You Should Know about Transposable Elements. *Genome Biology* 19(1): 199. doi:10.1186/s13059-018-1577-z

Charlesworth, B. & Charlesworth, D. (2010). *Elements of Evolutionary Genetics*. Greenwood Village, CO: Roberts & Company Publishers.

Currie, A. (2018). *Rock, Bone and Ruin: An Optimist's Guide to the Historical Sciences*. Cambridge, MA: MIT Press.

Deans, C. & Maggert, K. A. (2015). What Do You Mean, "Epigenetic"? *Genetics* 199(4): 887–896.

Dickins, T. E. & Rahman, Q. (2012). The Extended Evolutionary Synthesis and the Role of Soft Inheritance in Evolution. *Proceedings of the Royal Society B: Biological Sciences* 279(1740): 2913–2921.

Excoffier, L., Foll, M., & Petit, R. J. (2009). Genetic Consequences of Range Expansions. *Annual Review of Ecology Evolution and Systematics* 40: 481–501.

Fleck, L. (1979 [1935]). *Genesis and Development of a Scientific Fact*. Chicago: Chicago University Press.

Griffiths, P. E. & Stotz, K. (2013). *Genetics and Philosophy: An Introduction*. Cambridge: Cambridge University Press.

Illiari, P. & Russo, F. (2014). *Causality: Philosophical Theory Meets Scientific Practice*. New York: Oxford University Press.

Kampourakis, K. (2017). *Making Sense of Genes*. Cambridge: Cambridge University Press.

Keller, E .F. (2000). *The Century of the Gene*. Cambridge, MA: Harvard University Press.

Kuhn, T. S. (1996) [1962]. *The Structure of Scientific Revolutions (3rd edn.)*. Chicago: University of Chicago Press.

Lakatos, I. (1978). *The Methodology of Scientific Research Programmes: Philosophical Papers Volume 1*. Cambridge: Cambridge University Press.

Laland, K., Uller, T., Feldman, M., Sterelny, K., Muller, G. B., Moczek, A., et al. (2014). Does Evolutionary Theory Need a Rethink? – POINT Yes, Urgently. *Nature* 514 (7521): 161–164.

Laland, K. N., Sterelny, K., Odling-Smee, J., Hoppitt, W., & Uller, T. (2011). Cause and Effect in Biology Revisited: Is Mayr's Proximate–Ultimate Dichotomy Still Useful? *Science*, 334(6062): 1512–1516.

Laland, K. N., Uller, T., Fellman, M. W., Sterelny, K., Muller, G. B., Moczek, A., et al. (2015). The Extended Evolutionary Synthesis: Its Structure, Assumptions and Predictions. *Proceedings of the Royal Society B: Biological Sciences* 282(1813). doi:10.1098/rspb.2015.1019

Levis, N. A. & Pfennig, D. W. (2016). Evaluating "Plasticity-First" Evolution in Nature: Key Criteria and Empirical Approaches. *Trends in Ecology & Evolution* 31(7): 563–574.

Lewens, T. (2015). *The Meaning of Science*. Milton Keynes, UK: Penguin Random House.

Love, Alan C. (2008). Explaining Evolutionary Innovations and Novelties: Criteria of Explanatory Adequacy and Epistemological Prerequisites. *Philosophy of Science* 75(5): 874–886.

Mayr, E. (1961). Cause and Effect in Biology – Kinds of Causes, Predictability, and Teleology Are Viewed by a Practicing Biologist. *Science* 134(348): 1501–1506.

Mill, J. & Heijmans, B. T. (2013). From Promises to Practical Strategies in Epigenetic Epidemiology. *Nature Reviews Genetics* 14(8): 585–594.

Muschick, M., Indermaur, A., & Salzburger, W. (2012). Convergent Evolution within an Adaptive Radiation of Cichlid Fishes. *Current Biology* 22(24): 2362–2368.

Noble, D. (2013). Physiology Is Rocking the Foundations of Evolutionary Biology. *Experimental Physiology* 98(8): 1235–1243.

Oyama, S. (2000). *The Onotgeny of Information*. Durham, NC: Duke University Press.

Pocheville, A. (2019). A Darwinian Dream: On Time, Levels, and Processes in Evolution. In T. Uller & K. N. Laland (eds.), *Evolutionary Causation: Biological and Philosophical Reflections*, pp. 265–298. Cambridge, MA: MIT Press.

Potochnik, A. (2017). *Idealization and the Aims of Science*. Chicago: Chicago University Press.

Rheinberger, H-J. (2000). Gene Concepts: Fragments from the Perspective of Molecular Biology. In P. Beurton, R. Falk, & H-J. Rhienberger (eds.), *The Concept of the Genes in Development and Evolution*, pp. 219–239. Cambridge: Cambridge University Press.

Roughgarden, J. (2009). *The Genial Gene: Deconstructing Darwinian Selfishness*. Los Angeles: University of California Press.

Schneider, R. F. & Meyer, A. (2017). How Plasticity, Genetic Assimilation and Cryptic Genetic Variation May Contribute to Adaptive Radiations. *Molecular Ecology* 26(1): 330–350.

Seehausen, O. (2006). African Cichlid Fish: A Model System in Adaptive Radiation Research. *Proceedings of the Royal Society B: Biological Sciences* 273(1597): 1987–1998.

Strevens, M. (2008). *Depth: An Account of Scientific Explanation*. Cambridge, MA: Harvard University Press.

Tobi, E. W., van den Heuvel, J., Zwaan, B. J., Lumey, L. H., Heijmans, B. T., & Uller, T. (2018). Selective Survival of Embryos Can Explain DNA Methylation Signatures of Adverse Prenatal Environments. *Cell Reports* 25(10): 2660–2667.

Uller, T., Feiner, N., Radersma, R., Jackson, I. S. C., & Rago, A. (2019). Developmental Plasticity and Evolutionary Explanations. *Evolution & Development*, 22(1–2): 47–55.

Uller, T. & Helanterä, H. (2019). Niche Construction and Conceptual Change in Evolutionary Biology. *The British Journal for the Philosophy of Science* 70(2): 351–375.

Waddington, C. H. 1957. *The Strategy of the Genes: A Discussion of Some Aspects of Theoretical Biology*. London: George Allen & Unwin Ltd.

Walsh, D. M. (2015). *Organism, Agency and Evolution*. Cambridge: Cambridge University Press.

West-Eberhard, M. J. (2003). *Developmental Plasticity and Evolution*. New York: Oxford University Press.

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

Woodward, James, "Scientific Explanation," *The Stanford Encyclopedia of Philosophy* (Fall 2017 Edition), Edward N. Zalta (ed.), https://plato.stanford.edu/archives/fall2017/entries/scientific-explanation/

Wray, G. A., Hoekstra, H. E., Futuyma, D. J., Lenski, R. E., Mackay, T. F. C., Schluter, D., & Strassmann, J. E. (2014). Does Evolutionary Theory Need a Rethink? COUNTERPOINT No, All Is Well. *Nature* 514(7521): 161–164.

Ylikoski, P. & Kuorikoski, J. (2010). Dissecting Explanatory Power. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 148(2): 201–219.

2 What Constitutes an Explanation in Biology?

ANGELA POTOCHNIK

2.1 Introduction

"Explaining" and "explanation" are words that tend to feature prominently in even the most basic descriptions of science. This is a big part of what science is about: generating explanations of our world (see also Chapter 3). It seems to be a big part of what biology is all about as well. Research in biology undoubtedly leads to practical applications in pursuits from medicine to agriculture to conservation, but one of its fundamental aims is to generate *understanding* of the living world around – and within – us.

The centrality of explanation to the scientific enterprise is matched by philosophers' enthusiasm for debating the nature of explanation in science. Philosophers of science have been up to our elbows in debates about the nature of scientific explanations since at least the middle of the twentieth century. Pet theories abound, but some basic insights into the broad contours of scientific explanation have also emerged.

In this chapter, I aim to provide a relatively nonpartisan discussion of the nature of explanation in biology, grounded in widely shared philosophical views about scientific explanation. At the same time, this discussion reflects what I think is important for philosophers and biologists alike to appreciate about successful scientific explanations. So, some points will be controversial, at least among philosophers. Along the way, I indicate which ideas are controversial and say something about the nature of controversy. I make three main points: (1) causal relationships and broad patterns have often been granted importance to scientific explanations, and they are in fact both important; (2) some explanations in biology cite the components of or processes in systems that account for the systems' features, whereas other explanations causal information. Likening this to the source of the flamingo's pink, the cardinal's red, and the goldfinch's yellow indicates something about the scope of this causal pattern. The scarlet ibis's coloration is due to one main form of avian pigmentation, a process that can create red, pink, orange, and yellow coloration. It can also be enlightening to point out a difference between the scarlet ibis and white ibis: The former but not the latter has a substantial volume of a carotenoid carrier protein in its blood (Trams 1969). This indicates why the scarlet ibis is able to absorb and transport carotenoids in this way while the related white ibis is not – also indicating something about the scope of the explanatory causal pattern.

The idea that causal patterns explain is a causal approach to explanation, but it is not merely a causal approach. On this view, simply providing some causal information is not sufficient for explanation. An explanation also needs to indicate the scope of the explanatory causal dependence, the range of circumstances in which similar causal dependence obtains. The explanatory value of this relates to the insight that explanations should unify by showing how disparate phenomena arise for the same reason or fit the same pattern. With the explanation of scarlet ibis coloration mentioned previously, insight is provided into conditions in which similar coloration occurs. Such explanations show how (potentially disparate) phenomena fit the same pattern. One might even say that causal pattern explanations are akin to the original deductive-nomological view of explanation, for they show how phenomena result from regularities in our world, though the regularities are not universal laws but limited in scope and may have exceptions.

I have suggested that the idea that causal patterns explain relates comfortably to a range of other ideas philosophers have had about explanation. But why think it is true? That is, why think that causal patterns are the sort of things that help us explain our world? Grasping the nature of a causal dependence and the scope in which that dependence holds is key to determining the causal structure of the world, which Gopnik (1998) has influentially argued is the endpoint of explanation. Grasping what I call causal patterns is also, as Woodward (2003) and others have argued, key to effective action. This indicates how and in what conditions we can act to bring about or prevent the focal phenomenon. Additionally, research in cognitive science suggests that causal information and broad generalizations are both the kinds of information that strike our intellects as explanatory (Lombrozo & Carey 2006) and that we learn more through the act of explaining (Lombrozo 2011). Uncovering causal patterns thus seems to do exactly the tasks we expect from our scientific explanations, and grasping causal patterns seems to look a lot like explaining.

Perhaps in some explanations, the causal dependence does more of the explanatory lifting than information about the scope of dependence, and vice versa in some other explanations. Explaining the production of ATP in anaerobic respiration by detailing the steps of glycolysis seems to get more explanatory "oomph" from the detailed causal information about the chemical reactions involved than from indications of the scope of this pattern, that is, the conditions in which anaerobic respiration occurs. In contrast, explaining a trait as a product of natural selection assimilates it to a broad range of phenomena – all physical and behavioral traits that have been positively selected – while giving relatively little causal detail. Natural selection, after all, has myriad ecological sources and leads to an astonishing variety of outcomes. Kim Sterelny (1996) calls these two approaches "actual sequence" and "robust process" explanations, respectively. It may be that some explanatory dependence patterns aren't even causal in nature. The statistical pattern of regression to the mean has a broad scope of applicability; this pattern can explain a number of phenomena, such as why the outliers in some quantitative trait tend to have offspring with less extreme values for that trait. But this pattern does not seem to be causal. The insight that causal patterns are explanatory can accommodate this variety. I say more later about which causal patterns are explanatory in which circumstances, but I suspect most scientific explanations occur between these extremes. Science generates understanding by depicting causal dependence and the patterns of when that dependence holds. That is, scientists by and large explain phenomena by depicting causal patterns.

2.3 From Mechanisms to Large-Scale Causes

My first point about explanation was the idea that information about causal dependence and information about the scope of that dependence are both important to scientific explanations. Scientific explanations feature causal patterns. The second point is that not all explanations cite components and processes. There is sometimes a tendency among philosophers and biologists alike to expect that information about causal dependence boils down to information about the parts of an entity or system and the processes they collectively carry out. Attending to the role of not just causes but also patterns in causal action makes clear this is not so. Biological explanations vary from component- and process-based to those that feature large-scale or structural causes. By "large-scale causes," I mean influence from some spatial or temporal distance, and by "structural causes," I mean contextual influences that shape a phenomenon but that do not change to precipitate the phenomenon. These varieties of causal influence are just as important, just as causal, and just as explanatory as components of a system and the processes they carry out.

Let us begin this discussion by returning to the idea, mentioned at the beginning of the previous section, that there is regularly a call in biology to identify the cause or mechanism responsible for something. Some philosophers of science take very seriously such appeals to mechanisms. The so-called "new mechanists" have put a lot of work into defining exactly what a mechanism is, and they think explanation in biology and related disciplines consists in describing mechanisms. These philosophers disagree about some details regarding the nature of mechanisms and their role in explanation, but the general picture is that mechanisms are integrated networks of components that carry out certain activities, thereby bringing about predictable outcomes. Paradigmatic examples of mechanisms are the ATP cycle (Bechtel & Richardson 1993), protein synthesis (Darden 2006), and the action potential (Craver 2006).

This view of mechanisms emphasizes the explanatory value of identifying processes carried out by the components of an entity. This approach fits very well with some areas of research in biology, but advocates of mechanisms tend to go further, expecting processes carried out by the components of a given entity to be central to *any* explanation, at least in biology. For example, Connolly et al. (2017) suggest that ecology needs to focus exclusively on component- and process-based models, which depict the causal roles of components of a system or the processes that precipitate some outcome. These authors claim that the value of such models is that they can capture the causal structure of a system. Philosophers who emphasize the explanatory value of mechanisms tend to equate a lack of detail about causal processes with a failure to explain (see, e.g., Craver 2006). This is related to McGill & Nekola's (2010) characterization of ecologists sometimes justifying the value of their work by appealing to it being more mechanistic, resulting in "ideological squabbles" about what qualifies as a mechanism.

I urge a much broader interpretation of the call to identify causes. Not all explanatory causal information regards processes carried out by components. Some explanations feature causal patterns regarding the environmental context, such as optimal foraging models for evolved food preferences, which cite ecological factors such as patterns of food distribution that give rise to certain selection pressures. These cite structural causes, by which I mean contextual factors that may not have changed to precipitate the phenomenon. The distribution of food need not have changed to bring about selection for given food preferences, but had this distribution been different in certain respects, the evolved food preferences would have been as well. Other explanations cite large-scale causes that are distant in space or time from the phenomenon, such as evolutionary or phylogenetic explanations for traits (temporally distant causes) and some ecological explanations, such as wetlands suffering due to decreased snowpack in the mountains (spatially distant cause). Finally, still other explanations describe highly general causal patterns to which focal phenomena cohere, such as appealing to the second law of thermodynamics to explain ice melting, cooling a beverage in the process; indicating that some trait is a product of natural selection; or indicating how scarlet ibis coloration is an instance of carotenoid pigmentation in birds. All of these are causal pattern explanations – explanations that feature information about causal dependence and information about the scope of that dependence – but none are naturally described as processes carried out by the components of the system. (Of course, as I have already suggested, some other causal pattern explanations *are* naturally characterized as such.)

Philosophers who advocate mechanistic explanation disagree with me on this. For any of these cases of structural or large-scale causes or highly general causal patterns, most mechanists will either claim that the causal pattern is aptly characterized as a mechanism or call into question whether there is a genuine causal explanation. But I think causal patterns featuring contextual factors, large-scale causes, and highly general regularities significantly shape the phenomena of our world. This leads to their explanatory significance.

The expectation that all explanatory causes are local, component-based processes thus inhibits our recognition of a range of important causal patterns. These include, among many others, how ecological features shape selection pressures, phylogenetic influences on traits, and the highly general pattern of carotenoid pigmentation in birds. In my view, focusing on a narrow sense of mechanism both results from and contributes to an inaccurately reductionist view of the world, where causal significance is expected to be local and component-based. Such an expectation renders large-scale causal patterns less visible or even "spooky" seeming.

In this section, I have suggested that the call in biology for causes, process, or mechanism should be interpreted as a call for information about causal patterns, wherever they are found. This is a broader interpretation than that of philosophers of science who emphasize the significance of mechanisms, understood roughly as processes carried out by components. Yet this broader interpretation of the call for mechanism goes beyond establishing correlation, or the existence of a "mere" pattern. Causal patterns are more than just patterns: they are regularities in how causes exert their effects. This gives information about what to expect in different circumstances and about how to intervene on a phenomenon to bring about a desired effect. And this is so whether the causal pattern in question is local and component-based or large-scale. In the next section, I give reason to think that the very same phenomena will sometimes be explained by citing local, mechanistic causes and other times large-scale causes.

Much more argument would be needed for me to provide full support for this broad conception of explanatory causal patterns. But I hope this brief discussion is sufficient for two purposes. First, to provide some initial motivation for the idea that biologists should look beyond the local components of a system in their hunt for causes, explicitly including consideration of the significance of large-scale causal factors. And, second, simply to highlight that a philosophical question about scientific explanation is the degree to which explanations must be component- and process-based.

2.4 A Variety of Explanations Without Conflict

Scientists are regularly in the position of trying to discern whether a proffered explanation is right. This is a challenging task, and the details example, such a focus may be the evolution of carotenoid metabolism and selective transport in birds, or whether and in what ways sexual selection is responsible for carotenoid coloration in the scarlet ibis, or how parasite load influences carotenoid metabolism and selective transport, etc.

The relationship among multiple explanations of a given phenomenon has come up in a few different contexts in philosophical discussions of scientific explanation. A handful of influential philosophers have suggested that any given explanation is partial, so there are inevitably multiple different explanations of any one phenomenon (Railton 1981; Lewis 1986). Philosophers have disagreed about whether and in what ways such multiple explanations should relate to one another, with some anticipating integration of these explanations (e.g., Mitchell 2003) and others arguing, as I have also suggested here, that multiple different explanations of a phenomenon remain independent from one another. From my perspective, this explanatory independence, as I have called it elsewhere (Potochnik 2010), arises due to varying research interests even among those investigating the same phenomenon and to different causal patterns being explanatory in light of those various interests.

In some cases, these different interests are obvious. Other times, biologists may take themselves to disagree about causal facts, but the disagreement is also motivated by different research priorities and thus different aims for the explanations they are developing. To illustrate this point, let's return to Mayr's proximate–ultimate distinction but this time for a different purpose. Laland et al. (2011) challenged an implication this distinction has often been taken to have, namely that developmental processes are evolutionarily unimportant (one might say, merely proximate). These researchers emphasize that, to the contrary, feedback loops exist by which developmental processes influence evolution. They conclude: "It is now vital to recognize that developmental processes frequently play some role in explaining why characters possess the properties that they do, as well as in accounts of the historical processes that explain their current state" (p. 1516). This is an important observation that has significant implications for evolutionary theory. But, I do not think the significance of developmental processes to evolution is a reason to replace selective explanations of traits with selective-developmental explanations. Rather, in my view, this is better interpreted as the identification of a neglected kind of causal pattern, namely patterns in how development influences potential evolutionary outcomes. These different causal patterns are explanatory in light of different research questions. Sometimes a classic evolutionary explanation suffices, and the influence of development on evolution can be ignored; other times the latter is central to what biologists aim to understand.

I have suggested that there can be multiple noncompeting causal pattern explanations for any given phenomenon. In light of this idea, I urge biologists to take seriously the possibility that the apparent conflict among different research programs arises not due to different competing explanations of the same phenomenon, but rather due to different research agendas that lead to emphasizing different causal patterns. Sometimes a breakthrough in understanding warrants revisiting what we thought we knew, including the nature of the causal patterns we had posited in our explanations. It may be that, for some traits, there is no evolutionary explanation – no causal patterns to be found – without taking both selection and development into account. Other times, breakthroughs in understanding bring to light new causal patterns but do not undermine our existing explanations.

2.5 Conclusion

Explanation is taken to be an important aim, if not the central aim, of science. In this chapter, I have motivated three ideas about the nature of scientific explanations. These ideas are grounded in philosophical debates about explanation, even as they also reflect my particular views. First, I have suggested that philosophical debates about the definitive features of explanation support the idea that both causal dependence and the scope of that dependence – that is, causal patterns – are important to explanation. If this is so, biologists might explicitly think about both the causal content and the generality or scope of the explanations they develop. It is not always more explanatory to build in more detail. When scarlet ibis coloration is investigated in the context of explaining carotenoid metabolism and selective transport in birds, any reference to, say, sexual selection for the ibis's coloration is mere distraction. Omissions and simplifying assumptions are ways of signaling that those details don't matter given the present research aims – that the pattern in question is independent of them.

Second, I have suggested that while some explanations focus on components and processes, others focus on large-scale causes, including contextual features, distant causes, and highly general patterns. I urge biologists to look beyond the local components of a system in their hunt for causes, and to explicitly include consideration of the significance of large-scale causal factors. I am inclined to think that all of us – scientists, philosophers, and the public alike – share certain reductionist tendencies. Among these is a tendency to consider the large-scale to be a fixed background and the local and tiny to be where the causal action is. Across science, again and again, this expectation has been revealed to be incorrect. And yet the tendency persists.

Third and finally, I have suggested that phenomena of interest in biology may have multiple explanations, each occasioned by different research agendas and featuring different causal patterns. In my view, some disputes about research strategies and methods are, at root, disagreements about which explanations are most interesting – which causal patterns enlightening – rather than disagreements about which explanations are accurate. This idea also relates to the idea I motivated about the explanatory value of large-scale causal patterns. One way in which large-scale causes have been rendered invisible is by pointing out that there is already an explanation in terms of components or other local factors. But if phenomena have multiple explanations, the recognition of local, small-scale influences shouldn't lead us to expect the absence of large-scale explanations. Cancer has genomic causes, but it also has developmental, environmental, and socioeconomic causes. And yet our research dollars seem to go disproportionately to studying the tiny molecular bits residing inside us.

This reveals the error in what I take to be another reductionist tendency: an implicit expectation that events have just one or a few causes. To the contrary, complex causal relations abound, with any event bearing the influence of many causes, and causal interaction and feedback common (Love 2017). Recognizing and emphasizing that biological phenomena embody multiple causal patterns, and that different causal patterns can figure into explanations tailored to different questions, is one step toward counteracting these reductionist tendencies.

References

Bechtel, W. & Richardson, R. (1993). *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Cambridge, MA: MIT Press.

Connolly, S. R., Keith, S. A., Colwell, R. K., & Rahbek, C. (2017). Process, Mechanism, and Modeling in Macroecology. *Trends in Ecology & Evolution* 32(11): 835–844.

Craver, C. (2006). When Mechanistic Models Explain. *Synthese* 153(3): 355–376.

Darden, L. (2006). *Reasoning in Biological Discoveries*. New York: Cambridge University Press.

Friedman, M. (1974). Explanation and Scientific Understanding. *The Journal of Philosophy* 71(1): 5–19.

Gopnik, A. (1998). Explanation as Orgasm. *Minds and Machines* 8: 101–118.

Hempel, C. (1965). *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York: Free Press.

Kitcher, P. (1981). Explanatory Unification. *Philosophy of Science* 48(4): 507–531.

Laland, K., Sterelny, K. Odling-Smee, J., Hoppitt, W., & Uller, T. (2011). Cause and Effect in Biology Revisited: Is Mayr's Proximate-Ultimate Dichotomy Still Useful? *Science* 334(6062): 1512–1516.

Lewis, D. (1986). Causal Explanation. in *Philosophical Papers*, Vol. II. Oxford: Oxford UP.

Lombrozo, T. (2011). The Instrumental Value of Explanations. *Philosophy Compass* 6(8): 539–551.

Lombrozo, T. & Carey, S. (2006). Functional Explanation and the Function of Explanation. *Cognition* 99(2): 167–204.

Love, A. (2017). Building Integrated Explanatory Models of Complex Biological Phenomena: From Mill's Methods to a Causal Mosaic. In M. Massimi et al. (eds.), *EPSA15 Selected Papers, European Studies in Philosophy of Science* 5: 221–232.

Mayr, E. (1961). Cause and Effect in Biology. *Science* 134(3489): 1501.

McGill, B. & Nekola, J. (2010). Mechanisms in Macroecology: AWOL or Purloined Letter? Towards a Pragmatic View of Mechanism. *Oikos* 119(4): 591–603.

Miller, J. H., Druckenmiller, P., & Bahn, V. (2013). Antlers of the Arctic Refuge: Capturing Multi-Generational Patterns of Calving Ground Use from Bones on the Landscape. *Proceedings of the Royal Society B: Biological Sciences* 280(1759): 20130275.

Mitchell, S. (2003). *Biological Complexity and Integrative Pluralism*. Cambridge: Cambridge University Press.

Potochnik, A. (2010). Explanatory Independence and Epistemic Interdependence: A Case Study of the Optimality Approach. *The British Journal for Philosophy of Science* 61(1): 213–233.

Potochnik, A. (2013). Defusing Ideological Defenses in Biology. *BioScience* 63(1): 118–123.

Potochnik, A. (2015). Causal Patterns and Adequate Explanations. *Philosophical Studies* 172(5): 1163–1182.

Potochnik, A. (2017). *Idealization and the Aims of Science*. Chicago: University of Chicago Press.

Railton, P. (1981). Probability, Explanation, and Information. *Synthese* 48(2): 233–256.

Reutlinger, A. & Saatsi, J. (eds.) (2018). *Explanation Beyond Causation*. Oxford: Oxford University Press.

Sterelny, K. (1996). Explanatory Pluralism in Evolutionary Biology. *Biology and Philosophy* 11: 193–214.

Strevens, M. (2004). The Causal and Unification Approaches to Explanation Unified – Causally. *Noûs* 38(1): 154–179.
Strevens, M. (2008). *Depth: An Account of Scientific Explanation*. Cambridge: Harvard University Press.
Trams, Eberhard G. (1969). Carotenoid Transport in the Plasma of the Scarlet Ibis. *Comparative Biochemistry and Physiology* 28(3): 1177–1184.
Woodward, J. (2003). *Making Things Happen*. Oxford: Oxford University Press.

know-how to do something one must have the ability to do it. For example, one knows how to swim, if one can swim. If one cannot swim, one does not know-how to do it. Admittedly, one might have a lot of *propositional* knowledge about how to swim, but one does not really *know-how* without being able to actually swim.[4] Knowledge-how to do something amounts to having certain abilities – namely, the ability to do that thing.

Propositional knowledge is a bit more complex to spell out. Fortunately, it has been a key area of study for many philosophers, so the ground has been covered fairly well. Although there is disagreement about the specifics, it is widely accepted that three primary requirements of knowledge are truth, belief, and justification. This is sometimes referred to as the "traditional account of knowledge" or the "justified true belief account of knowledge" (McCain 2016). The general idea is that knowledge is made up of, or at least entails, justified true belief.[5] That is to say, when you know a proposition is true, you believe the proposition is true, it is true, and your belief is justified (it is based on sufficiently strong evidence).

Why think that knowledge requires justified true belief? To perhaps oversimplify, knowledge is information that we can go on – the things that we know can be relied upon when deciding how we should act (Fantl & McGrath 2009, ch. 3). If someone wants orange juice and she knows that there is orange juice in the fridge, she can use that knowledge to guide her actions. This entails that knowledge requires belief. Only if the person holds the belief that there is orange juice in the fridge can this belief serve as knowledge that can guide this person to open the fridge to get the orange juice when she wants to drink it.

[4] There are complications here. One might know-how to swim but recently broke a leg. In such a case, it seems that one still knows how to swim even though one cannot at this time swim. This sort of general concern will be set aside as the primary focus here is on typical cases of know-how. For an excellent discussion of the current state of scholarship on knowledge how, see Carter & Poston 2018.

[5] The reason for this qualification is that while the traditional (and still predominate) view is that knowledge can be analyzed into components which include justified true belief, following Williamson (2000) a significant number of philosophers take knowledge to be an unanalyzable primitive concept. That being said, those who, like Williamson, think that knowledge is primitive in this sense are often willing to grant that knowledge entails justified true belief – i.e., if you know some claim, that entails that the claim is true, you believe the claim, and that you have sufficiently strong evidence for the claim.

Continuing with the thought that knowledge is information that we can rely upon, knowledge requires truth. We cannot really rely upon false information – it is very likely to lead us astray. While at times things are said that suggest that knowledge can be false – "people knew the earth was flat," for instance – this is best understood as simply loose talk. People did not know that the earth was flat because it was not, and is not, flat. Rather, people *thought* they knew, or even had good reasons to think that they knew, that the earth was flat without actually knowing this. Sometimes we think we know, and we do; other times we think we know, but we are mistaken. Holding beliefs of various kinds is possible, but only those that are true are candidates for knowledge.

Finally, knowledge requires justification because this is what separates knowledge from lucky guesses. There is a very big difference between someone who sees the result of a coin toss and someone who simply guesses that the coin landed "heads" up without seeing it. Even if the coin did in fact land "heads" up, only the first person has knowledge. What is the difference? The person who sees the coin is justified in believing this – she has good reasons/evidence in support of the claim that the coin landed "heads" up. The second person does not know how the coin landed. True, she guessed correctly, but she simply guessed. Getting the right answer by luck is not sufficient for knowledge. For these sorts of considerations, it is widely held that in order to know that some claim is true one must have a justified true belief that that claim is true.

Biological knowledge is a species of knowledge. Consequently, biological knowledge has these same features. Know-how in biology (e.g., using a population model to make predictions) requires having certain skills and abilities. Propositional biological knowledge involves believing a true biological proposition on the basis of sufficiently strong evidence. Hence, when it comes to the general nature of knowledge there is nothing particularly special about biological knowledge as such. Nonetheless, there *are* features of how biological knowledge is generated and transmitted as well as additional distinctions between kinds of biological knowledge that make biological knowledge a special kind of knowledge.

3.2 Kinds of Propositional Biological Knowledge

Although both know-how and propositional knowledge are important features of biological knowledge and play key roles in biological practice, from this point on the focus will be on propositional knowledge. When it comes to propositional biological knowledge there is an important distinction that should be drawn. It is useful to distinguish between direct/observational knowledge on the one hand and theoretical knowledge on the other. Direct/observational knowledge is knowledge gained by way of observations. For example, one learns the size of the population of lions in a particular area of Africa, say, by observing the area and keeping track of the number of lions. Such direct/observational knowledge can be readily transmitted, of course. Once someone has observational knowledge of the population size, they can easily share this knowledge with others via testimony – they can tell others or publish the results in a journal, for example.

Theoretical knowledge is different. It is not gained simply by way of making observations. Instead, theoretical knowledge often starts with observational knowledge and builds from there. Theoretical knowledge arises from the attempt to explain the observational knowledge, which is descriptive. As Kampourakis and Niebert (2018, p. 237) have aptly put the point, "While a description [observational knowledge] aims to answer *what* has happened, explanations try to give us an answer to *why* it took place." We gain theoretical knowledge by explaining the observations we make both in and out of the lab. Similar to observational knowledge, we can readily share theoretical knowledge via testimony – common mediums for this are academic publications and professional conference presentations. But, how do we come to have knowledge through the act of attempting to explain what has been observed? This is achieved through inference to the best explanation (IBE).

3.3 Theoretical Knowledge and Inference to the Best Explanation

Theoretical biological knowledge arises via inference to the best explanation (IBE). This prompts three questions. The first is simply what exactly is an explanation? Many answers to this question have been proposed, and it is still a matter of philosophical debate (McCain 2016, ch. 9). For our purposes here, it is enough to think of an explanation as an answer to the question of why (or how) a particular phenomenon occurred (see Chapter 2 for more in-depth discussion). The other two questions, which will be the focus of the rest of this section, are: How can we gain knowledge via IBE? And, is IBE a reliable way to generate knowledge?

3.3.1 Getting Theoretical Knowledge from Explanations[6]

Inference to the best explanation (IBE) is familiar and ubiquitous in our lives. When a veterinarian determines what is wrong with the family pet, she uses IBE. She considers the symptoms the pet has displayed and the potential explanations of those symptoms. She then infers that the explanation that best explains those symptoms is correct. Similarly, when you try to determine why your car will not start, you employ IBE. It was running fine earlier today, the battery is new, it has plenty of fuel, and the alternator has been making a sound for weeks. A reasonable inference to draw here is that the alternator is bad. The bad alternator best explains the data you have about your car. Philosophers have argued that not only do we use IBE to come to knowledge via inferences of this sort, we also employ it any time we gain knowledge via testimony (Fricker 1994; Lipton 1998). Their thinking is that we can reasonably accept what someone tells us (whether this testimony comes orally or via some written medium) only if the best explanation of why they are telling us what they say is that they know what they are talking about.

Given the ubiquity of IBE, it may not be surprising that scientific reasoning employs it in a careful and refined form. It will be helpful to make the general form of IBE more precise here. McCain & Poston (2019) have formalized it this way (though I have changed the phrasing slightly here):

Inference to the Best Explanation (IBE)

(1) There is some data, *d*, and some background evidence, *k*.
(2) E explains *d* better than any available competing potential explanation.
(3) E is a good explanation given *k*.
(4) Therefore, E is true.[7]

[6] For more on how explanatory reasoning leads to knowledge in science in general, see McCain 2016, 2019.

[7] A key difference between McCain & Poston's formulation and others is their inclusion of (3). In essence, the inclusion of (3) means that IBE in terms of a slogan should be "infer the best sufficiently good explanation" rather than, as it is typically understood,

The most famous instance of IBE in biology comes from Charles Darwin. In *The Origin of Species* he justified natural selection on the grounds that it provides the best explanation of a wide variety of biological facts:

> It can hardly be supposed that a false theory would explain, in so satisfactory a manner as does the theory of natural selection, the several large classes of facts above specified. It has recently been objected that this is an unsafe method of arguing; but it is a method used in judging of the common events of life and has often been used by the greatest natural philosophers. (1872/1962, p. 476)

Before moving on, there are three points about IBE that need clarification. The first is that IBE involves inferring that the best explanation of a field of competing explanations is the one that is true. Since at most one of these competing (mutually exclusive) explanations can be true, it is important to realize that the explanations being compared are "potential explanations" – they are each such that *if* true, they would explain the relevant data. Hence, IBE involves inferring that the best *potential* explanation is the *actual* explanation of the data (or phenomenon).

The second point needing clarification is what makes an explanation the "best." A large number of explanatory virtues have been proposed by both scientists and philosophers of science. Common explanatory virtues include: *simplicity* (as Newton 1999, p. 794 put it, "No more causes of natural things should be admitted than are both true and sufficient to explain their phenomena ... For nature is simple and does not indulge in the luxury of superfluous causes"); *explanatory power* (the amount of data explained); *conservatism* (consistency with background knowledge); and *predictive power* (making accurate predictions). Of course, much more could be said about these various virtues. For instance, explanatory power concerns not only the individual points of data explained but also the kinds of data explained. So, an explanation that explains seemingly disparate phenomena is more explanatorily powerful, all else equal, than an explanation that only explains one kind of phenomena. Also, there are many additional explanatory virtues that have been proposed (McMullin 2008; Beebe 2009). Finally, there are a

"infer the best explanation." Lipton (2004) and Musgrave (1988) have also argued in support of including similar restrictions on IBE.