Through the Fractured Looking Glass

Sandra D. Mitchell*†

I argue that diversity and pluralism are valuable not just for science but for philosophy of science. Given the partiality and perspectivism of representation, pluralism preserving integration can increase accuracy. Perspectivism is often supported by appeal to visual representation. I draw further insights from multimodel sensory integration for understanding experiment-based predictions of protein structure. The epistemic lessons learned from the scientific case also apply to philosophy of science itself. Finally, I suggest that a critical, nuanced philosophical view of legitimate sources of pluralism in science has an important role to play in public discourse.

Introduction. The theme of my address is that diversity and pluralism in science are not just ethically and politically valuable, but they are required for science to be epistemically effective. Moreover, there are lessons I will draw from these arguments for best practices for philosophers of science. There are many studies of the importance of increasing the diversity of our social institutions. Providing opportunities for previously underrepresented members of the population to participate more fully in scientific, academic, and other forms of social life meets the demands of fairness and equal opportunity. Multiple forms of diversity also constitute an epistemic good. Integrating multiple scientific perspectives does not merely help us know more,

*To contact the author, please write to: Department of History and Philosophy of Science, 1101 Cathedral of Learning, University of Pittsburgh, Pittsburgh, PA 15260; e-mail: smitchel@pitt.edu.

†This article is a revised version of the presidential address I delivered at the PSA 2018 biennial meeting in Seattle, Washington. The presented version can be viewed at https://spark.adobe.com/page/AiXfAUmLTaEbB/. It takes the contributions of many to ensure a successful meeting and an effective organization. I want to thank PSA 2018 program chairs Kevin Elliott and Doreen Fraser and the members of the PSA Governing Board, PSA committees, and PSA members for their commitment and skill at navigating the complexities and coordination problems involved. I especially want to thank our extraordinary executive director, Jessica Pfeifer, for her wise and inspired leadership.

Philosophy of Science, 87 (December 2020) pp. 771–792. 0031-8248/2020/8705-0001$10.00
Copyright 2020 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/710787 Published online by Cambridge University Press

which in itself advances science, but it can increase the accuracy of scientific models. I will argue that there are situations in which perspectival diversity and model pluralism are required for making science an effective epistemic enterprise.[1]

In the history of philosophy of science, diversity as an epistemic good has sometimes been overlooked in pursuits that emphasize unification as the goal of the best scientific practices. Newton's theoretical unification of celestial and terrestrial motions stood as an icon for nineteenth-century philosophers of science, and in the 1930s the Vienna Circle founded the *International Encyclopedia of Unified Science*. Unifying efforts are clearly valuable in science, but that is not the whole story or the most common story. Science actually and properly produces multiple research efforts and associated products, so today it would be more appropriate to publish an *International Encyclopedia of Diversified Science*.

Philosophy of science has at different times and in different ways attempted to understand how we, mere mortals, can have knowledge of what nature is like, of how it works, of why things happen when they do. We look to the practices of science from bench and field to computer models and theoretical hypotheses to answer normative epistemological, metaphysical, and ethical questions. Collectively we have dug deep to unearth the presuppositions of experimental inference and mathematical representation, and we have looked wide to expose the roles of social values and historical context in what constitutes scientific knowledge. In doing these things, philosophers of science have brought a diversity of questions and perspectives to the study of science, just as scientists bring a diversity of questions and perspectives to the study of nature.

Naturally, we critique the different approaches others take, judging them to be more or less valuable for both describing and offering normative advice about the practice of science. While such critiques and ensuing controversies are part of the life blood of our efforts to get things right, I maintain that philosophers of science should embrace epistemic diversity for the same

1. This work has been deeply inspired by what I have learned from Nancy Cartwright, Ronald Giere, and Helen Longino as well as from many more colleagues, teachers, and students. I am indebted to the many scientists who have shared their understanding of science with me and especially those with whom I have had the opportunity to collaborate: Robert E. Page Jr., Wolfgang Streeck, and Angela Gronenborn. This article, while based on my 2018 PSA presidential address, includes additional explications and arguments. The discussion of experimental convergence was presented in a keynote lecture at the International Congress of Logic, Methodology and Philosophy of Science and Technology in Prague in August 2019. I am grateful to the University of Hamburg for the opportunity to try out some of the arguments in the Carl Friedrich von Weizsäcker Lectures in 2016. My sincere thanks to Joel M. Smith and Kevin Elliott for their extremely helpful comments and suggestions.

https://doi.org/10.1086/710787 Published online by Cambridge University Press

reasons that apply to scientists. By making explicit what each of our necessarily partial perspectives can provide, specifying how different approaches are tuned to different questions or goals, and determining how different perspectives can be mutually enhancing, we can be more successful at answering the questions that unite philosophy of science.

In Defense of Pluralism. For many years I have defended a form of pluralism, integrative pluralism (Mitchell 1992, 2003, 2009). In developing this view I have taken seriously both the complexity of nature and the complexity of the human enterprise of representing, explaining, and intervening in nature. Expanding the field of vision in the study of science beyond its foundations and logic to include situated practices makes clear the dynamic, provisional, and iterated character of how scientists learn about nature. My attempts to understand the multiple processes of learning by reasoning, experimenting, and engineering and the strategies for using what is learned to direct individual choices and social policies has led me to reject a number of dichotomous frameworks that are too limiting for representing what and how scientists know. I am not alone in this. I join a substantial number of philosophers (Cartwright 1983; Dupré 1993; Wylie 1999; Woodward 2005, among others) whose work shatters long-held overly rigid images of science such as the requirement of strict laws for explanation, the direct role of experience and experiment in confirmation, and a mirror-of-nature account of the relationship of scientific models to the parts of the world they represent. While many of us reject rigid versions of these views, we do not always agree on the details or alternatives to them.

This address considers some insights about scientific practice that the integrative pluralism approach provides. I will discuss features of ontological complexity, and of representational partiality, and present a view of science that emphasizes its dynamic character. Finally, I will suggest that by providing a sophisticated, nuanced understanding of science, philosophers of science can contribute in very important ways to public discourse about the authority of science to speak truth to power, which is much needed in current political contexts around the world.

My title, "Through the Fractured Looking Glass," pays homage less to Lewis Carroll's *Alice in Wonderland* than it does to the failed mirror-of-nature picture of how scientific knowledge represents the world (Rorty 1979). The "mirror of nature," a metaphor introduced in the seventeenth century as capturing the goal of both art and science, implies both veridicality and completeness in which every feature of a phenomenon is precisely described in a one-to-one mapping into a model. Indeed, no representational model can attain this type of completeness without it being an exact, to-scale copy of the target phenomenon and its context, since we know scale matters, context matters, and small differences in conditions can have large consequences

https://doi.org/10.1086/710787 Published online by Cambridge University Press

in behavior. Cluttering the lab with duplicates of the phenomenon studied would fail to bring science much closer to explaining and predicting natural behavior than directly engaging with the original. Complete models, in this sense of completeness, would be of little scientific use. To be usable, a representational model has to leave some things out by abstraction and often knowingly distort other features by idealization. Thus, scientific models are always partial and often, in the strictest sense, false. Veridicality and completeness both fail, and thus the mirror-of-nature assumption cannot provide grounds for the epistemic warrant of scientific claims. That this is the case will not be a surprise to most philosophers of science. Cartwright (1983) exposed the ways that the laws of physic lie, Teller (2001) described the collapse of the Perfect Model Model, and Giere (1999, 2010) explained how the relationship between models and the world is expected to be short of a perfect fit. I have long argued that some significant implications of the partiality of scientific representation are overlooked to the detriment of both scientific practice and an adequate understanding of those practices. The mirror of nature has been shattered and needs to be replaced by a plurality of partial yet accurate and, for some purposes, adequate models of what nature affords.

Contemporary science studies complex structures and behaviors through a fractured mirror, and at a variety of levels of organization from the most basic component parts to the emergent properties of entire systems, and uses representations of different kinds and degrees of precision from fine- to coarse-grained. The relationships among these multiple scientific models continue to be in dispute. Reductionists argue that all explanations and predictions can in principle, if not in practice, be crafted from descriptions of the properties and behaviors at the most fundamental level of physical substance. This argument presumes that since all more complex arrangements are composed from the basic physical building blocks, knowledge of the basic physical level could in principle (or, at the so-called end of science, will in fact) tell us everything about composite structures and their behaviors. But that is not the case. Some coarse-grained explanations or predictions that appeal to higher-level structures are required to explain at least some features of complex systems, like biological robustness (Mitchell 2012).

I have argued (Mitchell 2003) that this ontological reductionist argument confuses compositional physicalism with "physics-ism." The first, compositional physicalism, is a metaphysical presumption that is constitutive of scientific methodology. Science looks to material causes to explain nature and thus does not appeal to miracles. This assumption has been taken to support physics-ism, an epistemic doctrine, that scientific representations of what we take to be the most basic matter, typically representations in physics, can recover all the information that is contained in representations of composite matter of chemical, biological, psychological, or social structures. But physics-ism does not follow from physicalism because of the fact that all representations

https://doi.org/10.1086/710787 Published online by Cambridge University Press

are partial, whether they target the most basic constituents of nature or more complex composite ones. Furthermore, the ways in which they are partial do not necessarily, or arguably often, stand in neat, nested, intertranslatable relationships. Different models often tell us different things, not the same things in different ways. I will argue that effective science integrates perspectives encoded in models to provide more and better representations than the sum of what is provided by each.

Given that representational models are partial, in order to learn about natural systems, science must employ a diversity of models, methods, and representations. Scientific representations cannot veridically and completely mirror either the basic or the compositional relations of natural systems; rather, they are nonreductively partial. The flip side of partiality is that by leaving some features out, every method, model, and representation "selects" features to be included. This "selection" encodes a perspective.

My approach to perspectivism has been structured by the seminal work of Ronald Giere (2006b). In what follows I build on Giere's views about perspectivism. Attending to integrative strategies in scientific practice, I apply and extend his picture in ways that serve as a friendly correction to his original account.

There are many sources of the perspectival diversity of scientific representations: we can start by sorting them by what is investigated, why it is investigated and how it is investigated. Models may be variously tuned to answer different questions or target different levels of organization or contexts in which phenomena occur. There are different purposes that models are designed to serve. For example, having the goal to produce an effect is importantly different from trying to prevent an effect. This contrast can clearly shape the choices of variables. For the first case, sufficient but not all the necessary factors must be included in an effective model, while for the second only necessary factors are required. And there are a host of different methods for constructing a representational model. On Giere's (2010) hierarchical account, this includes derivation from theoretical principles and abstraction from experimentally produced data. In addition, I have argued, there are strategies that violate the hierarchy, namely, semi-empirical model construction (Mitchell and Gronenborn 2017). Models are accepted as accurate and adequate by logical and practical feedback between theories and empirical data. Scientific models are not entailed by the phenomena; rather, on my perspectival view, the phenomena afford diverse, correct representations.

Different Models for Different Purposes. Take the example of maps that represent the same natural phenomenon for different purposes. Maps, as Giere (2006b) and many others have pointed out, are nonpropositional, graphic modes of representation that encode relational, typically spatial, information that can be isomorphic, homeomorphic, or otherwise similar to the relations

https://doi.org/10.1086/710787 Published online by Cambridge University Press

in nature it targets.[2] Consider multiple maps of the London underground. One map, designed by Harry Beck in 1933, replaced the more geographically accurate 1908 map (see fig. 1).

The Beck map idealized the geographical distances, to optimize for navigation function. The scale and degree of precision differs in the two maps. Which is correct? Both. Which should you use? It depends on what you want to do. Purpose determines the adequacy of the accurate representations. The 1908 map is more accurate at representing the geographical distances and locations of the stations, but it was difficult for a rider to use it. The Beck map retained the relative locations of the stations, but idealized the geographical distances between the stations to optimize navigation function. Both maps are partial and perspectival. Both are accurate with respect to the features selected to represent. However, they are not both adequate to the same purpose. The 1908 geographical distance map would be of use to engineers for determining, for example, how much track needs to be laid or repaired. The 1933 map's usefulness to riders of the underground was evident by how quickly demand for the new map replaced the older maps. Difference in goals is one way diversity of representations provides greater resources than trying to unify or reduce all the information into one model. There is no one-to-one mapping between a representation and the possible goals it serves. One map may serve multiple end goals, and one goal may have multiple maps. Thus, the relationships are many to many.

Different Models for the Same Purpose. Partiality and perspectivism necessitate model pluralism. A single model cannot and should not deliver a complete, maximally precise representation for all possible purposes. But what about situations in which both the explanatory or predictive goals and the target phenomena of different models are the same? Goals, of course, can be expressed with more or less specificity. One could shrink the goal and the target to limit the relevance of other approaches and models or expand it to include almost all investigations in a field. For example, the goal of *ab initio* models of protein structure, which search the space of conformations for the minimal Gibbs energy states, might be characterized as just that—a search algorithm for finding thermodynamically stable states. Or it could be expressed, which I believe describes most of the scientific practice in this field, as a way to predict the functional states of proteins in organic bodies. Which

2. Accounts of how scientific theories and models represent the world by appealing to maps has a very long history. Among many who explore the map metaphor are Toulmin (1953), Kitcher (2001), Sismondo and Chrisman (2001), and Contessa (2007).

https://doi.org/10.1086/710787 Published online by Cambridge University Press

a



b

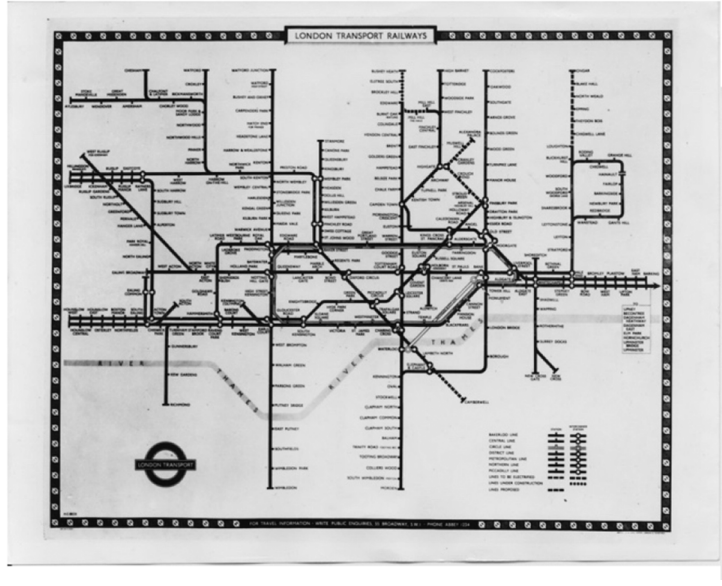


Figure 1. *a*, 1908 London underground map; *b*, 1933 Harry Beck London underground map (https://londonist.com/2016/05/the-historyof-the-tube-map).

https://doi.org/10.1086/710787 Published online by Cambridge University Press

alternative perspectival models are taken to be relevant to judge the success of an *ab initio* model will depend on the goal. In the set of alternatives within the scope of a goal, what one representation abstracts out could be, and often is, represented by other perspectival models. *In silico* protein models idealize water in their calculations, while *in vitro* nuclear magnetic resonance (NMR) experiments intervene on a protein in an environment that includes water to predict structure, although NMR typically excludes the presence of other molecules in the setup and inferences. *In vivo* predictions of protein structure experimentally and theoretically investigate proteins in a cell and include water plus numerous other molecules. If the features that are left out in one model but included in another are causally independent, partitionable into distinct subfeatures, or neatly mereologically nested, then the multiple models might be simply combined to form a single, more complete model of the phenomenon. But there are many cases in which a neat, linear combination of models is not possible. In those cases, a plurality of models is ineliminably required to represent and reason about a phenomenon. When a single model cannot deliver all the causally relevant aspects of a given phenomenon with complete precision for all purposes, using multiple models will be required to achieve explanatory or predictive goals. Thus, the question of how such models should be integrated becomes both philosophically and scientifically important. When the enduring plurality of models cannot be reduced or unified to produce a single model, sometimes they can be interactively integrated, yielding increased accuracy while preserving perspectival pluralism.

There are multiple sources of diversity of scientific models that are both of the same phenomenon and for the same purpose. Different methods for causally generating experimental data engage different features of a given phenomenon and employ different theories, assumptions, and sources of error in translating the causal effects of experimental intervention into a representational model from the data.

Clearly not all cases of multiple models of a single phenomenon stand in the same relationship to each other. In addition to the simple additive relation among models discussed above, they might stand in a reductive relationship, if in fact all the causal content of one can be represented, without loss, by the other. Some have argued (although others have disputed) that reduction correctly describes the relationship between thermodynamic models and statistical mechanical ones (see Sklar 1999, 2015). While reductive relationships do hold between some models, not all can be described in that way.

It is also possible that two models of the same phenomenon are in competition, where at best only one should be retained and the other eliminated. These have been the focus of accounts of the value of model diversity for the scientific community as whole as a way of hedging our bets until sufficient

https://doi.org/10.1086/710787 Published online by Cambridge University Press

evidence tilts the community to adopt one competitor and reject the other (Kitcher 1990). I will explore the use and value of a different kind of relationship among models, integration, for which models are not additively unified, reductive, or in direct competition. In these cases the plurality and diversity of models does not go away but endures.

One example of scientific practice in which this occurs is in multiple experimental sources for producing data, generating measurements, and predicting representational models. Different experiments are perspectival, as Giere (2006b) has suggested, similar to different visual systems, in that they causally interact with different features of a phenomenon in generating a representation. In both cases, how a phenomenon "looks" from one perspective is different from the way it "looks" from another. Consider, as many philosophers have, how the world looks to beings that have different visual systems. Bees and humans are trichromatic; that is, they each have three photoreceptors within the visual system by which color perception is constructed. Humans base color combinations on red, blue, and green wavelengths, while bees base all their colors on ultraviolet, blue, and green. Thus, the same flower from the perspective of a bee and the perspective of a human looks very different (Giere 2006a, 2006b). It is the same flower made of the same material in the same environment, but the visual apparatus of a bee and the visual apparatus of a human access different visual signals that are afforded by the same flower. There is one phenomenon, the flower, and one goal, representing its color, but two diverging visual models that are generated by different modalities.[3] Both of the representations are correct.

This example also exposes the interactive nature of the generation of a model from a causal detection process. Giere uses color vision as a window into understanding the diversity of models generated by scientific experimentation. He claims that "as with color vision, the compatibility of different observational perspectives is easily understood as a consequence of all observations being observations of one world" (Giere 2006a, 31). I suggest that, perhaps, it is not so easily understood. The details of experimental practice expose the complications, assumptions, and strategies used in determining the compatibility of

3. This presumes that the target phenomenon and the goal are the same for both perceivers. That is, there are features of the flower that afford two different interactions and representations. That feature we call the color of the flower, not the color detectable by a bee, which is different from the color detectable by a human. The shared goal is to represent those features of the flower. Perspectival realism, a view I endorse (see Massimi 2016, 2018), posits a source external to the perceiver or experimental devise that is causally relevant to the generation of data that are the basis of the representational models. That still leaves it to be decided how to characterize the scope of the phenomenon, how narrow to shrink it, and how wide to expand it, which will structure the space of alternatives.

https://doi.org/10.1086/710787 Published online by Cambridge University Press

different perspectives. However, I do want to extend the appeal to the visual system to multiple sensory modalities for insight into how to think about experimental perspectivism. In so doing, we see that a new resource becomes available for understanding not just model compatibility but the process by which a plurality of perspectives can increase the accuracy of representational models. Here color is not the best case; other properties in the world, like relative position and distance, have greater stability across variations of perspective and yet are equally subject to a perspectival analysis. These also afford a variety of causal interactions yielding multiple "correct" yet diverse representations.

When different experimental results agree, this type of robustness is taken as a defense of the reality of the shared result. How do we accommodate models that disagree or diverge? Consider the situation in which the target phenomenon is the same, namely, a specific protein, and the goal is the same, characterizing the three-dimensional atomic structure of a protein molecule, but the representational models produced by different experimental protocols diverge. Prima facie this appears to be a case of conflict. But I will argue that for some divergent models, perspectivism offers new insights on how they can be used together. Different experimental setups are used to detect the structure of a protein. They target different features of the protein and appeal to different background assumptions and theories, in order to interpret the detected signals to generate a prediction of the protein structure.

X-ray crystallography and NMR spectroscopy are the most used and trusted experimental methods to generate predictive models at the atomic level of resolution of protein structure. Yet when x-ray and NMR target the same protein, the predicted structures that are inferred can diverge. If we take those representations to be propositional descriptions of the actual protein in the world, then they make contradictory claims, and we are faced with saying that at most one of the representational models is correct, the one generating the true proposition. Perspectivism allows a different resolution to divergence, one that more accurately captures scientific practice.

X-ray crystallography and NMR spectroscopy involve setting up experiments to causally interact and measure features of a protein that are accessible to their different instruments. For x-ray experiments, a protein must be crystallized and then a beam of incident x-rays is diffracted by the electron clouds of the atoms of the protein in many specific directions. The angles and intensities of the diffracted beams are measured. From those, using theories of photon-electron interactions, a three-dimensional electron density map of the protein is constructed. This permits a calculation of the mean positions of the atoms in the crystal.

In contrast, for NMR, the target protein is prepared in solution and placed inside a very powerful magnet. By so doing, the normal spin of the nuclei of

https://doi.org/10.1086/710787 Published online by Cambridge University Press

the different atoms in the protein is realigned by the magnetic field of the instrument. A radio frequency pulse is introduced to disrupt this new equilibrium, and the return to equilibrium is tracked, revealing information about the effects not just of the experimental magnet but of other nearby atoms. Measurements along with electromagnetic theory provide a map of how the different atoms are chemically linked and thus how close they are spatially.

In short, with experimental interventions, the x-ray perspective detects electron cloud diffraction patterns, while the NMR perspective detects nuclear magnetic resonance spectra. Each has a partial perspective on the target atomic structure of the protein. X-ray experiments allow an inference to the relative positions of atoms, and NMR experiments permit inference to the relative distance between atoms. Clearly the scientific models engage different features of the phenomenon; x-ray treats the protein as composed of structures of electron clouds of atoms that causally scatter x-rays, while NMR treats it as structures of nuclei in atoms whose relative resonant frequencies are causally influenced by variant magnetic environments. Each is partial and perspectival, and they generate different predictions by using different parts of physical theory and different methods of interacting with their samples. When the two experimental protocols are operating as they should and yet deliver different predictions of a protein's structure, which one is correct? I suggest that, just as in the case of the two maps of the London underground and the color detected by bees and humans, the answer is both. Each experimental modality correctly detects the features it targets and correctly represents the protein from that perspective.

Convergence and Divergence. First, consider the situation in which the two experiments determined the exact same structure for the same protein. This would constitute robustness of evidence that the theoretical and experimental models generating the prediction were compatible and that the structure predicted was more likely to be correct than if obtained from repeated or replicated instances of only one of the types of experiments. Many philosophers have written about the epistemic value of convergence, including Salmon (1984, 221), citing the "remarkable agreement" of Perrin's 15 experiments to determine Avogadro's Number as an iconic example. In Perrin's own description: "Our wonder is aroused at the very remarkable agreement found between values derived from the consideration of such widely different phenomena. Seeing that not only is the same magnitude obtained by each method when the conditions under which it is applied are varied as much as possible, but that the numbers thus established also agree among themselves, without discrepancy, for all the methods employed, the real existence of the molecule is given a probability bordering on certainty" (1923, 215–16).

https://doi.org/10.1086/710787 Published online by Cambridge University Press

Philosophical analyses of the type of reasoning from experimental convergence to realism vary, from Salmon's (1984) appeal to common causes in the background conditions, to no miracle arguments, to more recently Kuorikoski and Marchionni's (2016) appeal to error and bias reduction. Although there are some who deny the importance of convergence in scientific practice (see Hudson 2013), most philosophers agree that convergence of multiple experimental results methodologically increases the confirmation for claims about a phenomenon. Does that tell us anything about how to think about divergence?

In recent work, Stegenga (2009) raises challenges for the epistemic role of convergence on the basis of the ubiquity of divergence. Having relegated convergence to only ideal situations that are rarely met, he laments that "it is unclear what support is provided to a hypothesis in the more common situation in which multiple techniques give results that are inconsistent or incongruent" (654).

Below I detail how in the case of protein structure prediction there are ways in which nonconvergent data from different experimental protocols can in fact increase the support for a model inferred from data. First, I explore why Stegenga is skeptical. He characterizes divergence in terms of dichotomous concepts, namely, inconsistency or incongruence. This static interpretation presumes that a prediction yields either X or not-X. However, this framework obscures the interactive processes that constitute effective reasoning about divergent results. For Stegenga, "Inconsistency is straightforward: petri dishes suggest x and test tubes suggest ¬x" (2009, 654). But models from experiments do not generally entail simple propositional outcomes. In the case of experimental protein prediction, a set of atomic structures are consistent with generated data, as there are uncertainties associated with every experimental protocol. Thus, comparing the results of different experiments is not best represented as strict dichotomies. The perspectival approach makes that apparent.

Stegenga's second source of skepticism involves incongruency. By this he means something close to incommensurability. Here Stegenga claims, given that different experiments have different vocabularies and metrics, there is no easy translation of the divergent data and inferences of one to the other. This is true. It is not easy. But that does not mean it is not done or that it makes the information gleaned from different experimental sources not integrable. Stegenga's definition of incongruence applies to results of protein x-ray crystallography and NMR because there is no direct translation of primary data of electron-scattering into a vocabulary of magnetization transfer. Why is this the case? An obvious difference is that the sample protein studied is in different states. Proteins are in a fixed crystalline state for x-ray crystallography but are in motion in solution for NMR. Even so, there are reasons that

https://doi.org/10.1086/710787 Published online by Cambridge University Press

suggest the results are compatible, and even complementary, and thus not incongruent.

There are some basic assumptions about molecular integrity that are shared by the theories used in designing and implementing the two experiments and in the inferences to the predicted structures that result. One observed atom cannot be in two places at the same time. There are constraints on the physically permitted torsion or dihedral angles of twists in forming secondary plane structures of a protein. There are other chemical and physical constraints that limit the possibilities for the relative positions and distances of the atoms in a protein as it folds. These apply equally to the atoms in the structures derived from electron cloud diffraction and nuclear magnetization transfer. Each experiment has additional constraints that are driven by optical theory and theories of magnetization, respectively. However, there are some shared assumptions that support the claim that the data generated from the two experiments could be compatible. One includes the identification of the same protein type by its primary structure, the unique sequence of amino acids linked together by covalent peptide bonds, which can then acquire a large number of two- and three-dimensional conformational states. Identifying which conformation is the stable, functional structure is the goal of protein structure prediction whether it is done *in silico*, *in vitro*, or *in vivo*.[4] I suggest that the process of determining congruence of the results of two experiments, pace Stegenga, is not by means of a static one-one translation of x-ray data into NMR data or by a propositional rendering of the results.

Integration: Sensory and Experimental. I suggest scientists take a different approach to divergent predictions. Rather than listing the terms, assumptions, and results of each experimental protocol side by side to see whether the evidence of different experimental data about the same phenomenon are convergent, inconsistent, or incongruent, scientists engage in an interactive and integrative process of mutual error correction. The systematic errors of one means of experimental detection can be brought to light by comparison with another means of detection. Extending the appeal to visual perception for understanding the perspectival character of representational models, I suggest there is another feature of sensory perception that suggests a way to understand the model divergence case of protein structure prediction. That is the case of multimodel sensory integration.

4. See Mitchell (2020) for arguments against the reductionist assumption that protein sequence determines structure that determines function. In that paper and in Mitchell and Gronenborn (2017), we consider functional complexity that can invoke the involvement of multiple molecules and cellular structures in the stabilization of protein structure and in the performance of function.

https://doi.org/10.1086/710787 Published online by Cambridge University Press

Our five senses, sight, sound, taste, hearing, and touch permit us to acquire information about nature. These different modalities detect different aspects of a given phenomenon. Again, perspectivism and partiality are evident, but here they are not choices we make; rather, they are what types of signals are accessible by a given sensory apparatus. Consider sight and sound. Visually we can detect reflectance, saturation, color, light reflection, and refraction. The human eye has a cornea and a lens that focuses light onto the retina. The retina includes millions of light-sensitive cells called rods and cones. When light hits the rods and cones it is converted into an electrical signal that is relayed to the brain's visual cortex via the optic nerve. What about sound? The basilar membrane in the inner ear detects frequencies of sound waves by vibration. Different frequencies activate different groups of neurons on this membrane. In addition to detecting what tone is being emitted by the target source, hearing also can locate the source of the sound by using the difference in loudness and timing between the two ears. As the eye and ear illustrate, the same organism can have multiple apparatuses to detect different features or different aspects of the same feature of one phenomenon. The auditory and visual models of the relative location of an entity, for example, may well be different as a result of the kind of causal interactions of signals from the source with the two detecting instruments.

For most tasks in which we engage something in the external world, we employ multiple senses. The representational resources vary with the sensory modality. An object affords different, but accurate, representations through these different sensory modalities just as it affords different visual representations by bees or humans. The individual senses may have optimal usefulness in different circumstances, but each is partial and perspectival, getting at some features but not others, and getting at them by different protocols each with its own range of precision.

Not surprisingly, using two senses is better than just using one. Experiments have shown that neural processes of integrating multiple modalities, like vision and audition, produces an increase in the likelihood of detecting and identifying events or objects over a single modality. What might be surprising is the degree to which multisensory integration is better; it is superadditive. Multimodal integration in the brain has been studied both at the neuronal level and by experiments on task completion: "the integrated product reveals more about the nature of the external event and does so faster and better than would be predicted from the sum of its individual contributors" (Stein and Stanford 2008, 255). So not only does using different perspectival senses provide more information than using a single modality, the cross-modal interaction of stimuli can lead to multisensory integration that yields a nonlinear, amplified neuronal response and more than additive reduction of time to task completion. Using compatible input from two senses generates more accurate models of external phenomena.

https://doi.org/10.1086/710787 Published online by Cambridge University Press

I suggest something like this happens in the case of x-ray and NMR experimental models of proteins. This is the process called joint refinement. Joint refinement is a method of mutual error correction that recognizes the different systematic biases present in each of the experimental approaches. A joint refinement procedure investigates the compatibility of the different data generated by the two methods and, when possible, attributes differences to the instrumental biases. For example, proteins that are crystallized are in a different physical state from those in solution in the NMR experiments, and that difference may account for some of the divergence in the predicted structures. In the fixed, crystallized protein, the same signal can be given by atoms in elements that are on one sheet of a protein as atoms on elements that are on distinct sheets. In NMR those sheets move independently, and thus a discrimination between atoms on one instead of two sheets is possible. Other sources of divergence between the two protein models include different degrees of error in data retrieval and in the ranges of uncertainty in the inferential algorithms used by the two experimental protocols (Carlon et al. 2016). When the known systematic biases have been resolved, then the refined data from the two experiments can be compared to see whether there are regions of overlap. The two instrumental perspectives each produce an ensemble of structures that are consistent with their individual data. While NMR scientists publish the ensemble, it is interesting that x-ray scientists publish a single optimal protein structure. However, in conjunction with the single structure, they include a measure of the quality of the atomic model obtained from the crystallographic data and, hence, de facto infer a range of consistent structures. The lower the quality, the greater the uncertainty, thus the larger the set of underdetermined structures from the x-ray data.

If there is no overlap in the NMR and x-ray ensembles of data-supported structures, then they would be determined to be incompatible or incongruent. The joint conclusion would be that the target structure is within either the NMR predictions or the x-ray predictions, so integrating both experiments would generate a prediction that is more uncertain than either one individually.[5]

If there is overlap, then the data from each can be refined to include only the overlapped structures that will be smaller than the original sets from each initial experiment. This reduces the underdetermination, and hence integrating both experiments allows the prediction of a protein structure that is more accurate than what could be obtained by either method alone. The blind spots of x-ray crystallography and NMR cannot be removed, but when they are system relative, then they can be exposed by the mutual analysis of joint

5. Of course one could also decide that one or both of the experiments was not reliable, so that the target structure would not be found in one or either of the predicted ensembles.

https://doi.org/10.1086/710787 Published online by Cambridge University Press

refinement. In the words of one protein scientist: "joint structural refinement using both NMR and X-ray data provides a method to obtain a more reliable structural model, which may disclose additional relevant information on its functional mechanisms" (Carlon et al. 2016, 1601).

Notice that there is no third perspective, no view from somewhere of the objective color of the flower or of the objective three-dimensional structure of the protein from which to judge the accuracy of sensory modalities or x-ray and NMR perspectives. There is no access outside of the different experimental modalities to reveal the "one world" that Giere posited would ensure the compatibility of the different instruments. Instead, in order to judge the accuracy of any instrument or inferred model, scientists have only the measured data from observations by the multiple detecting devices we have, interpreted by current best scientific theories. When you have data from one perspective, that is used to provide justification for a predictive structure of a protein. When you have data from a plurality of perspectives, that can provide stronger justification for a predictive structure. Joint refinement is a form of integration that is pluralism preserving. At the end of joint refinement, there is neither global unification of two theoretical perspectives or the associated methods into one nor reduction or elimination of two theoretical perspectives or the associated methods to a single perspectival modality. Each is partial, and we need them both. The representational models inferred by different perspectives from the causal interactions afforded by the targeted phenomenon may converge or diverge, may be compatible or incompatible. This example is presented to suggest that there is another option: they may be mutually informative.[6]

This parallels the story of sensory integration. Each of the modalities, vision and audition, can be improved to make better discriminations and increase precision. However, you cannot teach the eyes to hear or the ears to see. So too with experimental perspectives. We can increase the precision of detecting reflectance patterns (get better crystallization processes) or magnetization transfer (get a bigger magnet), but you cannot teach an x-ray crystallography spectrometer to detect magnetization transfer or vice versa. The information afforded by a phenomenon in its interaction with diverse modalities, when jointly integrated, can yield more accurate information than any one sense or one experiment could ever yield. In these cases, the preservation of diversity is the means to more accurate representations of nature through a dynamic, interactive, and iterative process.

Lessons for Diversity for Philosophy of Science. I want to suggest that this view has implications for how we view our own various philosophical

6. The view of pluralism I am defending shares a number of features and concerns with Chang's account of "active realism" (2012, 2018).

https://doi.org/10.1086/710787 Published online by Cambridge University Press

perspectives. We all agree that doing philosophy of science is not identical to doing science. But are there sufficient parallels to make the lesson of the value of integrative, perspectival pluralism from science apply to our own practice?

A well-rehearsed trope about philosophy of science, and one I use when responding to the puzzled look when I tell someone what it is I do for a living, is that philosophy and science are somewhat artificial disciplinary divisions; that William Whewell introduced the term "scientist" in 1833; that before that all would have been identified as "natural philosophy" or "moral philosophy"; that nowadays some scholars sit on the boundary, especially when engaged in theoretical work; that some cross the boundary, publishing and speaking to both communities; and so on. We have all been there. Whereas scientists study nature, philosophers of science step back to understand how the very enterprise of science does what it does. We study the practices and products of science.

At least for many these days, we have adopted naturalistic norms for judging our philosophical claims that include not just analytical and logical precision, conceptual clarity, and argumentative acumen but also empirical capture. That is, we study the philosophy of *science*, and so if scientific practice defies to a significant degree the picture we paint, then it seems likely that we have missed something important. We take seriously representing what scientists in practice do, as well as what they must do under certain assumptions or what they should do given the underlying norms. How do we know we accurately represent science, that is, that we are successful in the empirical part of our enterprise? I have my data, you have your data. I draw my inferences, you draw yours. We can be studying different sciences, we could be asking different questions, thus each getting a piece of the story. But often we are studying the same parts of science, trying to answer the same question, broadly speaking, but with different methodologies and getting different answers. Which of us is correct? My answer, again, is that both can be.

I propose that we extrapolate the lessons from the empirical and representational constraints on science to the empirical and representational constraints on philosophy of science. The various perspectives (i.e., the models, methods, and representations) that we as philosophers of science deploy will focus our attention on different features of science, highlighting some and ignoring others. Each choice of models, methods, and representations by a philosopher of science encodes a partial perspective about a scientific belief or practice; no one account can capture it all for all purposes to answer all questions. Thus, diversity and integration may also be tools for increased accuracy in our own profession. If, as many have documented, the ways in which social values are incorporated into science shape its content, not just to deform it but also to inform it, then looking at both the social context

https://doi.org/10.1086/710787 Published online by Cambridge University Press

and internal epistemic context together will be better at explaining the resulting scientific practices.[7] Hull (1988) exemplified this type of integrated approach. He attempted to dissolve the dichotomy between "internalist" and "externalist" accounts of science, which were then, and still are, taken as diametrically opposed alternatives. Hull, who offered an alternative, evolutionary approach to the "science of science," recognized the conceptual and social perspectivism embedded in all approaches, even his.

Philosophy of science should embrace not just social diversity for ethical and political reasons but philosophical diversity for epistemic reasons. To appreciate one way in which perspectival diversity is epistemically valuable, I suggest we embrace the features of scientific knowledge I have described in the case of protein structure prediction. Science is in the business of learning things about nature but rarely in the position of knowing propositionally what nature is like, once and for all time. Learning is a process that is ongoing within any one perspective or when integrating multiple perspectives such that any claim that "x is y" will be provisional, a snapshot of a moment. Nature affords a plurality of accounts, developed through experimental interaction and theoretical abstraction, such that many, although not all, can be accurate, that is, can capture the most reliable data. Integrating strategies, like joint refinement, that engage multiple perspectives are pluralism-preserving practices. By embracing diversity, science can increase the accuracy of models of nature and thus their adequacy in solving the problems we pose. That is not because science has delivered the true model or converges on a single account. It is because by actively integrating—searching for stabilities that go beyond one context and beyond one perspective—science is learning more about what nature is like. Philosophers can and should adopt the same integrative stance.

Further Lessons for Social Engagement. The picture of science I have painted, which embraces both the complexity of nature that scientists study and the diversity of the ways in which science investigates, represents, and uses the knowledge of nature gained, is one of enduring diversity. I have also extrapolated that picture to characterize a way of doing philosophy of science, which embraces the complexity of the nature of science and the diversity of ways in which philosophers investigate, represent, and use the knowledge of science gained by their investigations.

7. The work of many feminists and social epistemologists has exposed the problem of taking a narrow view of the factors that shape scientific knowledge. Harding's standpoint theory (1993, 2009) and Longino's situated knowledge framework (1990, 2001) are foundational to the many new developments in philosophy of science providing more sophisticated accounts of science.

https://doi.org/10.1086/710787 Published online by Cambridge University Press

While we may disagree about the details, I venture to assert that philosophers of science agree that science is a complicated, nuanced, value-influenced practice that exceeds the grasp of the limited types of logical analysis of confirmation, explanation, and prediction that many of us were trained on. The logical analyses get at important aspects of the reasoning that is central to science, but there are additional important features that we have learned are essential for accurately modeling scientific practice. Overly simple views of confirmation or falsification as ways to support the authority of scientific results as forms of proof or generating eternal facts have been superseded by more sophisticated views of evidential support and defeasible warrant. Exclusively social accounts of scientific authority also suffer from the dichotomous framework in which they have been proposed and evaluated. In short, both logical analysis and social embeddedness may be necessary to account for scientific practice, but rarely is only one sufficient (Downes 1993; Solomon 2001; Reisch 2005; Douglas 2014). More nuanced accounts of science are important not just for contributing to more accurate models of actual scientific practice but also as a corrective to a dangerous strategy for engaging in public discourse on the authority of science. This is especially significant in the current climate in which relying on science for individual and social decisions is increasingly being challenged. But while pluralism is enduring, not all proposed explanatory or predictive models are legitimate.

In order to enact the best personal decisions and public policies we can, given what we know at this time, we should let climate scientists be the source of what we know about global climate change and epidemiologists and biologists be the source of what we know about the efficacy and potential harms of viruses or vaccines. I suggest that philosophers of science are the experts to speak to what we know about evidence and warrant in science and thus have an important role to play in the debates about what counts as the facts that serve to direct individual decisions and collective policies (Fehr and Plaisance 2010). In the current context, some scientists are reluctant to publicly discuss the provisional nature of current findings or the ongoing internal debates that are crucial for scientific development.

A strategy of oversimplification, claiming that science has the eternal facts, that it gets them by proof, while seemingly wielding a powerful sword against so-called alternative facts and interest- and ideology-driven reasoning, can lead to suspicion and rejection of all scientific claims. There are two ways in which this can happen. One is "Yesterday scientists showed that X is bad for you, but today it is reported that X is good for you." Fill in the X with wine, aspirin, butter, and so on. The second is that the lack of complete consensus in a scientific community is taken to indicate that we do not know anything. Until there is proof "more research is needed" and thus no policy should be informed by what, say in climate science, is mistaken to be a lack of consensus (Cook et al. 2016). There is confusion on what the

https://doi.org/10.1086/710787 Published online by Cambridge University Press

consensus is about (Oreskes 2014), which has allowed climate deniers to elide the distinction between skepticism and denialism (the deniers call themselves climate skeptics). So any scientific disagreement on any component of the complex set of factors relating to the human role in global climate change is equated with denial of climate change. Complete consensus on all related claims is not the indicator of scientific knowledge.

If a view of good science as proof, eternal truth, and complete consensus is publicly promoted and accepted, then any chink in that fictional armor can be used as grounds for disregarding what scientists actually can and do know about nature. We, as philosophers of science, have a responsibility to promote a more accurate account of science. While some might think it is risky to promote an account of science that acknowledges the provisional nature of current knowledge, and the roles of disagreement and consensus (see John 2018), failing to explain these complexities may well contribute to the far greater danger that the public ceases to accept the authority of scientific findings. I believe we can provide a better, more accurate understanding of science to the public, to prevent the rejection of science driven by oversimplified accounts. Of course, which strategy to adopt will depend on the social and political context in which we act.

We saw the consequences of failure to understand evidence tragically in November 2018 in my community in Pittsburgh when an anti-Semitic gunman killed 11 individuals at the Tree of Life Synagogue. The terrorist's explicit motivation was his belief that Jews were responsible for the migrant invaders in a convoy in Mexico coming to the US to kill his (meaning white, male, Christian) people. In an earlier post the shooter accused the "evil Jews" of bringing in the "filthy, evil Muslims" (Gessen 2018). Why would anyone think the migrants are dangerous, that they are invaders rather than asylum seekers, and that Jews were responsible? What evidence is there that would support these claims? With respect to the claim that there are ISIS terrorists in the caravan, President Trump replied to the question from a reporter, "is there proof of them being in the caravan?" with "you've heard it before" and "it happens all the time." When the question was repeated, "so there is no proof then?" about ISIS being in the caravan, President Trump claimed, there is "no proof of anything, but they could very well be" (Hutchinson, Phelps, and Gutman 2018).

I interpret this exchange as illustrating a distorted view of evidence by presuming that if there is no proof then anything can be asserted. If the choice is dichotomous, proof or anything goes, then claims based on ideology or economic interests or religious doctrine are just as good as those based on unproved but strongly warranted scientific evidence. This is a dire consequence of misunderstanding the nature of science's epistemic authority to ground good decisions and policies. Reason will not sway those who willfully distort, or those who are unable to reason, but there need to be

https://doi.org/10.1086/710787 Published online by Cambridge University Press

more voices in the public sphere providing more sophisticated and accurate accounts of science. There need to be the voices of philosophers of science.

Conclusion. The lessons I have learned from studying the sciences of evolutionary biology, social insect complexity, behavioral genetics, and protein structure have led me to develop a pluralistic, pragmatic, and dynamic understanding of science: integrative pluralism. This is a view that highlights and explicates the epistemic value of diversity. I have suggested that those lessons can also be applied to our own practices as philosophers of science. Furthermore, I suggest that the diversity of perspectives we have developed regarding the methods that ground the authority of science, nuanced and provisional though that authority is, gives us an important role to play in public discussions of how science can inform and improve individual lives and public policies.

REFERENCES

Carlon, Azzura, Enrico Ravera, Janosch Hennig, Giacomo Parigi, Michael Sattler, and Claudio Luchinat. 2016. "Improved Accuracy from Joint X-Ray and NMR Refinement of a Protein-RNA Complex Structure." *Journal of the American Chemical Society* 138:1601–10.

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

Chang, Hasok. 2012. *Is Water H_2O? Evidence, Pluralism and Realism*. Dordrecht: Springer

———. 2018. "Realism for Realistic People." *Spontaneous Generations* 9:31–34.

Contessa, Gabriele. 2007. "Scientific Representation, Interpretation, and Surrogative Reasoning." *Philosophy of Science* 74:48–68.

Cook, John, et al. 2016. "Consensus on Consensus: A Synthesis of Consensus Estimates on Human-Caused Global Warming." *Environmental Research Letters* 11 (4). https://doi:10.1088/1748-9326/11/4/048002.

Douglas, Heather. 2014. "The Moral Terrain of Science." *Erkenntnis* 79:961–79.

Downes, Stephen M. 1993. "Socializing Naturalized Philosophy of Science." *Philosophy of Science* 60:452–68.

Dupré, John. 1993. *The Disorder of Things: Metaphysical Foundations of the Disunity of Science*. Cambridge, MA: Harvard University Press.

Fehr, Carla, and Kathryn S. Plaisance. 2010. "Socially Relevant Philosophy of Science: An Introduction." *Synthese* 177:301–16.

Gessen, Masha. 2018. "Why the Tree of Life Shooter Was Fixated on the Hebrew Immigrant Aid Society." *New Yorker*, October 27. https://www.newyorker.com/news/our-columnists/why-the-tree-of-life-shooter-was-fixated-on-the-hebrew-immigrant-aid-society.

Giere, Ronald N. 1999. *Science without Laws*. Chicago: University of Chicago Press.

———. 2006a. "Perspectival Pluralism in Scientific Pluralism." In *Scientific Pluralism*, ed. Stephen H. Kellert, Helen E. Longino, and C. Kenneth Waters, 26–41. Minneapolis: University of Minnesota Press.

———. 2006b. *Scientific Perspectivism*. Chicago: University of Chicago Press.

———. 2010. "An Agent-Based Conception of Models and Scientific Representation." *Synthese* 172:269–81.

Harding, Sandra. 1993. "Rethinking Standpoint Epistemology: What Is Strong Objectivity?" In *Feminist Epistemologies*, ed. Linda Alcoff and Elizabeth Potter. New York: Routledge.

———. 2009. "Standpoint Theories: Productively Controversial." *Hypatia* 24:192–200.

Hudson, Robert. 2013. *Seeing Things: The Philosophy of Reliable Observation*. Oxford: Oxford University Press.

https://doi.org/10.1086/710787 Published online by Cambridge University Press

Hull, David. 1988. *Science as a Process: An Evolutionary Account of the Social and Conceptional Development of Science*. Chicago: University of Chicago Press.
Hutchinson, Bill, Jordyn Phelps, and Matt Gutman. 2018. "Trump Admits He Has 'No Proof' of 'Middle Easterners' in Caravan, 'But There Could Very Well Be.'" ABC news, October 23. https://abcnews.go.com/US/trump-admits-proof-middle-easterners-caravan/story?id=58686056.
John, Stephen. 2018. "Epistemic Trust and the Ethics of Science Communication: Against Transparency, Openness, Sincerity and Honesty." *Social Epistemology* 32 (2): 75–87.
Kitcher, Philip. 1990. "The Division of Cognitive Labor." *Journal of Philosophy* 87:5–22.
———. 2001. *Science, Truth, and Democracy*. New York: Oxford University Press.
Kuorikoski, Jaakko, and Caterina Marchionni. 2016. "Evidential Diversity and the Triangulation of Phenomena." *Philosophy of Science* 83:227–47.
Longino, Helen. 1990. *Science as Social Knowledge*. Princeton, NJ: Princeton University Press.
———. 2001. *The Fate of Knowledge*. Princeton, NJ: Princeton University Press.
Massimi, Michela. 2016. "Four Kinds of Perspectival Truth." *Philosophy and Phenomenological Research* 96 (2): 342–59.
———. 2018. "Perspectival Modeling." *Philosophy of Science* 85:335–59.
Mitchell, Sandra D. 1992. "On Pluralism and Competition in Evolutionary Explanations." *American Zoologist* 32 (1): 135–44.
———. 2003. *Biological Complexity and Integrative Pluralism*. Cambridge: Cambridge University Press.
———. 2009. *Unsimple Truths: Science, Complexity and Policy*. Chicago: University of Chicago Press.
———. 2012. "Emergence: Logical, Functional and Dynamical Accounts." *Synthese* 185:171–86.
———. 2020. "Perspectives, Representation and Integration." In *Understanding Perspectivism: Scientific Challenges and Methodological Prospects*, ed. Michela Massimi and Casey D. McCoy. Milton Park: Taylor & Francis.
Mitchell, Sandra D., and Angela M. Gronenborn. 2017. "After Fifty Years, Why Are Protein X-Ray Crystallographers Still in Business?" *British Journal for the Philosophy of Science* 68:703–23.
Oreskes, Naomi. 2014. "The Scientific Consensus on Climate Change: How Do We Know We're Not Wrong?" In *Climate Change: What It Means for Us, Our Children, and Our Grandchildren*, ed. Joseph F. C. DiMento and Pamela Doughman, 105–48. Cambridge, MA: MIT Press.
Perrin, Jean. 1923. *Atoms*. Trans. D. L. Hammic. New York: Van Nostrand. Originally published as *Les Atomes* (Paris: Alcan, 1913).
Reisch, George A. 2005. *How the Cold War Transformed Philosophy of Science: To the Icy Slopes of Logic*. Cambridge: Cambridge University Press.
Rorty, Richard. 1979. *Philosophy and the Mirror of Nature*. Princeton, NJ: Princeton University Press.
Salmon, Wesley C. 1984. *Scientific Explanation and Causal Structure of the World*. Princeton, NJ: Princeton University Press.
Sismondo, Sergio, and Nicholas Chrisman. 2001. "Deflationary Metaphysics and the Nature of Maps." *Philosophy of Science* 68 (Proceedings): S38–S49.
Sklar, Lawrence. 1999. "The Reduction(?) of Thermodynamics to Statistical Mechanics." *Philosophical Studies* 95:187–202.
———. 2015. "Philosophy of Statistical Mechanics." In *Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta. Stanford, CA: Stanford University. https://plato.stanford.edu/archives/fall2015/entries/statphys-statmech/.
Solomon, Miriam. 2001. *Social Empiricism*. Cambridge, MA: MIT Press.
Stegenga, Jacob. 2009. "Robustness, Discordance, and Relevance." *Philosophy of Science* 76:650–61.
Stein, Barry E., and Terrence R. Stanford. 2008. "Multisensory Integration: Current Issues from the Perspective of the Single Neuron." *Nature Review Neuroscience* 9:255–66.
Teller, Paul. 2001. "Twilight of the Perfect Model Model." *Erkenntnis* 55:393–415.
Toulmin, Stephen. 1953. *The Philosophy of Science: An Introduction*. London: Hutchinson.
Woodward, James F. 2005. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.
Wylie, Alison H. 1999. "Rethinking Unity as a 'Working Hypothesis' for Philosophy of Science: How Archaeologists Exploit the Disunities of Science." *Perspectives on Science* 7:293–317.

https://doi.org/10.1086/710787 Published online by Cambridge University Press