Hypothetical Pattern Idealization and Explanatory Models

Yasha Rohwer and Collin Rice*†

Highly idealized models, such as the Hawk-Dove game, are pervasive in biological theorizing. We argue that the process and motivation that leads to the introduction of various idealizations into these models is not adequately captured by Michael Weisberg's taxonomy of three kinds of idealization. Consequently, a fourth kind of idealization is required, which we call *hypothetical pattern idealization*. This kind of idealization is used to construct models that aim to be explanatory but do not aim to be explanations.

1. Introduction. Drawing on previous work by Leszek Nowak (1972), Nancy Cartwright (1983, 1989), Ernan McMullin (1985), William Wimsatt (2007), and Michael Strevens (2009), Michael Weisberg (2007) has recently characterized three kinds of idealization: multiple models, Galilean, and minimalist. In this article, we argue that the process and motivation that leads to the introduction of various idealizations into many biological models is not sufficiently captured by Weisberg's categories. Although we will focus on Weisberg's account because of its clarity and influence, the previous accounts of idealization he discusses also appear to overlook the dis-

Received November 2012; revised April 2013.

*To contact the authors, please write to: Yasha Rohwer, Department of Humanities and Social Sciences, Oregon Institute of Technology, 3201 Campus Dr., Klamath Falls, OR 97601; e-mail: ynrg24@mail.missouri.edu. Collin Rice, Department of Philosophy, Lycoming College, 700 College Place, Williamsport, PA 17701; e-mail: crdv3@mail.missouri.edu.

†Both authors contributed equally to this work. We would like to thank Chris Pincock, André Ariew, Paul Weirich, Michael Weisberg, William Wimsatt, and several anonymous reviewers for comments and feedback on previous versions of this article. We would also like to thank the members of the University of Pittsburgh's Center for the Philosophy of Science for helpful discussion, comments, and support. This research was partially supported by the Al and Mary Agnes McQuinn Chair in Entrepreneurial Leadership.

Philosophy of Science, 80 (July 2013) pp. 334–355. 0031-8248/2013/8003-0008$10.00
Copyright 2013 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/671399 Published online by Cambridge University Press

tinctive features of the cases we present.[1] Therefore, we argue that there is at least a fourth kind of idealization that we call *hypothetical pattern idealization*. This fourth kind of idealization is needed to adequately characterize these highly idealized models that are especially important to evolutionary biology.[2] These models do not aim for veridical representation of any target system(s) but instead investigate a highly idealized hypothetical situation to justify background beliefs, resolve seeming inconsistencies between theories and our observations, or explore how-possibly questions.[3] In addition, our analysis of this fourth kind of idealization raises important questions about the status of these highly idealized models as scientific explanations. We go on to clarify the representational goals of hypothetical pattern idealization by drawing a distinction between a model being a stand-alone explanation versus merely being explanatory. An explanatory model is one that produces scientific understanding relevant to answering a why question, but the model need not provide an accurate enough representation to provide an explanation (or even a partial explanation). This distinction helps explicate the nature of and motivation behind hypothetical pattern idealization since models with these idealizations aim to be explanatory but do not aim to be explanations. We also investigate the scope of hypothetical pattern idealization and argue that it is likely central to a kind of modeling that is widely used in biology and economics. As a result, incorporating this fourth kind of idealization is essential to providing an adequate account of the use of idealization in science.

2. Weisberg's Three Kinds of Idealization. In "Three Kinds of Idealization" (2007), Michael Weisberg presents a unified framework for thinking about idealization.[4] Weisberg begins by noting that idealization is an activity wherein theories or models are distorted. He then identifies three kinds of idealization: multiple models, Galilean, and minimalist. Although there is some overlap among the categories, Weisberg distinguishes these kinds of idealization via their representational ideals (i.e., their ultimate representational goals) and the justifications used to motivate them.

1. While some of these authors—e.g., Cartwright (1983) and Batterman (2002)—hint at some of the features we identify, previous accounts have not treated these unique cases as a distinctive kind of idealization, nor have they given a detailed analysis of its justification and representational goals.

2. We will discuss the potential scope of this kind of idealization in more detail in sec. 5.

3. Of course, game-theoretic models are used in various disciplines and in various ways. However, we will focus on one way of using a game-theoretic model that is important within biology.

4. It is important to note that Weisberg never claims that his classification is exhaustive. Therefore, our arguments are not so much a criticism of Weisberg as they are a call to expand on his original analysis.

https://doi.org/10.1086/671399 Published online by Cambridge University Press

Multiple-models idealization is "the practice of building multiple related but incompatible models, each of which makes distinct claims about the nature and causal structure giving rise to a phenomenon" (Weisberg 2007, 645). This kind of idealization is distinguished by "not expecting a single best model to be generated" (646). This building of multiple models with incompatible assumptions is justified because modelers often have multiple representational goals when modeling natural systems. Multiple models are needed because no single model can provide every representational goal at its highest possible resolution. The reason for this is that many representational goals trade off (Levins 1966). For example, increasing precision decreases generality. As a result, modelers will sometimes require multiple incompatible models to maximize the representational goals they are interested in achieving. Weisberg's description of multiple-models idealization draws heavily on Richard Levins's (1966) work on the trade-offs faced by biological modelers.

An example Weisberg mentions is the National Weather Service's use of multiple incompatible models to predict the weather. Each of the models used makes different idealizing assumptions about the target system (i.e., each model is inaccurate). The National Weather Service's goal is to make the most accurate predictions, and they have determined that using three inconsistent and idealized models is the best way to maximize predictive accuracy. In this case, it appears that in order to achieve maximal predictive accuracy the modeler must sacrifice representational accuracy and use multiple incompatible models.

Galilean idealization is "the practice of introducing distortions into theories with the goal of simplifying theories in order to make them computationally tractable" (Weisberg 2007, 640). The justification for this activity is pragmatic; the activity is performed because of present computational limitations. The ultimate representational goal of Galilean idealization is a completely accurate representation of all the properties of the target phenomenon (655). Galilean idealization is important for the study of complex systems, but ultimately, with advances in computational power and mathematical techniques, Galilean idealizations should be systematically removed (641).[5] Weisberg's account of Galilean idealization comes from McMullin, who described a process in which we aim to "grasp the real world from which the idealization takes its origin" by making the problem more tractable. However, ultimately, "models can be made more specific by eliminating simplifying assumptions and 'deidealization' as it were" (McMullin 1985, 248, 261). While McMullin's account does not appear to require that the model actually be deidealized, it does suggest that it should be possible to deidealize the

5. While the representational ideal of Galilean idealization is to provide a completely accurate representation, in practice perhaps no model can ever be completely deidealized.

https://doi.org/10.1086/671399 Published online by Cambridge University Press

model in some principled (i.e., theory driven) way, in order to show that the idealized model is approximately true (264). Furthermore, although actual deidealization may not be required by all accounts of Galilean idealization, once scientific progress overcomes the tractability limitations that motivated the introduction of the Galilean idealization, the justification for the idealization vanishes. Hence, Weisberg's description of Galilean idealization as the "introduction of distortion to make the problem more tractable, then systematic removal of the distorting factors" plausibly includes deidealization as a distinguishing feature of this kind of idealization (2007, 641). Indeed, Weisberg is not alone in characterizing Galilean idealization in this way (Suárez 1999; Wayne 2011).

One of Weisberg's examples of Galilean idealization comes from Galileo himself, who was investigating gravitational acceleration in a medium devoid of resistance. However, since he lacked such a medium, Galileo instead investigated a system with the least resistance available and assumed that it functioned in (approximately) the same way (McMullin 1985, 267). Yet, as science has progressed, we can now create a vacuum in space—thereby doing away with the idealization.[6]

Minimalist idealization is "the practice of constructing and studying theoretical models that include only the core causal factors which gave rise to the phenomenon," where the core causal factors are "those factors that *make a difference* to the occurrence and essential character of the phenomenon in question" (Weisberg 2007, 642).[7] A model is minimalist whenever it "*accurately* captures the core causal factors" (643; emphasis added). The justification for this activity is that it aids in scientific explanation. According to Weisberg, "The key to explanation is a special set of explanatorily privileged causal factors. Minimalist idealization is what isolates these causes and thus plays a crucial role for explanation" (645). The ideal representation for this kind of idealization is one that accurately represents only causes that make a difference. Since the practice is key to scientific explanation, we should not expect these idealizations to be removed as science progresses. Among others, Weisberg cites Michael Strevens's (2009) account of idealization, Nancy

6. Of course, we cannot create a perfect vacuum, but we can investigate the kind of baseline case that Galileo hoped to approximate. As another example, in his investigation of pendulums Galileo assumed that the pendulum is not subject to air resistance, that the wire is massless and inelastic, and that there are no other influences or imperfections. However, Wayne (2011) has argued that this example is not a case of Galilean idealization.

7. In this article, we follow Weisberg (2007) and Strevens (2009) in calling the explanatorily relevant causal factors "difference-makers." Of course, according to other accounts, all causes make a difference. The key point, for our purposes, is that minimalist idealizations aim for accurate representation of causes in the target system in order to provide an explanation or a partial explanation.

https://doi.org/10.1086/671399 Published online by Cambridge University Press

Cartwright's (1989) description of abstraction, and Robert Batterman's (2002) account of asymptotic explanation as examples of minimalist idealization. However, since we have some doubts about whether Batterman's and Cartwright's accounts are best characterized as instances of minimalist idealization (see n. 8), we will focus on Strevens's account.[8]

Strevens explains his account of idealized models in this way: "The content of an idealized model, then, can be divided into two parts. The first part contains the difference-makers for the explanatory target. . . . The second part is all idealization; its overt claims are false but its role is to point to parts of the actual world that do not make a difference to the explanatory target. The overlap between an idealized model and reality . . . is a standalone set of difference-makers for the target" (2009, 318). In other words, the goal of a minimalist model is to accurately represent difference-makers and use idealizations to indicate those causal factors that are irrelevant.

An example of minimalist idealization cited by Weisberg and Strevens is Boyle's gas law. One idealizing assumption in the model is that gas molecules do not collide with each other. Although this assumption is false, in low-pressure gases these collisions make no difference to the behavior of the system since taking them into account would not change the model's predictions. Therefore, these collisions can be safely ignored when the goal is to represent only causes that make a difference.

3. Idealized Models in Evolutionary Biology. Weisberg's framework is a much-needed tool to synthesize previous work on idealizations and to characterize the distinct kinds of idealization in science. Unfortunately, we think it fails to adequately capture a kind of idealization that is widely used in evolutionary biology. In this section, we will present the Hawk-Dove game and argue that the justification and representational ideals for the idealizations in this instance of biological modeling are not adequately captured by Weisberg's taxonomy. While we present only one biological model, our anal-

8. One important difference is that Batterman explicitly denies that idealizations are used to isolate causal factors. Rather, for Batterman idealizations play an essential role in showing why various heterogeneous (causal) details of particular systems are irrelevant to the overall behavior of those systems (Batterman 2002, 73). Moreover, Batterman focuses on various other roles that idealizations play—e.g., in understanding intertheoretic relationships. Finally, contrary to Weisberg, Batterman uses the term 'minimal model' to refer to a model within the same universality class as the actual system. In addition, while Cartwright does discuss the isolation of causal powers, she also discusses the role of idealizations in making a model mathematically tractable and describes idealized models as nonveridical fictions (Cartwright 1983, 1989). Given the diverse and complicated threads in these authors' work, contra Weisberg's presentation, it is unclear precisely which category (or set of categories)—if any—best characterize(s) these authors' views. However, we do think the distinctions drawn by Weisberg's taxonomy are useful for sorting out some importantly different aspects of idealization mentioned by these previous accounts.

https://doi.org/10.1086/671399 Published online by Cambridge University Press

ysis applies equally well to the use of several models that employ similar idealizations (e.g., the Divide the Cake game or the Prisoner's Dilemma). Indeed, this kind of idealization is extremely important to a particular use of game-theoretic models that is central to various kinds of theorizing in evolutionary biology.

3.1. The Hawk-Dove Game. In the natural world, when conspecifics compete for a resource they often exercise restraint in combat instead of fighting to the death. This observed pattern was puzzling for adaptationists since individual-level selection would presumably favor fierce physical combat, given that the winner would gain obvious benefits (e.g., mates, desirable territory) that would translate into transmitting its genes to future generations at higher frequencies than the loser (Maynard Smith and Price 1973). However, the pattern we observe across numerous species is that animals often exercise restraint in combat—what Maynard Smith and Price refer to as "limited war" strategies. At the time the Hawk-Dove model was introduced, this observation was commonly explained by appealing to "group selection" (15). Group selection was believed to be required to explain the observed pattern because fierce physical combat would be detrimental to the survival of the species. The Hawk-Dove game is intended to demonstrate how individual selection acting alone is consistent with the observation of this behavior in a wide range of populations (Maynard Smith and Price 1973; Maynard Smith 1982).[9]

In the basic Hawk-Dove game there are two strategies: be a Hawk or be a Dove. Hawks escalate until injured or until the opponent retreats. Doves display and then retreat if their opponent escalates.[10] There are three kinds of interactions: (1) *Hawk vs. Hawk*, where each player has a 50% chance of obtaining the resource and a 50% chance of being injured; (2) *Hawk vs. Dove*, where the Hawk obtains the resource and the Dove retreats; and (3) *Dove vs. Dove*, where the resource is shared equally. These interactions lead to the payoff matrix in table 1, where H is Hawk, D is Dove, V is the fitness obtained by winning the resource, C is the cost of being injured, and $V > V/2 > 0 > \frac{1}{2}(V - C)$. The preceding inequality is crucial to the dynamics of the game. The value of the resource must be positive, and the cost of injury must be greater than the benefit of the resource.

9. In support of our interpretation, in another paper Maynard Smith and Parker explicitly claim they are not interested in building a model that explains any particular behavior but only in investigating a necessity claim (Maynard Smith and Parker 1976, 159).

10. There are, of course, more complicated formulations of the Hawk-Dove game. We use the simplest version to illustrate Maynard Smith and Price's use of the model and because it is sufficient to play the explanatory role we want to analyze.

https://doi.org/10.1086/671399 Published online by Cambridge University Press

TABLE 1. Hawk-Dove Game Payoff Matrix

	H	D
H	½(V − C), ½(V − C)	V, 0
D	0, V	V/2, V/2

Note.—H = Hawk; D = Dove; V = the fitness obtained by winning the resource; C = the cost of being injured.

The model assumes that this game is played over and over again by individuals in an infinitely large population that reproduces asexually. In a pure Dove population, any mutant Hawk would do very well since $V > V/2$. In a pure Hawk population, any mutant Dove would do very well since $0 > \frac{1}{2}(V - C)$. A particular mixture of Hawks and Doves, however, leads to stability in the population. Such a stable equilibrium occurs when the average payoffs for Hawks are equal to the average payoffs for Doves. This stable state of the population could occur in two ways. First, the population could consist of a mixture of individuals who played "pure" strategies. Alternatively, the population could consist of individuals who all adopt a "mixed" strategy of playing Hawk with probability x and Dove with probability $1 - x$. Either way, the model shows how individual selection could give rise to restraint in combat.

Several key assumptions underlie the Hawk-Dove model and its use, many of which are idealizations. These idealizations include (1) infinite population size, (2) random pairing of players, (3) asexual reproduction, (4) symmetric contests, (5) pair-wise contests, (6) constant payoff structure across individuals and across iterations of the game, and (7) perfect correlation between winning the resource and reproductive success (Maynard Smith 1982). While no actual population is infinite, this idealization is introduced because, for Maynard Smith and Price, the primary goal of the model is only to investigate an instance of individual selection acting alone in a highly idealized case to see if restraint could evolve. Assuming that players pair off randomly within the population is another idealization. In real-world systems, individuals will be more likely to compete against local individuals than others in the population, but the model idealizes these details away. Asexual reproduction is also an idealization when considering the many populations that reproduce sexually. In addition, the model assumes that the game is symmetric: individuals have the same available strategies, and the payoff matrix is the same for both individuals in the game. This, of course, idealizes away many features of real-world populations, such as differences between animals' abilities to fight, their ability to play alternative strategies, or cases in which animals occupy different roles (e.g., when one is the current owner of the resource). Next, the model makes the idealizing assumption that when

https://doi.org/10.1086/671399 Published online by Cambridge University Press

this game is iterated the payoffs do not change. Yet, the fitness payoffs to individuals will be different at different times, given that the individuals as well as the resources they are competing for may change. Finally, winning the resource is assumed to be perfectly correlated with the fitnesses of different strategies. Yet, most resources over which animals actually compete (e.g., food, mates, or territory) will not lead directly to increased fitness.

Despite all of these idealizations, the Hawk-Dove game still produces some understanding of how individual selection could possibly lead to restraint in situations of animal conflict. This is presumably because although the model fails to accurately represent the selection dynamics of any real-world population, it does tell us something about how individual selection could lead to the trait in a wide range of possible systems, by investigating a hypothetical scenario. Therefore, the model is explanatorily valuable, even though it uses several idealizations that make it an inaccurate representation of how individual selection actually occurred in any given real-world population.

3.2. Why These Idealizations Do Not Fit Weisberg's Taxonomy. The idealizations used in the Hawk-Dove game are common assumptions in many biological models. While we have distinguished seven idealizations in the above model, we can give a unified analysis that shows why the motivation behind and representational goals of this kind of idealized modeling are not adequately captured by Weisberg's categories. As a result, we will argue that a fourth kind of idealization is needed to adequately characterize these instances in evolutionary biology.[11]

To begin, these idealizations are not multiple-models idealizations since the goals of the modeler will not be better achieved through the use of multiple conflicting models; rather, a particular kind of highly idealized model is sufficient. This is because the goal of the modeler is to address a how-possibly question concerning a pattern that is highly general: Could individual-level selection lead to restraint in a wide variety of populations? No additional models with incompatible idealizations are required to achieve the goals of the modeler. Indeed, as Maynard Smith and Price explain in their conclusion, "The analysis is . . . sufficient to show that individual selection *can* explain why potentially dangerous offensive weapons are rarely used in intraspecific contests" (Maynard Smith and Price 1973, 17; italics added). Although it may be possible to achieve these goals with multiple conflicting models, this would be an unnecessarily complicated process for achieving the modeler's goals. That multiple incompatible models are not required is precisely why these are not instances of multiple-models idealizations. In-

11. However, as is the case with Weisberg's own categories, there will be some commonalities between this fourth kind of idealization and those distinguished by Weisberg.

https://doi.org/10.1086/671399 Published online by Cambridge University Press

stead, the goals of the modeler can be better achieved by the use of a single highly idealized model.

All seven idealizations may seem to be instances of Galilean idealization since each makes the model more computationally tractable. However, unlike Galilean idealizations, these idealizations should not be eliminated from the model (nor will they become unjustified) with computational progress. This is because the motivation behind them is not merely computational tractability. The goal of the model builder is to investigate the consistency of animal conflict with individual selection in a wide range of systems that are heterogeneous in the particular details eliminated by these idealizations. For example, assuming that obtaining the resource correlates perfectly with fitness allows the model to investigate this consistency claim across a wide range of possible resources. In other words, rather than focusing solely on computational tractability, these idealizations play an important role in providing the computational flexibility to investigate a range of possible systems in which restraint in combat might arise.[12] In this way, the idealizations play an important role in allowing the model to capture the generality of the pattern of interest. Eliminating these idealizations would, essentially, eliminate the model's ability to achieve the goals of the modeler: showing how individual selection could lead to the phenomena in a wide variety of possible situations. Consequently, the observed biological pattern is no longer puzzling within an individual selectionist framework. Without these idealizations, the model would be unable to address the highly general observation that is so puzzling for the adaptationist. In addition, whereas Galilean idealizations ultimately aim for a completely accurate description of some target system(s) through a process of deidealization, the Hawk-Dove model does not aim to accurately represent the features of any target system(s). Indeed, no process of deidealization would improve (or is required for) the model's ability to accomplish the goals of the model builder. Therefore, these idealizations are not Galilean.

Given these considerations, it might seem that these idealizations are minimalist since they appear to allow us to ignore various causal factors in real-world populations. Indeed, Weisberg's minimalist idealization involves including "only those [causal] factors that *make a difference* to the occurrence and essential character of the phenomenon" (2007, 462). Furthermore, a possible interpretation of the goal of these game-theoretic models in biology is to accurately represent the causal contribution of natural selection while ignoring other evolutionary factors. However, the modeler who uses the Hawk-Dove game does not aim to accurately represent the core causal factors of any real-world system(s) but instead investigates a hypothetical case in order to

12. Thanks to an anonymous reviewer for suggesting this as a way to clarify how these idealizations can contribute to capturing the generality of the pattern of interest.

https://doi.org/10.1086/671399 Published online by Cambridge University Press

demonstrate how individual selection could have produced this pattern across a range of causally heterogeneous systems. As Maynard Smith and Price repeatedly emphasize, "real animal conflicts are vastly more complex than our simulated conflicts" (1973, 17). There is no target system(s) whose difference-makers the model aims to accurately represent, nor does the model aim to establish any claims about what actually caused the phenomenon to occur.

In response, one suggestion might be that the model aims to accurately represent a particular kind of causal process (Strevens 2009). That is, the Hawk-Dove game might be thought to be using minimalist idealizations since it accurately represents the causal contribution of natural selection at the level of types, which might be the only difference-maker at this level of abstraction. However, given the extreme heterogeneity and complexity of the populations in which restraint in combat has evolved, it is unlikely that there is a set of core causal factors that is common to each of them, even if we move to more abstract levels of description. More importantly, however, in order for the model to cite type-level causes to try and explain a coarse-grained event, it must accurately represent causal types within real-world systems such that what is true of the types will be true of the tokens within particular populations. However, the Hawk-Dove model is so idealized that it fails to accurately represent the individual-level selection processes of any real-world population, even at the level of types—that is, the type of selection process described within the model is never instantiated by any real-world system. Indeed, as Sugden (2009) notes, "The Hawk-Dove model does make use of some accepted principles of biology. . . . But the workings of those principles are explored in a counterfactual world created by [Maynard Smith and Parker]. Many of the features that have been built into that world—for example, asexual reproduction and the entirely genetic determination of behavior—seem to be modeling conventions rather than accepted principles. . . . This makes it hard to make sense of the idea that the model isolates an other-things-being-equal tendency that is at work in real-world cases" (21).

Minimalist models aim to give a veridical representation of core causal factors within real-world systems. The Hawk-Dove model, in contrast, does not aim to provide an accurate representation of any actual causal factors that could be used to provide an explanation (or a partial explanation). Instead, the central aim of Maynard Smith and Price's original use of the Hawk-Dove model is to show that individual selection is compatible with the observed biological pattern. Indeed, Maynard Smith and Price state this aim explicitly when they claim: "A main reason for using computer simulation was to test whether it is *possible even in theory* for individual selection to account for 'limited war' behaviour" (1973, 15; italics added). Yet, the model is able to provide this insight without accurately representing any causes within real-world systems—and this remains true even if we

https://doi.org/10.1086/671399 Published online by Cambridge University Press

consider causal processes at the level of types. Therefore, the motivation behind and representational goals of the Hawk-Dove game are importantly different from those of minimalist idealization.

Given that the idealizations of the Hawk-Dove game do not nicely fit into any of Weisberg's categories, a fourth kind of idealization is needed—a kind of idealization that is extremely important to biological modeling. We call this kind of idealization *hypothetical pattern idealization*.[13] These idealizations are used to construct models of hypothetical scenarios that need not be instantiated by any real-world system—indeed sometimes they will present impossible scenarios. In addition, we call this kind of idealization hypothetical *pattern* idealization because it is most likely to be fruitful when the phenomenon of interest is a general pattern that ranges over extremely heterogeneous and complex systems (although this kind of idealization may be used in other modeling contexts as well). Consequently, the motivation behind hypothetical pattern idealization is to construct models of hypothetical scenarios that, even though they may not accurately describe any core causal factors of a real-world system(s), are able to aid in the investigation of general patterns across extremely heterogeneous and complex systems. This situation is common in modeling biological systems since a model is often intended to investigate patterns that range across systems composed of individuals as diverse as humans and viruses (e.g., models of biological altruism). Models that use hypothetical pattern idealizations aim to show how such a pattern could arise in a (or perhaps a wide range of) hypothetical scenario(s) but do not aim to represent how it did arise in any actual case.

4. Idealization and Explanation. The nature of and motivation for hypothetical pattern idealization can be clarified by looking at the relationship between idealization and explanation. Since Hempel (1965), it has been widely accepted that a satisfactory explanation must be true.[14] When considering models, this truth requirement amounts to the requirement that the model provide a veridical representation of the explanatorily relevant features of the target system(s). A version of this requirement is present in contemporary causal theories of explanation (Woodward 2003; Strevens 2009). For these theories, a satisfactory explanation must accurately represent the relevant causes of the event to be explained.

Many philosophers claim that highly idealized biological models are explanations (e.g., Potochnik 2007; Rice 2012). In fact, when discussing the Hawk-Dove model, Elliott Sober claims that "the model is a plausible explanation of why the population exhibits the configuration it does" (2000,

13. Thanks to Christopher Pincock for suggesting this name.

14. One notable exception is Cartwright, who has consistently argued that the prominence of idealization in science shows that some of our best explanations are false.

https://doi.org/10.1086/671399 Published online by Cambridge University Press

142). However, models such as the Hawk-Dove game are so highly idealized that they drastically misrepresent all real-world systems. Indeed, the model describes a selection process that does not (and could not) occur in any real-world system. Claiming that all such highly idealized models are explanations appears to conflict with the requirement that model explanations provide an accurate representation of the explanatorily relevant features of their target system(s). One way to resolve this tension is to identify a unique explanatory role for many highly idealized models—one that avoids equating all such models with explanations. This explanatory role helps explicate the motivation behind the use of hypothetical pattern idealizations.

4.1. The Difference between X as an Explanation and X as Explanatory. Explanations are usually thought of in terms of two components: the explanandum (a proposition that represents the phenomenon to be explained) and the explanans (a set of propositions that does the explaining). When philosophers analyze the nature of explanation, they attempt to understand what properties and relationships these components must have.

One widely accepted feature of scientific explanations is that they produce scientific understanding in agents who grasp them (Friedman 1974; Kitcher 1981; Achinstein 1983; Lewis 1986; Salmon 1998; Grimm 2006; Lipton 2009; Strevens 2009). In fact, Michael Friedman once argued that our theory of explanation "should somehow connect explanation and understanding—it should tell us what kind of understanding scientific explanations provide and how they provide it" (1974, 14). Understanding is a cognitive achievement (Lipton 2009). A representation, such as the propositions of an explanation, is what allows an agent to attain this achievement. This relationship is not between propositions; rather, it is a relationship between something representational and an agent who gains understanding by grasping it.

However, being an explanation is not the only way to produce understanding. For instance, many philosophers claim that if something produces understanding, then it is explanatory (e.g., Sober 1983; Lipton 2009). In addition, rather than requiring an explanatory representation to actually produce understanding for an agent, on our account being explanatory is an objective property of a representation that is merely capable of producing understanding.[15] Whether a representation actually produces understanding in any particular case will, of course, depend on various features of the context (e.g., the agents involved and their current beliefs). The idea here, however, is that a representation can be capable of producing understanding in various contexts, even if an agent fails to gain any understanding from the representation in a particular case.

15. The objective nature of being explanatory distinguishes our account from those involved in pragmatic theories of explanation (e.g., Van Fraassen 1980).

https://doi.org/10.1086/671399 Published online by Cambridge University Press

Friedman (1974) discusses this kind of understanding and its relationship to explanation in the following way: "If one concentrates only on the local aspects of explanation—the phenomenon being explained, the phenomenon doing the explaining, and the relation (deductive or otherwise) between them—one ends up trying to find some special epistemological status . . . for the phenomenon doing the explaining. . . . However, attention to the global aspects of explanation—the relation of the phenomena in question to the total set of accepted phenomena—allows one to dispense with any special epistemological status for the phenomenon doing the explaining" (18). Although we disagree with several parts of Friedman's account, we think his view captures some important features of scientific understanding. First, understanding need not be about a particular target system. Second, understanding is sometimes produced by incorporating new (presumably justified) beliefs into our entire body of scientific knowledge (see also Rohwer, forthcoming). Viewing understanding in this way allows us to dispense with the requirement that the representation that produces understanding must veridically represent a particular target system(s).

In the epistemology literature, it is commonly assumed that understanding is factive (or quasi factive) and requires some form of justification (Kvanvig 2003; Depaul and Grimm 2007; Mizrahi 2012).[16] If this is so, then an explanatory representation will be one that is capable of justifying true beliefs in an agent who understands. However, a representation will not count as explanatory merely by being capable of justifying some random true beliefs. In order to count as explanatory, the content of these true beliefs must be relevant to providing an explanation. At the very least, the true beliefs will need to be about the phenomenon of interest and be relevant to answering a why question that we are interested in. Furthermore, in the epistemology literature, understanding requires not only justified true beliefs but also the grasping of certain relations between those justified true beliefs. Indeed, in the case of explanatory representations, there also appears to be a minimum requirement that agents who understand must grasp the connection between their justified true beliefs and the particular why question under consideration.

So in order to be explanatory, a representation must (1) be capable of justifying true beliefs about the phenomenon of interest for the agent and (2) enable the agent to grasp how those beliefs are relevant to answering a why question he or she is interested in. These criteria, however, fall short of requiring that the explanatory representation be a sufficiently veridical representation of any (or any part of a) real-world system(s) that could provide an explanation or partial explanation. Therefore, there is an asymmetric re-

16. To be quasi factive requires that all the central propositions in one's understanding be true but allows for some of the peripheral propositions to be false (Mizrahi 2012).

https://doi.org/10.1086/671399 Published online by Cambridge University Press

lationship between that which is an explanation and that which is explanatory. If something is an explanation, then it is explanatory. However, something that is explanatory is not necessarily an explanation. A representation can be capable of justifying true beliefs that are relevant to a why question without accurately representing (the explanatorily relevant features of) any real-world system. A standard view about models claims that "if we want to use models to learn about the world, the model needs to map onto the real world" (Morgan 1999, 366). In contrast, we argue that many highly idealized models are capable of producing scientific understanding without having to meet any additional veridical representation requirements—indeed, these models need not have a "target system" at all.[17]

At this point, one might object that surely some connection to the real world is required for a model to provide understanding of real-world systems (e.g., see Pincock 2012). To see why some connection is required, consider that one could construct a model in which restraint in combat persists, but if that model were completely disconnected from the real world, then no one would view the model as alleviating the inconsistency Maynard Smith and Price initially sought to investigate.[18] However, although some minimal relationship of "relevance" or "connectedness" to the real world is certainly necessary to use models to support various explanatory claims, we maintain that no veridical representation of the features of a target system(s) is required for a model to produce understanding. In other words, an explanatory model will need to present a scenario that is relevant or related to the real world in some way, but this minimal connection does not require the veridical representation of any explanatorily relevant features (e.g., causal factors) of a target system(s). For instance, the Hawk-Dove game is similar to real-world systems in that it includes a population of organisms that reproduce and interact in ways that affect their fitness. The model is able to be explanatory by presenting a relevant hypothetical scenario, but it does not aim to accurately represent the explanatorily relevant (e.g., causal) features of any real-world system(s). The important point is that the truth or accuracy of a model with respect to a target system(s) is often irrelevant to its ability to play this explanatory role within scientific inquiry. In addition, precisely what the necessary connection is between the real world and an explanatory model will depend on the particular goals (or purposes) of the model builder. For example, in certain modeling contexts, the model may need to represent a credible parallel world (i.e., a situation that is a possi-

17. By "target system" we mean a physical system in the world that a model might aim to represent. While some models may have other kinds of targets—e.g., a theory or other models—a treatment of these is beyond the scope of this article.

18. Thanks to an anonymous reviewer for pointing out this sort of case and emphasizing the need for some connection between an explanatory model and the real world.

https://doi.org/10.1086/671399 Published online by Cambridge University Press

ble candidate for truth; Sugden 2000, 2009; Grüne-Yanoff 2009). Alternatively, some explanatory models may be fictions that are only relevantly similar to real-world systems (Cartwright 1983; Godfrey-Smith 2006, 2009; Suárez 2009). Another possibility is that an explanatory model might be related to the real world via some mathematical transformation, for example, being in the same universality class as the actual system (Batterman 2002). Finally, a model can be explanatory by providing a veridical representation that is an explanation or partial explanation (see below). What these examples illustrate is that the relevant connection between an explanatory model and the real world will depend on the modeling context (Bokulich 2012)—that is, no particular relation is likely to work for every case. So while various minimal connections are likely to be sufficient for a model to be explanatory, none of them is necessary across all modeling contexts. Instead, explanatory models are unified by their ability to provide understanding that is relevant to answering a why question of interest to the modeler—not by any requirements concerning their truth or accuracy with respect to a target system(s) or any particular connection to the real world. In the end, being relevantly related to the real world does not require veridically representing a set of features of some real-world target system(s).

This analysis allows us to identify a key distinguishing feature of hypothetical pattern idealization. Whereas models that use Galilean and minimalist idealizations ultimately aim at a partially veridical representation of their target system(s) that is an explanation (or a partial explanation), models that use hypothetical pattern idealization do not. Models that use hypothetical pattern idealizations aim at producing understanding via a single representation that may never be instantiated. In other words, the primary motivation behind the introduction of hypothetical pattern idealizations is to build models that are explanatory but that may fail to be explanations.[19]

4.2. Some Ways of Being Explanatory. The Hawk-Dove model, although filled with falsehoods, allows us to understand that individual-level selection is not incompatible with restraint in combat in a wide range of populations. If we take seriously that a necessary condition for a model to be an explanation is that it veridically represent the explanatorily relevant features of its target system(s), then this highly idealized model cannot be an explanation. After all, the dynamics described by the model do not—indeed could not—occur in any real-world system. However, the Hawk-Dove model

19. Of course, it could turn out that, despite the presence of hypothetical pattern idealizations, the model still accurately represents the features required to give an explanation. In such a case, the motivation behind and representational goals of the idealizations make them instances of hypothetical pattern idealization, but the model may still be able to provide an explanation.

https://doi.org/10.1086/671399 Published online by Cambridge University Press

is capable of justifying true beliefs about the pattern of interest for agents who grasp the workings of the model, for example, the belief that individual-level selection is not incompatible with restraint in combat in a wide range of populations. And this true belief is certainly relevant to answering the question, "why do animals exercise such restraint in so many natural populations?" since citing individual selection will likely be part of that explanation. Even if individual selection is not part of the explanation, the model is still able to answer certain key how-possibly questions (Resnik 1991; Forber 2010; Reydon 2012). Establishing this consistency claim is relevant to the why question of interest, regardless of what actually explains the phenomenon in any given real-world population. Indeed, this understanding is relevant to the why question of interest because it justifies important background assumptions that aid in formulating the explanation. Therefore, one way the Hawk-Dove model is explanatory is that it answers an important how-possibly question that is related to the why question of interest, even though it fails to be an explanation.

As a result, the Hawk-Dove game illustrates how models that use hypothetical pattern idealization can often be essential stepping-stones en route to building models that are explanations or partial explanations. Therefore, although distinguished by their motivation and representational goals, models that use hypothetical pattern idealization and models that use minimalist idealization are nonetheless sometimes related. Models that involve hypothetical pattern idealizations are sometimes used to justify background beliefs or answer how-possibly questions that later may be important for the construction of minimalist (or Galilean) models that aim to be veridical representations.[20]

Another way the Hawk-Dove game is explanatory is that it isolates a single factor and investigates a hypothetical case, in order to better understand its potential contribution to the overall behavior of various possible systems.[21] The model tells us about how, given a particular idealized payoff structure, natural selection might lead to the evolution of restraint in combat in a wide range of possible systems. Understanding these potential contributions can be relevant to answering the why question of interest, even if the

20. In addition, models that use hypothetical pattern idealizations can contribute to constructing models that aim for accurate representations, by providing a template that can be made more realistic by introducing various complexities or more realistic assumptions. For example, later modelers sought to make the Hawk-Dove game more realistic. This illustrates how when interests change—e.g., if a modeler is interested in providing an explanation—a modeler may attempt to make the model more realistic. The important point is that the original model is able to be explanatory without requiring the (possible or actual) construction of more realistic versions of the model.

21. Thanks to William Wimsatt for suggesting this kind of case in his comments on an earlier version of this article.

https://doi.org/10.1086/671399 Published online by Cambridge University Press

model does not accurately represent how the isolated factor functions in any real-world system. So the Hawk-Dove model also produces understanding by isolating a factor that is relevant to the pattern of interest and investigating its influence within a hypothetical scenario. However, understanding the possible contributions of this single factor in this hypothetical case is insufficient to provide a satisfactory explanation.

Although models that use hypothetical pattern idealizations are distinguished by their aiming at being explanatory without aiming for veridical representation of a target system(s), this does not mean that other kinds of idealized models cannot be explanatory. For example, models that use minimalist idealizations might produce understanding by providing a veridical representation of some of the core causal factors of the system. Veridically representing the contributions of a single (or a few) causal factor(s) can be thought of as providing a partial explanation of the phenomenon to be explained. A model that provides this kind of partial explanation can later be used in tandem with other models that capture the contributions of other factors to demonstrate how taken together they might provide an explanation or be used to formulate a different model that is an explanation. So, in some instances, a model can be explanatory by providing a veridical representation of an essential component of the actual explanation.

5. The Scope of Hypothetical Pattern Idealization. In this section, we investigate the scope of hypothetical pattern idealization. In addition to the biological cases discussed above, we argue that hypothetical pattern idealizations are likely widely used in economic modeling.

Economic and biological modelers often face similar challenges in attempting to investigate patterns that range across systems that are causally complex and very heterogeneous. Given this modeling context, it is likely that hypothetical pattern idealizations will also be useful in many instances of economic modeling. Indeed, in line with our analysis of hypothetical pattern idealization, several authors have recently argued that many (if not most) economic models do not aim for veridical representations of mechanisms or causes in any particular target system(s) (Grüne-Yanoff 2009; Knuuttila 2009; Sugden 2009).

For instance, Knuuttila describes economic modeling in this way: "Instead of directly trying to represent some selected aspects of a given target system—as has conventionally been assumed—modelers proceed in a roundabout way, seeking to build hypothetical systems in the light of the anticipated results or of certain general features of the phenomena they are supposed to exhibit" (2009, 74). In other words, many economic models present hypothetical scenarios that do not aim to veridically represent some aspects of a target system(s). Furthermore, when these models are successful they can be used as a "starting point" for further theorizing that may aim to in-

https://doi.org/10.1086/671399 Published online by Cambridge University Press

vestigate the "real mechanisms" (74). Presumably, one reason these economic models are sometimes able to aid in the formulation of veridical models is because they are able to produce understanding that is relevant to answering why questions. Therefore, it seems likely that at least some economic modelers appear to be using a process similar to that involved in hypothetical pattern idealization.

What is more, similar to the Hawk-Dove model, one way economic models produce understanding is by aiding in the investigation of inconsistency claims (or answering a how-possibly question). For example, Grüne-Yanoff (2009) analyzes the use of economic models to "learn" about inconsistency and necessity claims.[22] He begins his account by describing the representational goals of many economic modelers: "Theoretical modelers . . . think of the model as a concrete situation—yet not a situation in the real world. They interpret formal structures not as descriptions of the real world but as describing 'parallel worlds' (Sugden 2000, p. 25), which exhibit familiar features of the real world but may not be identifiable with any of its particulars. Thus, economic models consist of both a formal structure and an interpretation of this structure as an imaginary scenario or world . . . from which modellers hope to learn about their ultimate target" (84). Although these economic models do not aim at satisfying various 'world-linking' properties, Grüne-Yanoff argues that one can still use such "imaginary" economic models to "learn" about necessity or impossibility hypotheses. Furthermore, he argues that economic models sometimes produce "new beliefs (about something being possible). This role of eliciting beliefs does not depend on the imaginer's believing something to be true or probable. On the contrary, [these] judgments . . . are often elicited solely through consideration of imaginary worlds" (Grüne-Yanoff 2009, 94). One of his main examples is Schelling's checkerboard model (Schelling 1978).

In Schelling's agent-based model, dimes and nickels are used to represent two types of individuals, *A* and *B*. The set of nine adjacent squares on a chessboard is used to represent an individual's 'neighborhood'. The model assumes that individuals prefer to have neighbors that are at least 30% of the same type—for example, *A*s want at least 30% of their neighbors to be *A*s. Using these preferences, the agents take turns determining whether they have at least three neighbors of the same type. If so, the agent remains in the same location; if not, the agent moves to the nearest unoccupied location. The model is run until all agents are satisfied with their location. Schelling's results showed that, given a small group of dissatisfied individuals, segregation is the equilibrium point of the model. Moreover,

22. Of course, the understanding produced by models that use hypothetical pattern idealizations need not be about inconsistency claims. We merely want to point out this parallel between Grüne-Yanoff's examples and our own.

https://doi.org/10.1086/671399 Published online by Cambridge University Press

this result continues to hold across a wide range of changes to the dynamics of the model, including the use of different utility functions, rules for updating, neighborhood sizes, and spatial configurations. Consequently, Schelling showed how minor preferences for like neighbors make it extremely hard to avoid segregation.

Grüne-Yanoff argues that Schelling's checkerboard model allows us to learn something important about an inconsistency hypothesis since "before the models' publication, it seems, many people believed that segregation was necessarily a consequence of explicitly racist preferences. Schelling's model showed that there were plausible settings in which this was not so" (2009, 96). Indeed, Schelling was interested in investigating how it is possible for segregation to occur in a wide range of systems without explicit collective preferences for segregation.[23] The answer he discovered through the use of a highly idealized model is that this is possible when every individual acts on a small preference for similar neighbors. This use of a highly idealized model to investigate an apparent inconsistency claim parallels the use of the Hawk-Dove game to investigate the compatibility of individual selection with the observation of restraint in combat in a wide range of populations.[24] In both cases, the modeler uses a highly idealized model that does not aim to accurately represent any real-world system(s), in order to investigate the observed pattern of interest. By showing that a particular kind of process could give rise to the phenomenon in a (or a wide range of) hypothetical case(s), these models were able to produce understanding relevant to answering a particular why question. Namely, in Schelling's case, the model produces understanding if you want to answer the question, "why does segregation exist in so many real-world populations?" Schelling's model is able to justify the true belief that a neighborhood can become segregated even if there are no explicitly racist preferences. However, Schelling's model is able to provide this insight, even though the model describes a segregation process that never occurs in any real-world system—that is, the model's ability to be explanatory is independent of the truth or accuracy of the model with respect to any target system(s). Therefore, the checkerboard model (like the Hawk-Dove model) is explanatory but does not aim to be an explanation.

Given our analysis of these authors' accounts of modeling in economics, it seems likely that many of the idealizations used within economic modeling will also be instances of hypothetical pattern idealization. This suggests

23. In fact, Schelling himself says his goal is to investigate "some of the individual incentives and individual perceptions of difference that *can* lead collectively to segregation" (1978, 138).

24. Sugden (2009) discusses several other ways that Maynard Smith's use of the Hawk-Dove model parallels the use of many economic models, e.g., Schelling's checkerboard model.

https://doi.org/10.1086/671399 Published online by Cambridge University Press

that the scope of hypothetical pattern idealization is much broader than its use within biological modeling. In fact, Godfrey-Smith has described something like this process as the general "strategy of model-based science" wherein scientists use models as a "deliberate detour through merely hypothetical systems" (2006, 734).[25] Therefore, in the end, hypothetical pattern idealizations may play a central role in a particular form of scientific theorizing that aims to construct explanatory models that do not aim to be explanations.

Consequently, not only has our analysis shown that a new category of idealization is needed, but when we investigate the scope of hypothetical pattern idealization, we find that it applies more generally. It is also plausible that this new category of idealization will apply to several cases that may previously have been thought to be instances of Galilean or minimalist idealization (or some combination of the two).[26] Indeed, many idealizations that were originally thought to be eliminable simplifications or isolations of causal factors appear to be closer to hypothetical pattern idealizations (e.g., see Batterman 2002; Wayne 2011; Rice 2012). However, the actual scope of hypothetical pattern idealization will have to be determined on a case-by-case basis, by examining particular instances of idealized models and paying close attention to the goals of the model builders.

6. Conclusion. Building on several previous accounts of idealization, Michael Weisberg has provided an extremely useful taxonomy for thinking about the use of idealizations in science. Unfortunately, Weisberg's taxonomy is unable to adequately characterize a kind of idealization that is widely used in biological modeling. In order to sufficiently characterize this kind of modeling, a fourth kind of idealization is required. We have named this kind of idealization *hypothetical pattern idealization* and have provided an analysis of its justification and representational goals. This kind of idealization is most often used to produce understanding relevant to patterns across extremely complex and heterogeneous systems. When these models are capable of justifying true beliefs that are relevant to answering a why question, they are explanatory. However, a model may be explanatory but fail to provide a veridical representation of the features of any real-world system necessary to be an explanation (or a partial explanation). Although we have focused on the use of hypothetical pattern idealizations within evolutionary

25. Although some authors have suggested that all models are fictional scenarios, we do not intend to endorse this stronger claim. We only want to point out that certain features emphasized by our account of hypothetical pattern idealization are widespread.

26. Thanks to an anonymous reviewer for suggesting that we be clearer about the difference between these two claims concerning the intended scope of hypothetical pattern idealization.

https://doi.org/10.1086/671399 Published online by Cambridge University Press

biology, this kind of idealization is likely pervasive in many sciences; for example, we have discussed some plausible examples in economics. As a result, incorporating hypothetical pattern idealizations is essential to providing an adequate account of the use of idealization in science.

REFERENCES

Achinstein, Peter. 1983. *The Nature of Explanation*. New York: Oxford University Press.
Batterman, Robert W. 2002. *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*. Oxford: Oxford University Press.
Bokulich, Alisa. 2012. "Distinguishing Explanatory and Nonexplanatory Fictions." *Philosophy of Science* 79:725–37.
Cartwright, Nancy. 1983. *How the Laws of Physics Lie*. New York: Oxford University Press.
———. 1989. *Nature's Capacities and Their Measurement*. New York: Oxford University Press.
Depaul, Michael, and Stephen Grimm. 2007. Review of *The Value of Knowledge and the Pursuit of Understanding*, by Jonathan Kvanvig. *Philosophy and Phenomenological Research* 74: 498–514.
Forber, Patrick. 2010. "Confirmation and Explaining How Possible." *Studies in History and Philosophy of Biological and Biomedical Sciences* 41:32–40.
Friedman, Michael. 1974. "Explanation and Scientific Understanding." *Journal of Philosophy* 71: 5–19.
Godfrey-Smith, Peter. 2006. "The Strategy of Model-Based Science." *Biology and Philosophy* 21: 725–40.
———. 2009. "Models and Fictions in Science." *Philosophical Studies* 143:101–16.
Grimm, Stephen. 2006. "Is Understanding a Species of Knowledge?" *British Journal for the Philosophy of Science* 57:515–35.
Grüne-Yanoff, Till. 2009. "Learning from Minimal Economic Models." *Erkenntnis* 70:81–99.
Hempel, Carl G. 1965. *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York: Free Press.
Kitcher, Philip. 1981. "Explanatory Unification." *Philosophy of Science* 48:507–31.
Knuuttila, Tarja. 2009. "Isolation Representations versus Credible Constructions? Economic Modeling in Theory and Practice." *Erkenntnis* 70:59–80.
Kvanvig, Jonathan. 2003. *The Value of Knowledge and the Pursuit of Understanding*. New York: Cambridge University Press.
Levins, Richard. 1966. "The Strategy of Model Building in Population Biology." *American Scientist* 54:421–31.
Lewis, David. 1986. *Philosophical Papers*. Pt. 2. New York: Oxford University Press.
Lipton, Peter. 2009. "Understanding without Explanation." In *Scientific Understanding: Philosophical Perspectives*, ed. Henk W. de Regt, Sabina Leonelli, and Kai Eigner, 43–63. Pittsburgh: University of Pittsburgh Press.
Maynard Smith, John. 1982. *Evolution and the Theory of Games*. New York: Cambridge University Press.
Maynard Smith, John, and Geoffery A. Parker. 1976. "The Logic of Asymmetric Contests." *Animal Behavior* 24:159–75.
Maynard Smith, John, and G. R. Price. 1973. "The Logic of Animal Conflict." *Nature* 246:15–18.
McMullin, Ernan. 1985. "Galilean Idealization." *Studies in History and Philosophy of Science* 16:247–73.
Mizrahi, Moti. 2012. "Idealizations and Scientific Understanding." *Philosophical Studies* 160: 237–52.
Morgan, M. S. 1999. "Learning from Models." In *Models as Mediators: Perspectives on Natural and Social Science*, ed. M. Morgan and M. Morrison, 347–88. Cambridge: Cambridge University Press.

https://doi.org/10.1086/671399 Published online by Cambridge University Press

Nowak, Leszek. 1972. "Laws of Science, Theories, Measurement." *Philosophy of Science* 39 (4): 533–48.
Pincock, Christopher. 2012. "Mathematical Models of Biological Patterns: Lessons from Hamilton's Selfish Herd." *Biology and Philosophy* 27:481–96.
Potochnik, Angela. 2007. "Optimality Modeling and Explanatory Generality." *Philosophy of Science* 74:680–91.
Resnik, David B. 1991. "How-Possibly Explanations in Biology." *Acta Biotheoretica* 39:141–49.
Reydon, Thomas A. C. 2012. "How-Possibly Explanations as Genuine Explanations and Helpful Heuristics: A Comment on Forber." *Studies in History and Philosophy of Biological and Biomedical Sciences* 43:302–10.
Rice, Collin. 2012. "Optimality Explanations: A Plea for an Alternative Approach." *Biology and Philosophy* 27 (5): 685–703.
Rohwer, Yasha. Forthcoming. "Lucky Understanding without Knowledge." *Synthese*.
Salmon, Wesley. 1998. *Causality and Explanation*. New York: Oxford University Press.
Schelling, Thomas. 1978. *Micromotives and Macrobehavior*. New York: Norton.
Sober, Elliott. 1983. "Equilibrium Explanation." *Philosophical Studies* 43:201–10.
———. 2000. *Philosophy of Biology.* 2nd ed. Boulder, CO: Westview.
Strevens, Michael. 2009. *Depth: An Account of Scientific Explanation*. Cambridge, MA: Harvard University Press.
Suárez, Mauricio. 1999. "The Role of Models in the Application of Scientific Theories: Epistemological Implications." In *Models as Mediators: Perspectives on Natural and Social Science*, ed. M. Morgan and M. Morrison, 168–96. Cambridge: Cambridge University Press.
———. 2009. *Fictions in Science: Philosophical Essays on Modeling and Idealization*. New York: Routledge.
Sugden, R. 2000. "Credible Worlds: The Status of Theoretical Models in Economics." *Journal of Economic Methodology* 7:1–31.
———. 2009. "Credible Worlds, Capacities and Mechanisms." *Erkenntnis* 670:3–27.
Van Fraassen, Bas C. 1980. *The Scientific Image*. Oxford: Oxford University Press.
Wayne, Andrew. 2011. "Expanding the Scope of Explanatory Idealization." *Philosophy of Science* 78:830–41.
Weisberg, Michael. 2007. "Three Kinds of Idealization." *Journal of Philosophy* 104:639–59.
Wimsatt, William. 2007. *Re-engineering Philosophy for Limited Beings: Piecewise Approximations of Reality*. Cambridge, MA: Harvard University Press.
Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.

https://doi.org/10.1086/671399 Published online by Cambridge University Press