Truth and Reality: How to be a scientific realist without believing scientific theories should be true

Angela Potochnik

Forthcoming in K. Khalifa, I. Lawler & E. Shech (eds.), *Scientific Understanding and Representation: Modeling in the Physical Sciences*, Routledge

Abstract

Scientific realism is a thesis about the success of science. Most traditionally: science has been so successful at prediction and guiding action because its best theories are true (or approximately true or increasing in their degree of truth). If science is in the business of doing its best to generate true theories, then we should turn to those theories for explanatory knowledge, predictions, and guidance of our actions and decisions. Views that are popular in contemporary philosophy of science about scientific modeling and the centrality of idealization create several challenges for this traditional form of scientific realism. Yet the basic idea behind scientific realism that science has been and will continue to be epistemically successful is deeply appealing. This chapter explores the challenges posed by idealization and scientific modeling to motivate a scientific realism fully divorced from the idea that science is in the business of generating true theories. On the resulting view, the objects of scientific knowledge are causal patterns, so this knowledge only ever provides partial, simplified accounts of a complex reality. This variety of selective realism better accommodates the nature of our present-day scientific successes and offers an interpretation of scientific progress that resists the antirealist's pessimism.

Scientific realism is a thesis about the success of science. Most traditionally: science has been so successful at prediction and guiding action because its theories are true (or approximately true or increasing in their degree of truth). Here's Anjan Chakravartty in the Stanford Encyclopedia of Philosophy:

> Scientific realism is a positive epistemic attitude toward the content of our best theories and models, recommending belief in both observable and unobservable aspects of the world described by the sciences… most people define scientific realism in terms of the truth or approximate truth of scientific theories or certain aspects of theories.

If science is in the business of doing its best to generate true theories, then we should turn to those theories for explanatory knowledge, predictions, and guidance of our actions and decisions.

At least at first glance, views that are popular in contemporary philosophy of science create challenges for this traditional form of scientific realism. The embrace of model-based science introduces questions about whether theories are even the main epistemic currency in science. And though in the quote above Chakravartty includes models in his formulation of realism, it's not clear how traditional realism applies to them. Models are not definitive like theories may be: it's common to employ a variety of different models to the very same phenomena. Further, the occurrence of tradeoffs among modeling priorities such as predictive accuracy vs. generality calls into question whether the same scientific accounts can even deliver explanations, predictions, and policy-guidance. Perhaps we need to employ a variety of models to accomplish all our goals. Finally, the widespread importance of idealization in both models and theories is taken by some to cast doubt on the idea that the best scientific accounts are true, approximately true, or increasing in degree of truth over time.

Yet the basic idea behind scientific realism that science has been and will continue to be epistemically successful is deeply appealing. In this chapter, I use the challenges of modeling and widespread idealization to motivate the view that scientific realism should be fully divorced from the idea that science is in the business of generating true theories. In section 1, I'll say more about these challenges to a traditional realism as I understand them, trying to base this in views broadly endorsed in our field. Then, in section 2, I'll motivate the possibility of a realism that is consistent with a rejection of the idea that science is in pursuit of truth. Again here, my aim is to carve out a general place in philosophical space, but I will also say a little about the version of this view that I have advocated. Finally, in section 3, I will address why this deserves to be called scientific realism and gesture toward what I think may be some advantages of this way of going about being a realist. Realism without a strict commitment to truth better accommodates the nature of our present-day scientific achievements, and it opens up an interpretation of cumulative scientific progress that resists the antirealist's pessimism.

1. The Trouble with Truth

I will begin by describing how views endorsed in philosophy of science put pressure on the idea, bound up with traditional realism, that science aims for truth. I'll start in the literature on scientific modeling. I consider models and theories to be two forms of what we might generically call scientific accounts. I will skirt the issue of how models and theories relate to one another, other than rejecting the idea associated with the semantic view of theories that models just are interpretations of scientific theories. Most agree that scientific models should not be judged by the metric of truth.

First of all, models have reduced ontological commitments compared to theories. As a result, there's more of an air of convenience in choice of modeling approach than in theory selection. According to Morrison (2016), for instance, while scientific theories have truth conditions that indicate what must be the case in order for them to apply, models incorporate idealized assumptions to aid in their application to specific types of systems, with the result that the models do not have straightforward truth conditions. In many fields, it is common practice to employ multiple different models to the same phenomena (Weisberg, 2013), and the same models are sometimes redeployed to very different kinds of phenomena. (See Potochnik, 2012, for example, regarding how game theory has been reinterpreted and redeployed across evolutionary biology and behavioral sciences.) Wimsatt (1987) outlines several ways in which even models recognized to be innaccurate (false, as he puts it) can nonetheless lead to truer theories, such as by suggesting refinements to existing theory, enabling scientists to localize errors to particular aspects of the model, or as a limiting model to define an extreme for the phenomenon of interest.

Second, following from the point that it is common to employ multiple models to investigate a single phenomenon, philosophers have also observed that there are often or always *tradeoffs* among the advantages of different modeling approaches. The biologist Richard Levins wrote the seminal paper on this point (1966), arguing that scientists must choose the extent to which their models prioritize accuracy versus generality and that it is often decided to sacrifice some accuracy in favor of greater generality of application. Since, Odenbaugh (2003), Matthewson and Weisberg (2009), and others have more fully worked through these ideas. I want to make two points about this in the present context. First, this provides reason to think that a single model won't accomplish all epistemic ends for a given investigation; in (Potochnik,

2015) I develop this idea with an eye to the diverse aims of science. Second, scientists regularly prioritize other values for their models other than accuracy. Some degree of accuracy might be sacrificed for other aims. Compromises in truth or accuracy are regularly made, at least in model-based science, to further other scientific goals, including applicability across related phenomena. This sacrifice in accuracy seems, at least at first consideration, to contribute to epistemic success rather than to compromise it.

Third, as Bokulich (see especially 2009, 2011) among others has emphasized, model-based explanations proceed with substantial assistance from idealizations or even fictions. If false posits or fictions facilitate at least some varieties of scientific explanations, this role is another form of sacrificing some truth to facilitate other aims. Philosophers working on scientific explanation have developed a variety of views about how to accommodate idealizations' contribution to explanation, many of which aim to maintain a commitment to the truth of successful explanations. But this contribution poses a potential challenge—even if the challenge is solvable—to the idea that scientific accounts should be (wholly) true. I will say more below to motivate the particular challenge idealizations pose to a traditional truth-centric realism.

These points about scientific models are well appreciated in philosophy of science, and they make articulating epistemic success in terms of truth less obviously the right path forward for scientific modeling compared to scientific theorizing. Insofar as model-based science is an important form of science, it seems truth, approximate truth, or increasing truth may not always be the best descriptions of the accomplishments—or even the aim—of science. Insofar as scientific models can explain, successful scientific explanations may not always be true or approximately true. Indeed, some of the philosophers with the most influential early accounts of scientific models also suggested attenuated versions of realism on that basis, including Giere's

(1988) constructive realism and Suppe's (1989) quasi-realism. Others have suggested that the discussion about scientific realism can proceed independently from considering such details of scientific modeling. Chakravartty (2001) argues that model-based science inherits the same traditional challenges facing scientific realism. Psillos (2011), in turn, argues that scientific realism is fully consistent with the use of models in science and, in particular, with models' apparent reliance on idealizations. Finally, Chakravartty, in the *Stanford Encyclopedia of Philosophy* article I quoted from above, last revised in 2017, formulates all the statements about scientific realism as claims regarding "theories and models," generically. The article does not address any promise or challenges for realism specifically related to modeling.

In the rest of this initial discussion, I will focus on the challenge posed by idealizations in particular. By idealizations, I mean assumptions made without regard for whether they are true and often with full knowledge they are false. Familiar examples include the assumptions that a gas is ideal, that a phenotypic trait reproduces to its degree of success, and that humans are perfectly rational agents. Though idealizations have perhaps received the most attention in discussions of scientific modeling, they are not just used to facilitate models but also support the application of a host of scientific laws and theories. Idealizations are widely recognized to contribute to scientific understanding; many of our most heralded scientific explanations involve idealizations. This goes beyond the model-based explanations that are Bokulich's primary focus to include theory- and law-based explanations—consider, for instance, Cartwright's (1983) decades-earlier discussion of how the laws of physics rely on idealizations. Chakravartty (2001) mostly puts the question of idealization's significance for scientific realism to one side, but he does say in passing, "Ultimately, the question of how to make sense of idealization may pose the greatest challenge to the realist" (329).[1]

Here is how the worry might be developed. Idealizations are common in our scientific accounts—including not just models but also theories and laws. As several philosophers have emphasized, idealizations persist in at least some of our best scientific accounts and contribute positively to the explanatory value of those accounts. Examples include Cartwright's (1983) seminal account of idealizations in the laws of physics, Elgin's (2004, 2017) account of how scientific understanding benefits from compromising truth, and Potochnik's (2015, 2017) argument for the widespread use of idealizations to further diverse scientific aims, including understanding different features of the same phenomena. From the broadly recognized role of idealizations in scientific explanations, these philosophers have drawn the conclusion that explanations need not be true—that the understanding they produce is non-factive. These philosophers thus reject *veritism* about scientific explanation. Of course, none of these accounts posits that we are indifferent to the truth value of our scientific accounts or prefer falsehoods to truths. The claims made are subtler, such as that idealizations contribute directly to the epistemic success of our scientific explanations, and that less accurate explanations are sometimes better than more accurate alternatives.

Put in terms of (successful) scientific explanations, these are my grounds for rejecting veritism:

- An explanation Y can explain the target phenomenon without being entirely true of it.

[1] Levy (2017) explores the prospects for realism in light of challenges from modeling and idealization. The approach I develop here is most like his characterization of potential responses to perspectivism.

- An explanation Y can better explain the target phenomenon than does an explanation Z that is more accurate of it.
- A posit P in explanation Y can be crucial for Y explaining X without being remotely true.

Here are very brief illustrations of each idea, using the tried-and-true example of the ideal gas law. An application of the ideal gas law can explain the approximate behavior of a real gas with molecules exhibiting attractive or repulsive interactions, despite the real gas's properties deviating from the ideal gas law's predictions. And, in many circumstances, the ideal gas law explanation is preferrable to a more accurate explanation with the van der Waals equation. Finally, the posit that a gas is ideal, i.e. composed of non-interacting point particles, is crucial for an ideal gas law explanation but may nonetheless deviate significantly from the properties of the real gas. (A note about example choice: I take it a toy example suffices insofar as disagreement about veritism does not hinge on whether there are examples with these features, but rather whether these are sufficient reasons to reject veritism.)

Other philosophers accept the idea that idealizations contribute positively to scientific explanation but have developed accounts that seek to retain veritism about explanation. Such accounts have been developed by Lawler (2019), Sullivan and Khalifa (2019), and Pincock (2021), among others. One motivation for developing a veritism about explanation that accommodates idealization is the perception that this will enable scientific realism to be retained (e.g. Rice, 2019). Nonetheless, for the sake of argument in the current investigation, let's take seriously the view that these features of idealization warrant reconsidering whether or to what extent truth is a requirement or potentially even an aim of our scientific accounts. There may be a way to rehabilitate a central role for truth. But on the face of it, the important roles for idealizations—false assumptions—in our scientific accounts of the world, as well as the features of model-based science discussed earlier in this section, seem to support looking for a scientific realism that does not require a commitment to the truth, approximate truth, or increasing truth of our scientific accounts. At the very least, such a realism can remain noncommittal on these questions of the extent to which truth is a benchmark for epistemic success in science, thus sidestepping the complications of scientific modeling and idealization that have received significant attention of late in philosophy of science.

2. Realism About What?

In section 1, I used scientific modeling and idealizations to motivate the idea that it might be beneficial to develop a form of scientific realism that does not require strict truth, increasing truth, or potentially even the pursuit of true scientific accounts. Now I want to outline the form such a realism might take. As above, I am aiming for carving out a general place in philosophical space, but I will also develop a specific version of the view that I would like to advocate, influenced especially by the use of idealizations in science described above.

The first task is to find a generic formulation of scientific realism that doesn't already presume a commitment to truth, then we can work out the specifics of how to make good on a commitment to such a realism. Here's Anjan Chakravartty again, a bit later in the SEP entry I drew from above: "What all of these approaches [of defining realism in terms of epistemic achievement] have in common is a commitment to the idea that our best theories have a certain epistemic status: they yield knowledge of aspects of the world, including unobservable aspects." Without too much of a stretch, we can broaden this formulation to apply to any scientific accounts used to explain our world, which might include models and laws as well as theories, and perhaps still other varieties of accounts. Then, at its most basic, realism is the idea that *our best scientific accounts qualify as epistemic achievements and yield knowledge of the world*. This idea strikes me as at the heart of a commitment to realism. This formulation also creates room to maneuver: if there is a natural way to interpret scientific accounts as epistemic achievements that yield knowledge, despite any sacrifices of truth, then that can enable a form of realism to be preserved. This approach is akin to Asay (2013), who argues against 'truth-mongering' approaches to scientific realism that locate the debate in specifics about the theory of truth embraced. Asay instead advocates a 'metaphysical' approach to realism that locates the true debate in questions of ontological commitments. Though Asay does not identify a need for distancing scientific realism from the commitment to the truth of our scientific accounts, his articulation of what should be at issue for realists versus antirealists accommodates such a distancing.

So, in what way—if any—do our scientific accounts yield knowledge bearing on observable and unobservable entities (despite potential sacrifices of truth)? On this initial question, insight can be gained from a convergence of views among recent accounts of idealization in explanation. Bokulich (2011), Batterman and Rice (2014), and I (Potochnik, 2015, 2017), among others, emphasize patterns. Here's Bokulich: "The model explains the explanandum by showing how elements of the model correctly capture the pattern of counterfactual dependence of the target system" (2011, 39). Batterman and Rice articulate the role of models in "[explaining] universal patterns across diverse real systems" (2014, 350). And, my take: "scientific knowledge consists of truths about causal patterns. Grasping those truths about causal patterns comprises understanding of the phenomena embodying the patterns. … [T]ruths of causal patterns are by and large partial untruths about phenomena, accomplished with the use of idealizations" (2017, 119). Though Elgin (2004, 2017) focuses less explicitly on patterns, her emphasis on the aim of exemplification suggests implicit reliance on pattern or a similar concept as well—instantiation of a property or relation across some range of circumstances. This suggests a potential way forward: such patterns might be the objects of our scientific knowledge. This would be a form of selective realism, perhaps akin to a kind of structural realism: realism about the relations depicted in our scientific accounts, not taken as descriptions of the targets of investigation directly but of the patterns those targets of investigation embody.[2]

The key move here is allowing non-identity between the target of investigation and the object of knowledge. You may study/learn about one thing, while the object of knowledge is technically something else. This non-identity creates two issues to address. The first issue regards the difference between the target of investigation and the object of knowledge. If the latter is not the same as the former, then what is it? Is the object of knowledge structure, patterns, or something else? Implicit in this question is the need to justify why this, whatever we settle on, should be taken to be the intended object of scientific knowledge. The second issue regards the connection between the target of investigation and the object of knowledge if we accept that this is not a relation of identity.[3] What makes this knowledge qualify as an epistemic achievement *about* the target of investigation? There must be some special relationship—if not identity—for this to be the case. So, to summarize, the two questions are:

1. **The Difference:** What is the object of knowledge if not the phenomena targeted in scientific investigation?

[2] This is particularly akin to Saatsi's (forthcoming) proposed selective realism about counterfactual dependence, as these relationships are at the heart of causal patterns.
[3] Or, at least, if we accept that this is not *always* a relation of identity; this approach only requires that objects of knowledge at least sometimes be distinct from targets of investigation.

2. **The Connection:** What makes this knowledge qualify as an epistemic achievement about the target of investigation?

In the remainder of this section, I will indicate my suggestions for how to address these two questions. Along the way, I will gesture toward allied positions.

2.1 The Difference

Above I suggested on the basis of some recent views about the explanatory value of idealizations that the object of scientific knowledge might be *patterns*. In particular, following Bokulich's and my formulations, I propose as the object of scientific knowledge patterns in counterfactual dependence—by and large, in the manipulability relations constituting causal dependence (Woodward, 2003). In (Potochnik, 2015, 2017) I termed these *causal patterns*. Postulating causal patterns rather than causal relationships or processes as the object of scientific knowledge reflects the scientific practice of strategically isolating difference-makers to indicate the scope of the dependence, i.e. the scope of what Woodward calls the invariance of the causal relationship.

Thus, what I have in mind for what I called 'the difference' is causal patterns. Idealized scientific accounts cannot generate knowledge of the targets of investigation directly insofar as they are not true of those targets. But they can and do generate knowledge of causal patterns embodied by those targets. Beyond according with the insights of several recent accounts of idealization, this also captures well science's frequent focus on general laws, patterns, and tendencies. Many causal patterns, like the pattern depicted by the ideal gas law, are broad regularities that are limited in scope and can have exceptions. On the other hand, sometimes scientists are interested in accounting for the interplay of multiple causes, which motivates a focus on a more specific causal pattern with a narrower scope, that is, present in a smaller set of phenomena. The van der Waals equation, with additional parameters specifying molecular size and attraction, depicts a more specific pattern with a narrower scope. As this illustrates, any target of investigation—such as the behavior of some real gas—embodies multiple, perhaps countless, causal patterns. The strategic isolation of a set (small or large) of causal factors gives rise to a focus on just one causal pattern among many embodied in the target of investigation. Just as it is a matter of choice which phenomena are targeted in scientific investigation, it is a matter of choice which causal pattern is targeted in the investigation of a given phenomenon.

This many-one relationship between causal patterns and target phenomena is, in my view, of central importance. Insofar as our scientific knowledge regards causal patterns, the resulting knowledge is inevitably an incomplete perspective on the target of investigation—and, when idealizations facilitate the knowledge, a perspective enabled by false posits of features of the target that are incidental to the focal pattern. Positing causal patterns as the object of scientific knowledge introduces a variability in our scientific knowledge that goes beyond what phenomena our investigations target. The nature of our scientific knowledge is, on this view, relational to the sets of causes deemed interesting by investigators or their audience.[4] This may be considered a form of perspectival realism. The view is in some respects similar to Massimi's (2019) perspectival realism, particularly her emphasis on realism regarding robust phenomena. On the other hand, for my view, the status of qualifying as knowledge is not indexed to research context, which is often thought to be entailed by perspectivism. On my approach, the role for perspective is in what knowledge is prioritized, i.e., knowledge of which causal pattern(s).

[4] This relationality is, I gather, a distinctive feature of my view (Potochnik, 2017) in comparison with the other accounts advocating a focus on patterns. See Potochnik (2020) for a discussion of this difference from Elgin's (2017) view in particular.

So, in my view, the difference between the objects of scientific knowledge and targets of investigation is that of causal patterns versus the phenomena embodying those (and other) patterns. The key difference regarding truth and knowledge is that idealizations (such as the posit that a gas is ideal) can contribute directly to knowledge of a causal pattern, even as they introduce falsehoods of phenomena. Why create this gap, what I've called the difference, at all? First, based on considerations raised in section 1, to accommodate or even take our lead from scientific practices that appear to be driven by considerations other than maximizing the truth of our scientific accounts about the phenomena they target. Second, the landing point—a form of realism that accommodates how science has developed many different accounts of the very same phenomena—is, in my view, a welcome one. And third, I'll offer a promissory note: I will argue for some further philosophical advantages to embracing this kind of approach to realism in section 3 below.

2.2 The Connection

Now that I have outlined a view for the difference between targets of investigation and objects of scientific knowledge, what is the connection? That is, why take knowledge of causal patterns to constitute an epistemic achievement *about* the phenomena targeted in scientific investigations? The brief answer is that, in my view, when a causal pattern is embodied in a phenomenon, knowledge of the causal pattern suffices to explain the phenomenon. But, given the many-one relationship between causal patterns and target phenomena, there is more that must be said about the conditions in which knowledge of the former suffices to explain the latter. Here I will sketch an approach, but it requires moves that I won't be able to make fully precise or adequately defend in this chapter. For fuller articulation and defense in another philosophical context, see (Potochnik, 2017).

Let's presume a given scientific account qualifies as knowledge of some causal pattern, since what is at issue here is how that knowledge relates to the target of investigation such that it constitutes an epistemic achievement about that target. In my view, to explain the target of investigation, such an account must relate to the target phenomenon, how that phenomenon is characterized in the explanandum, and the audience in the proper ways. First, the pattern must be embodied by the target phenomenon. This entails that the causal claim(s) are approximately true—importantly, they need not be a comprehensive account of the phenomenon but merely depict the influence of at least one causal factor. Even the same causal factor(s) may be depicted differently in characterizations of distinct causal patterns, but the claims themselves still must be approximately true (tout court, not merely within a specific research context). Second, the causal pattern must account for the explanandum, that is, the features of the target phenomenon specified in the call for explanation. Roughly put, this means that what the depiction of the causal pattern, including any idealizations, would entail is specified by the (approximately true) explanandum as having occurred. Any given causal pattern embodied in a phenomenon can only explain features conforming to that pattern, not unrelated features of the phenomenon or deviations from the pattern. Third and finally, the depicted causal pattern must address the cognitive needs of those seeking explanation. Given the many-one relationship between causal patterns and target phenomena and the inevitable relationality of depicting one causal pattern rather than another, as described in 2.1, there is an ambiguity to which causal pattern is explanatory that needs to be resolved with reference to the agenda of those seeking explanation. This is influenced by their interests and background knowledge, as well as by incidental circumstances, potential biases, and blind spots. Specifying the explanandum cannot resolve the ambiguity (Potochnik, 2016).

Our simple ideal gas law example can show what these requirements amount to. The ideal gas law succeeds as an explanation of the pressure in a balloon doubling when the volume was halved (the explanandum) when: (1) the depicted relationship between pressure and volume is roughly true; (2) this application of the ideal gas law entails the pressure roughly doubling; (3) the ideal gas law application addresses the explanation-seekers' cognitive needs. This would be the case, for example, for someone with a rough understanding of the relevant physics but who's forgotten the exact relationship between pressure and volume or what other variables, like temperature, that relationship depends on. This would not be the case for someone who well remembers the ideal gas law and how it applies here but thought the temperature had also increased to an extent that the balloon's pressure would have caused it to pop. That requires a different causal pattern, featuring different causal influences, to explain. A background assumption here is that all scientific explanations are relative to the audience's cognitive needs in this way. The shared background knowledge of a scientific community may constrain the variety of causal patterns that are explanatory, but variability will persist due to which specific questions are posed, which other phenomena are of interest, which methods are employed, and so on.

This, briefly, is what I take to be necessary for knowledge of a causal pattern to constitute an epistemic achievement—in the form of an explanation—of the target of investigation. The best explanation of a given explanandum features a causal pattern (1) embodied by the phenomenon, (2) accounting for the explanandum, that (3) best addresses the audience's cognitive needs. If increased accuracy of the phenomenon contributes to this, it does so only incidentally. More development of this view is needed, which you can find in (Potochnik, 2015, 2016). Here I am simply sketching a view with the potential to connect target of investigation to object of scientific knowledge to flesh out one candidate for the selective realism I seek to motivate in this chapter. This view creates the desired space between objects of scientific knowledge and targets of investigation. The former are multiple while the latter—empirical reality—is single; knowledge of the former inevitably yields only partial and perspectival insight into the latter.

3. A Realism Worth Having

Let's take stock. I have suggested that science achieving or even aiming for true accounts (of target phenomena) should not be a requirement for scientific realism. There may be a deep relationship between the objects of scientific understanding and the objects of scientific knowledge that is not identity. In particular, perhaps science explains phenomena by providing knowledge of (some of) the laws, patterns, and tendencies they embody. This kind of view requires specifying what the objects of scientific knowledge are, what I called the difference, and how that qualifies as an epistemic achievement *about* the targets of investigation, what I called the connection. I've sketched one way of working out those components of this form of realism: We come to understand phenomena by generating knowledge of causal patterns they embody.

This approach suggests that the question of realism might be reformulated from 'does science uncover legitimate knowledge (even about unobservables)' to 'about what, if anything, does science uncover legitimate knowledge?' What, if anything, should we be realists about on the basis of our scientific findings? I have proposed causal patterns as a candidate—in the epistemic sense of claiming that we have scientific knowledge of them and thus also in the ontic sense of positing their existence. This is one version of a selective realism. Causal patterns might well not be the *only* kind of thing we should be realists about on the basis of scientific findings; that isn't a requirement of the view I have outlined. For instance, I think some existence claims about individual events also likely qualify as legitimate scientific knowledge, such as 'life on Earth evolved from one or a few common ancestors' and 'adding the starch will turn this iodine solution blue.' In some cases, scientific explanations may also provide knowledge of target phenomena themselves. I have pointed to idealization and features of scientific modeling as reasons to doubt that, by and large, this is so; explanations that do not involve models or idealizations may not face the attendant complications. The selective realism introduced here is meant to replace a realism predicated on the requirement that scientific accounts aim for or achieve truth about the targets of investigation directly. It is meant to replace realism about the nature of phenomena depicted in science with a realism about some of the causal patterns embodied by those phenomena.

In the remainder of this last section, I want to explore some advantages to this approach to realism in contrast to a realism that posits target phenomena as the objects of scientific knowledge. To start, this realism about causal patterns better accommodates some features of our present-day scientific achievements. The motivation for this approach is accommodating scientific practices bound up with the production of scientific knowledge that don't seem to be in pursuit of truth (of target phenomena) or at least not truth alone, as detailed in section 1. Beyond that, this approach predicts a continuing plurality of scientific accounts of the very same phenomena whenever scientists are interested in different features of those phenomena rather than movement toward a single, integrated account. I take it this well describes the present situation in many fields, as evidenced by the philosophical popularity of various pluralisms. This approach reconciles this continuing plurality of accounts with both a realist perspective and the expectation of a unitary empirical reality, both of which I, at least, find to be compelling ideas, worth maintaining if possible.

I wonder whether a causal pattern realism such as I have sketched here might not also open up an interpretation of cumulative scientific progress that resists the antirealist's pessimism. If we are traditional realists, then we must conclude that science has fallen short of knowledge again and again, and that theory change doesn't by and large seem to be getting us closer to the truth. Many superseded scientific theories posit radically different views of the world; they thus must have been radically incorrect. Thus the pessimistic meta-induction: Why think our current accounts are true, when their many predecessors were not? However, if we are realists about knowledge of causal patterns, the outlook is perhaps rosier. Taking our scientific achievements to constitute objective understanding of phenomena based in knowledge of some among many causal patterns enables a subtler accounting of what past scientific achievements have had right. The relationality of knowledge of causal patterns to the cognitive needs of the investigators and their audiences enables a possible explanation for radically different accounts. Though all cannot be true, they might all achieve or at least pursue knowledge of different causal patterns. Further, the sensitivity of our stores of scientific knowledge to the cognitive circumstances of scientists and their audiences can help legitimize the social influences on past scientific findings that other realisms may take to be threats.

Of course, some scientific conjectures are simply wrong: they posit causal patterns that do not exist. But science arguably *has* latched onto some of the many causal patterns embodied in phenomena of interest. Many previously accepted theories captured causal patterns that are embodied in phenomena, even if they were later replaced by other theories that depicted other patterns, including sometimes more refined relatives of the original pattern. Causal pattern realism may be able to account for theory change resembling reduction of prior theories to their replacements in terms of related causal patterns and radical theory change in terms of transforming priorities, questions, and background assumptions. Of course, one would need to motivate the plausibility of this approach for specific instances of theory change, which is a project that is, unfortunately, not only beyond what I can accomplish in this chapter but also more or less beyond my expertise. So, my articulation of how a causal pattern realism can accommodate theory change will by necessity remain suggestive.

Stanford (2003) and Tulodziecki (2017), among others, have argued against selective realism's ability to rebut the pessimistic meta-induction. Both highlight how jettisoned theories fare poorly with regards to successful reference and approximate truth from the perspective of our current theories, which are presumed by the realist to be true. A causal pattern realism such as I have described, however, does not posit the truth of our current best theories—at least not their truth of the targets of investigation. Because different accounts of the same phenomena may target distinct and even unrelated causal patterns, misconstruing other features of the target phenomena along the way, these accounts may emphasize different aspects and even posit different entities. This does not give us grounds for claiming that a past theory's successes will be conserved from the perspective of current theory but rather grounds for suspecting that past scientific accounts succeeded at generating scientific knowledge of causal patterns apt for the epistemic and social locations in which they were formulated. There can be different and even incompatible ways to define causal variables that give rise to patterns of counterfactual dependence invariant over some range of conditions. Note that this does leave intact the ability to diagnose shortcomings—empirical and programmatic—in past research. One may still complain that this suggests too low a bar for a proper realism. I intend attenuated commitments to truth or attempted truth to be a feature of the view, so I accept that the resulting realism—and its strategy for interpreting past scientific successes—might strike some as insufficiently yoked to accuracy.

I propose that the same approach I have suggested for accommodating past scientific change suffices also as a basis for confidence in the epistemic success of our present scientific accounts. Our current best theories are probably not true. Some of our current best theories even contradict one another. All of those theories, along with scientific models and perhaps other forms of scientific accounts, stand a good chance of being jettisoned or significantly altered sometime in the future. But even if that comes to pass, it need not undermine their contribution to human understanding. If humans now use these accounts to yield knowledge of real causal patterns that satisfy our cognitive needs, then it does not matter if they aren't strictly true, or turn out to not be true, of the phenomena they target. Some postulated causal patterns may turn out not to obtain at all, but many more are real—revealing important aspects of the world, even as they are incomplete and shaped by our present epistemic and social locations.

Whatever knowledge science generates is in response to our cognitive needs. That much is uncontroversial, I think. And our cognitive needs regularly motivate simplified accounts of our complex world. Simplified accounts privilege some aspects of phenomena to draw connections with other related phenomena, while neglecting other aspects. This is one reason for the widespread value of idealizations in scientific accounts. In this chapter, I have argued that these simplified accounts can qualify as knowledge even if they are not strictly true of phenomena insofar as they are knowledge of something else. It has often been said that science focuses on laws, regularities, and repeat phenomena rather than one-off events. It's not a far stretch to suggest that science achieves an understanding of phenomena—the events and occurrences of our world—by generating knowledge of some of the laws, patterns, and tendencies these phenomena embody. On the view I have developed, the resulting scientific knowledge is not technically of the targets of scientific investigation but of some of the causal patterns they embody. This causal pattern realism enables scientific realism to be maintained regardless of scientists' habit of sacrificing truth of phenomena for other epistemic priorities, and it holds promise for a realism-friendly interpretation of past scientific change and present scientific knowledge. Science is always in revision, but along the way, it is amassing ever more understanding via knowledge of causal patterns.

Acknowledgements

Thanks to Kareem Khalifa, Insa Lawler, and Elay Shech for organizing this volume; to Chris Pincock for fruitful exchanges and helpful feedback on this chapter; and to Insa Lawler, Kareem Khalifa, Michela Massimi, Aaron Novick, and the Caltech philosophy of physics reading group as well for their feedback and suggestions.

References

Asay, J. 2013. "Three Paradigms of Scientific Realism: A Truthmaking Account." *International Studies in the Philosophy of Science*, 27: 1-21. https://doi.org/10.1080/02698595.2013.783971

Bokulich, Alyssa. 2009. "Explanatory Fictions," in *Fictions in Science: Philosophical Essays on Modeling and Idealization*, Mauricio Suárez Ed. Routledge Studies in the Philosophy of Science. New York: Routledge.

Bokulich, Alyssa. 2011. "How Scientific Models Can Explain," *Synthese* 180:33-45. https://doi.org/10.1007/s11229-009-9565-1

Cartwright, Nancy. 1983. *How the Laws of Physics Lie*. Oxford: Oxford UP. https://doi.org/10.1093/0198247044.001.0001

Chakravartty, Anjan. 2001. "The Semantic or Model-Theoretic View of Theories and Scientific Realism," *Synthese* 127:325-345.

Elgin, Catherine. 2004. "True Enough," *Philosophical Issues*. 14: 113-131. https://doi.org/10.1111/j.1533-6077.2004.00023.x

Elgin, Catherine. 2017. True Enough. Boston: MIT Press. https://doi.org/10.7551/mitpress/9780262036535.001.0001

Giere, Ronald N. 1988. *Explaining Science - A Cognitive Approach.* Chicago: University of Chicago Press. https://doi.org/10.7208/chicago/9780226292038.001.0001

Lawler, Insa. 2021. "Scientific Understanding and Felicitous Legitimate Falsehoods," *Synthese* 198 (7): 6859-6887. https://doi.org/10.1007/s11229-019-02495-0

Levy, Arnon. 2017. "Modeling and Realism: Strange Bedfellows?," in *Routledge Handbook of Scientific Realism*, Juha Saatsi Ed. New York: Routledge. https://doi.org/10.4324/9780203712498-20

Massimi, Michela. 2019. "Realism, Perspectivism, and Disagreement in Science," *Synthese*. https://doi.org/10.1007/s11229-019-02500-6

Matthewson, John and Michael Weisberg. 2009. "The Structure of Tradeoffs in Model Building." *Synthese* 170: 169-190. https://doi.org/10.1007/s11229-008-9366-y

Morrison, Margaret. 2016. "Models and Theories", in *Oxford Handbook to Philosophy of Science*, P. Humphreys (ed.), Oxford: Oxford UP. https://doi.org/10.1093/oxfordhb/9780199368815.013.32

Odenbaugh, Jay. 2003. "Complex Systems, Trade-offs, and Theoretical Population Biology: Richard Levin's ‘Strategy of Model Building in Population Biology’ revisited," *Philosophy of Science* 70: 1496-1507. https://doi.org/10.1086/377425

Pincock, Christopher. 2021. "A Defense of Truth as a Necessary Condition on Scientific Explanation." *Erkenntnis*. https://doi.org/10.1007/s10670-020-00371-9

Potochnik, Angela. 2012. "Modeling Social and Evolutionary Games." *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences* 43:202-208. https://doi.org/10.1016/j.shpsc.2011.10.035

Potochnik, Angela. 2015. "The Diverse Aims of Science." *Studies in History and Philosophy of Science* 53: 71-80. https://doi.org/10.1016/j.shpsa.2015.05.008

Potochnik, Angela. 2015. "Causal Patterns and Adequate Explanations." *Philosophical Studies* 172: 1163-1182. https://doi.org/10.1007/s11098-014-0342-8

Potochnik, Angela. 2016. "Scientific Explanation: Putting Communication First." *Philosophy of Science* 83: 721-732. https://doi.org/10.1086/687858

Potochnik, Angela. 2017. *Idealization and the Aims of Science.* Chicago: University of Chicago Press. https://doi.org/10.7208/chicago/9780226507194.001.0001

Potochnik, Angela. 2020. "Review of Catherine Elgin's *True Enough*." *Philosophical Review* 128: 363-366. https://doi.org/10.1215/00318108-7537361

Psillos, Stathis. 2011. "Living with Models," *Synthese* 180: 3-17. https://doi.org/10.1007/s11229-009-9563-3

Rice, Collin. 2019. "Understanding Realism," *Synthese* 198 (5): 4097-4121. https://doi.org/10.1007/s11229-019-02331-5

Saatsi, Juha. forthcoming. "Realism and Explanatory Perspectivism," in *Understanding Perspectivism: Scientific Challenges and Methodological Prospects*, M. Massimi & C.D. McCoy (Eds.), New York: Routledge.

Stanford, Kyle. 2003. "Pyrrhic Victories for Scientific Realism," *The Journal of Philosophy* 100: 553-572.

Sullivan, Emily and Kareem Khalifa. 2019. "Idealizations and Understanding: Much Ado About Nothing?," *Australasian Journal of Philosophy* 97: 673-689. https://doi.org/10.1080/00048402.2018.1564337

Suppe, Frederick. 1989. *The Semantic Conception of Theories and Scientific Realism.* Chicago: University of Illinois Press.

Tulodziecki, Dana. 2017. "Against Selective Realism(s)," *Philosophy of Science* 84. https://doi.org/10.1086/694004

Weisberg, Michael. 2013. *Simulation and Similarity: Using Models to Understand the World*. Oxford: Oxford UP. https://doi.org/10.1093/acprof:oso/9780199933662.001.0001

Wimsatt, William. 1987. "False Models as Means to Truer Theories," in *Neutral Models in Biology*. Oxford: Oxford UP.

Woodward, J. 2003. *Making Things Happen: A Theory of Causal Explanation.* Oxford: Oxford UP. https://doi.org/10.1093/0195155270.001.0001