

The Diverse Aims of Science

Angela Potochnik*

Abstract

There is increasing attention to the centrality of idealization in science. One common view is that models and other idealized representations are important to science, but that they fall short in one or more ways. On this view, there must be an intermediary step between idealized representation and the traditional aims of science, including truth, explanation, and prediction. Here I develop an alternative interpretation of the relationship between idealized representation and the aims of science. I suggest that continuing, widespread idealization calls into question the idea that science aims for truth. If instead science aims to produce understanding, this would enable idealizations to directly contribute to science's epistemic success. I also use the fact of widespread idealization to motivate the idea that science's wide variety aims, epistemic and non-epistemic, are best served by different kinds of scientific products. Finally, I show how these diverse aims—most rather distant from truth—result in the expanded influence of social values on science.

Keywords: idealization, understanding, tradeoffs, social values

The important role that models play in science has, in the past decades, been increasingly appreciated by philosophers. This attention to scientific modeling has, in turn, led to an

*Department of Philosophy, University of Cincinnati, 206 McMicken Hall, Cincinnati, OH, 45221-0374.
Email: angela.potochnik@uc.edu; Telephone: 513-556-6324

emphasis on the centrality of idealizations. As Wimsatt (1987, 2007) says, “Any model implicitly or explicitly makes simplifications, ignores variables, and simplifies or ignores interactions among the variables in the models and among possibly relevant variables not included in the model (p.96).”¹ These are all idealizations. Most broadly, idealizations are features of representations that misconstrue the represented systems. Examples include the common assumption in physics of frictionless planes and the common assumption in economics that humans are perfectly rational agents. These assumptions are false of every real system: every plane has friction, and no human is perfectly rational. Assimilating several views about the nature of idealization, including Wimsatt’s, Weisberg (2007, 2013) identifies three distinct purposes to which idealizations are put. These include Galilean idealizations, which are simplifications needed to secure computational tractability, to be eliminated if and when it proves possible; minimalist idealization, which is the elimination of all but the most significant causal influences on a phenomenon; and multiple-models idealization, which is the use of several distinct models that together shed light on a phenomenon. Rohwer and Rice (2013) argue that the roles of idealization are even more varied still.

One common view is that all of this idealization may be necessary, but it results in models that are lacking in various ways. Accordingly, the view goes, we must look for a subsequent step, a way to connect these idealized models to the successful pursuit of the aims of science, whether the specific aim is prediction, empirical confirmation, explanation, accurate representation, etc. The textbook version of this view would hold that science aims for truth, and so idealized models must be de-idealized in order to be useful. It seems that Odenbaugh and Alexandrova (2011) assume something like this view, for they argue that without the removal of all idealizations—complete de-idealization—we have “no ground, beyond that of our background knowledge that informed the model, for claiming that the

¹(Wimsatt, 1987) is republished as Chapter 6 in (Wimsatt, 2007). Page numbers here refer to the latter publication.

model specifies a causal relation” (p.765). Odenbaugh and Alexandrova conclude that even the use of multiple models with different idealizations—i.e. robustness analysis—cannot yield the description of a causal mechanism. Thus, they claim, analysis based on multiple, idealized models does not allow for the confirmation of models, nor can it generate explanations.

Other versions of this view do not hold de-idealization to be necessary but still anticipate the need to bridge the gap between idealized models and the traditional aims of science. Wimsatt (2007), for instance, argues that idealized, “false” models can be used to produce “truer” theories without recourse to de-idealization. Similar to Odenbaugh and Alexandrova’s concern with causal description and explanation, Rohwer and Rice (2013) argue that at least one purpose of idealizations, the investigation of general patterns across heterogenous systems, prevents the accurate description of causal factors and, thus, prevents the formulation of explanations (though they hold that resultant models may still be explanatory in some weaker sense). This style of view endorses the continuing practice of idealization, but also holds idealized models to be somewhat distant from the traditional aims of science. These authors accordingly explicitly or implicitly commit themselves to an intermediary step of some kind between idealized representation and achieving the aims of science. On this strategy, even though idealized models are of scientific value, they are not sufficient to provide adequate explanations, trustworthy predictions, causal representations, etc.—or at least not by themselves.

One could instead take a very different approach to reconciling idealization with the aims of science. The observation of widespread idealization in science, and the distance between idealized models and traditional articulations of the aims of science, might be seen as grounds for concluding that those traditional articulations of the aims of science are incorrect. On this alternative approach, nothing has gone wrong with or is lacking from idealized models, and no intermediary step is needed for idealized models to achieve the aims of science. Those aims just stand in need of clarification. This is the tack I take in this paper. In §1 I develop

a conception of science's epistemic aim to which idealized models can directly contribute. In particular, I suggest that science does not aim to provide truth, but instead to provide understanding. In §2 I outline a second consequence a positive conception of widespread idealization might be seen to have for the aims of science. Science's diverse aims, both epistemic and non-epistemic, often conflict and thus motivate different kinds of scientific products. Finally, in §3 I demonstrate that this alternative conception of science's epistemic aim and of the relationships among science's various aims creates new room for the influence of social values.

1 Understanding at the Expense of Truth

Wimsatt (2007) points out, regarding idealized models, that “unless they could help us do something in the task of investigating natural phenomena, there would be no reason for choosing model building over astrology or mystic revelation as a source of knowledge of the natural world” (p.101). This must be right. Idealized models, even though they are false in some regards, must get us somewhere that mystic revelation does not. At issue is what exactly idealized models are helping us accomplish, and in particular, the nature of their epistemic value. Here I will explore the idea that false models are not a means to truer theories, as Wimsatt believes, but instead themselves accomplish the end goals of science, including its epistemic success. I do not provide a conclusive argument in favor of understanding and against truth as science's epistemic aim. Instead, in what follows, I distinguish the aim of understanding from the aim of truth; motivate the former; and show how this enables widespread idealization to directly contribute to science's epistemic success. Because idealizations are patently untrue, their continued presence in models cannot be justified by their contribution to the truth of those models. Accordingly, if idealizations directly contribute to science's epistemic success, then this suggests the epistemic aim is

something other than truth.

A first step toward a conception of the epistemic aim of science to which idealized models can directly contribute is provided by Elgin (2004). Elgin is also impressed by how many scientific laws, models, and theories diverge from the truth in various ways. Her aim is thus to show how these scientific products can be epistemically acceptable without being entirely true. She says,

I take it that science provides an understanding of the natural order. By this I do not mean merely that an ideal science *would* provide such an understanding or that in the end of inquiry science *will* provide one, but that much actual science has done so and continues to do so (p.114, emphasis in original).

Elgin's strategy is to accept today's actual science as a successful venture, then look to see what this science accomplishes. What she finds is that science regularly produces understanding, even as it falls short of producing truths. Accordingly, rather than make excuses for the myriad ways in which our science fails to produce truth, Elgin redefines science's epistemic success to consist of understanding, not (necessarily) truth.

For this approach to have promise, it must be possible for the achievement of understanding to occur without the possession of complete truth, but understanding must still qualify as an epistemic success. A key feature of the concept of understanding enables it to play this role: it has a dual nature. Understanding is at once a cognitive state and an epistemic achievement. Because it is an epistemic achievement, understanding is not simply an "aha" moment. A felt sense of understanding is not sufficient for the possession of understanding; understanding requires successful mastery, in some sense, of the target of understanding. Both Grimm (2010, 2012) and Strevens (2013) describe this mastery as a form of grasping. As Grimm (2012) stresses, "grasping" is a success term. And so, "the mind of someone who understands mirrors or reflects reality" (Grimm, 2012, p.109). There is some debate over whether understanding is a species of knowledge (see, e.g., Grimm,

2006), but it is widely regarded as an epistemic success of some kind, and I take it to be as well. On the other hand, because understanding is also a cognitive state, its achievement partly depends on the psychological characteristics of those who seek to understand. This distinguishes understanding from truth, for whether a proposition is true in no way depends on the psychology of one who entertains or believes that proposition. Understanding's characteristics qua cognitive state accounts for how idealized representations can be well-positioned to provide understanding.

For Grimm (2012), you cannot have objective understanding without possessing truth (though in his view this is not understanding's full epistemic value). In contrast, Elgin (2004) suggests that understanding may be furthered by some departures from the truth. In her view, "felicitous falsehoods," or idealizations, can facilitate understanding insofar as they "impose an order on things, highlight certain aspects of the phenomena, reveal connections, patterns and discrepancies, and make possible insights that we could not otherwise obtain" (p.127). Elgin gives the example of drawing a smooth curve and treating the data's deviation from the curve as error or noise. This idea—that departures from truth can contribute to understanding by revealing patterns and enabling insights that would otherwise be inaccessible—is corroborated by research in psychology. Williams et al. (2013), for instance, provide empirical support for the idea that human understanding is furthered by generalizations about broad patterns.² Such generalizations are facilitated by employing idealizations. The idealization that humans are perfectly rational agents enables generalizations across diverse decision-making processes. The common idealization in population biology that a population is infinite in size plays a similar role. This enables a model to represent populations of different sizes, and to ignore the causal role of that

²More specifically, according to this research, seeking and producing explanations furthers understanding in virtue of an emphasis on generalizations about broad patterns, and this same process can sometimes lead to overgeneralization. Here I set aside the connection between understanding and explanation in order to narrow the discussion, although the considerations raised here also have interesting implications for accounts of scientific explanation.

population size, i.e., genetic drift, a kind of sampling error.

Though allowing some departures from the truth is a natural attribute for understanding qua psychological state, this creates trouble for the other part of understanding's dual nature, understanding qua epistemic success. Elgin (2004) has a solution on offer. She proposes that truth is in fact still a requirement on understanding, but it functions as a threshold concept. A claim must be "true enough" in order to be epistemically acceptable; that is, any divergence from the truth must be negligible, or safely neglected. Whether a claim is epistemically acceptable depends on the function it plays in an argument, explanation, or theory—or, one might further generalize, in a model or other representation. Elgin appeals to the example of Snell's law, which governs light's angle of refraction when it passes from one medium to another. Snell's law is only true of optically isotropic media, but it is true *enough* of media that are nearly isotropic, which include a wide range of media in which physicists are interested.

I follow Elgin in taking truth to be a threshold requirement for epistemic acceptability, but in my view, one amendment is needed. Whether a claim is true enough depends not only on its function in an argument, explanation, theory, or representation, as Elgin holds. It also depends on the purpose of the research. Elgin points out that Snell's law applied to anisotropic media is of limited use "if we are interested only in the path of a particular light ray," but is useful "if we are interested in optical refraction in general" (p.118). This divergence in whether or not Snell's law is true enough cannot be accounted for solely by referencing the role the assumption of isotropic media plays in the law. It also requires reference to the researchers' different aims, that is, what they intend to get out of an application of the law. This is a minor expansion of the recognized influences on a claim's epistemic acceptability, but it hints at a more significant shortcoming of Elgin's conception of understanding that will be addressed below.

When the achievement of understanding is distanced from truth, it must also be distanced

from belief; otherwise one would be in the position of endorsing the epistemic value of false beliefs. Indeed, idealizations are generally regarded to be false. No biologist actually *believes* she is dealing with an infinite population of organisms. Accordingly, as Elgin (2004) puts it, “understanding is often couched in and conveyed by symbols that are not, and do not purport to be, true” (p.116). Likewise Winsberg (2010) argues that the proper attitude to have toward “successful model-building principles” (what I would call idealizations) is not holding them to be true or even approximately true (p.134).³ Elliott (2013) and Elliott and Willmes (2013) argue that scientists adopt a number of different cognitive attitudes other than belief toward the products of science. They follow Cohen (1992) in distinguishing between accepting and believing a body of content, such as a hypothesis, theory, model, or other representation. On this view, contributors to understanding, including idealizations, must be accepted—and must be epistemically acceptable—but need not be believed.⁴

I have characterized understanding as both an epistemic achievement and a cognitive state, with significant features deriving from each of these attributes. I also have adopted Elgin’s view that understanding requires only some variable threshold of approximate truth be met, and I have suggested that this threshold depends in part on scientists’ immediate interests. This characterization distinguishes understanding from truth, thereby setting it up as a proper alternative for the epistemic aim of science. It also demonstrates how idealizations can directly contribute to understanding. Idealizations facilitate understanding when they enable the representation of cognitively valuable connections and patterns that more accurate portrayals would miss.

There are two limitations of the conception of idealization’s value developed so far that

³Winsberg argues that these model-building principles with which he is concerned go beyond idealization and are best described as fictions.

⁴There is a question regarding the relationship between this broader concept of epistemic acceptability and knowledge. In particular, one may well wonder whether the epistemic acceptance required for scientific understanding is sufficient to produce scientific knowledge. I sidestep this difficult question here. Let me say, though, that I do not expect this account of science’s epistemic aim to constitute a commitment to scientific antirealism.

must be addressed. First, so far I have exclusively focused on science's purely epistemic role, that is, the production of understanding. But science plays a wide range of roles. Some are epistemic, such as understanding; others non-epistemic, such as action within a short timespan; and still others seem to involve both epistemic and non-epistemic elements, such as accurate prediction (cf. Elliott, 2013). Instead of a single successor aim for truth, we should expect a variety of scientific aims, which suit science to the range of roles it plays. Indeed, the list of science's aims must be open-ended, for science is a continually creative process. Its procedures as well as its products are always in development. Surely idealized representation contributes to these diverse aims of science in a variety of ways, not just in virtue of its contribution to the narrowly epistemic goal of creating understanding. I say more about how idealized representation relates to science's other aims in the next section.

The second regard in which the view I have sketched is lacking is that, following Elgin (2004), it preserves a role for truth that is both too central and too uniform. As a result, this view does not fully accommodate the ways in which idealizations facilitate understanding. Consider what Elgin says about the epistemic value of the ideal gas law:

The model is illuminating though, because we understand the properties of real gases in terms of their deviation from the ideal. In such cases, understanding involves a pattern of schema and correction. We represent the phenomena with a schematic model, and introduce corrections as needed to closer accord with the facts (p.127).

Phrasing this as schema and correction sounds awfully like the traditional view of idealization's contribution that I outlined above, according to which idealizations are distortions, to be overcome or circumvented in the pursuit of truth. This is quite similar to Elgin's description of an idealized schema, followed by corrections in order to "closer accord with the facts." But the use of idealized representations like the ideal gas law and Snell's law (discussed above) does not conform to this picture of schema and correction.

In many circumstances the idealized representation alone, without any corrections, is employed. For example, the ideal gas law remains central, even though there is a more accurate alternative—van der Waal’s equation, which takes into account molecular size and intermolecular attraction, properties the ideal gas law simply idealizes away.

Part of the difficulty here stems from Elgin’s presupposition that all of science generates claims that figure into arguments, explanations, or theories. This is wrong. Modeling in science may proceed largely independent of theory, and there is a diverse array of scientific products, many of which have little or no relationship to theory or explanation. But a broadened conception of scientific products requires that Elgin’s definition of “true enough” be significantly revamped. Her proposed standard of negligible divergence from truth, taking into account a claim’s role in an argument, explanation, or theory, is still too truth-conservative. This can be seen from Elgin’s treatment of the ideal gas law, as quoted above. For many scientific projects, an uncorrected, schematic—that is, heavily idealized—representation suffices. Such projects would be hindered, not furthered, by closer accordance with the facts. Elgin also says, “if, for example, evidence shows that friction plays a major role in collisions between gas molecules, then unless compensating adjustments are made elsewhere, theories that model collisions as perfectly elastic spheres will be discredited” (p.129). But this too is wrong. There may be perfectly good reasons, including epistemic reasons, to continue to model gas-molecule collisions as if they were collisions among perfectly elastic spheres, even if the *theory* that they are similar to perfectly elastic spheres is discredited. This would be like biologists modeling relatively small populations as if they were infinite, which enables the neglect of genetic drift when other causal influences are focal. Significant causal influences are often idealized in just this way.

This suggests that what best facilitates understanding is not determined solely by the relationship between a representation and the world. Instead, what representations best facilitate understanding depends also on a range of considerations about scientists

themselves. Those considerations include, prominently, scientists' particular research interests. The idealization of a population as infinite in size or a gas as comprising non-interacting point particles is appropriate when the focus of research is elsewhere. These idealizations are inappropriate, that is, unhelpful, when researchers want to investigate the role of genetic drift or intermolecular attraction. A range of other considerations about scientists themselves are also potentially relevant, including their cognitive faculties, psychological characteristics, temporal and spatial location, etc. Because understanding is a cognitive state, its achievement depends in part on the characteristics of those who seek to understand. I have suggested that idealizations facilitate understanding in virtue of the contributions broad patterns make to human understanding. It is also the case that a human scientist must be in the right position for a representation to enable her to understand some phenomenon. She must have the proper background information, the proper type of question, and she must have the proper aims.

Noting these two limitations demonstrates the features that a reconstrual of the aims of science must possess in order to accommodate the positive contributions of widespread idealization. First, it must be acknowledged that science has a variety of aims. Understanding may replace truth as the purely epistemic aim of science, but it is not the ultimate aim of science simpliciter. Accordingly, the aim of understanding does not alone account for all uses of idealizations, just their narrowly epistemic value. Second, even within the epistemic aim of understanding, the research purpose to which an idealized representation is put must be taken into account in the determination of whether an idealization promotes understanding. Features of scientists themselves, including their interests and intentions, influence what generates understanding. There is no lower limit on when an idealization qualifies as "true enough." Idealizations, no matter how little they resemble the systems they represent, may facilitate an understanding of those systems in the proper circumstances.

This last point deserves a bit more discussion. It is significant that an idealization may be

radically untrue, that is, quite different from the true state of affairs, but nonetheless facilitate understanding. A radically untrue idealization may actually facilitate understanding *in virtue of* its distance from the truth. As Strevens (2008) discusses, this style of idealization can advertise the causal irrelevance of the idealized factor. Elgin (2004) seems to have something like this in mind as well when she discusses idealizations' use in "a fortiori arguments from limiting cases" (p.119). The assumption of an infinite population size can be a way to show that the actual size of the population is causally irrelevant, that is, that the population is large enough—close enough to infinite—for its actual size to be completely ignored. But this does not exhaust the contribution to understanding of radically untrue idealizations. I suggested above that an infinite population is sometimes assumed when the population size is small, and the role of genetic drift is significant. This occurs when research interests focus on the causal role of other factors, like natural selection, to the exclusion of a focus on genetic drift. Put generally, the point is that idealizations that are far from the actual state of affairs are warranted, even when they stand in for significant causes, when the idealized features are not central to what scientists seek to understand. These idealizations do not facilitate understanding in virtue of their relationship to the world, that is, to the causal facts, but in virtue of their relationship to the research interests of those seeking to understand.

In summary, I have suggested that idealizations directly contribute to the production of scientific understanding. If understanding is taken to be the epistemic aim of science, this would vindicate the continuing prominence of idealized models in science. Positing understanding as the epistemic aim of science is, then, a version of the approach I adopted at the outset, according to which widespread idealization is used to motivate a conception of the aims of science to which idealizations can directly contribute. Because idealizations are patently untrue, I follow Elgin (2004) in distancing the achievement of understanding from the achievement of truth. As Elgin says, representations, including their idealized

assumptions, must simply be “true enough,” where this can be quite far from actual truth. I have argued that features of human scientists, including their specific research interests in particular instances, also influence what counts as true enough. On this view idealizations are not preparatory to truer theories; instead they are valuable for their direct contribution to understanding. As a result, many of the immediate products of science are not things we believe to be true.

I conclude this section by addressing a few possible concerns with this view. First, notice that truth does not disappear entirely from this picture. An idealization may fail to qualify as true enough—that is, it may not be accurate enough of the world to serve its intended purpose. In this case, the idealization will fail to facilitate understanding. If the causal role of genetic drift is under investigation, the assumption that a medium-sized population is infinite is inaccurate in a way that undermines the aim of research. Accordingly, this assumption is simply not made in that research context. This is one regard (among many, I expect) in which mystic revelation and astrology fall short. Consider astrology. No matter what specific research project is pursued, information about celestial bodies simply does not yield insight into human traits and actions. Astrology asserts that there are such causal relationships, and those assertions have been empirically demonstrated to be false. Falsity is inconsistent with any intended purpose of an assertion (of fact). In contrast, as we saw above, idealizations such as an infinite population are not believed nor asserted, but assumed contrary to fact. They are, then, epistemically acceptable even though untrue.

There is another regard in which truth remains relevant. I have suggested that understanding is a candidate for the purely epistemic aim of science, but that science has a range of other aims, some non-epistemic and some partly epistemic. Some of those aims might best be served by true claims. This is particularly plausible for the aim of accurate prediction. It seems that accurate predictions are best provided by true assertions about

states of affairs.⁵ In this case, though, the true claim is in service of a distinct scientific aim. Perhaps the aim of prediction is an epistemic aim in its own right? I take it that prediction is traditionally thought to have derived epistemic value (if any), and so it seems strange to consider prediction to be an epistemic aim on its own. But nothing is at stake for me in that judgment. If one holds prediction to constitute an epistemic aim, then the position developed here is that understanding is *an* epistemic aim of science, alongside prediction (but not truth more generally).

But is truth-in-general not also the epistemic aim of science, at least sometimes? Let us distinguish between three possibilities. First, perhaps sometimes straightforwardly true claims also are best at generating understanding. In these cases, it is impossible to distinguish the aim of truth from the aim of understanding, for the same scientific products accomplish both. Examples might include straightforward existence claims, like that Saturn has a moon larger than the planet Mercury, or that life on Earth has evolved from one or a few common ancestors. Notice, though, that it is still possible to describe the prevailing epistemic aim as understanding. If this interpretation is adopted, then these cases are naturally described as instances in which departures from truth do not further understanding. Second, I have said that general patterns are valuable for human understanding, but perhaps sometimes scientists value more accurate representations over representations of general patterns. This must be right: sometimes the exceptions or the individual instances are central to research interests. But this too accords with the view of understanding as science's epistemic aim I have put forward. Whether a representation, or individual assumption, facilitates understanding depends in part on the nature of the understanding sought. This builds in variability to accommodate the wide range of focuses for scientific understanding. Third and lastly, consider the possibility that sometimes true scientific products are developed at the expense of understanding, that is, that produces *less* understanding of a focal

⁵Thanks to an anonymous referee for pointing this out.

phenomenon in order to be more veridical. I can imagine this when there is some other aim at play, such as prediction, but not otherwise. Of course, this might simply be due to my impoverished imagination; I would happily entertain possible counterexamples. But I would wager that, when the world's complexity and variability inhibits representations that are both accurate and comprehensible, comprehension (or predictability, or etc.) always trumps simple accuracy.

One final concern is that this approach seems to provide a terribly subjective standard for science's epistemic success. There is one subjective element: what best facilitates understanding depends in part on the features of the practitioners and consumers of science. But this element of subjectivity is unproblematic. Indeed, it is to be expected for an account of a scientific enterprise developed by limited and historically located human beings. What best meets human requirements—including our epistemic requirements—depends not just on features of the world we investigate but also, in part, on the features of human investigators. Other forms of subjectivity that would be problematic do not apply. One need not rely on the subjective experience of an “aha!” moment, viz., the felt *sense* of understanding (Trout, 2002), in order to judge whether an idealization facilitates understanding. Instead, objective reasons guide that judgment. These reasons might include limited computational power or absolute computational limits; the limitations of human powers of cognition or the limitations in one researcher's training; the existence of a causal pattern or the existence of specific research interests. All are objective features of the world under investigation or of the scientists leading that investigation that together determine what best generates human understanding.

2 The Aims of Science in Tension

If the position developed above is right, the epistemic aim of science is not truth, but understanding. Yet, as I have already acknowledged, there are a variety of aims of science, both epistemic and non-epistemic. Traditionally appreciated aims include (at least) accurate prediction, explanation, and representation. Other aims of science have recently received increasing attention. These include providing information to guide policymaking (Douglas, 2009b); action within a short timespan (Elliott, 2013); and facilitating the public uptake of scientific knowledge (Elliott, 2011). There are surely many other aims and, indeed, many other existing articulations of these and other aims. My project here is not to delineate the range of scientific aims; above I suggested that we should expect an open-ended list of aims. My goal is instead to examine the relationship among the various aims of science. In particular, I suggest that it is the norm for the pursuit of one aim to occur at the expense of others. Successful pursuit of one among the various aims of science generally inhibits success with other aims. Accurate prediction is achieved by tools poorly suited to explain; the aim of quick action is at odds with full causal representation; etc. At root, this is because the different aims of science are furthered by different means. This too traces back to the centrality of idealizations in science.

Notice first that the diversity of scientific aims is linked also to a diversity of cognitive attitudes toward the products of science. Recall from above that embracing understanding, instead of truth, as the epistemic aim of science requires shifting from a focus solely on belief to the broader concept of acceptance. This is because many aids to understanding are not things we believe to be true. I followed Elgin in phrasing the alternative, broader concept “epistemic acceptance.” But acceptance, like the aims of science themselves, comes in many varieties. Elliott and Willmes (2013) define acceptance as follows: “*S* accepts that *h*, iff *S* presupposes *h* for specific reasons in her deliberation” (p.5). This definition can

yield the subspecies of epistemic acceptance by restricting the relevant specific reasons to the furtherance of understanding. In this way, one can generate other varieties of acceptance by focusing on other, specific aims of science. For instance, acceptance as predictively useful amounts to presupposing for purposes of predictive value (alone). It is clear that one form of acceptance need not entail another. Epistemic acceptance does not entail acceptance as true (i.e. belief), nor does it entail acceptance as predictable useful. The value of an assumption in the production of understanding does not entail that it is true, nor does it entail its predictive usefulness.

Just as one form of acceptance need not entail another, the pursuit of one aim of science need not contribute to other aims. Believing a true claim might be useful for a variety of other aims, such as prediction or guidance of action. In contrast, accepting a claim for its contribution to a specific purpose may not yield similar success with other purposes. Indeed, something stronger is true: success with one aim often *inhibits* success with other aims. Science as a whole employs a variety of tools to achieve, e.g., predictions, explanations, causal representations, and the basis for action. What suits a tool to further one of these aims does not well suit it for the other aims. For example, one method used to generate predictions is the analysis of a variety of models with competing assumptions, called robustness analysis. This is a common method in climate modeling (Parker, 2011). But none of those models are expected to accurately represent the causal influences on climate, nor to explain climate change. Indeed, accurate representation of specific causal influences is impossible, for the models represent influences in multiple, incompatible ways. Here the tool of robustness analysis helps achieve one aim, but there are many other aims to which it cannot contribute. A second example of conflicting aims in climate modeling is that the aim of accurate prediction calls for assumptions that are expected to deliver the most likely outcomes, whereas the aim of policy guidance might instead motivate placing particular value on the riskiest possible outcomes.

This division in the pursuit of different aims of science is due in part to the widespread use of idealizations and the variety of purposes idealizations serve. Recall from above that whether an idealization furthers understanding depends on the specific goals of the research. This is why, e.g., Snell's law is appropriate to apply to anisotropic media when researchers want to understand optical refraction in general, but not when they want to understand the path of a specific light ray. To generalize, whether an idealization furthers any particular aim of science depends on the specifics of that aim. Snell's law may help us understand optical refraction, but it is too idealized to give precise predictions of light's path of travel in anisotropic media. For that we need a different tool. Consider also the example just above of making climate predictions with the use of several idealized models, models that do not help us understand climate change.

The idea that different aims of science generally must be pursued separately is evocative of Cartwright's (1983) study of how laws are either false and explanatory or else predictive (but not both). Elsewhere I have defended the view that the aims of explanation and confirmation push in different directions (Potochnik, 2010). Isaac (2013) argues that models have specific functions, such as prediction and use in policymaking, and that many of these functions are best satisfied by abandoning the goal of realistic representation. The separate pursuit of scientific aims is also related to the view that there are tradeoffs among the desirable features of models, such as their generality, precision, and realism or accuracy (Levins, 1966; Odenbaugh, 2003; Matthewson and Weisberg, 2009). This is the idea that increasing a model's generality, for example, must be achieved by decreasing its precision or accuracy. Different features of models, such as greater generality or greater precision, will differently position those models to contribute to particular representational or predictive aims. In my view, the nature of the selected tradeoff reflects the specific purpose to which a model or other scientific product is put.

This conception of the aims of science in conflict requires one amendment. Not only

do the aims of science often conflict, but so too can deployments of a single aim, including even the narrowly epistemic aim of understanding. Recall from above that idealizations can facilitate understanding of a phenomenon by demonstrating that some of its features—including causally important features—are nonetheless unimportant to the current research focus. Which features are to be understood is, thus, relative to a specific research focus. This means that an understanding of some features of a phenomenon can be purchased at the cost of misunderstanding, i.e. misrepresenting, other features. So, for example, an evolutionary game theory model may demonstrate the role of natural selection in producing cooperative behavior, while occluding the role of non-selective and non-evolutionary influences. These might include specific genetic influences and alternative, non-evolutionary influences like learning, to name just two possibilities. In the context of other research programs, such as those with a population genetic focus or an epigenetic focus, understanding other influences on cooperative behavior will move to center stage. This shows how the epistemic aim of understanding can, in itself, motivate different scientific products that are appropriate for different research focuses. The same is true for other scientific aims. Different models or sets of models of a single phenomenon often are suited to different predictive tasks. Different explanations of the same phenomenon are called for in different research contexts (Lewis, 1986; Potochnik, 2015).

There are two primary reasons for the tension among different aims of science and different deployments of a single aim—reasons, also, for the variety of idealizations’ contributions to the aims of science. These are the complexity of phenomena of scientific interest coupled with the limited powers of human cognition and action. The complexity of phenomena investigated in science is by now well appreciated; see, for instance, (Dupré, 1993; Cartwright, 1999; Mitchell, 2003; Strevens, 2006; Wimsatt, 2007). All of the scientific aims discussed in this paper are valuable for their ability to further human comprehension and control of this complex world. Their furtherance is, thus, relative to the limitations of human

cognition and action. This leads to the tension among the successful pursuit of different aims, and even among versions of the same aim. Guiding policy and action are obvious examples of scientific aims influenced by human limitations, and these aims can motivate sacrificing, for example, some predictive accuracy or depth of causal analysis in order to inform policy in a timely manner or focus especially on the riskiest scenarios. The aim of prediction is similarly shaped by human limitations. We have already seen how robustness analysis can generate predictions, without accurately representing causal influences or, for that matter, providing understanding of the system in question. Additionally, when faced with the complex phenomena that are the norm in scientific investigation, scientists must even choose which features of phenomena to focus on successfully predicting. So, human limitations in the face of complexity even lead to tension among different specific predictive aims.

The pursuit of scientific explanations is also influenced by human limitations. Explanation is in furtherance of human cognition, and its features are crucially shaped by this goal (Potochnik, 2011). This is why, as I suggested above, the aims of explanation and confirmation motivate opposed scientific practices, and why different explanations of the same phenomenon are called for in different research contexts. Finally, if I am right that the epistemic aim of science is understanding, then this too is in service to, and shaped by, the particularities of human powers of comprehension. The limited powers of human beings and especially of our cognition, when faced with incredibly complex phenomena, are accommodated with the help of idealizations, but idealizations that are specific to their purposes. Idealizations that further understanding of one feature of a phenomenon tend to obscure other features, and they cannot be expected to serve other aims, such as accurate prediction, equally well. This scenario results in focus on one particular scientific aim (at a time) to the exclusion of others and results in a divide-and-conquer approach, where scientific products are tailored to their precise aims. In order to successfully predict, or to represent

a certain element of the causal structure, or to provide quick guidance for policy, one often must sacrifice other aims.

It is instructive to consider the nature of the position most directly opposed to my view of the conflicting aims of science. The opposed position is that the same scientific products individually further all aims of science—at least accurate representation, explanation, and prediction, and possibly also including more human-centric aims like providing grounds for policymaking. If this were the case, a single, best model or theory would simultaneously offer the best causal representation, the best explanation, and the best predictions. I suspect this kind of view is motivated by an implicit commitment to the idea that scientific products are true depictions of the world. If that were generally the case, then one could expect individual scientific products to play all these roles. Representations that were true in all important regards would provide a sufficient causal representation, be explanatorily unimpeachable, and ground accurate predictions. But science does not aim for such truth, or so I have argued here. The aims of science are accordingly in tension.

This view of opposition among scientific aims also conflicts, though less directly, with views that relate confirmation and prediction to explanation. One example of such a view is inference to the best explanation, where explanatory considerations are used to guide beliefs. Another is Douglas's (2009a) argument that focusing on the aim of prediction sheds light on approaches to explanation. A similar kind of view is suggested by appeals to the simultaneous pursuit of multiple scientific aims—often prediction and explanation—that some philosophical works make in passing. Consider, as just one example, Odenbaugh and Alexandrova (2011) who, in the process of criticizing robustness analysis, emphasize the need to test “empirical hypotheses that later figure in explanations of particular. . . phenomena” (p.758). They see the confirmation of hypotheses—presumably by the accuracy of the predictions those hypotheses ground—to be preparatory for explanatory work. In contrast to this kind of view, I have suggested that these aims of science are in tension. If this is right,

then one should not expect successful prediction to facilitate explanation, nor vice versa.

The conception of the relationship among scientific aims developed here helps make sense of a feature of science that might otherwise be puzzling. Across science there is a proliferation of different approaches to studying the same phenomenon, approaches that apparently are in conflict with one another. For instance, Longino (2013) surveys a wide variety of approaches to the study of human behavior, and in particular, human sexuality and aggressivity. Similar to my criticism in this section of the idea that individual scientific products further *all* scientific aims, Longino criticizes the supposition “that ‘pure’ (a.k.a. ‘basic’) research can provide comprehensive knowledge of a phenomenon that can then be applied to or drawn on for the solution to practical problems” (p.149). Instead she takes a pluralist stance, according to which different approaches provide knowledge, knowledge that is incommensurable across approaches. According to Longino, multiple, incommensurate items of knowledge are possible because each approach differently defines the exact phenomenon under investigation and differently parses the space of causal influences. Thus, on her view “there are many truths,” each useful to a different purpose (p.149). Despite this, Longino allows that “we may remain metaphysically committed to the view that multiple factors or processes are interacting in our one world” (p.146, 147).

The tension among the various scientific aims (and specific deployments of individual aims) grounds one possible interpretation of this proliferation of approaches and the apparent conflict among approaches. Consider the case of investigations of human aggressivity. If accurate representation of the full suite of causes of aggressivity were the uniform aim of those investigations, and if one is metaphysically committed to one world with a single causal structure, then we should expect different research results to be reconcilable. That is, there would be something to say about the respects in which each does better or worse at representing the causes of human aggressivity. In contrast to this expectation, Longino makes the case that “each approach offers a partial understanding framed by distinctive

questions and distinctive parsings of the causal space” (p.135). This suggests that there is no unitary aim for all these different research approaches. On my interpretation, the various aims guiding this research may include a number of distinct representational goals, predictive goals, use in policymaking, etc. A variety of such different aims would generate incommensurable findings of the sort Longino identifies, findings that nonetheless are about the same phenomenon and may equally qualify as successful science.

Conflict among different aims of science also accounts for the persistence of deep disagreements about fundamental principles within otherwise functional fields of research. Such disagreements are found, for instance, in a variety of debates within the field of population biology (?). Distinct scientific products, utilizing different idealizations, are helpful in the pursuit of different particular representational aims, and other aims entirely. This not only results in incompatible approaches, as discussed just above; it also can result in scientists adhering to different, incompatible approaches. On the face of it, this can appear to constitute a deep theoretical divide, and may even be construed as such by some scientists. But subscribing to the view set out in this section enables an alternative interpretation. Each approach may be successful given its specific aim, and no approach succeeds in addressing all the relevant aims.

In summary, in this section I have developed the view that the different aims of science and even specific deployments of individual aims are often best pursued separately. Different scientific products are generally best suited to different aims of representation, prediction, explanation, guidance of policy or action, etc. This is due in part to the variety of purposes served by idealizations, and so the variety of idealizations incorporated into models and other representations. I have suggested that this outcome is ultimately due to the limited powers of human cognition and action, when faced with the exceedingly complex phenomena under scientific investigation.

In light of this variety of scientific aims and multiple deployments of individual aims,

one might wonder how to discern which aim is served by each individual scientific project. If there are many different, potentially applicable scientific aims, by what standard should we judge the success of individual scientific products? Some generalizations can be made about what determines the applicable aims. First, as outlined above, the particular research focus directly influences which features of a phenomenon are to be understood. The research program also influences the broader determination of the importance of prediction, explanation, guidance of action, etc. Some research agendas in behavioral ecology focus on the role of natural selection in producing a behavior, whereas research on epigenetics will foreground environmental influence on genetic factors. For both of these focuses, prediction tends to be much less central than it is in climate science. A second type of influence on what aim is applicable is the features of the practitioners of science and their target audience. Much climate research aims to furnish us with predictions in part because climate change is expected to have vast humanitarian, economic, political, and other consequences for humans. In turn, whether a prediction is intended for other researchers or for policy-makers may lead to an emphasis on, respectively, accuracy or risk-minimization.

When it comes to determining the aim by which a particular scientific product should be judged, the applicable aim can often be ascertained by charitable interpretation of the research conducted, including its setup, the conclusions drawn, and its research and social contexts. There may nonetheless be ambiguity regarding what aim is pursued and, thus, by what standard the work should be judged. Even more common is a mismatch between the pursued aim and conclusions drawn from the research by other researchers or, especially, by the popular media. A first step toward eliminating such confusions is to explicitly recognize that any given scientific product is in service of some particular aim, and may well not contribute to any other aims of science.

3 The Expanded Influence of Social Values

In §1 I motivated an alternative construal of science’s epistemic aim, based on the observation of continuing widespread idealization in science. In particular, I argued that if the epistemic aim of science is not truth but understanding, this would enable idealizations to directly contribute to science’s epistemic success. I claimed there that the aim of understanding is not problematically subjective, because what furthers human understanding depends on objective features of the world under investigation and of the scientists leading that investigation. Then, in §2 I suggested that there are many diverse aims of science—some epistemic, others non-epistemic, and still others somewhere in between—and that these diverse aims motivate different types of scientific products. We should expect an open-ended list of scientific aims. Yet there is an overarching theme that unifies those aims. All the aims of science further cognition or action; indeed, they further *human* cognition or action, as a science by and for human beings should. Moreover, I have suggested that in order to determine the aims by which a scientific product’s success is to be judged, one must consider features of the particular researchers and their target audience. The account of the aims of science developed in this paper thus explicitly links those aims to particularities about human scientists and human consumers of science.

In this final section I consider two implications this human-centric conception of science’s aims has for the role of social values in science. First, there are several avenues by which rather obvious and mundane effects of social values turn out to have epistemic significance, in virtue of shaping the requirements for scientific understanding. The epistemic influence of social values is kept in check, though, with the threshold requirement of “true enough”—a loose and variable tether to reality. Second, the diversity of science’s aims, and the diverse scientific products this results in, also extends the influence of social values on science’s other aims. Which aim, and which variety of an aim, is pursued depends also upon social values.

Social values are of obvious relevance to at least some of the purely non-epistemic aims of science. For instance, Douglas (2009b) argues that social values are relevant to determining the evidential standards behind the use of science in policymaking. But, as Elliott and Willmes (2013) point out, there seems to be a relevant difference between this and the influence of social values on science's epistemic aims. As those authors indicate, ethical values do not seem to be relevant to determining whether a hypothesis is true. Construing science's epistemic aim as truth thus seems to tightly constrain social values' epistemic influence. Our values may influence what truths we seek, and how we employ those truths, but not what is and is not true.

In contrast, the alternative construal of science's epistemic aim as understanding creates additional inroads for the epistemic influence of values. One obvious role for the influence of social and political concerns is in the choice of which research questions are pursued and how those research programs are subsequently developed. The possible directions for research are vast, and human concerns and values shape which of these possible directions are in fact pursued. Considerations introduced in this paper elevate this obvious role for social values to epistemic significance. I have argued that the precise research agenda helps guide the requirements for understanding. If this is right, then the values that shape the research agenda have direct epistemic implications. Many different aims of representation and understanding (among other aims entirely) can be brought to the investigation of any given phenomenon. This is illustrated by Longino's (2013) study of the wide variety of scientific research into human sexuality and aggression, as discussed in the prior section. Different research agendas lead to different emphases and parsings of the causal space and, as a result, vastly different varieties of understanding. Additionally, as I showed above, particularities of a research program partially determine whether a claim or other representational element contributes positively to understanding and, accordingly, whether it is true enough.

I have also claimed that, besides the research program, features of researchers themselves and of the target audience influences what variety of understanding is sought and, thus, what representational elements are true enough. These are, then, other avenues for the epistemic influence of social values. Our values are reflected not only in the nature of the research programs pursued in science, but also in who pursues that research, and in whom it is pursued for. These seemingly humble influences of social values on science often have been ignored because of their apparent irrelevance to epistemic concerns. But there is now a well-developed literature, especially in feminist philosophy of science, about how features of the practitioners of science influence the content of scientific knowledge. The view developed here suggests one way in which that influence is exerted. Features of researchers and their audience have epistemic influence in virtue of how they shape the specific requirements for scientific understanding.

A scientist's features can, of course, influence the nature of the research program she pursues. But even beyond that, what a scientist brings to her work can influence which causal influences are salient, what patterns are seen in the phenomena, and thus, the variety of understanding produced. Consider, as an example of the epistemic influence of researchers' characteristics, Hrdy's (1986) analysis of how a shift in the gender ratio of primatologists led to significantly different types of observations about the behavior of male and female primates. Particular features of the intended audience—such as their background knowledge, and their position to act—can also shape which aspects of phenomena scientists should aim to provide them with an understanding of. Recall that whether any given posit facilitates understanding is what determines whether it is true enough. As a result, while social and ethical concerns may not directly influence whether a claim is true, they do influence whether a claim is true *enough* via their influence on research program, researchers, and the audience for the research.

I also claimed in §1 that taking understanding to be the ultimate epistemic aim of science

is not problematically subjective. The avenues I have suggested for social values' epistemic influences correspond to the ways in which epistemic success *is* properly subjective, viz., that what best facilitates understanding depends in part on the features of the practitioners and consumers of science. Just as our values cannot determine whether a hypothesis is true, similar limitations obtain for understanding, taken as a form of epistemic success. There are regards in which social concerns cannot legitimately influence the content of our understanding. Once the focal phenomenon and research agenda are set, we cannot choose whether an idealized assumption facilitates understanding, nor can we choose whether a representation can convey understanding. These are shaped entirely by features of the phenomenon under investigation, the scientific tools available, and human psychology. Accordingly, the requirement of "true enough" cannot be influenced by our social concerns. The significantly weakened version of this requirement I have motivated provides only a loose and variable tether to reality, but it is a tether nonetheless.

The idea that science's diverse aims motivate different types of scientific products also expands the space for the influence of social values. In §2 I opposed that idea to the view, often left implicit, that true (or approximately true) representations are used to accurately predict, explain phenomena, guide policy decisions, etc. If this were so, then the influence of social values on these latter scientific activities would be constrained by their inability to influence what is and is not true. On that view, what we aim to predict or our evidential standards for policymaking, e.g., may be influenced by our values, but what representations satisfy these aims cannot be. Embracing instead the idea that success with one aim inhibits successful pursuit of other aims eliminates this constraint. The most predictively accurate representations are not very explanatory, our best explanations tend to be somewhat inaccurate representations, etc. Values can find their way into which aim is pursued and what specific deployment of an aim is intended. This in turn influences what characteristics the resulting scientific product should have. This idea is very similar to the

conclusion reached by Elliott and McKaughan (2014) that nonepistemic considerations can override epistemic considerations in virtue of the different goals scientists can have, when coupled with tradeoffs among the desirable features scientific representations might possess. It is also reminiscent of Longino's suggestion that even the epistemic value of empirical adequacy, which might seem basic to all empirical pursuits, is negotiable, for some accuracy can be sacrificed for other desirable epistemic features, such as simplicity or generality (2001, p.185,186).

Consider two brief examples of how values might influence the specific research aim. First, the political commitments of researchers or funding agencies might influence whether predictive research on climate change places greater emphasis on the riskiest scenarios, the least risky scenarios, or weights scenarios simply by their likelihood. None of these strategies is epistemically more sound than the others, yet they are predictive aims distinct from one another that will likely result in different results. As a second, less conjectural example, consider Joan Roughgarden's research in behavioral ecology. According to Roughgarden (2009), her values lead her to emphasize the role of "kindness and cooperation" in animals' traits and their evolutionary history. The view I have motivated here would construe this emphasis as a commitment to a particular type of representational project, namely, representing the causal role(s) of cooperative behavior and mutually beneficial outcomes in the evolution of animal behavior. Here too there are limits on values' proper role. Social values cannot influence the accuracy of our predictions. The desires of researchers and their funders cannot determine the likelihood of the riskiest scenarios, only what emphasis those likelihoods receive in scientific research. One cannot represent a pattern that does not obtain. Roughgarden may pursue a representational project that has been largely neglected, but the pursuit of that project is necessary but not sufficient for its success. Any scientific project may, in the end, suffer a fate similar to astrology's.

I have not here attempted a full analysis of influences on the aims of scientific research

and the specific deployments of those aims. There may be many other inroads for the epistemic influence of social values beyond those I have identified here. My purpose in this section has been to show how the views developed above elevate simple ways in which particular humans and their values influence science to great significance, including epistemic significance. Two individuals cannot know contradictory things about the same phenomenon, defined in the same way. This is so even if those individuals have different backgrounds, values, interests, etc. But all of those differences among the individual practitioners and consumers of science *are* relevant to what phenomena are under investigation; how those phenomena are construed; which causal influences and patterns are of primary interest; what connections are drawn to other phenomena; and other elements of scientific research. All of these differences affect the variety of understanding that is sought or, more generally, the specific scientific aim pursued. This provides a means for our individual characteristics and values to legitimately influence the nature of our scientific understanding in a way that they could not influence scientific truth.

Acknowledgments

Thanks to Kevin Elliott, Dan McKaughan, and three anonymous referees for helpful comments on this paper. I also received feedback on some of these ideas from Anthony Chemero, Thomas Polger, Robert Skipper, the participants in the 2013 Missouri Philosophy of Science Workshop, and the audience of a talk at the 2014 Pacific Division Meeting of the American Philosophical Association.

References

Cartwright, Nancy (1983), *How the Laws of Physics Lie*, Oxford: Oxford University Press.

- (1999), *The Dappled World: A Study of the Boundaries of Science*, Cambridge: Cambridge University Press.
- Cohen, L. Jonathan (1992), *An Essay on Belief and Acceptance*, New York: Clarendon Press.
- Douglas, Heather (2009a), “Reintroducing Prediction to Explanation”, *Philosophy of Science* 76: 444–463.
- (2009b), *Science, Policy, and the Value-Free Ideal*, University of Pittsburgh Press.
- Dupré, John (1993), *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*, Cambridge: Harvard University Press.
- Elgin, Catherine Z. (2004), “True Enough”, *Philosophical Issues* 14: 113–131.
- Elliott, Kevin C. (2011), *Is a Little Pollution Good for You?: Incorporating Societal Values in Environmental Research*, Oxford University Press.
- (2013), “Douglas on Values: From Indirect Roles to Multiple Goals”, *Studies in History and Philosophy of Science* 44: 375–383.
- Elliott, Kevin C., and Daniel J. McKaughan (2014), “Nonepistemic Values and the Multiple Goals of Science”, *Philosophy of Science* 81: 1–21.
- Elliott, Kevin C., and David Willmes (2013), “Cognitive Attitudes and Values in Science”, *Philosophy of Science* 80.
- Grimm, Stephen R. (2006), “Is Understanding a Species of Knowledge?”, *British Journal for the Philosophy of Science* 57: 515–535.
- (2010), “Understanding”, in Duncan Pritchard and Sven Berneker, eds., *The Routledge Companion to Epistemology*, New York: Routledge.

- (2012), “The Value of Understanding”, *Philosophy Compass* 7: 103–117.
- Hrdy, Sarah Blaffer (1986), “Empathy, Polyandry, and the Myth of the Coy Female”, in R. Bleier, ed., *Feminist Approaches to Science*, The Athene Series, Pergamon Press, 119–146.
- Isaac, Alistair M.C. (2013), “Modeling Without Representation”, *Synthese* 190: 3611–3623.
- Levins, Richard (1966), “The strategy of model building in population biology”, *American Scientist* 54: 421–431.
- Lewis, David (1986), “Causal Explanation”, in David Lewis, ed., *Philosophical Papers*, Oxford University Press, vol. II.
- Longino, Helen E. (2001), *The Fate of Knowledge*, Princeton: Princeton University Press.
- (2013), *Studying Human Behavior: How Scientists Investigate Aggression and Sexuality*, Chicago: University of Chicago Press.
- Matthewson, John, and Michael Weisberg (2009), “The structure of tradeoffs in model building”, *Synthese* 170.
- Mitchell, Sandra D. (2003), *Biological Complexity and Integrative Pluralism*, Cambridge Studies in Philosophy and Biology, Cambridge: Cambridge University Press.
- Odenbaugh, Jay (2003), “Complex systems, trade-offs, and theoretical population biology: Richard Levins’ ‘The strategy of model building in population biology’ revisited”, *Philosophy of Science* 70: 1496–1507.
- Odenbaugh, Jay, and Anna Alexandrova (2011), “Buyer Beware: Robustness Analyses in Economics and Biology”, *Biology and Philosophy* 26: 757–771.

- Parker, Wendy S. (2011), “When Climate Models Agree: The Significance of Robust Model Predictions”, *Philosophy of Science* 78: 579–600.
- Potochnik, Angela (2010), “Explanatory Independence and Epistemic Interdependence: A Case Study of the Optimality Approach”, *The British Journal for the Philosophy of Science* 61: 213–233.
- (2011), “Explanation and Understanding: An Alternative to Strevens’ *Depth*”, *European Journal for the Philosophy of Science* 1: 29–38.
- (2015), “Causal Patterns and Adequate Explanations”, *Philosophical Studies* 172: 1163–1182.
- Rohwer, Yasha, and Collin Rice (2013), “Hypothetical Pattern Idealization and Explanatory Models”, *Philosophy of Science* 80: 334–355.
- Roughgarden, Joan (2009), *The Genial Gene: Deconstructing Darwinian Selfishness*, Berkeley: University of California Press.
- Strevens, Michael (2006), *Bigger than Chaos: Understanding Complexity through Probability*, Cambridge: Harvard University Press.
- (2008), *Depth: An Account of Scientific Explanation*, Cambridge: Harvard University Press.
- (2013), “No Understanding Without Explanation”, *Studies in History and Philosophy of Science* 44: 510–515.
- Trout, J.D. (2002), “Scientific Explanation and the Sense of Understanding”, *Philosophy of Science* 69: 212–233.

- Weisberg, Michael (2007), “Three Kinds of Idealization”, *The Journal of Philosophy* 104: 639–659.
- (2013), *Simulation and Similarity: Using Models to Understand the World*, Oxford University Press.
- Williams, Joseph J., Tania Lombrozo, and Bob Rehnder (2013), “The Hazards of Explanation: Overgeneralization in the Face of Exceptions”, *Journal of Experimental Psychology: General* 142: 1006–1014.
- Wimsatt, William C. (1987), “False Models as Means to Truer Theories”, in N. Nitecki and A. Hoffman, eds., *Neutral Models in Biology*, Oxford: Oxford University Press, 23–55.
- (2007), *Re-Engineering Philosophy for Limited Beings*, Cambridge: Harvard University Press.
- Winsberg, Eric B. (2010), *Science in the Age of Computer Simulation*, The University of Chicago Press.