

Levels of Explanation Reconceived

Angela Potochnik

Abstract

A common argument against explanatory reductionism is that higher-level explanations are sometimes or always preferable because they are more general than reductive explanations. Here I challenge two basic assumptions that are needed for that argument to succeed. It cannot be assumed that higher-level explanations are more general than their lower-level alternatives, nor that higher-level explanations are general in the right way to be explanatory. I suggest a novel form of pluralism regarding levels of explanation, according to which explanations at different levels are preferable in different circumstances because they offer different types of generality, which are appropriate in different circumstances of explanation.

1 Introduction

According to the classic reductionist conception of science, the existence of different fields of science may now be a practical necessity, but an ultimate aim of science is to show how all events follow from microphysical laws and events (Oppenheim and Putnam, 1958). On this view, the explanation of an event is always improved by giving information about the lower-level determiners of the event—ideally, the microphysical determiners. A common response to this reductionism about scientific explanation is to defend the value of non-reductive explanations on account of their generality (Fodor, 1974; Putnam, 1975; Garfinkel, 1981;

Kitcher, 1984; Sober, 1984; Jackson and Pettit, 1992; Sober, 1999). The basic argument is that, whereas a microphysical explanation applies only to systems fitting that microphysical description, a higher-level explanation has a much broader range of applicability. This is because higher-level descriptions can be realized by a variety of microphysical systems. This argument has been used to defend the value of a variety of high-level explanations over their lower-level competitors. The idea is that higher-level explanations are sometimes or always preferable because they apply in a broader range of circumstances.¹

Consider Putnam's (1975) familiar example of the peg and the board. There are two holes in a board: each has a diameter of one inch, but one is circular and the other is square. The peg is a cube with sides slightly shorter than one inch. If we are to explain why the peg passes through the square hole but not the round hole, we do so by citing the dimensions of the peg and the two holes, along with geometrical information about the relationships among these dimensions. This higher-level explanation is much more general—applies in a much broader range of circumstances—than a lower-level explanation that specifies the exact atomic structure of the peg and the board and computes all possible trajectories of these atoms in order to deduce that the peg never passes through the round hole but does pass through the square hole.

This style of defense for higher-level explanations depends on the idea that generality is of explanatory worth, but *why* general explanations should be taken to be valuable often remains unarticulated.² Examining the motivation for this idea is the first step to assessing this popular line of argument against explanatory reductionism. Garfinkel (1981) suggests

¹Jackson and Pettit (1992) and Sober (1999) do not argue that general explanations are always best; they are pluralists regarding levels of explanation. They claim that higher-level explanations are valuable because they are more general, but in their view whether generality or specificity is of explanatory value varies. I discuss this view in Section 4.

²On the unification account of explanation (Friedman, 1974; Kitcher, 1981), generality is constitutive of explanation. Generality is also an important consideration for Strevens's (2009) causal account of explanation. However, the antireductionist arguments considered here do not rely on a specific theory of explanation. I similarly avoid committing to a particular account of explanation in my analysis.

one possible motivation. He argues that higher-level explanations are preferable because they focus on the causes that critically affect the outcome to be explained, whereas the alternative lower-level explanations cite overly specific causal details. Perhaps, then, general explanations are valuable because they identify the critical causal factors and exclude causal information that is in some sense unimportant.

This idea can be put into broader terms that are compatible with a range of views regarding the nature of explanation. The sense of generality at issue is the breadth of an explanation's applicability to different systems, or equivalently, in different circumstances. An explanation's generality is inversely related to the amount of information it provides about the system in question. Including more information more finely specifies the process, which reduces the explanation's breadth of applicability, whereas removing information eliminates specifications of the process, which extends the explanation's breadth of applicability.³ A general explanation thus includes fewer details about the target system, thereby focusing attention on the limited information it does contain. This information is more perspicuous than it would be if lots of other features of the system were also specified. Moreover, the explanation advertises the unimportance of other features of the target system by neglecting them. For both of these reasons, if the information featured in an explanation regards the explanatorily important factors, then the explanation benefits from its generality. Accounts vary on what factors are explanatorily important. For Garfinkel (1981) they are causes critical to the outcome to be explained. Others who defend the value of general explanations think the explanatorily important factors are the difference-makers (Strevens, 2009), or the factors that identify the system as fitting into a unified pattern (Friedman, 1974; Kitcher, 1981).

³This assumes that other features of the explanations in question are held constant. The conception of generality as breadth of applicability may not include all senses in which the generality of explanations vary, but this is the sense of generality that is at work in the defenses of higher-level explanation under consideration. All claims about generality in this paper regard this particular sense of the word.

This is one argument for the idea that explanations benefit from being general. A general explanation focuses attention on the information that is properly explanatory and advertises the unimportance of other features of the system by neglecting them. I assume that those who defend the value of higher-level explanations on the grounds of their generality would be amenable to a justification along these lines. The examples they offer suggest so. The square peg fails to travel through the round hole because of a geometric relationship regardless of the particular atomic structure of the peg and board (Putnam, 1975); Gresham's Law governs economic exchanges regardless of whether the exchanges involve wampum or dollars (Fodor, 1974); and modeling the evolution of altruism in terms of fitness provides "a useful generality" lacking from physical descriptions of various instances of altruism and selfishness (Sober, 1999).

Two claims are needed in order to use this argument for the explanatory worth of generality to defend higher-level explanation. First, it must be the case that higher-level explanations are more general than lower-level explanations. Otherwise, the explanatory value of generality would not translate into the value of higher-level explanations. Fodor (1974) and others appeal to the idea of multiple realizability to motivate this claim. Second, it must be the case that the information included in a higher-level explanation is important to the outcome to be explained, whereas the neglected information is unimportant to the outcome. According to the argument sketched above, the value of general explanations is that they focus attention on the important information. If neglected lower-level information is explanatorily important, then the generality of a higher-level explanation does not improve the explanation. Both of these claims turn out to be problematic. The problems stem from the fact that the relationship between higher- and lower-level explanations often cannot be accurately described in terms of supervenience.

2 Supervenience and Levels of Explanation

The successful defense of nonreductive explanations on the grounds of their generality relies on the idea that higher-level explanations are more general than their lower-level alternatives. This idea is widely accepted (Fodor, 1974; Putnam, 1975; Garfinkel, 1981; Jackson and Pettit, 1992; Sober, 1999; Strevens, 2009) and is motivated by the belief that higher-level properties are multiply realized by lower-level properties—that is, that many lower-level properties individually suffice as supervenience bases for a given higher-level property. If a lower-level explanation cites properties that are only one way among many to realize the properties cited in a higher-level explanation, then the higher-level explanation is more general than the lower-level explanation. That is, it applies to a broader range of systems. This is the situation in Putnam’s (1975) peg and board example. The lower-level explanation cites the atomic structure of the peg and board to account for why the square peg cannot travel through the round hole, but this exact atomic structure is only one way among many for the geometric relationship between the peg and the board to be realized. A higher-level explanation that cites only the geometric relationship is more general than the alternative lower-level explanation.

Yet not all higher-level explanations are more general than their lower-level alternatives. Let us grant that all higher-level properties supervene on lower-level properties, and that multiple realization is common. Even so, this relationship among higher- and lower-level *properties* need not result in the predicted relationship between higher- and lower-level *explanations*. Establishing that implication requires an additional premise: that lower-level explanations are formulated in terms of the supervenience bases of the competing higher-level explanations. I see no reason to think that this premise is always or even often true.

Aside from being multiply realized, higher-level properties are often *complexly* realized. That is, higher-level properties are realized by a complicated combination of lower-level

properties, even in individual (token) occurrences. Fodor (1974) points out that the physical description of an economic exchange will sometimes involve dollar bills, sometimes strings of wampum, sometimes signing one's name to a check; this is multiple realization. But consider the physical description of an individual event that has the property of being an economic exchange—say, you putting a dollar bill into a vending machine yesterday and getting a candy bar in return. What physical property qualifies this event as an economic exchange? It's not enough that you gave up a dollar bill and received a candy bar: a marooned explorer with a penchant for dramatics might throw her cash to the wind just as her sole companion hands her the last candy bar. The physical properties of your vending machine encounter that jointly qualify it as an economic exchange must include (at least) the relinquishing of the bill, the acquisition of the candy bar, the circulation of relevantly similar bills in relevantly similar circumstances such that the bill qualifies as valuable, the vending machine specifications such that the encounter counts as its ordinary functioning, specifications of an economic entity on behalf of which the machine acts, and your specifications such that candy is of potential value to you and agreements are things you can enter into.

Fields of science that are equipped to provide lower-level explanations formulate those explanations for their own purposes, as part of investigations into phenomena of interest in those fields. Consequently, lower-level explanations feature properties that are relevant to those investigations, not the complex, highly particular conjunctions of properties that serve as supervenience bases. This is so even when the target of the explanation is a higher-level event. Consider the relationship between an organism's genotype (genetic composition) and phenotype (observable characteristics). Philosophers tend to think of genetic properties as occupying a lower-level of organization than phenotypic properties (e.g., Dupré, 1993; Rosenberg, 1994), and evolutionary change is modeled by focusing on one or the other. Yet phenotypes do not supervene on genes or genotypes, but on a complex combination of properties including many other properties of the organism,

properties of the environment, and sometimes even properties of other organisms. The lower-level properties under investigation—genes—are not proper candidates for the supervenience bases of phenotypes. As a result, explanations featuring genes or genotypes are not related by supervenience to explanations featuring higher-level, phenotypic properties, even when these properties are used to explain the same evolutionary outcomes.

This example is not atypical. Because different fields of science formulate explanations for different purposes, lower-level explanations and higher-level explanations of the same phenomena are often not related by supervenience. This is so even for some examples used to defend higher-level explanation. Above I suggested this is true for Fodor's example of monetary exchanges. Garfinkel (1981) provides a biology example that also illustrates my point: competing explanations for the death of a rabbit. Garfinkel considers a lower-level explanation for the death that cites the presence of the unfortunate rabbit in the capture-space of a particular fox, and he considers a higher-level explanation that cites the high concentration of foxes in the area. Yet the property of high fox-concentration clearly does not supervene upon the property of the rabbit being located in the capture-space of one fox, but on a combination of the spatial location of all foxes in the area. The only relationship apparent between these higher-level and lower-level properties is that the high concentration of foxes makes it more likely that the rabbit is in the capture-space of some fox or other. This bears no resemblance to supervenience.

Explanations that are not related by supervenience may still qualify as lower- and higher-level. The explanations in Garfinkel's rabbit example are formulated at lower and higher levels in the sense that one deals with an individual and the other with a population of individuals. In general, a lower-level explanation cites properties of objects that stand in a part-whole relationship to objects referenced in the competing higher-level explanation. Genetic and phenotypic explanations are also related in this way: genes are parts of organisms. Often, the properties cited in a lower-level explanation are part of the complex

group of properties on which the properties cited in a higher-level explanation supervenes, and occasionally, lower-level properties are proper supervenience bases.

This lower/higher distinction may be less crucial when explanations are not related by supervenience, but it remains significant that different fields formulate competing explanations for a single event. Evolutionary explanations provide a ready example of this. Population genetics and evolutionary ecology provide competing explanations for why many traits evolve. Genetic explanations show how the distribution of genotypes related to the trait changed generation by generation, whereas phenotypic explanations show how the environment selectively advantaged the trait in question. These explanations need not make competing claims about the world: the genetic explanation is silent regarding the environmental influences on selection, and the phenotypic explanation is silent regarding the genetic causes of the trait (Potochnik, forthcoming). The explanations *do* compete, though, in a way important to the debate about levels of explanation. The two explanations cite different properties, at different levels of organization, to explain the same event.

When higher- and lower-level explanations are not related by supervenience, nothing ensures that higher-level explanations are more general than their lower-level competitors. If a lower-level explanation does not cite the supervenience bases of properties cited in a higher-level explanation, then the former may not have a more limited range of applicability than the latter. Without supervenience and multiple realization, competing lower- and higher-level explanations may simply have different ranges of applicability. Such explanations are general in different ways; they offer different types of generality. By this I do not mean that they satisfy different *definitions* of generality, but that each applies to some range of systems that the other does not.

So it is for genetic and phenotypic explanations of evolutionary phenomena (Potochnik, 2007). A phenotypic explanation of a trait is applicable whenever the represented fitness relationship obtains, even if the genetic details vary. For example, a single phenotypic model

explains the trait of when in the lifecycle reproduction begins for a vast array of species, even though the genes that influence that trait differ (Bull et al., 2004). Genetic explanations are not general in this way. In contrast, a genetic explanation of a trait is applicable whenever the represented combination of selection and transmission obtains, even if the environmental sources of fitness are different. For example, a single genetic model explains how a selectively advantaged trait influenced by two genes evolves, regardless of the environmental sources of fitness (e.g. Curnow and Ayres, 2007). Phenotypic explanations represent the sources of fitness, so they are not general in this way.

3 Contextual Explanation

In Section 1 I distinguished two claims that must be true in order to use the explanatory worth of generality to defend nonreductive explanations. I dealt with the first in Section 2. The second claim is that higher-level explanations provide information that is important to the explanation and omit only unimportant information. General explanations are valuable if they focus attention on explanatorily important information. If an explanation purchases greater generality by omitting information that is important to the explanation, then the increased generality impoverishes rather than improves the explanation. Suppose that a higher-level explanation is more general than its lower-level competitor (though I argued in Section 2 that this is uncommon). This greater generality only results in a better explanation if it is the right type of generality; the information omitted must be unimportant to the explanation.

That more generality is not always better is demonstrated by taking this idea to the extreme. Rigorously applying the precept of maximizing generality would result in explanations that cite every possible circumstance that would lead to the event to be explained. So, e.g., the most general explanation of a room getting dark is:

Either the light was turned off, or the lightbulb burned out, or it is twilight, or someone drew the curtains, or a storm is gathering outside, or...

This vast disjunction is maximally general, but it is not very informative.⁴ Expansive disjunctions like this surely cannot be the preferable form of explanation, so it seems explanations should not simply be maximally general. Other arguments to this effect have been made for different accounts of explanation. For instance, Strevens (2009) claims that a causal account of explanation must specify which aspects of a causal process belong in an event's explanation and which aspects should be omitted for the sake of generality. Kitcher (1981) claims that the generality of an explanatory pattern should be limited by considerations of the pattern's stringency—the commonalities among its instances. Explanations should not simply be maximally general. If generality is of explanatory worth, the trick is to find a more nuanced way to express this.

Consider evolutionary explanations again. A vast number of factors influence individual evolutionary outcomes. For instance, *Homo sapiens* have evolved large brains relative to our hominid kin. This is the result of (at least) genetic mutations, developmental processes, environmental conditions, other evolutionary events, and phylogenetic history. An explanation that neglects some of these many influences focuses attention on the influences that are cited and demonstrates that small variations in other factors would not have interfered with the evolution of large brains. For these reasons, a general explanation is preferable. Yet some ways of increasing the generality of the explanation would be detrimental, e.g., an explanation that describes how large brains develop but does not provide information about genetic factors and other evolutionary influences would be ill-suited as an *evolutionary* explanation.

Generality may be of explanatory worth, but explanations can be too general, or general

⁴This is a version of the disjunction problem, which traces back to discussions of Hempel's theory of explanation; see Cover and Curd (1998, p. 786).

in the wrong way. An exhaustive disjunction of all possible causes of the room getting dark is too general to be a good explanation. Citing only information about the development of humans' large brains is general in the wrong way to be a good evolutionary explanation. The degree and type of generality that benefits an explanation must be specified. In the discussion that follows, I assume a broadly causal view of explanation. An alternate formulation of my point is consistent with a unification approach to explanation, but providing both formulations here would be overly cumbersome. On the assumption that explanations cite causal information, generality gained by leaving out unimportant causal information is good, whereas generality gained by leaving out important causal information is bad. The key question is thus: what causal information is important to an explanation?

What information about a causal process is important to explaining the resulting event depends upon which of the many causal relationships involved are explanatory. This in turn depends upon the role the explanation is intended to play. Each causal factor leading to an event is potentially relevant to explaining that event, or alternately, each difference-maker is of potential relevance (Woodward, 2003; Strevens, 2009). Explanatory relevance follows from causal or difference-making relevance. Yet every event results from innumerable causal factors; causal processes extend indefinitely back in time and often involve several interacting influences at any given point in time. The potential explanatory relevance of many causes, coupled with the explanatory value of generality, makes it so that different causal relationships are explanatorily important in different circumstances. Which causal relationship is explanatory, and consequently what causal information is explanatorily important, depends on the context of the call for explanation, that is, on the research program to which the explanation contributes.⁵

Consider the different contexts in which an explanation may be sought for why humans

⁵For the sake of brevity, I do not develop an account of this context-dependence here. See Potochnik (forthcoming) for discussion of what governs this pragmatic feature of explanation.

have large brains relative to other hominids. In a research program in evolutionary ecology, the causal relationship between environmental influences and the resultant selective advantage for large-brained humans is explanatorily important. In a developmental biology research program focused on the human nervous system, the causal relationship of note is instead how certain developmental pathways in humans lead to the fabrication of big brains. Both of these causal relationships—and many, many more besides—influence the event to be explained. Which causal relationship is explanatorily important depends on the role that the explanation is intended to play in scientific investigation.

Those who want to resist this move to contextual explanation might argue that different contexts lead to different explananda. For the above example, one might argue that in the first context one explains the evolution of humans' large brains, whereas in the second context one explains the development of humans' large brains. But individuating explananda in this manner is simply a way of covertly specifying which causal relationship is of interest. The explanandum is that humans have large brains, or if you like, that humans evolved to develop large brains. Depending on research interests—that is, on context—this explanandum might be expressed alternately as (a) humans evolved (to develop) large brains, or (b) humans (evolved to) develop large brains. Both regard a single evolutionary/developmental event, resulting from a single complex causal process. The alternate expressions of the explanandum result from singling out a distal part of this causal process and from singling out a proximate part of this causal process, respectively (Mayr, 1961).⁶ This is simply the context of the explanation directing attention toward one part of the causal process, which leads to bracketing other parts of the process.

Another way to resist the move to contextual explanation is to claim that the role

⁶Developmental systems theorists maintain that the evolutionary/developmental process often cannot be prized apart in this way, that is, that it is not possible to adequately answer evolutionary questions while bracketing development. If this is right, then in contexts favoring evolutionary information, the explanandum would be the evolution of the developmental process leading to large brains.

of context is captured by specifying the contrast class of the explanandum. Contrastive explanation is the idea that an explanation is in part determined by what one considers to be the alternatives to the event to be explained (Dretske, 1972; van Fraassen, 1980; Garfinkel, 1981). For instance, one might wonder why humans have large brains rather than small brains, or why humans have large brains rather than chimps having large brains. Although specifying different contrast classes may lead to different explanations, this does not capture the full relevance of context. Consider explanations for why humans have large brains rather than smaller brains. Research interests in development will still result in an explanation that focuses on developmental pathways: why it is that a given genotype leads to large brains rather than small brains, whereas research interests in evolution result in an explanation that focuses on natural selection: why it is that large brains were selected for rather than small brains. Specifying the contrast class does not fully account for the relevance of context. On this, I agree with van Fraassen (1980).⁷

The idea that the explanatory importance of causal information is context-dependent undermines the claim that higher-level explanations always focus on important information and only neglect unimportant information. Which causal relationships are explanatorily important and which should be neglected for the sake of generality depends on the research program in which the explanation is developed. Even if a higher-level explanation is more general than its lower-level competitor, there are likely to be contexts in which the causal relationships that are focal in the lower-level explanation are explanatorily important. When this is so, the generality of the higher-level explanation results in the omission of information that is important to the explanation, thereby diminishing the value of the explanation. The force of this point only increases when one takes into account that higher- and lower-level

⁷Much of what I say here about context-dependence fits with van Fraassen's views on explanation. According to van Fraassen, why-questions do not even arise without specific interests—“legal, medical, economic; or just an interest in optics or thermodynamics rather than chemistry”—and without abstracting, “[thinking] of Caesar's death *qua* multiple stabbing, or *qua* assassination” (van Fraassen, 1980, p. 130).

explanations often are not related by supervenience. Oftentimes, lower-level explanations and higher-level explanations simply differ in which causal relationships they represent and which they neglect, and so are general in different ways. Whether the higher-level explanation or the lower-level explanation offers the right type of generality—i.e. captures the important causal information and neglects the unimportant information—depends on the circumstances.

4 New Reasons for Explanatory Pluralism

I have argued that, despite common assumptions to the contrary, explanations formulated at different levels are often general in different ways, and those different forms of generality are valuable in different contexts of explanation. These points undermine the traditional argument against explanatory reductionism and suggest an alternate conception of the relationship among levels of explanation. The result is a new form of pluralism regarding levels of explanation, one that is sensitive to the complexity of real-world causal processes and to the ways in which science grapples with this complexity.

The antireductionist argument considered here is right about at least one thing: generality is always valuable to explanation. The generality of an explanation is what enables it to apply to a range of systems, thereby grouping together systems that bring about similar outcomes via similar means. An explanation highlights the distinctive features of the range of systems to which it applies and demonstrates the causal regularities that are responsible for the outcome to be explained.⁸ The way in which an explanation is general—what information it includes and what it neglects—determines the range of systems to which the explanation applies. Relatively general explanations apply to a broad range of systems and, consequently, pick out a broad causal pattern. More specific explanations pick out a more limited causal

⁸This suggests how the present ideas are compatible with both causal and unification approaches to explanation, though the idea that citing causal regularities is explanatory is probably most closely associated with the difference-making approach to explanation (Woodward, 2003; Strevens, 2009).

pattern, but the explanation still applies to a range of systems that behave similarly to one another. The explanation is still general in a way that makes perspicuous a causal regularity, and that indicates the range of circumstances in which that regularity holds. Explanations that offer different types of generality, that apply to different ranges of systems, offer information about different causal regularities that are involved in bringing about the event in question.

The type of generality that is appropriate depends on the explanation's context. Of course, which causal relationships have the potential to be explanatory in the first place is constrained by the nature of the causal process that leads to the event to be explained. A causal relationship cannot be explanatory if it is not involved in bringing about the event in question. But which among the potentially explanatory causal relationships belongs in a particular explanation is decided by the context in which the explanation is formulated. This context is determined by the role to be played by the explanation, that is, the focus of the research program in which the explanation is formulated. Because explanations formulated on different levels are often not related via supervenience, they are often general in different ways. That is, higher- and lower-level explanations generalize to different ranges of systems, thereby highlighting different causal regularities. Whether the higher-level or the lower-level explanation is better depends upon which of these causal regularities is relevant to the guiding research interests.

Consider once more the distinct genetic and phenotypic explanations for evolved traits. A genetic explanation provides insight into how the dynamics of genetic transmission increased the frequency of the trait in question. A phenotypic explanation provides insight into how the environment selectively advantaged the trait in question and caused it to predominate. These genetic and phenotypic explanations compete in the sense that they are distinct explanations offered for a single event. Moreover, they cite properties that are reasonably described as lower- and higher-level, respectively. Yet, as demonstrated in Section 2, they

apply to different ranges of systems; in other words, they offer different types of generality. By doing so, they highlight different causal regularities governing the processes by which traits evolve. Each type of explanation captures one aspect of why, e.g., humans have large brains. Which explanation is better depends on the research interests that guide the investigation. Accordingly, genetic explanations are typically developed in population genetics, while phenotypic explanations are typically developed in evolutionary ecology.

Jackson and Pettit (1992) and Sober (1999) develop versions of the classic antireductionist response that I criticize, but they are also pluralists about levels of explanation. They grant that higher-level explanations are more general than lower-level explanations, for they accept that higher- and lower-level explanations cite properties related by supervenience and that multiple realization is common. They are pluralists because they think that explanations sometimes should be general (higher-level) and other times specific (lower-level). This misidentifies the value of multiple levels of explanation. Maximal generality of the right type always improves explanations. Pluralism about explanation results from variance only in what *type* of generality is called for, i.e., what range of systems the explanation should apply to. This is determined by contextual factors—in particular, by what causal relationship is focal given the research program at hand. I have shown that higher-level explanations need not be more general than lower-level explanations, and that the preferable level of explanation depends on the type of generality that best captures the causal relationship of explanatory importance in the context at hand.

This conception of the explanatory role of generality and its context-dependence grounds a rich pluralism of approaches to explanation that reflects the varied projects of science. Which of the many causal relationships involved in bringing about an event belongs in an explanation of that event depends on the research interests that occasion the call for explanation, and determines the type of generality the explanation should have. The rarity of supervenience relationships among competing explanations undermines explanatory

reductionism, as well as the common antireductionist rejoinder. Indeed, the framework of “levels” of explanation emerges as somewhat forced. It is more natural to distinguish among alternate approaches to explanation that simply vary as to which entities and properties are cited, and thus which causal relationships are focal.

Acknowledgments

This paper has benefitted from comments by Michael Friedman, Peter Godfrey-Smith, Elliott Sober, Michael Strevens, and Michael Weisberg, as well as two anonymous referees for *Philosophy of Science*.

References

- Bull, J., Pfennig, D. W. and Wang, I. N. [2004]: ‘Genetic details, optimization and phage life histories’, *Trends in Ecology and Evolution*, **19**(2), pp. 76–82.
- Cover, J. A. and Curd, M. (eds.) [1998]: *Philosophy of Science: The Central Issues*, W.W. Norton and Company.
- Curnow, R. and Ayres, K. [2007]: ‘Population genetic models can be used to study the evolution of the interacting behaviors of parents and their progeny’, *Theoretical Population Biology*, **72**, pp. 67–76.
- Dretske, F. [1972]: ‘Contrastive statements’, *The Philosophical Review*, **81**(4), pp. 411–437.
- Dupré, J. [1993]: *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*, Cambridge, MA: Harvard University Press.
- Fodor, J. [1974]: ‘Special sciences: The disunity of science as a working hypothesis’, *Synthese*, **28**, pp. 97–115.
- Friedman, M. [1974]: ‘Explanation and scientific understanding’, *The Journal of Philosophy*, **71**(1), pp. 5–19.

- Garfinkel, A. [1981]: *Forms of Explanation: Rethinking the Questions in Social Theory*, New Haven, CT: Yale University Press.
- Jackson, F. and Pettit, P. [1992]: ‘In defense of explanatory ecumenism’, *Economics and Philosophy*, **8**, pp. 1–21.
- Kitcher, P. [1981]: ‘Explanatory unification’, *Philosophy of Science*, **48**(4), pp. 507–531.
- Kitcher, P. [1984]: ‘1953 and all that: A tale of two sciences’, *Philosophical Review*, **93**, pp. 335–373.
- Mayr, E. [1961]: ‘Cause and effect in biology’, *Science*, **134**, pp. 1501–1506, reprinted in E. Mayr, *Towards a New Philosophy of Biology*. Harvard University Press, 1988.
- Oppenheim, P. and Putnam, H. [1958]: ‘Unity of science as a working hypothesis’, in H. Feigl, M. Scriven and G. Maxwell (eds.), *Minnesota Studies in the Philosophy of Science*, volume 2, Minneapolis: University of Minnesota Press, pp. 3–36.
- Potochnik, A. [2007]: ‘Optimality modeling and explanatory generality’, *Philosophy of Science*, **74**(5), pp. 680–691.
- Potochnik, A. [forthcoming]: ‘Explanatory independence and epistemic interdependence: A case study of the optimality approach’, *British Journal of the Philosophy of Science*.
- Putnam, H. [1975]: *Philosophy and our Mental Life*, volume 2 of *Philosophical Papers*, chapter 14, Cambridge, UK: Cambridge University Press, pp. 291–303.
- Rosenberg, A. [1994]: *Instrumental Biology or the Disunity of Science*, Chicago, IL: University of Chicago Press.
- Sober, E. [1984]: *The Nature of Selection*, Cambridge, MA: MIT Press.
- Sober, E. [1999]: ‘The multiple realizability argument against reduction’, *Philosophy of Science*, **66**, pp. 542–564.
- Strevens, M. [2009]: *Depth: An Account of Scientific Explanation*, Cambridge, MA: Harvard University Press.
- van Fraassen, B. C. [1980]: *The Scientific Image*, Oxford: Clarendon Press.
- Woodward, J. [2003]: *Making Things Happen: A Theory of Causal Explanation*, Oxford, UK: Oxford University Press.