

2016

Factive Scientific Understanding Without Accurate Representation

Collin Rice

Bryn Mawr College, crice3@brynmawr.edu

[Let us know how access to this document benefits you.](#)

Follow this and additional works at: https://repository.brynmawr.edu/philosophy_pubs

Part of the [Philosophy Commons](#)

Custom Citation

Rice, Collin. 2016. "Factive Scientific Understanding Without Accurate Representation." *Biology and Philosophy*, 31.1: 81-102.

This paper is posted at Scholarship, Research, and Creative Work at Bryn Mawr College. https://repository.brynmawr.edu/philosophy_pubs/19

For more information, please contact repository@brynmawr.edu.

Factive Scientific Understanding Without Accurate Representation

Author Information:

Collin C. Rice
Department of Philosophy
Lycoming College
700 College Place
Williamsport, PA 17701
Email: rice@lycoming.edu
Phone: (913) 961-2018

Abstract: This paper analyzes two ways idealized biological models produce factive scientific understanding. I then argue that models can provide factive scientific understanding of a phenomenon without providing an accurate representation of the (difference-making) features of their real-world target system(s). My analysis of these cases also suggests that the debate over scientific realism needs to investigate the factive scientific understanding produced by scientists' use of idealized models rather than the accuracy of scientific models themselves.

Keywords: Modeling, Scientific Understanding, Idealization, Optimality, Scientific Realism

Factive Scientific Understanding Without Accurate Representation

Abstract: This paper analyzes two ways idealized biological models produce factive scientific understanding. I then argue that models can provide factive scientific understanding of a phenomenon without providing an accurate representation of the (difference-making) features of their real-world target system(s). My analysis of these cases also suggests that the debate over scientific realism needs to investigate the factive scientific understanding produced by scientists' use of idealized models rather than the accuracy of scientific models themselves.

1. Introduction

In this paper, I investigate two ways highly idealized models can produce the cognitive achievement of factive scientific understanding. I then argue that models can produce factive scientific understanding of a phenomenon without providing an accurate representation of the (difference-making) features of any real-world target system. My analysis also suggests that the debate over scientific realism needs to investigate the factive scientific understanding produced by scientists' use of idealized models (or theories) rather than the (approximate) accuracy of scientific models (or theories) themselves.

According to most accounts of explanation, a necessary condition for something to explain is that it be, at least in some sense, true. Hempel (1965) originally distinguished between *true* explanations and *potential* explanations, which would be adequate if they were true. In modeling terms, this truth requirement claims that in order for a model to explain it must *accurately represent* (at least some of) the relevant features of the physical target system(s). Many contemporary accounts of how models explain make this accuracy requirement explicit. For example, Michael Strevens asserts, “no causal account of explanation—certainly not the

kairetic account—allows non-factive models to explain” (Strevens, 2009, p. 320).¹ More generally, for causal and mechanistic accounts, in order for a model to explain it must provide an accurate representation of the difference-making causal relationships or causal mechanisms that led to the target explanandum (Craver 2006; Kaplan and Craver 2011; Kaplan 2011; Strevens 2009; Woodward 2003).²

In addition, several philosophers have emphasized the connection between providing an explanation and the cognitive achievement of understanding (Achinstein 1983; Friedman 1974; Grimm 2006; Kitcher 1981; Lewis 1986; Salmon 1984, 1998; Strevens 2009). For example, Wesley Salmon writes, “understanding results from our ability to fashion scientific explanations” (Salmon 1984, p. 259). In addition, Michael Freidman argues that our theory of explanation, “should tell us what kind of understanding scientific explanations provide and how they provide it” (Freidman 1974, p. 14).

What is more, several philosophers have recently claimed that the *only* way to achieve scientific understanding is by grasping a *correct* explanation (de Regt 2009b; Khalifa 2012; Strevens 2009, 2013; Trout 2007). For example, J.D. Trout states that, “scientific understanding is the state produced, and only produced, by grasping a *true* explanation” (Trout 2007, p. 585–586). In addition, Michael Strevens argues: “An individual has scientific understanding of a phenomenon just in case they grasp a *correct* scientific explanation of that phenomenon” (Strevens 2013, p. 1). Given that most accounts—including Strevens’s kairetic account—require some form of accurate representation for a model to explain, these views seem to imply that

¹ While Strevens’s account does allow some idealized models to explain, accurate representation continues to play a key role since, “the overlap between an idealized model and reality...is a standalone set of difference-makers for the target” (Strevens 2009, p. 318).

² Indeed, mechanistic accounts of explanation typically include strong accurate representation requirements in order for a model to explain (Craver 2006; Kaplan and Craver 2011). For example, David Kaplan and Carl Craver argue for a model-to-mechanism-mapping (3M) requirement that involves various kinds of “correspondence” between the model and the actual causal mechanisms in the model’s target system (Kaplan 2011, p. 347).

accurate representation of difference-making features is necessary for producing scientific understanding.

Understanding is a cognitive achievement that involves grasping a fairly comprehensive body of information (Elgin 2007; Grimm 2006, 2012, Kvanvig 2003). Moreover, it is widely agreed that scientific understanding is *factive* in the sense that what one understands must have some connection to the way things really are (Grimm 2006, 2012; Mizrahi 2012; Kvanvig 2003). In general, when it comes to scientific understanding, “what we are trying to understand is how things actually stand in the world” (Grimm 2006, p. 518).

However, as many philosophers have noted, these accuracy (or truth) requirements for explanation and understanding are in tension with a widely recognized fact: that *idealization* is an essential and pervasive aspect of scientific theorizing (Cartwright 1983; Elgin 2007; Godfrey-Smith 2009; Mäki, 2011a, 2011b; Odenbaugh 2011; Psillos 2011; Rohwer and Rice 2013; Weisberg 2007a, 2007b, 2013; Wimsatt 2007). Generally, idealizations are known to be false assumptions—they deliberately misrepresent or distort the features of real-world systems. Consequently, several philosophers have suggested that the pervasiveness of idealized models (and theories) raises a serious challenge to scientific realism (Cartwright 1983; Levy 2012; McMullin 1985; Odenbaugh 2011; Psillos 2011; Saatsi 2014; Suárez 1999). The general idea is that, given that our best models and theories contain known to be false assumptions, even if they make accurate predictions we have little reason to believe they are true. For example, Jay Odenbaugh argues that, “if idealizations are generally ineliminable, we are rarely justified in believing our models” (Odenbaugh 2011, p. 1187). This has led many authors to suggest that, characterizing models as highly idealized fictions, “is incompatible with the most basic tent of

realism—namely, that attaining truth is the central aim of scientific investigation” (Levy, 2012, p. 741).³

This paper investigates two case studies from biology to illustrate how highly idealized models can produce factive scientific understanding despite the accuracy requirements outlined above. After showing how both cases produce factive scientific understanding, I will argue that models can produce factive scientific understanding of a phenomenon without providing an accurate representation of the (difference-making) features of a real-world target system(s). Finally, where Cartwright (1983), Odenbaugh (2011), Mäki (2011a, 2011b), Suarez (1999), Ladyman et al. (2007), Peters (2014), Worrall (1989) and others all focus on whether scientific realism can be defended based on the partial or approximate truth of our models or theories, I will suggest that a promising (but unexplored) approach to defending realism is to investigate the factive scientific understanding produced by scientists’ use of highly idealized models.

The paper will proceed as follows. In the next section, I outline an account of factive scientific understanding. Then, in Section 3, I analyze two cases in which a highly idealized model is used to produce factive scientific understanding in biology. Next, Section 4 argues that models can produce factive scientific understanding of a phenomenon without providing an accurate representation of the (difference-making) features of any real-world target system(s). Finally, Section 5 suggests that the debate over scientific realism needs to include an investigation of the factive scientific understanding produced by scientists’ use of idealized models.

2. What is Factive Scientific Understanding?

To begin, in order for something to count as an instance of *scientific* understanding, it needs to be the case that the understanding is a product of scientific inquiry; rather than say resulting from

³ However, Levy (2012) goes on to argue that if we understand fictions in a particular way we can maintain realism.

history, art, or literature. In addition, I am specifically interested in the understanding scientific models provide *of natural phenomena* rather than an individual's understanding of a model, theory, or subject matter; e.g. my understanding of relativity theory (de Regt 2009a). Jonathan Kvanvig (2009) discusses this distinction:

The issue here concerns the object of understanding. One might understand the model or theory itself, as when one understands phlogiston theory. One does not thereby understand combustion, however. Understanding the world scientifically is not simply a matter of understanding the given model but involves, rather, some relationship between the model and reality. (p. 342)

In light of this distinction, I will be assuming that scientific understanding is concerned with understanding “real phenomena” in the world (Schurz and Lambert 1994, p. 68).

It is widely accepted that the understanding science provides of natural phenomena involves “grasping something further” than what is involved in merely accepting a set of beliefs about the phenomenon (Elgin 2007; Grimm 2006, 2012, Kvanvig 2003). In general, understanding is thought to involve the ability to grasp some important information about, “how the various parts of the world [are] systematically related” (Grimm 2012, p. 103). This is typically thought to require that one who understands must grasp certain relations among the components (e.g. propositions or beliefs) of a larger body of information. Catherine Elgin puts the point this way:

Understanding is primarily a cognitive relation to a fairly comprehensive coherent body of information. The understanding encapsulated in individual propositions derives from an understanding of larger bodies of information that include those propositions (Elgin 2007, p. 35)

So, according to Elgin, understanding involves incorporating information into a comprehensive and coherent network of information. Philosophers of science echo this idea, but in slightly different terms: “to understand a phenomenon *P* is to know how *P* fits into one's background knowledge” (Schurz and Lambert 1994, p. 67). Following these views, on the account I present here, scientific understanding of a phenomenon requires that what one understands about the phenomenon must be systematically integrated into a wider body of information about the

phenomenon (or systems) of interest. This systematic integration can take multiple forms; e.g., grasping the kinds of functional, logical, modal, causal, prototypical, or exemplar relationships emphasized in the cognitive science literature on conceptual information (Machery 2009; Gopnik and Meltzoff 1997; Rice 2014; Weiskopf 2009). Many of these same relations—e.g. causal, modal, and theoretical relationships—have also been emphasized in the literature on explanation (Bokulich 2011, 2012; Rice 2015; Salmon 1984; Woodward 2003). While this list is probably not exhaustive, I will leave it to future cognitive science research to provide more details about the kinds of relationships human beings use to integrate new information into their existing background knowledge. Moreover, my focus here will be exclusively on modal (i.e. counterfactual) information, which is central to several accounts of conceptual information, explanation and understanding (Bokulich 2011, 2012; Grimm 2006, 2008; Rice 2014; Woodward 2003).

Additionally, in order to genuinely understand, an agent *must grasp how* the new information fits into this larger body of information about the phenomenon of interest. In other words, how the new information can be systematically integrated into this kind of larger cognitive corpus is the “something further” that must be grasped in order for the agent to genuinely understand.

Finally, what does it mean to claim that this kind of scientific understanding must be *factive*? Among philosophers of science, it is widely accepted that understanding is *factive* in the sense that (at least some of) the beliefs (or propositions) within one’s understanding must be true (Grimm 2006, 2012; Mizrahi 2012; Strevens 2009, 2013). For example, young earth creationists believe a great flood formed the Grand Canyon in about a year, but they do not understand because their story is incorrect (Strevens 2013). Indeed, there is a strong intuitive pull to say that

understanding natural phenomena cannot involve believing falsehoods since, “what we are trying to understand is how things actually stand in the world” (Grimm 2006, p. 518).⁴

One might think this requires that all the beliefs (or propositions) involved in one’s understanding must be true. However, this requirement is far too strong (Elgin 2007; Kvanvig 2003; Zagzebski 2001). Indeed, such a standard cannot do justice to the cognitive contributions of science since scientific understanding typically depends on the use of highly idealized models and theories. Moreover, since idealizations are so essential and pervasive in science, it is often difficult to see how the inferences used to produce scientific understanding can be isolated from the contributions idealizations make to those inferences (and understanding). As a result, the widespread use of idealizations within our best scientific models and theories suggests that requiring all of the beliefs (or propositions) that contribute to one’s understanding to be true is too high of a standard.

In response, Elgin (2007) contends that we should abandon a factive conception of scientific understanding. However, contrary to Elgin, I think that allowing that not all the beliefs within one’s understanding must be true does not require that we claim that scientific understanding is non-factive. For example, Kvanvig (2003) and Mizrahi (2012) both argue for what they call a “quasi-factive” account of understanding by distinguishing between central propositions and peripheral ones. On Kvanvig’s view, all of the central propositions of one’s understanding must be true, but a few false beliefs about peripheral propositions does not undermine one’s understanding. This allows some falsehoods to play a role, while still requiring a factive component to understanding.

⁴ In addition, Pritchard (2009), Grimm (2006), and Kvanvig (2003) all argue for a factive notion of understanding that requires at least some of the propositions in one’s understanding be true.

While this is a step in the right direction, one problem with Kvanvig's view is that idealizations are often *central* (or essential) to the understanding provided by many, many scientific models and theories (Elgin 2007; Mizrahi 2012). Moreover, as Elgin notes, there is often "no expectation that in the fullness of time idealizations will be eliminated from scientific theories...elimination of idealizations is not a desideratum. Nor is consigning them to the periphery of a theory" (Elgin 2007, p. 38). As a result, requiring that the idealizations within our best scientific models be removed or resigned to the periphery of our understanding is the wrong way to argue that scientific understanding can still be factive.

Instead, my view is that scientific understanding is factive because in order to genuinely understand a natural phenomenon *most* of what one believes about that phenomenon—especially about certain contextually salient propositions—must be true. To be clear, I am not claiming that simply having a majority of what one believes about the phenomenon be true is sufficient for understanding. Nor do I suggest that determining if an agent understands requires counting up all the beliefs within their understanding and determining whether the number of true beliefs meets some universally applicable threshold (e.g. sixty two percent). Rather than simple methods of proposition counting, I recommend a case-by-case approach that allows for a plurality of context-sensitive ways that one's understanding might meet this factive requirement. For one thing, in different contexts of scientific inquiry, some propositions will be more important (or salient) to one's understanding and so their truth-value will carry more weight in determining whether one's understanding meets this factive requirement. Another issue concerns cases where one's beliefs are only "approximately true" or are directly inferred from other false beliefs. I won't work through the details of each of these complications here, but I will suggest that the factive component of scientific understanding is somewhat flexible and is highly context-

sensitive. However, it is important to note that the context of scientific inquiry will typically establish a particular why question of interest, a contrast class, a set of features that are thought to be relevant and irrelevant, etc. As a result, merely having many true beliefs about the phenomenon of interests will typically be insufficient for understanding since the agent who understands will often be required to grasp certain truths that are made particularly salient by the context of inquiry. Precisely how these contextual factors interact with the factive requirements on scientific understanding will have to be uncovered by analysis of particular cases.

Therefore, while I maintain a factive requirement for understanding, I have no universal account to offer about how to determine precisely (across all cases) when most of what one believes about the phenomenon is true in a way that is sufficient for understanding. Still, although the factive requirement suggested above is admittedly vague, I think we can make clear judgments in many cases. For example, the Grand Canyon case above seems to clearly fail the requirement since these agents fail to grasp the particularly salient facts that the cause of the Grand Canyon was the Colorado River (not a great flood) and that the process took approximately 6 million years to complete. As a result, most of what these agents believe about the phenomenon and why it occurred is incorrect. Moreover, the facts the agents get incorrect are precisely those facts that are made most salient by the particular question being asked. We can contrast this case with physicists' understanding of the movement of planets which seems to clearly meet the factive requirement since—although it may contain some false beliefs—is mostly constituted by true beliefs and those true beliefs are those that are particularly salient in the context of inquiry; e.g. where the earth is, how planets rotate, which bodies orbit which others, and the approximately elliptical shape of planetary orbits.

What is more, I think this vagueness regarding the factive requirement for understanding is due to the inherent vagueness in our judgments about when the falsity of one's beliefs will undermine one's understanding—i.e. vagueness in how we apply the concept of understanding. Therefore, while “most of what one believes about the phenomenon is true” is admittedly vague and context sensitive, I suggest that this tracks the way we attribute understanding to individuals. Moreover, I see no reason not to call this a factive notion of understanding since truth continues to play a key role in our judgments about whether or not one genuinely understands.

In short, I claim that a scientific model produces factive scientific understanding of a natural phenomenon if it *enables an agent to grasp some true belief(s) about the phenomenon of interest and the agent grasps how that information can be systematically incorporated into a larger body of information in which most of what the agent believes about the phenomenon is true*. It is important to note, however, that the above account of factive scientific understanding leaves open the possibility that many (and perhaps central) propositions in one's understanding might be false. Indeed, in the next section I analyze two ways that highly idealized models in biology can produce this kind of factive scientific understanding.

3. Two Ways Highly Idealized Models Produce Factive Scientific Understanding

My case studies will both involve the use of optimization models in biology. The unifying feature of optimization models is the use of a mathematical technique called optimization theory. Optimization theory is widely applicable across the sciences since, “In engineering, as in evolution, the [best] attainable solution is often a compromise, owing to *constraints* on the feasible design options and *tradeoffs* among different benefits to be achieved by the design...Optimization is about constraints and tradeoffs” (Seger and Stubblefield 1996, p. 94). Indeed, these examples of optimality modeling serve as useful case studies because optimality

models are widely used in biology (Orzack and Sober 2001; Potochnik 2007, 2009; Rice 2012, 2015; Stephens and Krebs 1986), physics (Hartmann and Rieger 2002), economics (Pindyck and Rubinfeld, 2009; Rice and Rohwer 2013), cognitive science (Churchland 2013; Carruthers 2006), chemical engineering (Corsano et al. 2009), and various social sciences. Moreover, the cases I describe are similar to the uses of other kinds of models in other disciplines.

Consequently, the arguments that follow could easily be applied to other kinds of modeling and models outside of biology.

3.1. System-Specific Modeling

In the first kind of case, an optimality model produces factive scientific understanding by investigating a biological phenomenon that occurs within a particular target system. By showing how certain features contribute (i.e. are counterfactually relevant) to the occurrence of the phenomenon and showing why certain contextually salient features are irrelevant, an optimality model can enable us to understand why the phenomenon occurred in its target system.⁵ Precisely which features need to be shown to be relevant and irrelevant will depend on how we specify the target phenomenon, the context of inquiry, and the nature of the model's target system. However, once these features are taken into account, by providing the right set of modal information about the counterfactual relevance and irrelevance of various features of the target system, biologists can use a highly idealized optimization model to understand why the target phenomenon occurred (Orzack and Sober 2001; Potochnik 2007, 2009; Rice 2012, 2015; Sober 2000).⁶

⁵ In what follows I adopt something similar to counterfactual accounts of explanation that focus on a model's ability to answer what-if-things-had-been-different questions (Bokulich 2011, 2012; Rice 2015; Woodward 2003). However, my goal here is not to argue for a particular account of scientific explanation, but to investigate how the information provided by idealized models is able to produce factive scientific understanding.

⁶ Here I explicitly put things in terms of counterfactual (or modal) relevance and irrelevance in order to remain neutral on the need to reveal causal relations in order to provide understanding of a phenomenon. Indeed, it has been argued that optimization models provide noncausal explanations (Rice 2012, 2015), that natural selection is not a cause (Walsh et al. 2002; Matthen and Ariew 2009), and that only counterfactual information (but not causal information) is required for explanation (Bokulich 2011, 2012; Rice 2015). In order to remain as neutral as possible

Because the goal of these modelers is to understand a system-specific phenomenon, I will refer to this as *system-specific modeling*.

The goal of system-specific modeling is to provide accurate information about the counterfactual relevance and irrelevance of various contextually salient features within the model's target system. This information about counterfactual relevance and irrelevance then leads the modeler to acquire factive scientific understanding of why the phenomenon occurred in the target system. This understanding is constituted by a set of true beliefs about the counterfactual relevance and irrelevance of various contextually salient factors, which is incorporated into a larger body of mostly true scientific knowledge about the phenomenon of interest.

An example of system-specific modeling is Schmid-Hempel et al.'s (1985) use of optimality models to investigate honeybee foraging behavior.⁷ Honeybees forage for nectar in patches and carry their load back to the hive. The phenomenon of interest is that bees often leave food sources when their crops are only partially filled. Schmid-Hempel et al. attempted to understand this system-specific phenomenon by investigating two optimality models constructed from a detailed description of the steps in the honeybee foraging cycle (see Figure 1).

in this paper, I maintain the focus on counterfactual relevance, which can encompass both causal and noncausal kinds of relevance.

⁷ For considerations of space I do not work through all the detailed equations of the model (those interested should see Schmid-Hempel et al.'s original article).




Figure 1: Schmid-Hempel et al. (1985)'s depiction of the honeybee foraging cycle. From A to B the bee is traveling from the hive to the flower, from B to C the bee is collecting food from flowers within a patch, from C to D the bee returns to the hive, and from D to E the bee is flying within the hive.

The first optimality model assumed that selection would maximize the *net rate of energy delivery*, given by net energy divided by time. The second optimality model assumed that selection would maximize *energy efficiency*, given by net energy gain per unit of energy expended. Empirical data from other studies was then used to fill in the parameters of these models with precise values; e.g., values for the metabolic rate during flight for an unloaded bee and the linear increase in metabolic rate with load (Schmid-Hempel et al. 1985, p. 63). These modelers were then able to make detailed quantitative predictions that could be experimentally tested within real-world honeybee populations. The results showed that honeybees' foraging behaviors appear to be maximizing their energy efficiency, not the rate of energy intake.

Consequently, since energy efficiency is reduced with each additional flower, the optimality model that assumes that natural selection would favor strategies that maximize energy efficiency can be used to understand why members of the target population often leave food sources when their crops are only partially filled. Understanding this behavior involves understanding that the reduction of energy efficiency with each additional flower is the

counterfactually relevant feature in the evolution of the target phenomenon. This new information about counterfactual dependence is then incorporated into a larger corpus of background knowledge that includes the scientists' training, their understanding of the theory of natural selection, their understanding of honeybee foraging behavior, the information provided by previous studies, etc.

In addition, this optimization model uses several additional assumptions—many of which are idealizations. For example, the model assumes that natural selection will ultimately be able to overcome other evolutionary influences; e.g. drift or genetic recombination. In order to capture these adaptationist assumptions, these models typically make the idealizing assumptions that the population is infinite, phenotypes are passed on to offspring perfectly (i.e. like begets like), and there is no intergenerational overlap. These assumptions result in the optimal strategy—i.e. the strategy that maximizes energy efficiency—being the expected equilibrium of the model population. In other words, without these idealizations, the optimality model would not entail the prediction that current honeybee populations ought to leave food sources when their crops are only partially filled. However, as a result of these (and other) essential idealizations, these optimality models *fail to accurately represent the selection process of any real-world biological population*.

The challenge, then, is to see how these idealized models can still provide true modal information that allows scientists to understand the phenomenon of interest. For example, in this case, it is important that we can show (mathematically) that the actual size of the population is counterfactually irrelevant to the equilibrium point of its long-term evolution (as long as the population is sufficiently large). In addition, we can demonstrate that the results of evolution by natural selection, even with recombination and multilocus structures, will converge on the

equilibrium point of an optimality model (Eshel and Feldman 2001, p. 183). Because we are able to demonstrate that these features are counterfactually irrelevant to the equilibrium point of the evolving population, we can see how this highly idealized model can still provide true counterfactual information that enables us to understand the phenomenon of interest.⁸

More specifically, this optimality model shows us how the tradeoff between increased crop and increased energy expenditure is the counterfactually relevant feature in the evolution of honeybee's foraging behavior *and* the modelers can see how other features of the target system—e.g. the particular population size, the particular inheritance process, the initial conditions of the population, or the actual dynamical trajectory of the population—are counterfactually irrelevant to the ultimate equilibrium point of the evolving population. Therefore, despite the fact that it fails to accurately represent the selection process of its real-world target system, this optimality model allows the modeler to *acquire factive understanding of why* bees leave often leave food sources when their crops are only partially filled. The model accomplishes this task by allowing these modelers to grasp several true propositions about which features are counterfactually relevant and irrelevant to the phenomenon and the agent grasps how this new information can be systematically incorporated into a larger nexus of mostly true background information about evolving honeybee populations.⁹ Consequently, despite their

⁸ Of course, just how one understands what it means to “accurately represent” matters here. For example, one could argue that these models accurately represent the counterfactual structure of the phenomenon, but misrepresent (e.g. with a fictional ontology) the entities, relations, and processes of that phenomenon. If this could be maintained, then one might claim that these models produce this factive understanding only by accurately representing these counterfactual relations. However, I suggest that the distortion of the kinds of entities, relationships, and difference-making processes of their target phenomena entails that these models also distort the counterfactual structure of the target phenomena in some way. That is, this counterfactual structure cannot be easily isolated from the entities, relations, and processes of the phenomenon. Instead, that true counterfactual information typically has to be *inferred* by the scientific modelers who incorporate the results of the model into their larger set of background beliefs and assumptions. That is, the true counterfactual structure cannot be simply “read off” the idealized model that mirrors that structure.

⁹ In other cases, a model might be used to draw inferences about real-world systems even when the modeler is unsure whether the assumptions of the model are true or false. Alternatively, a model might lead to some false implications along with some true ones. All of these are important cases that require additional analysis. What I wish

being highly idealized and distorting the difference-making process(es) of natural selection, optimality models that investigate system-specific phenomenon can produce factive scientific understanding of the kind described in Section 2.

3.2. Modeling Hypothetical Scenarios

In the second kind of case, an optimality model is constructed to better understand the general behavior of a small set of related features that are believed to be present in many real systems, but the model's representation of those features is not intended to accurately represent the features of any particular real-world target system. That is, in these cases, the goal of the modeler is to *investigate the possible contributions of a few key features within a wide range of (actual or possible) systems*. This goal is typically accomplished by constructing a model of a *hypothetical scenario* that isolates the features of interest. A hypothetical scenario is one that is not intended to accurately represent any particular features of a real-world system—i.e. the model has no real-world “target system” whose (difference-making) features it aims to accurately represent. Because this kind of modeling involves the construction of a hypothetical scenario, I will refer to it as *hypothetical modeling*.¹⁰

Hypothetical modeling is suggested by many authors who view optimality modeling as a tool for discovering general interactions among some key variables. For instance, Seger and Stubblefield claim that: “[Optimality] models are intentionally caricatures whose purpose is to gain some insight about how a small number of key variables might interact” (Seger &

to focus on here, however, is that even when scientists know which assumptions of their models are false, when those idealizations play essential roles in the model and cannot be replaced with alternative assumptions, precisely how they are able to acquire factive understanding requires detailed analysis of particular cases. The various additional ways that this might occur will have to be analyzed in future work.

¹⁰ These hypothetical models are similar to what other authors have called “toy” models in that they are highly simplified and aim to provide insights concerning only a few key aspects of the target phenomenon. However, while many toy models (e.g. the Ising model) are claimed to capture the crucial features that produce the phenomenon, not all hypothetical models will achieve this goal. Therefore, while they are certainly related, hypothetical models are not the same as toy models. Toy models may, however, be a species of hypothetical model that accomplishes particular goals for the modeler.

Stubblefield 1996, p. 108). In addition, Potochnik (2009) identifies a “weak use” of optimality models in which, “the [optimality] model represents the role of natural selection in bringing about the evolutionary outcome”, but selection is only *one* important factor involved in the trait’s evolution (Potochnik 2009, p. 187). However, in contrast to Potochnik’s view, my hypothetical modeling does not require accurate representation of the process of natural selection since, as Potochnik notes, in some cases “the aim of optimality modeling is merely to represent *possible* selection dynamics” (Potochnik 2009, p. 188).¹¹ I will argue that modeling these possible systems via a hypothetical scenario is often sufficient to produce factive scientific understanding of a phenomenon without providing an accurate representation of the (difference-making) features of any real-world target system(s).

By building hypothetical models, scientists can investigate a particular set of features that is believed to function in a wide range of systems. Moreover, by building a related set of such models, scientists can begin to understand how those features might contribute to the overall behavior of a system in different contexts; e.g. in conjunction with different sets of assumptions. Understanding these features’ possible contributions to overall system behavior often enables the modeler to answer how-possibly questions (Forber 2010; Odenbaugh 2005; Resnik 1991), or justify background beliefs about what is necessary or possible (Rohwer and Rice, 2013). What is more, these background beliefs—e.g. about what is possible or necessary—are often *true* and can be incorporated into larger networks of knowledge concerning the phenomenon of interest. In other words, hypothetical models can produce factive scientific understanding of a real-world phenomenon by providing true modal information about that phenomenon—the same kind of information that enables system-specific models to produce factive scientific understanding of a

¹¹ Odenbaugh (2005) also discusses the use of idealized models to explore possibilities and answer how-possibly questions.

phenomenon. While there is a sense in which system-specific models aim to provide “how actually” information while hypothetical models typically aim to provide “how possibly” information, in both cases it is the model’s ability to provide true modal information about the space of possibilities that enables the model to produce factive scientific understanding.¹²

An example of hypothetical modeling is John Maynard Smith’s original use of the Hawk-Dove game (Maynard Smith 1978; Maynard Smith and Price 1973). In the natural world, organisms often exercise restraint in combat instead of fighting to the death. The Hawk-Dove game is intended to show how individual selection could *possibly* produce this behavior in a wide range of populations.

In the Hawk-Dove game, two organisms compete for a resource that will increase their fitness by V . The basic game allows only two strategies: Hawks (H) escalate until injured or until the opponent retreats; Doves (D) display, then retreat if their opponent escalates. This results in three kinds of interactions: (1) *Hawk vs. Hawk*, where each player has a 50% chance of obtaining the resource, V , and a 50% chance of receiving some cost, C , 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 following payoff matrix:

	H	D
H	$\frac{1}{2}(V - C), \frac{1}{2}(V - C)$	$V, 0$
D	$0, V$	$V/2, V/2$

Where $V > V/2 > 0 > \frac{1}{2}(V-C)$.

¹² In addition, one way to think about the difference is that system-specific models will typically yield *more* counterfactual information and that information will be focused on more specific changes to the actual features of the system. Hypothetical modeling, in contrast, will often explore the possibility space in which features differ more dramatically from those of the actual system. These differences may be important for certain kinds of understanding, but my goal here is to highlight the similarities in the ways that these two kinds of models produce a similar cognitive achievement by providing similar kinds of information.

In this game, neither Hawk nor Dove is an evolutionarily stable strategy (or ESS). However, a stable equilibrium does occur when the average payoffs for Hawks are equal to the average payoffs for Doves. This can occur in one of two ways. First, the population could consist of a mixture of some Hawks and some Doves. 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 predicts that individual selection will lead to restraint in combat in some instances. In this way, the Hawk-Dove game provides some factive scientific understanding by showing *how it is possible* for individual selection and a particular kind of payoff structure to produce the phenomenon of interest (Rohwer and Rice 2013).

However, the Hawk-Dove game's ability to provide this insight involves the use of several idealizations, including: (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; (7) perfect correlation between winning the resource and reproductive success (Maynard Smith 1982). In addition, the model represents the available strategies, interactions, and the payoffs in a highly distorted way—i.e. most of the features of real-world populations are left out or misrepresented. Indeed, as Maynard Smith and Price repeatedly emphasize, “real animal conflicts are vastly more complex than our simulated conflicts” (Maynard Smith and Price 1973, p. 17). As a result, the Hawk-Dove model fails to—and does not aim to—accurately represent the (difference-making) features of the selection process of any real-world system.

Still, despite its failure to accurately represent the (difference-making) features of any real-world target system, the model does enable us to *understand* that our observations could *possibly* be explained by individual-level selection. Indeed, as Maynard Smith and Price claim:

“A main reason for using [the model] was to test whether it is *possible even in theory* for individual selection to account for ‘limited war’ behaviour” (Maynard Smith and Price 1973, p. 15, my emphasis). Therefore, despite its failure to accurately represent the features of a particular target system, the Hawk-Dove game allows us to answer a key how-possibly question concerning the compatibility of individual selection with the observed behavior by investigating a hypothetical scenario. Moreover, the model helps us understand how certain key features might interact in a range of possible systems. This modal information about the phenomenon of interest *is true* and one can grasp how it fits into our larger body of knowledge about evolving biological populations in which restraint in combat occurs. Therefore, by providing true modal information about the phenomenon of interest, this hypothetical model produces factive scientific understanding of that phenomenon, despite the fact that it does not provide an accurate representation of the (difference-making) features of any particular real-world target system.

4. Factive Scientific Understanding Without Accurate Representation

As I noted above, one widely accepted feature of scientific explanations is that they produce scientific understanding (Achinstein 1983; Friedman 1974; Kitcher 1981; Kitcher and Salmon 1989; Salmon 1984; Strevens 2009). Hempel originally suggested that an explanation shows that “the phenomenon was to be expected; and it is in this sense that the explanation enables us to understand why the phenomenon occurred” (Hempel 1965, p. 337). In addition, Kitcher requires that our account of explanation “should show us *how* scientific explanation advances our understanding” (Kitcher and Salmon 1989, p. 168), and Salmon writes: “understanding results from our ability to fashion scientific explanations” (Salmon 1984, p. 259). More recently, both Strevens and Woodward have suggested that explanations produce understanding in agents who grasp them (Strevens 2009, p. 3; 2013, p. 1; Woodward 1993, p. 249). Indeed, by grasping the

explanans (or model), an agent can gain understanding about why the explanandum occurred.

While I agree that providing scientific explanations is a key way of producing factive scientific understanding, most accounts of explanation require that the model provide an accurate representation of at least some of the difference-making features of a real-world target system (Craver 2006; Kaplan and Craver 2011; Kaplan 2011; Strevens 2009; Woodward 2003).

However, I have argued that highly idealized models—such as hypothetical models—can produce factive scientific understanding of real-world phenomena *even if they fail to provide an accurate representation of the difference-making features of any real-world target system*. More specifically, for a model to produce factive scientific understanding requires the model to allow for a certain kind of cognitive achievement, but producing that kind of cognitive achievement *does not* require that the model, itself, be an accurate representation of the difference-making features of any real-world target system. Instead, a model can produce the cognitive achievement of factive scientific understanding of a phenomenon even if the model itself does not—and perhaps does not even attempt to—accurately represent the features of *any* real-world target system(s); e.g. by describing an impossible hypothetical scenario.

Of course, some kind of link is required between such a hypothetical model and the real world phenomenon, but providing an accurate representation of the difference-making features of the target system is not a necessary condition for producing understanding of a phenomenon. Instead, I maintain that there are numerous possible “links” between idealized models and real-world systems that can be sufficient for the model to produce factive scientific understanding in different contexts. For example, besides accurately representing difference-makers, one possible link is that the idealized model and the target system are in the same *universality class* (Batterman and Rice 2014). This entails that the model and the real-world system(s) will display

similar patterns of macroscale behavior even if the model drastically distorts the entities, relationships, and processes of its target system(s). Alternatively, the scientific modeler may play an essential role in establishing a sufficient link between the idealized model and its target system; e.g. by interpreting the assumptions of the model, interpreting the results obtained from the model, or connecting those results with those obtained from other models. In short, I contend that there are several ways that idealized models can be linked to their target systems—including but not limited to accurate representation—that can allow for the production of the kind of factive scientific understanding discussed in Section 2.

For example, the Hawk-Dove game allows scientists to grasp the true proposition that restraint in combat could *possibly* result from individual-level selection and the modeler can see how that information fits into their larger cognitive corpus of mostly true scientific knowledge concerning the phenomenon of interest. Moreover, the model is able to accomplish this despite the fact that it distorts the entities, interactions, relationships, and difference-making processes of the target phenomenon. It is only due to the modeler’s background knowledge about the kind of processes that count as “individual-level” selection and interpretation of how the model connects with real-world systems that the model is able to provide the true modal information about the target phenomenon.

These ideas build on Peter Lipton’s suggestion that merely potential explanations may produce understanding:

Potential explanations may provide actual understanding without approximating an actual explanation. . . The understanding involves a kind of cognitive gain about the actual phenomenon, even though the proffered explanation is not true of the actual phenomenon. (Lipton 2009, p. 52).

More specifically, in both of the examples given here, the cognitive gain about the target phenomenon is provided by *grasping true modal information about the phenomenon*. My focus on modal information also ties into a view originally proposed by Robert Nozick: “explanation

locates something in actuality...while *understanding* locates it in a network of possibility” (Nozick 1981, p. 12). More recently, Stephen Grimm has expanded on Nozick’s idea by suggesting that when we don’t have the understanding provided by an actual explanation we might have what he calls ‘proto-understanding’:

By an agent’s proto-understanding, I mean an agent’s convictions *about* the sorts of possibilities that are live or relevant, relative to the situation in question. [This is] a further specification of Nozick’s notion of a ‘network of possibility’; it is something like a person’s ‘modal sense’ of the various alternatives that might have obtained, relative to the fact in question. (Grimm 2008, p. 491).

I suggest that Grimm’s proto-understanding amounts to genuine scientific understanding that, although it may be provided by an accurate representation of the features of a real-world target system, may be provided in other ways as well. In the cases above, the modal information provided by these models produces beliefs about the network of possibilities and can answer how possibly questions (Forber 2010). Moreover, these beliefs about the network of possibility are often true of the phenomenon of interest and are incorporated into larger bodies of information containing mostly true beliefs about the phenomenon of interest. As a result, highly idealized models that fail to provide an accurate representation of the difference-making features of any real-world target system can still produce factive scientific understanding.

Consequently, my analysis of these cases raises a potential problem for several recent accounts that have claimed that the *only* way to produce scientific understanding is by providing a *correct* explanation (de Regt 2009b; Khalifa 2012, 2013; Strevens 2009, 2013; Trout 2007). For example, J. D. Trout claims that, “scientific understanding is the state produced, and only produced, by grasping a *true* explanation” (Trout 2007, p. 585–586). In addition, Michael Strevens argues that, “An individual has scientific understanding of a phenomenon just in case they grasp a *correct* scientific explanation of that phenomenon” (Strevens 2013, p. 1) and follows Trout in taking “scientific understanding to be that state produced, and only produced, by

grasping a *true* explanation” (Strevens 2009, p. 3). The problem with such claims is that according to most accounts of explanation—including Strevens’s own account—models will need to accurately represent difference-making features of a real-world target system in order to provide an explanation. If this is so, then Strevens’s and Trout’s claims require accurate representation of the difference-making features of a real-world target system in order for a model to produce scientific understanding.

In contrast to these accounts, I have argued that a highly idealized model can produce factive scientific understanding of a phenomenon even if it is an inaccurate representation of most (or perhaps even all) of the features of real-world systems. For example, in hypothetical modeling, the model provides factive scientific understanding by producing true beliefs about what is possible and these true beliefs can be systematically incorporated into a larger network of information in which most of what the modeler believes about the phenomenon is true. Recognizing that producing factive scientific understanding in this way is possible is important for characterizing the different epistemic contributions that idealized models make to scientific inquiry. Indeed, this distinction shows that providing an accurate representation of the difference-making features of a real-world system is only essential for *some* ways of producing factive scientific understanding. Consequently, accounts that require a model to provide an accurate representation of the (difference-making) features of real-world systems in order to produce understanding miss a large amount of scientific modeling that produces factive understanding of real-world phenomena—e.g. the factive understanding produced by hypothetical models.

5. Factive Scientific Understanding and Scientific Realism

Finally, my analysis of these cases also has interesting implications for the ongoing debate over scientific realism. Traditionally, scientific realism claims that science aims at truth and that we have reason to believe that our most successful scientific theories and models are true or approximately true. Several philosophers have suggested that the widespread use of idealizations in our models raises (at least potential) problems for the scientific realist (Cartwright 1983; McMullin 1985; Odenbaugh 2011; Psillos 2011; Suárez 1999). The central idea is that, given that we know our models include false assumptions, even if they make accurate predictions, we have no reason to believe that they, or the theories that employ them, are true (or accurate).

One potential response, which many philosophers have adopted, is for the realist to claim that highly idealized models can provide *partially* accurate representations (Bueno and Colyvan 2011; Kitcher 1993; Peters 2014; Psillos 1999; Pincock 2011; Strevens 2009; Weisberg 2007a, 2007b; Worrall 1989). Although models will typically (if not always) involve some level of idealization, this does not entail that they are incapable of accurately representing *some* features of real-world systems. Indeed, as many authors have noted, idealizations can play important roles in our best scientific models—e.g. by showing that certain features are irrelevant to the occurrence of some phenomenon (Batterman 2002; Strevens 2009; Wayne 2011; Weisberg 2007a, 2013). The key to this realist strategy is to point out that not every part of the model is required (or even intended) to accurately represent the target system(s). Rather, in many cases, highly idealized models can provide partially accurate representations of real system(s).

The greatest challenge for such a “selective confirmation” or “decompositional” strategy is to provide principled ways for distinguishing which parts, features, or aspects are actually required for the success of the model (or theory) from those that are idle or irrelevant (Stanford 2003, 2006). Much of the literature on idealization has focused either on how they can ultimately

be eliminated (McMullin 1985; Weisberg 2007a), or how they can be justified by playing a strictly periphery role in the explanations and understanding provided by models—i.e. by distorting only what is already known to be irrelevant (Elgin and Sober 2002; Weisberg 2007a, 2013; Strevens 2009).

However, although some idealizations can perhaps be justified in these ways, there are many instances in which *the idealizations are ineliminable and play central roles within our best scientific models* (Batterman 2002, 2009; Batterman and Rice, 2014; Rice 2012, 2015; Wayne 2011). Indeed, as I discussed above, idealizations are often *essential* assumptions within our best scientific models and there is often no expectation that the idealizations will eventually be removed or resigned to the periphery of our understanding (Batterman 2002, 2009; Elgin 2007; Morrison 2009; Rice 2015; Wayne 2011). For example, in both of the cases described above it is unclear how the idealizations could be removed from the model without consequently eliminating the model's ability to fulfill the purpose for which it was constructed. At the very least, eliminating the idealizations from these models would make the model much *worse* at fulfilling the goals of the model-builder. Moreover, removing (or correcting) the idealizations of these models would entail that the models no longer produce the target phenomenon; e.g. without the idealizations of infinite population size, asexual reproduction, isolation of the selected feature, no intergenerational overlap, etc. the optimal strategy would no longer be the expected outcome of the evolving population.

Furthermore, these models use a variety of idealizations that drastically distort the selection process(es) that led to the target phenomenon. Because they are adaptationist models, these models assume that natural selection is the most important—i.e. difference making—factor in the evolution of the trait in question. However, due to the numerous idealizations listed above,

these models distort the actual selection process(es) that led to the target phenomenon. Consequently, rather than distorting only irrelevant factors in order to isolate difference makers (as Strevens's view and other minimalist approaches would have it) these optimization models drastically distort the difference making processes (or factors) that led to the target explanandum. For example, selection in an infinitely large asexually reproducing population is an importantly different kind of process than what occurs in all real-world biological populations.

Indeed, the biggest problem for the selective confirmation strategy is that, in many cases, there simply is no clear mapping from each idealization onto a particular set of features that are irrelevant. Instead, in many cases, idealizations distort features that *do* make a difference to the target phenomenon (Batterman and Rice 2014; Levy 2011; Morrison 2009; Rice 2015). This makes it far more challenging to say precisely which aspects are essential and which are irrelevant to the success of the model. As a result, while this kind of selective confirmation strategy is possible, adequately defending it will require a far more complicated answer about how to distinguish the essential and irrelevant aspects of idealized models.

However, even if this selective confirmation strategy fails (or is exceedingly difficult), the above discussion suggests another more novel line of defense for the realist. The alternative realist strategy is for the realist to point out that both of the cases discussed here are capable of producing the cognitive achievement of *factive scientific understanding*—which involves *incorporating true beliefs into a cognitive corpus in which most of what one believes about the phenomenon of interest is true*. In system-specific modeling, the model will generate factive scientific understanding by providing information about the counterfactual relevance and irrelevance of various features within the model's target system. Perhaps more interestingly, in hypothetical modeling the model is still capable of producing factive scientific understanding

even when the model does not even aim to provide an accurate representation of the difference-making features, relations, or processes of a real-world target system.

Recognizing that idealized models can produce a large body of factive scientific understanding without having to accurately represent the (difference-making) features of their target system(s) reveals a large class of systematized true beliefs about the natural world. However, philosophical discussions of scientific realism frequently focus exclusively on the truth or accuracy of models (or theories) themselves rather than the truth, accuracy, or justification of the body of scientific understanding produced by scientists' use of those models. I contend that this focus is too narrow for adequately evaluating scientific realism. Philosophers must also look beyond the accuracy of the models (or theories) in question and *evaluate the body of factive scientific understanding* that can be produced by using idealized models in various ways.

Unfortunately, the debate over scientific realism has been so focused on the truth and continuity of parts of our theories or models themselves that it has failed to even consider alternative ways that the required truth and continuity might be achieved by scientific inquiry. As I noted above, in most cases the debate is put in terms of whether our best scientific theories or models are true or which parts of them are true. However, I hope to have shown that it is at least *possible* that our best scientific models are known to be wholly inaccurate and yet there can still be a kind of truth about the world that is achieved by scientific inquiry. Moreover, I think it is at least worth investigating whether the factive understanding produced by these models might be able to provide (at least some of) the continuity the realist seeks. Although our current models and theories will continually be replaced with other models and theories, perhaps the understanding we acquire from scientific models can “survive” those changes. The question then becomes *which* highly idealized models provide factive understanding and can this

understanding survive changes in the theories and models used to acquire it. In order to answer these questions, a case-by-case analysis will be required. However, it is important to note that this project is importantly different from what is traditionally focused on in the debate over realism.

Consequently, my analysis reveals a novel (although admittedly limited) form of scientific realism.¹³ This version of realism claims that even when they fail to accurately represent the features of real-world systems, highly idealized models are nonetheless capable of producing factive scientific understanding about the natural world. In other words, this version of realism claims that *science aims at factive understanding* of the natural world and that our best models (and theories) provide us with a large body of factive scientific understanding, even when they are known to be inaccurate representations of most, or perhaps even all, of the features of real-world systems.

The mistake of many accounts in the literature is to assume that the only way to argue for scientific realism is to show one of the following: that our models and theories (1) are literally true (or approximately true) as a whole (Cartwright 1983; Odenbaugh 2011); (2) are all partially true in the same way (Ladyman et al. 2007; Peters 2014; Worrall 1989); or (3) require accurate representation of the difference-making features of a target system(s) in order for us to have factive understanding of the natural world (Strevens 2009, 2013; Trout 2007). Instead, I recommend a more nuanced form of realism, which claims that highly idealized models often produce factive scientific understanding of how the natural world works—even when they do not (even aim to) accurately represent the difference-making features of their target system(s) (Bokulich 2012). Although far more will have to be said in order to adequately evaluate this view, recognizing it as a viable form of realism reveals a novel set of questions concerning a

¹³ However, this version of realism may only be applicable to model-based theorizing.

class of scientific truths that have been largely neglected by philosophical discussions of scientific realism.¹⁴ Most importantly, I contend that the realism debate needs to explore the body of factive scientific understanding that can be provided by the process of building, using, and investigating highly idealized models. The next step is to analyze additional ways that building idealized models can furnish factive scientific understanding across different contexts, modeling strategies, and disciplines.

6. Conclusion

Scientists often use multiple highly idealized models to investigate natural phenomena. This paper has investigated two ways that highly idealized models produce factive scientific understanding. In light of these cases, I have argued that models can provide factive scientific understanding of a phenomenon even when they fail to provide an accurate representation of the difference-making features of a real-world target system. In addition, I have suggested that the debate over scientific realism needs to investigate the factive understanding produced by scientists' use of multiple idealized models (and theories) rather than focusing exclusively on the accuracy of our best models (and theories) themselves. In the end, we need not believe our models are (approximately or partially) accurate representations of real-world systems in order to be justified in believing that the understanding we acquire from them accurately describes the world we live in.

¹⁴ While a complete remodeling of the debate over scientific realism may not be required, these additional questions ought to be part of the story or at least need to be investigated before we can adequately determine whether scientific inquiry can provide the kind of truth and continuity that is typically required of realism.

References

- Achinstein, P. (1983). *The Nature of Explanation*. New York: Oxford University Press.
- Batterman, R. W. (2002). *The Devil in the Details: Asymptotic Reasoning in Explanation, Reduction, and Emergence*. Oxford: Oxford University Press.
- Batterman, R. W. (2009). Idealization and modeling. *Synthese*, 169(3), 427-446.
- Batterman, R. W. and Rice, C. (2014). Minimal Model Explanations. *Philosophy of Science*, 81(3), 349-376.
- Bokulich, A. (2011). How scientific models can explain. *Synthese*, 180, 33-45.
- Bokulich, A. (2012). Distinguishing explanatory from nonexplanatory fictions. *Philosophy of Science*, 79, 725-737.
- Bueno, O., & Colyvan, M. (2011). An inferential conception of the application of mathematics. *Noûs*, 45(2), 345-374.
- Carruthers, P. (2006). *The Architecture of the Mind: Massive Modularity and the Flexibility of Thought*. Oxford: Oxford University Press.
- Cartwright, N. (1983). *How the Laws of Physics Lie*. New York: Oxford University Press.
- Churchland, P. (2013). *Touching a Nerve: Self as Brain*. New York: Norton.
- Corsano, G., Montagna, J. M., Iribarren, O., and Aguirre, P. (2009) *Mathematical Modeling Approaches for Optimization of Chemical Processes*. Nova Science Publishers: New York.
- Craver, C. F. (2006). When mechanistic models explain. *Synthese*, 153(3), 355-376.
- de Regt, H. W. (2009a). The epistemic value of understanding. *Philosophy of Science*, 76(5), 585-597.
- de Regt, H. W. (2009b). Understanding and scientific explanation. In H. W. de Regt, S. Leonelli, and K. Eigner (eds.) *Scientific Understanding*. Pittsburgh: University of Pittsburgh Press, pp. 21-42.
- Elgin, C. Z. (2007). Understanding and the facts. *Philosophical Studies*, 132, 33-42.
- Elgin, M., and Sober, E. (2002). Cartwright on explanation and idealization. *Erkenntnis*, 57, 441-450.
- Eshel I., Feldman M. W. (2001). Optimality and evolutionary stability under short-term and long-term selection. In Orzack, S.H. and Sober, E (eds.), *Adaptationism and Optimality*. Cambridge: Cambridge University Press, 114-160.
- Forber, P. (2010). Confirmation and explaining how possible. *Studies in the History and Philosophy of Biological and Biomedical Sciences*, 41, 32-40.
- Friedman, M. (1974). Explanation and scientific understanding. *The Journal of Philosophy*, 71(1), 5-19.
- Giere, R. (2004). How models are used to represent reality. *Philosophy of Science*, 64, S325-S336.
- Godfrey-Smith, P. (2006). The strategy of model-based science. *Biology and Philosophy*, 21, 725-740.
- Godfrey-Smith, P. (2009). Models and fictions in science. *Philosophical Studies*, 143, 101-116.
- Gopnik, A., and Meltzoff, A. N. (1997). *Words, Thoughts, and Theories*. Cambridge, MA: MIT Press.
- Grimm, S. (2006). Is understanding a species of knowledge? *British Journal for the Philosophy of Science*, 57, 515-535.

- Grimm, S. (2008). Explanatory inquiry and the need for explanation. *British Journal for the Philosophy of Science*, 59(3), 481-497.
- Grimm, S. (2012). The value of understanding. *Philosophy Compass*, 7(2), 103-117.
- Hartmann, A. K. and Rieger, H. (2002). *Optimization Algorithms in Physics*. Berlin: Wiley.
- Hempel, C. (1965). *Aspects of Scientific Explanation*. New York: Free Press.
- Kaplan, D. M. (2011). Explanation and description in computational neuroscience. *Synthese*, 183, 339-373.
- Kaplan, D. M. and Craver, C. F. (2011). The explanatory force of dynamical and mathematical models in neuroscience: A mechanistic perspective. *Philosophy of Science*, 78, 601-627.
- Khalifa, K. (2012). Inaugurating understanding or repackaging explanation? *Philosophy of Science*, 79, 15-37.
- Khalifa, K. (2013). The role of explanation in understanding. *The British Journal for the Philosophy of Science*, 64, 161-187.
- Kitcher, P. (1981). Explanatory unification. *Philosophy of Science*, 48, 507-531.
- Kitcher, P. (1993). *The Advancement of Science*. New York: OUP.
- Kitcher, P., & Salmon, W. C. (1989). *Scientific Explanation*. Minneapolis: University of Minnesota Press.
- Kvanvig, J. L. (2003). *The Value of Knowledge and the Pursuit of Understanding*. New York: Cambridge University Press.
- Kvanvig, J. L. (2009). Responses to critics. In A. Haddock, A. Millar, and D. Pritchard (eds.) *Epistemic value*. New York: Oxford University Press, pp. 339-353).
- Ladyman, J., Ross, D., Spurrett, D., and Collier, J. (2007). *Every Thing Must Go: Metaphysics Naturalised*. Oxford: Oxford University Press.
- Levins, R. (1966). The strategy of model building in population biology. *American Scientist*, 54, 421-431.
- Levy, A. (2011). Makes a difference. Review of Michael Strevens' *Depth*. *Biology and Philosophy*, 26, 459-467.
- Levy, A. (2012). Models, fictions, and realism: Two packages. *Philosophy of Science*, 79(5), 738-748.
- Lewis, D. (1986). *Philosophical Papers II*. New York: Oxford University Press.
- Lipton, P. (2009). Understanding without explanation. In *Scientific understanding: Philosophical perspectives*. Pittsburgh: University of Pittsburgh Press.
- Machery, E. (2009) *Doing Without Concepts*. Oxford: Oxford University Press.
- Mäki, U. (2011a). The truth of false idealizations in modelling. In Paul Humphreys and Cyrille Imbert (eds.) *Representations, Models and Simulations*, Routledge.
- Mäki, U. (2011b) Models and the locus of their truth, *Synthese*, 180, 47-63.
- Matthen, M. and Ariew, A. (2009). Selection and causation, *Philosophy of Science*, 76(2), 201-224.
- Maynard Smith, J. (1978). Optimization theory in evolution. *Annual Review of Ecology and Systematics*, 9(1), 31-56.
- Maynard Smith, J. (1982). *Evolution and the Theory of Games*. Cambridge: Cambridge University Press.
- Maynard Smith, J., & Price, G. A. (1973). The logic of animal conflict. *Nature*, 246, 15-18.
- McMullin, E. (1985). Galilean idealization. *Studies in History and Philosophy of Science*, 16, 247-273.

- Mizrahi, M. (2012). Idealizations and scientific understanding. *Philosophical Studies*, 160, 237-252.
- Morgan, M. S. and Morrison, M. (1999). *Models as Mediators: Perspectives on Natural and Social Science*. Cambridge: Cambridge University Press.
- Morrison, M. (2009). Understanding in physics and biology. In *Scientific Understanding: Philosophical Perspectives*. Henk W. de Regt, Sabina Leonelli, and Kai Eigner (eds.). Pittsburgh: Pittsburgh University Press.
- Nozick, R. (1981). *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Odenbaugh, J. (2005). Idealized, inaccurate, but successful: A pragmatic approach to evaluating models in theoretical ecology. *Biology and Philosophy*, 20, 231-255.
- Odenbaugh, J. (2011). True lies: Realism, robustness, and models. *Philosophy of Science*, 78, 1177-1188.
- Orzack, S. H., & Sober, E. (2001). *Adaptationism and Optimality*. Cambridge, UK: Cambridge University Press.
- Peters, D. (2014). What elements of successful scientific theories are the correct targets for “selective” scientific realism? *Philosophy of Science*, 81, 377-397.
- Pindyck, R. S., and Rubinfeld, D. L. (2009). *Microeconomics* (Seventh ed.). Upper Saddle River, New Jersey: Pearson Education.
- Pincock, C. (2011). Modeling reality. *Synthese*, 180, 19-32.
- Potochnik, A. (2007). Optimality modeling and explanatory generality. *Philosophy of Science*, 74(5), 680-691.
- Potochnik, A. (2009). Optimality modeling in a suboptimal world. *Biology and Philosophy*, 24(2), 183-197.
- Pritchard, D. (2009). Knowledge, understanding, and epistemic value. In A. O’Hear (ed.) *Epistemology*. Cambridge: Cambridge University Press, pp. 19-43.
- Psillos, S. (1999). *Scientific Realism: How Science Tracks Truth*. New York: Routledge.
- Psillos, S. (2011). Living with the abstract: realism and models. *Synthese*, 180, 3-17.
- Resnik, D.B. (1991). How-possibly explanations in biology. *Acta Biotheoretica*, 39, 141-149.
- Rohwer, Y. and Rice, C. (2013). Hypothetical pattern idealization and explanatory models. *Philosophy of Science*, 80, 334-355.
- Rice, C. (2012). Optimality explanations: A plea for an alternative approach. *Biology and Philosophy*, 27, 685-703.
- Rice, C. (2015). Moving beyond causes: Optimality models and scientific explanation. *Noûs*, 49(3), 589-615.
- Rice, C. (2014). Concepts as pluralistic hybrids. *Philosophy and Phenomenological Research*. DOI: 10.1111/phpr.12128.
- Saatsi, J. (2014). Inconsistency and scientific realism. *Synthese*, 91(13), 2941-2955.
- Salmon, W. C. (1984). *Scientific Explanation and the Causal Structure of the World*. Princeton, NJ: Princeton University Press.
- Salmon, W. C. (1998). *Causality and Explanation*. New York: Oxford University Press.
- Schmid-Hempel, P., Kacelnik, A., and Houston, A. (1985). Honeybees maximize efficiency by not filling their crop. *Behavioral Ecology and Sociobiology*, 17, 61-66.
- Schurz, G., and Lambert, K. (1994). Outline of a theory of scientific understanding. *Synthese*, 101(1), 65-120.
- Seger, J., & Stubblefield, J. W. (1996). Optimization and adaptation. *Adaptation*, 93-102.
- Sober, E. (2000). *The Philosophy of Biology* (Second ed.). Boulder: Westview.

- Stanford, K. (2003). No refuge for realism: Selective confirmation and the history of science. *Philosophy of Science*, 70, 913-925.
- Stanford, K., (2006). *Exceeding our Grasp: Science, History, and the Problem of Unconceived Alternatives*. Oxford: OUP.
- Stephens, D. W., & Krebs, J. R. (1986). *Foraging Theory*. Princeton: Princeton University Press.
- Strevens, M. (2009). *Depth: An Account of Scientific Explanation*. Cambridge, MA: Harvard University Press.
- Strevens, M. (2013). No understanding without explanation. *Studies in History and Philosophy of Science*, 44, 510-515.
- Suárez, M. (1999). The role of models in the application of scientific theories. In Mary S. Morgan and Margaret Morrison (eds.) *Models as Mediators: Perspectives on Natural and Social Science*. Cambridge: Cambridge University Press: pp. 168-195.
- Trout, J. D. (2007). The psychology of explanation. *Philosophy Compass*, 2, 564-596.
- Walsh, D., Lewens, T. and Ariew, A. (2002). Trials of life: Natural selection and random drift. *Philosophy of Science*, 69(3), 452-473.
- Wayne, A. (2011). Expanding the scope of explanatory idealization. *Philosophy of Science* 78: 83-841.
- Weisberg, M. (2006). Robustness Analysis” *Philosophy of Science*, 73, 730-742.
- Weisberg, M. (2007a). Three kinds of idealization. *Journal of Philosophy*, 104(12), 639-59.
- Weisberg, M. (2007b). Who is a modeler? *British Journal for the Philosophy of Science*, 58(2), 207-233.
- Weisberg, M. (2013). *Simulation and Similarity: Using Models to Understand the World*. Oxford: Oxford University Press.
- Weiskopf, D. A. (2009). The plurality of concepts. *Synthese*, 169, 145-173.
- Wimsatt, W. (2007). *Re-engineering Philosophy for Limited Beings: Piecewise Approximations of Reality*. Cambridge: Harvard University Press.
- Woodward, J. (2003). *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.
- Worrall, J. (1989). Structural realism: The best of both worlds?” *Dialectica*, 43:99-124.
- Zagzebski, L. (2001). Recovering understanding. In M. Steup (ed.) *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. New York: Oxford University Press.