



The tragedy of the canon; or, path dependence in the history and philosophy of science

Agnes Bolinska^a, Joseph D. Martin^{b,*}

^a Department of Philosophy, University of South Carolina, United States

^b Department of History, Durham University, United Kingdom

ARTICLE INFO

Keywords:

Canon formation
Case studies
Confrontation model
Integrated history and philosophy of science
Path dependence

ABSTRACT

We have previously argued that historical cases must be rendered canonical before they can plausibly serve as evidence for philosophical claims, where canonicity is established through a process of negotiation among historians and philosophers of science (Bolinska and Martin, 2020). Here, we extend this proposal by exploring how that negotiation might take place in practice. The working stock of historical examples that philosophers tend to employ has long been established informally, and, as a result, somewhat haphazardly. The composition of the historical canon of philosophy of science is therefore path dependent, and cases often become stock examples for reasons tangential to their appropriateness for the purposes at hand. We show how the lack of rigor around the canonization of case studies has muddied the waters in selected philosophical debates. This, in turn, lays the groundwork for proposing ways in which they can be improved.

1. Introduction

Philosophy of science assumes as its object of inquiry the entire edifice of human efforts to acquire and systematize natural knowledge. However, to the extent that they confront the actual practice of science, philosophers address themselves to only a sliver of all that occurs or has occurred under that umbrella. The legitimacy of philosophy of science as an empirical endeavor with aspirations to even local generality therefore requires that the minority of scientific activity it addresses is more representative, more influential, or otherwise more important than the vast majority it neglects (Mizrahi, 2015).

Many critics question whether the empirical foundation philosophers of science often use—composed of historical case studies—in fact exhibits these virtues. The latitude philosophers have when assembling, selecting, interpreting, and deploying their case studies, the worry goes, encourages them to develop examples that appear to support focused philosophical objectives, but which are riddled with bias that undermines their generality, and which might subvert their narrower aims as well. These challenges have drawn attention from those defending the utility of history for philosophy of science (e.g., Burian, 2001; Chang, 2012; Sauer and Scholl, 2016; Schickore, 2011; Mizrahi, 2020; Schindler & Scholl, 2020) as well as those advancing more skeptical accounts of the

use of empirical, especially historical, cases to evidence normative philosophical claims (e.g., Chakravarty, 2017; Giere, 1973; Kinzel, 2015a; Pitt, 2001; Tambolo, 2018).

In response to these issues, we have argued that historical cases can plausibly serve as evidence for philosophical claims once they are rendered *canonical*, where canonicity is established through iterative negotiation among historians and philosophers of science (Bolinska & Martin, 2020). Cases can be canonical with respect to a philosophical aim if, through that negotiation, they are shown to be sensitive to the right kinds of causal factors. A canonical case for questions about evidence and theory choice, for instance, should be one where careful historical scrutiny shows that scientists adopted theories for evidential reasons, rather than, say, ideological ones. Philosophical discourse often proceeds alongside revision and refinement of such cases.

Understanding case studies in terms of canonicity highlights a further question that has received little attention: how has the existing stock of canonical—or at least well-developed—cases been assembled?¹ We address this question by exploring the path dependencies that shape how historical cases become canonical for given philosophical aims, and by suggesting principles to guide more effective canonization practices. We first review the stakes of canonicity for recent discussions of how to integrate history and philosophy of science (HPS), aligning ourselves

* Corresponding author.

E-mail addresses: bolinska@mailbox.sc.edu (A. Bolinska), joseph.d.martin@durham.ac.uk (J.D. Martin).

¹ Schindler and Scholl (2020), in a notable exception, raise the question of why particular historical case studies are singled out for development in the course of comparing historical case studies in philosophy to model organisms in biology, an analogy in consonance with our view. See section 5.1.

with those who advocate iterative exchange within HPS. In section 3, we show how such iteration has in fact occurred productively in the past. The canon resulting from such informal iteration has nevertheless been assembled haphazardly. In section 4, we assess what we call *the tragedy of the canon*: the difficulties that emerge at the level of communal aims when we have a stock of well-developed case studies crafted to serve individual aims. In section 5, we enumerate key aspects of canonization, showing how attending to these aspects can help alleviate some of the unintended consequences of informal and haphazard canonization practices. In section 6, we close with some reflections on what our assessment implies for the professional labor of HPS scholars.

2. Canonicity and HPS

Exploring how canonicity functions within HPS contributes to larger discussions of the relationship between history of science and philosophy of science. These began in the Anglophone world when historicist philosophers like Thomas Kuhn and Norwood Russell Hanson brought the two perspectives together in the 1960s and it has evolved in response to shifting trends in HPS. Jutta Schickore (2018) traces thinking about the relationship between history of science and philosophy of science from these headwaters through to the recent resurgence of such discussions.

Schickore critiques what she calls the *confrontation model*, according to which philosophers devise theories about science and historians provide evidence with which to confront them. According to Larry Laudan (1989), who articulated this model as a foundation for discussions about rational theory choice, this resembles the common division of scientific labour between theoreticians and experimentalists. Schickore argues that this analogy misrepresents the nature of analyses of science (be they historical, sociological, anthropological, or what have you), which should instead “be characterised in terms of interpretation, clarification, and explication of scientific concepts and arguments” (2011, p. 471).

Hermeneutical analysis is Schickore’s preferred approach to the study of science, in which “initial case judgments—judgments that identify portions of the historical record as noteworthy—and provisional analytic concepts are gradually reconciled until they are brought into equilibrium” (2011, p. 471). A philosophical position is thus not established by surviving a confrontation with historical (or sociological, or anthropological) evidence; it rather emerges from an iterative process in which informational and cognitive resources from many disciplines might play a role.

We endorse this view, with some caveats. First, a hermeneutic approach is not so dissimilar to an appropriately nuanced confrontation model, which relies on a closer analogy with scientific practice. A naïve confrontation model, in which theoreticians and experimentalists work in isolation from each other, would fail as a model of most scientific theory construction.² Data must be interpreted before it can be brought to bear on theory; further, theory can be revised in small or significant ways in light of data. Theorists and experimentalists often collaborate closely, and individuals might act as both. Thus, the “confrontation” that occurs in science contains hermeneutic elements, so Schickore’s hermeneutic model is at least loosely aligned with scientific theory construction conceived in this way. Raphael Scholl (2018) makes a similar point when he describes confrontation between philosophical theory and historical evidence not as a one-time event, but as an iterative process through which abstract theory is refined upon repeated contact with concrete historical cases (see also Chang, 2012; Vickers, 2013; Kuukkanen, 2018).

² Similarly, Scholl (2018) argues that the confrontation model assumes scientific theory testing to proceed either via enumerative induction or the hypothetico-deductive method. If we adopt a broader perspective on how scientific theory can “confront” evidence—such as the way in which mechanistic hypotheses in biology are empirically founded—the model no longer faces the problems Schickore attributes to it.

These characterizations of the relationship between philosophy and historical cases emphasize the complex give-and-take that marks any productive interaction between bodies of knowledge and practice, but they are abstract and would benefit from more concrete specification. What does it mean for “preliminary concepts and points of view and initial case judgments [to be] brought together and mediated and adjusted until a cogent account is obtained” (Schickore, 2018, p. 472)? What does “improv[ing] our abstract models by refining, adding, or removing components” amount to, and what does it look like for “explanans and explanandum [to] evolve together over time” (Scholl, 2018, p. 226)?

We present the canonization of case studies as a key mechanism of iterative exchange in HPS. A case can become canonical for a philosophical aim only if its outcome is sensitive to the factors at issue in that question. For instance, the debate over why Arthur Stanley Eddington reported results validating Albert Einstein’s general theory of relativity following his 1919 eclipse expedition hinges on such questions of sensitivity (see Virtajoki, 2018). Some have argued that Eddington discarded portions of his data for ideological reasons: a pacifist and conscientious objector, he sought to reestablish ties with Germany following World War I and was favorably disposed to results supporting Einstein (e.g., Collins and Pinch, 1993; Waller, 2002). If that is so, then the case cannot help arbitrate disputes about the relationship between theory and data, such as those over novel prediction. A world in which Eddington would have validated general relativity *no matter the evidence* (or, at least, within suitably wide variance in the evidence) is a world in which that case can give us little salutary advice about reasoning from evidence. The case would thus be disqualified from the canon for such questions. Others, though, have argued that Eddington and his team discarded data for good methodological reasons—such as prudent mistrust of the behavior of well-understood instruments (Kenefick, 2019), or because being guided by theory when interpreting data should be regarded as methodologically sound (Schindler, 2013). If we are convinced that Eddington’s conclusions reflected sound practice, then we have reason to believe that different evidence would have compelled him to report different results, and so the case can be reinstated as a canonical one for philosophical aims germane to evidential reasoning (Bolinska & Martin, 2020).

The canonization of a case, that is, relies on just the sort of iteration that Schickore, Scholl, and others have called for. Such iterative scrutiny from historical and philosophical perspectives allows the many dimensions of a case to become well-enough understood to qualify the case for canonical status. Philosophical tools allow us to evaluate the salience of cases for well-defined purposes; incorporating these evaluations into our historical perspectives allows us to revise our understanding of the causal structure of those cases, often in light of their use for philosophical ends. The resulting back-and-forth—described in the next section—has the potential both to deepen our understanding of the relevant historical episode and to sharpen our sense of its salience to our philosophical aims.³

³ How should we delineate philosophical aims? At one extreme, we could discuss the aims of philosophy of science writ large—for example, to understand and account for the proper function of scientific reasoning. At the other extreme, philosophical aims might be sliced so thinly that they are distinctive to individual philosophers. It would be unreasonable to demand that case studies be canonical for goals so general as understanding the proper function of all scientific reasoning, and equally senseless to insist upon a new canon for each idiosyncratic take on an existing question. Our rough-and-ready guide is that, for practical reasons, we should seek canonical cases for intermediate-scale aims: broad enough to sustain discourse among multiple interlocutors but focused enough that a well-developed case study of genuine science might *not* be salient for discussing them. The reason, aside from the practical one, is that this is the scale at which we see community-level intellectual exchange of the type that constitutes iterative interaction among historians and philosophers of science. We thank Peter Vickers for raising this important consideration.

3. Iteration in action

Theoretical prediction of a novel empirical phenomenon seems like a great triumph, especially if that phenomenon is unexpected. The prediction that the shadow of a rigid disc would show a bright spot at its center, so the story goes, was a triumph for the wave theory of light, integral to its acceptance in the mid-nineteenth century. But does this imply that we should accord more epistemic weight to a theory that successfully predicts novel empirical phenomena than to one that merely accommodates existing evidence? After all, the wave theory also accommodated known phenomena, such as diffraction of light around straight edges; should its accommodation of straight-edge diffraction hold less sway than its prediction of diffraction by a circular disc?

We might approach this case in one of two ways. We might regard the epistemic value of a prediction as dependent strictly on the logical relationship between theory and evidence. From this perspective, historical examples do not tell us anything about the evidential value of novel prediction. At best, they can serve an illustrative function, making abstract claims more concrete. It would not matter whether scientists *in fact* accorded greater evidentiary weight to theories that made successful novel predictions. That is a descriptive point; the issue at hand is a normative one. But we could also approach the question of the value of novel prediction from the perspective of accounting for the success of past science. We might want to know whether past scientists accorded theories that made successful novel predictions greater evidentiary weight and, further, whether those theories stood the test of time—whether they continued to be well regarded. Here, the details of the examples matter because they bear on the substance of the philosophical argument. This approach, in which case studies can serve as *evidence*, has become more common in the wake of the postpositivist turn (Soler et al., 2014; Dresow, 2020).

Consider again the story of the bright spot, a classic case in the history of optics. In caricatured form, it stands as a striking testament to the power of novel prediction.⁴ In the early 1800s, disagreement reigned about whether light was a particle or a wave. Given the success of Isaac Newton's physics, his corpuscularist understanding of light informed the orthodox view. But challenging empirical phenomena, such as dispersion and diffraction, gave wave theories traction. With these debates in view, the French Académie des Sciences, in 1817, announced an essay competition on the subject of diffraction. Augustin-Jean Fresnel, his eyes on this coveted prize, penned a memoir showing the advantages of the wave nature of light for explaining diffraction phenomena. Siméon Denis Poisson, a member of the prize committee and a committed corpuscularist, remained skeptical and derived an apparently absurd consequence of Fresnel's view: an opaque disk illuminated by a point light source should cast a shadow with a bright spot at its center. François Arago, a member of the committee sympathetic to the wave theory, performed the experiment and found the bright spot. The wave theory carried the day.

For someone already convinced that novel prediction provides epistemic support for theories—and viewing the issue from the logical-relationship perspective outlined above—this offers a vivid illustration of the point; the case's particulars are merely academic. But the turn toward evidential use of case studies suggested that reading the historical particulars more closely was necessary to determine whether the case *supports* the philosophical claim it had hitherto been used to *illustrate*. Did scientists in fact accord greater weight to the wave theory because of its novel predictions? Did the prediction change anyone's mind, and why, or why not? The potted story above has rhetorical appeal, but can it withstand historical scrutiny?

John Worrall chastised Edmund Whittaker (1951), Thomas Kuhn (1962), and Ronald Giere (1984) for failing to look beyond the story's attractive veneer when deploying it in the course of arguments favoring

novel prediction. A closer look at the sources paints a less compelling picture. Fresnel competed for the prize against only one other mediocre submission. The prize citation studiously avoided mentioning the underlying theory of light at play; it pointed to practical rather than theoretical elements: the experimental techniques and calculational tools Fresnel introduced. Further, the corpuscularists on the prize committee remained steadfast in their metaphysical commitments, at least for some time thereafter. Most significantly, the prize committee was impressed by how straightforwardly Fresnel's theory was able to *explain* straight-edge diffraction. The bright spot, that is, appears to have played at best a minor role in the story of how the nineteenth-century physical community came to accept the wave theory of light. Partly on this basis, Worrall rejects *temporal* predictivism—the doctrine that predictions that are made chronologically prior to the observation of their predicted effects carry more epistemic weight (Worrall, 1989).

This is not to suggest that more careful historical work unproblematically motivates philosophical consensus. Reconstructions such as Worrall's, and similar interventions by Stephen Brush (2015), though they have made it unpopular to espouse temporal predictivism, have not settled the issue. They have, however, contributed to a series of shifts in the conversation. One is the shift toward the broader notion of use-novel prediction, which considers whether a theory accounts for evidence not used in its construction, irrespective of whether that evidence was available before or after the construction of the theory. Another is an increasing emphasis on the role of novel predictions—temporal or use-novel—in scientific practice, as opposed to prescriptive epistemology (see Douglas & Magnus, 2013).

We emphasize that Worrall did not single-handedly dethrone an old interpretation and usher in a new regime in the predictivism debate. Skepticism of temporal predictivism had been brewing for some time before Worrall's 1989 article (e.g. Zahar, 1973), as had support for use-novel prediction (e.g. Gardner, 1982; see also Barnes, 2018). Worrall's argument for the evidential inadequacy of the bright-spot case nevertheless sparked the realization that it is “ridiculously hard to find any historical evidence that would support the thesis of the special epistemic weight of temporally novel predictions” (Schindler, 2008b). We can therefore understand Worrall's intervention as an anchor point in an iterative communal process that shifted both our understanding of key historical cases and the terms of the philosophical debate around prediction and accommodation. Brush's historical research program, for instance, builds on Worrall's insights to show the impotency of novel predictive success for motivating theory choice in a range of examples (collated in Brush, 2015). And Hitchcock and Sober (2004) take Worrall's interpretation of the bright-spot case on board in the process of recovering a role for predictive success in theory assessment.

This episode clarifies how sustained attention to a case as deployed *evidentially* in a philosophical context benefits both history and philosophy. The evidential needs of a philosophical debate inspired closer historical analysis. What seemed an ideal illustration of a successful novel prediction straightforwardly supporting a new theory turned out to be more complicated. Careful historical scrutiny yielded a more nuanced and accurate historical case, and inspired a new historical approach to other cases. That scrutiny turned up surprising results from the perspective of the naïve predictivist. The results of historical investigation thereby moved the philosophical discussion in more productive directions.

We can discern another type of benefit in this example: to the robustness of the case for philosophical purposes. Worrall's historical work on the bright spot, and the episode's subsequent reinterpretation by supporters of novel prediction, made it *canonical* in the predictivism literature. That is, it was rendered suitable for use as evidence in the course of historical and philosophical interrogation, whereas the potted story, absent careful historical reconstruction, was no better than a thought experiment—valuable only insofar as it could illuminate arguments substantiated on other grounds. It is important, that is, to distinguish between reuse and development. A case study might be propagated

⁴ For a synoptic but nuanced account of the history of optics, see Darrigol (2012).

through a philosophical debate without broadening our understanding of that case.⁵ But iterative scrutiny from historical and philosophical perspectives renders the many dimensions of a case well-enough understood to merit its use for the philosophical purpose at hand, and only thereby can that case become canonical.

Examining the use of historical cases in the philosophy of science reveals that the productive iteration previous commentators endorse already takes place in a fruitful way, at least some of the time.⁶ We demonstrate in the next section that these iterative exchanges have an aggregate effect on what we will call the global canon—the set of all cases that are commonly used for philosophical purposes. Attending to potential pitfalls of this process can help us develop strategies to refine these exchanges.

4. The tragedy of the canon

The negotiation between historians and philosophers results in a case becoming canonical in one sense. But we must distinguish three distinct senses of the terms “canonical” and “canon.” These terms commonly refer to works belonging to an established corpus—for instance in literature—on one hand, and to works that are authoritative or accurate on the other. In line with this distinction, we can refer to the set of all the cases that have, in practice, been used (to a non-trivial extent) to address *some* philosophical problem or other as the *global canon*. Within the global canon, however, cases might have differing status. We will call a case study that corresponds to the “authoritative or accurate” sense of canonicity *saliently canonical*, in line our previous use of the term to designate a case that is appropriate for addressing a specific philosophical question (Bolinska & Martin, 2020). Other members of the global canon, however, might not clear this bar. These are *folk canonical* cases that have traditionally been applied to particular philosophical problems, yet have not undergone the scrutiny necessary to determine their salience, or which continue to be applied after they have, and have been found wanting. One goal for improving the canon is to ensure that as high a proportion as possible of the cases that compose it are saliently canonical; however, the characteristics of a global canon might have features unsuited to our aims even if salience is maximized.

This is so because canonization in any discipline is a social process of consensus formation; canons form in ways that reflect the contingencies of the communities assembling them (see Biagioli, 1996; Schaffer, 1996). An analogy to literary canons can clarify the point. Nations and linguistic communities have works considered canonical within them. We can interrogate the composition and mode of assembly of these canons, and debate whether individual works merit canonical status. Simultaneously, however, we might also evaluate characteristics of the collective body of works that enjoy canonical status within *any* national or linguistic tradition. Feminist literary critics, for instance, have warned of the consequences of a global canon of this sort that skews overwhelmingly male. Historical gender bias—locked in because canonical texts are entrenched—produces an emphasis on an unrepresentative set of perspectives and experiences within the literary canon (see Showalter, 1971).

Our objective is analogous. We are interested, first, in how cases can be made saliently canonical, and, second, in how saliently canonical and

folk canonical cases combine to form a global canon. In some cases, the global canon might have features amenable to our aims. However, we might find that—like in the overwhelmingly masculine literary canon—biases, lock-in effects, and differences of emphasis within it produce a global canon that is unrepresentative of key elements of scientific practice and so lacks the resources to address important philosophical aims. Recognizing such skews in the global canon is not necessarily an indictment of any specific slice of it. The fact that we find men problematically overrepresented in the literary canon does not mean we should think less of Shakespeare, or Tolstoy, or Borges, but it does point toward better ways to construct the literary canon moving forward. Similarly, difficulties with the global canon of HPS need not be an indictment of individual philosophers’ use of cases, but they do suggest practices that can better serve community needs.

In examining the global canon and the process by which it developed, we find a version of the tragedy of the commons—the tendency of groups of self-interested individuals to mismanage shared resources via rational custodianship of their personal resources. Individual philosophers, focused on problems of interest to them, select, deploy, and critique cases in ways suited to their aims. The result is a global canon that reflects the particularities of individual philosophical aims, but does not necessarily exhibit features that would make it most useful for the philosophy of science community. We call this effect *the tragedy of the canon*.

The tragedy of the canon manifests itself in two ways. First, the philosophical community is composed of individuals who choose examples to suit their aims, and choose which examples of others to critique in line with those aims. This process necessarily generates a canon the constitution of which is subject to some bias. It reflects the idiosyncratic interests and aims of the members of the community. As a result, it might be skewed in ways that make it suboptimal for addressing the needs of future philosophers with different (or even relevantly similar) aims. For instance, early efforts to understand features of scientific models focused overwhelmingly on simple physical models such as the model of the simple pendulum (e.g., Giere, 1988) or the billiard ball model of gases (e.g., Hesse, 1966).

This problem is exacerbated because emerging philosophical issues rarely build a stock of canonical cases *de novo*; new issues evolve from old ones, and so tend to draw from the global canon, as existing cases are adapted to new purposes. Such path dependence in the global canon can constrain philosophical inquiry. Consider the overwhelming focus on physics in the philosophy of science through the 1990s, which created a skew that made the global canon inapt for addressing questions distinctive to philosophy of biology and other special sciences. Philosophy of biology’s emergence as a major subspecialty therefore depended on cultivating new historical feedstocks (see Griffiths, 2008). We have reason to suspect, that is, that philosophers’ individual critical choices, even if independently methodologically sound, might not generate a global canon suited to the needs of the wider community, which often aims to answer different questions.

This difficulty extends to the internal structure of the philosophy of the special sciences. Philosophers of biology have long focused intensely on evolutionary biology, for instance, and still do to a great extent.⁷ This comes at the expense of subjects like developmental biology, ecology, and molecular biology (the last of which has received considerable attention, but largely in the context of evolutionary genetics). And because philosophers of biology—at least on the face of it—make claims that are representative of biology in general, we end up with a skewed understanding of that field. That is, we see biology as primarily concerned with evolution and we understand its central questions and methods to be those of evolutionary biology. This problem has seen some

⁵ For an illustration of the difficulties this can cause, see Potters (2019).

⁶ This is one illustration. Other cases include, *inter alia*: Eddington’s prediction of starlight bending (Collins and Pinch, 1993; Earman and Glymour, 1980; Kennefick, 2019; Schindler, 2013); the debate among Hasok Chang (2003), Kyle Stanford (2003), and Stathis Psillos (1994) about scientific realism and the caloric theory of heat; and the London brothers’ model of superconductivity (Bueno et al., 2012; Cartwright et al., 1995; Potters, 2019). See also the examples in Chang (2012). This is not to say that all cases are subject to such iterative scrutiny, but to observe that it has already proceeded productively in some cases.

⁷ Philosophy of biology textbooks typically represent central issues in the field; examining these gives us a sense of the overwhelming representation of evolutionary biology through the 1980s (e.g. Sober, 1984), and its continued prominence today (e.g. Godfrey-Smith, 2014).

redress in recent years, as more philosophers focus on other subfields of biology, such as those mentioned above.⁸ Moreover, related specialties such as the philosophy of neuroscience have come to prominence in their own right (see Bickle, 2009). But the point stands that whatever cases we have, we come to focus on, and these shape our philosophical understanding of the field.

The same can be said for physics, the history and philosophy of which has been shaped by a focus on high energy physics and cosmology. But these specialties are unrepresentative of much of physics. High energy physics, for instance, maintains an unusually sharp division of labor between theory and experiment. Its experiments are orders of magnitude larger and more complex than those to be found in other areas of physics. Its theoretical program is defined by a virulent reductionist ideology that finds little purchase elsewhere in the physics community. And it deals with phenomena far remote from our sensory experience. But areas such as solid state, condensed matter, materials, plasma, and chemical physics, which represent the majority of physical investigation over the past half century or so, differ dramatically in these critical respects (see Martin & Janssen, 2015; Weisel, 2017; Martin, 2018).

In philosophy of science, that is, we find a version of the “WEIRD” problem that has been identified in behavioral science. Most experimental studies that support general claims about human behavior are conducted on American university students, a population that is largely white, educated, industrialized, rich, and democratic. This combination of traits is not representative of most of the world, and so conclusions drawn from studies conducted on WEIRD populations are suspect if applied to others (Henrich et al., 2010). Similarly, in philosophy of science, our understanding of science writ large has been shaped by which cases we take to be canonical for particular questions—by the global canon—which is itself weird in some important respects.

A concrete example of the constraining effect the constitution of the global canon can have on an emerging philosophical debate can be found within the growing literature on the contingency/inevitability (C/I) issue, which focuses on whether the results of successful scientific inquiry are contingent or inevitable. It is of comparatively recent vintage, having been articulated first by Ian Hacking in *The Social Construction of What?* (Hacking, 1999), and subsequently expounded in a series of special issues, edited volumes, and stand-alone articles (Soler, 2008; Martin, 2013; Kinzel, 2015b; Soler, Trizio, and Pickering, 2015). Discussions of the C/I issue have clustered tightly around a few central cases, largely from quantum mechanics and high energy physics, which share peculiar features, such as being highly mathematical and involving the status of unobservable entities. As Catherine Allamel-Raffin and Jean-Luc Gangloff observe: “In our C-I corpus, the heterogeneity of the sciences does not appear, and remains widely to be explored” (Allamel-Raffin & Gangloff, 2015, p. 106).

At first blush, it might seem odd that a philosophical issue that emerged alongside the flourishing of new research into philosophy of biology, medicine, and the historical and social sciences would cleave so tightly to hoary cases from the history and philosophy of physics. But this is less mystifying when we consider that the issue also emerged in parallel with the turn toward the evidential use of historical cases. Early contributors to the C/I issue therefore looked for cases from the HPS literature that were well-developed enough to sustain evidential use, and fell upon discussions of interpretations of quantum mechanics (e.g., Cushing, 1994) and the existence of quarks (Pickering, 1984; Galison, 1987). Isolated authors have injected other examples into the discussion. Work on contingency in early genetics, for instance, has developed cases that offer promising insight into the consequences of the Mendelian paradigm’s success (Radick, 2005; Jamieson & Radick, 2013). The genetics case promises to push the C/I debate in new directions, but it has had to be constructed anew and so has not exerted the same influence on the debate

as older examples from physics. Few other cases have undergone the iterative scrutiny that would render them saliently canonical, and so the question of the extent to which dominant positions in the debate are general, or particular to certain types of physical inquiry, remains open.

This case illustrates how philosophers of science can hardly but be constrained by the global canon. Given that individuals are embedded in communities and that HPS students receive training that includes canonical cases, it is implausible that they would not be guided to an extent by existing community emphases and by familiarity with particular cases. This is not to say that the global canon alone restricts which problems philosophers address and how. We acknowledge that novel insights based, for instance, on attention to new disciplines, can emerge largely because philosophers’ interests shift from old, tried-and-tested cases to new ones. In other words, the tragedy of the canon implies the existence of constraining influences on the development of philosophical work, but these influences do not fully determine which cases or problems receive attention.

The second manifestation of the tragedy of the canon is a practical one. The development of our canon spans a transition across which philosophers of science moved from using examples largely for illustrative purposes to using them evidentially. The sort of examples one chooses for illustrative ends will often not be those that best serve evidential ends, and the standards for engagement with them will be different. We saw this in the case of the Fresnel bright spot. If the question of the value of novel prediction is conceived in a logical-relationship manner, then the potted story does the trick. However, if we want to use that story as evidence, we need to engage with and develop it more carefully. But despite this asymmetry, the particular examples used to argue for a wide range of philosophical positions have stayed relatively stable, even as the use of case studies as evidence has become more common. As a result, examples originally selected for their heuristic utility for illustrating logical relations have been asked to serve an evidential function for which they were not designed.

We saw the stakes of this practical problem in the case of the bright spot. Cases, when studied as evidence, often suggest conclusions different from those they were once used to illustrate heuristically. We might therefore expect an evidentially motivated approach to yield a different corpus of cases. This is a consequence for individual cases as well as for the global canon. When we shift from heuristic to evidential use, the evidential adequacy of a case study will have to be re-established if it was originally selected for its heuristic power. That is, we have to question anew whether it is saliently canonical for the problem. Moreover, we cannot be confident that folk-canonical cases, whatever their antiquity, have undergone the iterative scrutiny needed to render them salient for the issues to which they have been turned.

Returning to the prediction-accommodation debate can illustrate how these effects play out. David Harker (2008) identifies a folk canon for that debate: the phases of Venus falling out of Copernican theory; the discovery of new elements on the basis of Mendeleev’s periodic law; the Fresnel bright spot. These are the examples, Harker notes, that shape the predictivist instinct, and they have, for the large part, been present since the beginning of the debate. But, he objects, “examples motivate discussion, but don’t establish theses. In fact, the value of such historical case studies to the temporal predictivist appears to deteriorate under closer inspection of the pertinent historical details” (432). That is, as Worrall’s and Brush’s work demonstrates, historical scrutiny has tended to undermine, rather than uphold, predictivism. The conditions for the historiographical betrayal of predictivist instincts came about because the heuristic mode in which examples were once employed does not inspire the same iterative historical scrutiny that evidential use does.

Brush makes an observation that is telling of another difficulty with the global canon: even his broad survey of cases from physics, astronomy, chemistry, and biology is too limited to sustain broad generalizations. He notes that most examples people have chosen involve revolutionary change, not work-a-day theorizing, and so only reflect one aspect of scientific development. A more comprehensive picture would require cultivating examples that reflect other sorts of theoretical development.

⁸ For instance, see Laubichler and Maienschein (2007) on developmental biology, Inkpen (2017), O’Malley (2016), and Sagoff (2017) on ecology, and Ratti (2018) and Sarkar (2004) on molecular biology.

Indeed, this is one problematic feature of our informally assembled global canon: it favors flashy cases over a representative selection of cases. Worrall (1989) evocatively suggests that the story of Fresnel's prize could form the basis for a movie script, but most scientific practice exhibits somewhat less drama and intrigue. A representative global canon would include a large proportion of ordinary cases alongside the few flashy ones. Philosophical questions directed toward revolutionary change will need to rely on examples with some radical features; however, the overemphasis on these cases means that well developed examples are not as readily available to form a canon suited to other types of philosophical questions.

Several features of the hitherto informal mode of canon formation are worth emphasizing. The first is the strength of the founder effect. Those cases developed earliest in a debate exert disproportionate influence over its direction. The story of wave optics might very well be a good case, but its persistence owes more to its antiquity than to its fitness for purpose. When John Stuart Mill and William Whewell sparred over the value of novel predictions in the mid-1800s, they commonly invoked wave optics and the question of the existence of the luminiferous ether (Whewell, 1849; Mill, 1855 [1846]). The Fresnel bright spot remained a stock example of the evidential value of novel prediction for decades, despite the fact that better scrutiny of its historical details of the sort eventually supplied by Worrall would have undermined its apparent lesson. Some examples prove durable at least partly because they have become entrenched in the discourse.

Second, and relatedly, familiarity can cement a case study in the canon. Philosophers reach for cases they know, either from the philosophical literature or from ambient scientific discourse. Does the fact that a case has been discussed extensively imply that it is better suited than other cases for assessing a philosophical question? It is tempting to say "no." We should assess the appropriateness of an example on independent grounds, not by its popularity. We should not, however, overlook the advantages of working with examples that are easy to manipulate and likely to be understood. Familiarity can be a virtue in some circumstances, but we need to attend to the processes by which cases become familiar and the factors that shape the collective characteristics of the class of familiar cases.

Third, the distribution of case studies roughly corresponds to the prominence of the scientists featured in them. In some ways, this is to be expected. Philosophers want to use the most successful science when framing their convictions about its proper function; cases that appear as examples and are then probed in enough detail to attain saliently canonical status are thus often famous cases featuring prestigious scientists. But the prestige that a scientific community itself accords its members is a poor proxy for epistemic merit. Extensive literature has shown how the contributions of women and minority scientists have been systematically undervalued by contemporary reward systems (Kohlstedt, 1995; Martin, 2016), how intellectual contributions from routine technical labor and maintenance are elided (Shapin, 1989; Schaffer, 2011; Morus, 2016), and how otherwise similar areas of science enjoy vastly different professional capital and public approbation (Laubichler & Maienschein, 2007; Milam, 2010; Martin, 2017, 2018). A conscientiously constructed canon should be careful not to perpetuate prejudicial judgments of all varieties.

Finally, whimsy features prominently. Individuals choose cases that strike them. For Imre Lakatos, successful novel prediction was a key feature distinguishing progressive research programs from degenerative ones. Some notable examples illustrate this assertion, including Newtonian mechanics permitting Edmund Halley to foretell the return of his eponymous comet and relativity predicting the sun's gravitational effect on light. Lakatos, however, does not give any systematic reasons for selecting these particular cases. Rather, he asserts, the ability to "predict novel facts, facts which had been either undreamt of, or have indeed been contradicted by previous or rival programmes" is a shared feature of "the research programmes I admire" (Lakatos, 1978, p. 5).

The consequence of this state of affairs is that philosophers of science reason from a canon that is subject to idiosyncrasies and lock-in effects.

Given that cases used for one problem are often those repurposed for others, the global canon is consequently unrepresentative in important ways. Like other folk practices, this way of establishing canonicity often finds ways of working, as we saw in the iterative exchange around the case of the Fresnel bright spot. But canon construction can benefit from being more rigorously systematized. We take up this task in the next section.

5. Balancing the aspects of canonical case construction

The challenges described above concern the relationship between the features of case studies themselves and the philosophical purposes to which we put them. With that in mind, what aspects of that relationship should we consider when deploying case studies? How should a canon be assembled so that it reflects the collective consequences of the types of cases that compose it? How, that is, should historians and philosophers of science address the tragedy of the canon? Below, we outline some aspects for historians and philosophers to consider when introducing and developing case studies, which can be useful not only for their own purposes, but also for the HPS community at large. These aspects might at times be virtues we want to maximize, at other times not. They might come into tension depending on the philosophical aims we are addressing. But they all feature in one way or another in the intersection of historical cases and philosophical aims.

The following builds on previous literature that has advanced productive guidance for case selection with an eye to improving the evidential value of case studies, both individually and as they come together to form a canon. Scholl and Ráz (2016), in particular, have developed a taxonomy that distinguishes cases based on different aims they might serve: hard cases for challenging entrenched presumptions; paradigm cases for indicating typicality; big cases for revealing broadly influential shifts; and randomly selected cases as a bulwark against selection bias. We understand our goal to be consonant with such efforts, but it is somewhat different insofar as our primary interest is in the effect of our case selection practices on the canon, rather than on the evidential value of individual cases. We therefore encourage understanding the aspects of canonization discussed below as considerations to be taken alongside aspects of cases that bear on their evidential value.

5.1. The aspects of canonization

We identify six aspects of canonization, that is, of the historical case–philosophical problem relationship: salience, familiarity, typicality, representativeness, scope, and pertinence. Attending to these aspects is a prerequisite for developing more robust canonicity practices, ones that lead us toward a global canon that is more amenable to the collective aims of philosophers. We offer no prescriptive formula to balance these aspects, but we note that they will all be germane in some measure to any canonization process, though some might pertain more directly to cases, some to aims, and some to the relationship between them. Recognizing them, and noticing when they pull in different directions, can have salutary effects on the global canon.

Salience is the most critical aspect of canonization. It refers to sensitivity to the right causal factors, which determines a case's suitability for addressing the philosophical question at hand. This involves thinking carefully about contingency and the causal structure of history (Bolinska & Martin, 2020). Assessing the salience of a case for a philosophical aim requires close connections between historical and philosophical discourse. Philosophical aims will influence the sorts of causal questions historians ask and historians' answers to those questions will in turn help establish the salience of those cases for philosophical aims. Salience, unlike the aspects below, is *necessary* for a case to be saliently canonical for a philosophical aim.⁹ A case that is not sensitive to the right

⁹ But not sufficient: iterative scrutiny of a case by historians and/or philosophers is also necessary.

sorts of causal factors cannot be used to support philosophical claims targeting those factors.

Assessing salience requires careful historical work. As we saw with reference to general relativity and the Eddington expedition, our judgments of salience are delicately sensitive to historical scrutiny. A once straightforwardly salient case was called into question by interrogating Eddington's reasoning process. Its salience was restored, but only on the basis of further historical work leading to a deeper understanding of the case. That is, the salience of a case is not *established* so much as it is *maintained*. Because evolving historical understanding can bear on the philosophical salience of our cases, close connections between history and philosophy are necessary to ensure the understanding of case studies bearing evidential weight remains current.

Salience is a minimal standard for a case to qualify as the kind of thing that can play an evidential role in assessing a philosophical claim. But, as Scholl and R az point out, we can hold cases to different evidential standards. Hard cases, for instance, are those that seem at first glance to pose special challenges for a philosophical proposal. If a philosophical proposal can account for even the most *prima facie* difficult cases, then it is especially worthy of approbation, just as Darwin's theory of evolution by natural selection's ability to account for the evolution of something as complex as the eye is a more powerful testament to that theory than its accounting for the evolution of a giraffe's neck (Scholl & R az, 2016, §4.1).

Second, as we discussed above, **familiarity** often bears on case selection. The virtues of model organisms, in Samuel Schindler and Raphael Scholl's helpful analogy, apply also to case studies. That is, the standardization and reproducibility of, say, *Drosophila melanogaster* or maize, allows biologists to form a stable discourse, establish a shared set of problems, and work productively toward addressing them (Ankeny & Leonelli, 2011). Similarly, in HPS, familiar cases are easier to work with. They allow deeper and richer analysis than novel cases, which require a large investment to develop. They are also more likely to be easily understood, and therefore to gain traction in the community (Schindler & Scholl, 2020, §3.2). Scholl and R az (2016, §4.2) make a similar point about the value of paradigm cases for teaching and research.

Equally importantly, familiar cases are more likely to be understood in enough detail that we can consider them saliently canonical for certain questions. The risk with introducing new cases is that they have not undergone—and might never undergo—the sort of iterative scrutiny that would make them eligible for canonization, and so are subject to much greater uncertainty. As Philip Anderson, a former Bell Laboratories physicist, once remarked: “Any good engineer will tell you that the last thing in the world you want for your large engineering project to depend upon is truly innovative technology. Innovation in one piece of a project is permissible—you may only lose your deadline. But more than one truly innovative piece spells disaster. The delays and mistakes inevitable in new technology are just too common” (Anderson, 1987, p. 67). A philosopher might reasonably express a similar worry about novel case studies increasing the risk of breakage.

Nevertheless, familiarity has its limitations as a guide to case selection. HPS scholarship has emphasized the way in which choice of model organisms constrains research programs (e.g. Ankeny, 2010). Similarly, standard historical cases focus our attention on particular types of questions and often encode presumptions we might wish to revisit. The challenges of *overfamiliarity* are heightened in HPS when dealing with examples that began their lives as illustrative heuristics before evolving into evidential cases. The utility of familiarity therefore must be balanced against the other aspects of canonization.

The third of these is **typicality**. Does a case exhibit features that are common to other salient cases? Attending to typicality can defang objections to the use of historical evidence that rest on concerns about cherry-picking. Such concerns are founded on the worry that, even if one finds a historical case that genuinely supports a philosophical position, that case might be exceptional in some sense—it might be atypical of the part of scientific practice that position is about. But if a case is typical—if

it exhibits features that we would reasonably expect to, or actually do find in other cases, and those shared features are the source of those cases' philosophical salience—then this worry dissolves.

Scholl and R az's (2016, §4.2) discussion of paradigm cases notes that we can discern typicality in cases that are historically well understood. For instance, typicality can be established by demonstrating that a particular episode was influential. Scholl and R az give the example of Robert Boyle's air-pump experiments, which were widely imitated and helped shape early modern experimental practice. Given the influence of these experiments, we might reasonably expect conclusions made on their basis to be typical of experimental practice in this period. Typicality can, of course, be contested, but it can also be effectively and responsibly argued for, and such arguments can do a great deal of evidential heavy lifting that might otherwise require the elaboration of many additional cases.

Typicality pertains to individual cases, but we might ask similar questions about the canon itself. This is why, fourth, we should assess the collective **representativeness** of our cases, especially as they join to form the global canon. Is the global canon of philosophy of science exemplary of the ways in which scientists tend to work?¹⁰ Does a case contribute to or alleviate biases within it? We should also, though, attend to representativeness on a smaller scale, with respect to the set of cases salient for particular aims. A philosopher interested in modeling, for instance, who only has access to cases of ecological modeling will struggle to make claims applicable beyond ecology without cultivating a more representative set of cases. We should not, we stress, seek to maximize representativeness at all costs. Failure to increase the canon's representativeness is not cause to abandon a case study well suited to the purposes at hand. But it is appropriate to consider representativeness when choosing from among otherwise adequate cases, both to alleviate biases in the global canon, and because considering the representativeness of the global canon can be useful for opening new avenues for investigation.

For example, one common cognitive bias in both history and philosophy is that our attention falls more easily on exceptional cases, which we then sometimes treat as though they are exemplary. The overrepresentation of extraordinary, revolutionary, and high-prestige science in our global canon introduces significant distortions. Some worthwhile cases might be exceptional in revealing ways, of course, but collectively, emphasis on the exceptional contributes to philosophy's version of the WEIRD problem discussed in section 4. As Olivier Darrigol (2013) notes, for instance, a lack of suitably developed historical studies of fluid dynamics—the site of some of the stubbornest conceptual problems in physics—has impoverished several areas of philosophy of science. To be blunt, rectifying these specific biases (though we could of course identify others, which would demand other responses) requires more well-developed cases of *boring* science: cases that reflect the quotidian operation of the scientific enterprise, and cases that reflect the scientific subjects and labor that failed to grab contemporary headlines.

The problems of representativeness in the global canon can be addressed if philosophers commit to making arguments for the representativeness of their cases, relative to their aims, and to look for such arguments in the work of others.¹¹ The question of how to make such arguments is multifaceted, and it is beyond our scope to address it fully here. We note, however, that the literature offers apt resources for crafting them. Schindler and Scholl's (2020) model organisms analogy, for instance, suggests that cases fit into historical “phylogenies,” which

¹⁰ This is not to say that science has a unified method. Rather, we suggest that a global canon that fails to account for a large proportion of the ways in which scientists work by focusing disproportionately on others would be unrepresentative.

¹¹ Scholl and R az (2016, p. 12) make a similar point when they urge philosophers to argue for the applicability of philosophical claims made on the basis of “big cases” to scientific practice more generally.

warrants understanding them as representative, to a degree, of other members of that same historical lineage. Jürgen Renn (2020) identifies recurrent patterns of change in knowledge systems, such as the ubiquity of “challenging objects,” which push at a conceptual framework’s limits, and “borderline problems,” which compel different regimes of knowledge to integrate. A case study might be representative insofar as it instantiates a recurring phenomenon of this type. Whatever the mechanism, explicit arguments for the representativeness of case studies would both encourage greater sensitivity to biases in the global canon, and encourage philosophical claims with the appropriate scope.

Scope is, in fact, our fifth aspect of canonization. Whereas representativeness is a feature of the canon, scope pertains most directly to the aim for which a case is being canonized. What range of scientific activity does the philosophical question cover? To what extent can the answer to it be abstracted? The normative power of philosophical insight depends on that insight applying beyond individual cases, within well-defined parameters. As such, scope, like representativeness, is not an aspect we should always seek to maximize. More general insights are pursuit-worthy, if available, but we can expect that, in practice, the utility of a generalization will be constrained more strongly by how well we know the boundaries of its applicability. A broad generalization whose limits are fuzzy offers us much shakier ground than a local, but well constrained generalization.

The scope of a philosophical claim is again established through productive iteration between historians and philosophers of science. A philosophical claim that appears, at first blush, to have a very broad scope might be shown to apply more narrowly when tested against further historical cases. Morgan and Morrison (1999) argue that scientific models are partially independent from both theory and the world, thereby acting as mediators between them. This view has been widely accepted, and illustrated with diverse historical case studies in physics and economics. Nevertheless, Schindler (2008a) identifies a prominent historical case—the discovery of the DNA structure—in which *theory* mediates between model and data. This case can be understood as limiting the scope of the models-as-mediators view, in addition to giving us a more nuanced understanding of the roles that theory, models, and data each played in the discovery of the DNA structure.¹²

Finally, we should inquire after an aim’s **pertinence**. In short, is the philosophical question important? Pertinence is already an element of philosophical discourse, but it takes on additional significance when considered in the context of canonization. Cases do not become saliently canonical by virtue of their salience alone. Achieving that status requires investing considerable communal resources in developing the case to a sufficient degree. Given that the HPS community has limited resources, we would be best served to invest those resources in developing robust canonical cases for the problems that matter most. Of course, this is not to imply that pertinence is easy to establish, nor that there will always be agreement about which problems are pertinent. But to the extent that journal referees are asked about the importance of the topics of papers, job candidates are required to persuade committees of the value of their research, and keynote speakers are invited on the basis of their work’s significance, there is at least a minimal sense in which we can assess aims in terms of their pertinence.

Pertinence is therefore critical when considering which *new* case studies to develop. It is of comparatively lesser advantage to develop a new case study for a narrow and focused philosophical question—especially if existing case studies can be adapted—than it is to develop a new case study for an aim of more widely recognized pertinence. If the global canon becomes more representative, but only via the introduction of a broader range of examples directed toward issues of parochial interest, that does little to rectify the biases in the canon.

¹² See also Scholl and Rätz’s (2016) discussion of their (2013) research limiting the range of applicability of Weisberg’s (2007) notion of abstract direct representation.

Considering these aspects of canonization in the deployment and assessment of cases provides a mechanism for addressing the tragedy of the canon. The offer is, metaphorically, a more civic-minded HPS. When we concern ourselves from the outset with the salience of our cases to our aims, with the familiarity and typicality of our cases, with the scope and pertinence of our aims, and with the representativeness of our canon, we not only craft stronger arguments for ourselves, but we also contribute to the development and maintenance of more useful resources for others as well. We will now illustrate how these aspects of canonization might be considered using a familiar debate, in which prediction also figures prominently.

5.2. Attending to the aspects of canonization

Karl Popper (1963) thought successful novel prediction an impressive feat. He especially admired Einstein’s theory of general relativity, which “stuck its neck out,” predicting a result that would be very surprising from the perspective of classical mechanics: that starlight would bend to a particular degree in the presence of massive bodies, and that this bending would manifest through the star’s appearing in a different position in the night sky than we would expect when its light passed near a massive body. At the same time, Popper was wary of Marxist history, Freudian psychoanalysis, and Adlerian psychology.

Popper was troubled by a particular philosophical question: what distinguishes science from pseudoscience? In order to answer that question, he sought out some clear instances of science and some clear instances of pseudoscience, and he looked for differences between them. And Popper had just some cases in mind: Einstein’s theory of general relativity on the one hand, and Marxism, Freudian psychoanalysis, and Adlerian psychology on the other. With these cases in hand, Popper took the crucial next step: what was it about the pseudoscientific practices that made them importantly different from general relativity? His answer: these practices proposed theories that could—and often did—accommodate *any* evidence, even evidence that appeared to be exactly opposing those theories. In contrast, Einstein’s theory of general relativity predicted a phenomenon that would be very unlikely from the perspective of classical mechanics. Making specific predictions and subjecting them to tests was, for Popper, a mark of scientificity.

We might understand Popper as being a falsificationist about his own philosophical theory. He did not canvass all (or even a sizable minority) of science and every purported case of pseudoscience, gathering many instances from which to work inductively to a conclusion about science. Rather, he focused on one case, which he took to be a clear instance of science—one that nobody would dispute ought to qualify as such. He then compared it to some instances of what he took to be pseudosciences—practices which, although they had a scientific veneer, were actually something much more dubious in disguise. These, too, are cases put forth as being indisputable instances of pseudoscience. We can understand Popper, then, as issuing a challenge to philosophers: here is my theory about what makes science scientific; show me where it goes wrong.

Philosophers have taken up this challenge, and they have (as nearly everyone agrees) shown that Popper was wrong. Newtonian mechanics, a predecessor of general relativity, is also clearly a scientific theory. It, too, made predictions. But rather than being abandoned when those predictions were not borne out, it was instead modified in an *ad hoc* way. For instance, when John Adams and Urbain Le Verrier noticed deviations from the trajectory of the orbit of Uranus predicted by Newtonian mechanics, Newtonian mechanics was not scrapped, as Popper claimed that a genuinely scientific theory ought to be. Instead, it was retained, and the deviations were accommodated by the postulation of the existence of another planet. In hindsight, that kind of accommodation turned out to be appropriate, once the existence of Neptune was verified. Yet when the orbit of Mercury was similarly shown to deviate from the predicted Newtonian orbit, Newtonian mechanics again was not abandoned, even absent an explanation that could account for Mercury’s anomalous orbit.

We now understand these deviations as having been genuine falsifying instances of Newtonian mechanics, on Popper's view, since they became explicable only with the advent of Einstein's general theory of relativity, but Popper gives us no resources for distinguishing the cases where accommodating a falsifying instance is justified from those where it is not.

We can consider Popper's position on what distinguishes science from pseudoscience, and the cases he adduced in its support, with respect to the aspects of canonization described above. Was Popper's central case—the Eddington expedition's measurement of starlight bending to just the degree predicted by Einstein's theory—salient?¹³ To answer this question, we must determine whether the theory is a genuine example of science, and whether the expedition indeed tested it. General relativity is clearly a scientific theory. However, subsequent historical work on the Eddington expedition called into question whether that expedition really did put it to the test. If Eddington were ideologically motivated, as Collins and Pinch (1993) suggest, then the case would not have been sensitive to the right causal factors, those that are necessary for *bona fide* theory testing of the sort Popper had in mind. But if we accept more recent accounts of the expedition (Schindler, 2013; Kennefick, 2019), then the case of the Eddington expedition verifying general relativity can be understood as salient for the question of what distinguishes science from pseudoscience.

The question appears to be pertinent, given the use of “science” as an honorific and the worrying effects of skepticism about science, such as failure to stop anthropogenic climate change or to take measures to stop the spread of COVID-19. This might well stem from our (perhaps problematic) use of the term “science” as an honorific; nevertheless, we might plausibly advocate for maintaining this usage, but developing a better understanding of the difference between science and nonscience. Doing so would be a more straightforward route to our desired outcomes, at least pragmatically, if not in principle. We can argue about whether distinguishing science from pseudoscience is a pertinent question, but it seems to be *prima facie* promising in that respect.

The Eddington case is also familiar, and so it comes with the advantages (and drawbacks) of familiarity outlined above.¹⁴ It has garnered a good deal of historical and philosophical attention, and has been amply developed to serve Popper's philosophical aim—although of course much of this development has come after his time.

However, subsequent historical and philosophical work has shown that this case is not so broadly typical as Popper intended, given his deploying it to answer the question of what distinguishes science from pseudoscience. Popper (unlike many of today's philosophers) was trying to make a claim about *all of science*. And the examples from Newtonian mechanics show that his case is not even representative of our best examples of physics. Sometimes we ought not to discard our theories in the face of anomalous results; at other times, we should; and it is extraordinarily difficult to know when we should and when we should not without the benefit of hindsight. Moreover, Popper's case is an instance of revolutionary science, which itself is hardly typical of the kind of ordinary work that takes place in the lab and field during normal science, a difficulty that, as Brush (2015) emphasizes, persists in our present stock of well-understood cases.

Turning to additional cases, then, shows us that Popper erred in his assessment of what distinguishes science from pseudoscience. Although one example of science appears to fit his framework (other difficulties with that example notwithstanding), still further examples show that science is not only characterized by novel prediction, the failure of which renders the theory falsified. Other factors matter, too; in particular, knowing when a theory needs to be modified—and if so, how—after its predictions fail to be borne out is an important element of scientific

practice. And oftentimes novel prediction is not necessary, particularly if we look to evolutionary biology and historical sciences.

Given the difficulties with falsificationism, it is not clear that we can rescue it by narrowing its scope. Since Popper's aim was to distinguish science from pseudoscience, restricting scope is not a move that is available to him: modifying his view such that it applies only to some instances of science would undermine this motivation entirely. This is unlike the challenge to the models-as-mediators view issued by Schindler (2008a), since that view acknowledges the diversity of models and model-based reasoning, and was never intended as anything like providing a definition of models or demarcating them from theories by way of necessary or sufficient conditions.

This example illustrates two advantages of thinking about case studies with respect to these aspects of canonization. First, they can guide philosophers and historians of science in developing cases that function well as members of a more representative global canon, making them more useful to future work. Second, attending to these aspects facilitates critical reflection on how a case bears on the issue at hand. In particular, by asking whether the case is salient, we pay attention to the causal factors that are relevant to that question. And asking about representativeness enables us to gauge whether that answer is the right one. In determining the scope of the question, we can perhaps add further modifications either to the question or our answer. Thus, attending to the aspects of canonization is not just a tool for ensuring the health of the global canon, it also encodes many of the virtues of good philosophical argumentation and critique.

6. Coda: Affirmative action for case studies

For Garrett Hardin, who popularized the phrase, the tragedy of the commons was tragic insofar as it recalled Alfred North Whitehead's characterization of the theatrical genre: “The essence of dramatic tragedy is not unhappiness. It resides in the solemnity of the remorseless working of things” (Whitehead, 1926, p. 13; quoted in Hardin, 1968, p. 1244). Fittingly, for our purposes, Whitehead wrote this in the course of describing his attendance at the meeting of the Royal Society at which Eddington announced the results of his eclipse expedition. The occasion echoed Greek tragedy, with the inevitability of Newton's downfall pre-ordained by the remorseless workings of nature. Science, Whitehead observed, had internalized that sense of inevitability—fruitfully, but not always wisely—necessitating a need for science to embark upon “a thorough criticism of its own foundations.” The tragedy of the canon can also be understood as tragic in this sense. Our global canon displays shortcomings that emerge, unintended and unforeseen, from the ordinary and proper working of our professional machinery. That machinery might function fruitfully for a time. But if its foundations remain unexamined, we risk “the uncritical assumption of half-truths” (Whitehead, 1926, p. 21).

Hardin himself was not sanguine about the possibility of addressing tragedies of this type through individual action. His preferred response to the tragedy of the commons, to the extent that it threatened shared natural resources, involved coercive restrictions on human reproduction. He displayed the distinctive brand of Cold War reasoning that understood individual actions as driven by rational, if myopic, self-interest, while holding out the possibility of a regulative rational force to counteract individual short-sightedness. And he echoed the strains of authoritarianism, nativism, and racism that have long reverberated in American cultural discourse (Oakes, 2016).

Our view of the tragedy of the canon is both less pessimistic and less authoritarian. Rather than assuming that individuals will never act in a way that accounts for the health of commons, we should examine the conditions under which they have in fact done so. Indeed, while Hardin was espousing his view of cold, self-interested rationality, a contrasting way of thinking about collective responsibility was taking shape in the United States. In the wake of the Civil Rights Act of 1964, President Lyndon Johnson recognized that codifying civil rights in law did not itself

¹³ Of course, we would also have to apply the same procedure to the cases of pseudoscience that Popper relied upon. Because our purposes here are merely illustrative, we omit an analysis of these cases.

¹⁴ Indeed, familiarity guides our choice of Popper to illustrate our own position.

guarantee equality of opportunity, particularly for African Americans, on account of persistent structural inequalities. His policy response, first implemented by executive order and commonly known as “affirmative action,” provided that these structural inequalities form the background against which certain decisions should be made—initially government hiring decisions (Cravens, 2004). Although it began as an executive order targeting government hiring, the principle has been widely adopted as a voluntary measure to address systemic disadvantages, for example in university admissions. And powerful examples of natural resource management have also shown voluntary action to be equal to the task of productively addressing the problems of the commons (Ostrom, 1990).

Our proposal proceeds in a similar spirit. We need not, as some have suggested (e.g., Barseghyan, 2015), undertake dramatic disciplinary restructuring or introduce new layers of professional authority. Rather, we can address the tragedy of the canon with a subtle realignment of the norms and expectations to which practicing historians and philosophers of science respond by changing the background against which they select their evidence and make their arguments.

For historians, this means cultivating greater willingness to engage with generalities. The rise of microhistory began as a corrective but ended in an abdication. This point has been made forcefully of late, notably in *The History Manifesto* (Guldi & Armitage, 2014). Whatever the shortcomings of its positive program,¹⁵ *The History Manifesto* did some excellent diagnostic work. Concerns about overgeneralization have led many historians to disengage from questions beyond the narrow areas of their factual expertise, contributing to the erosion of history’s cultural and policy influence. But if historians are committed to their expertise having use—for philosophy and more—then they have a responsibility to look beyond narrow, local cases, not to construct simplistic grand narratives, but to explore how and to what extent the lessons of those cases might apply outside of them. The tools for constraining those claims appropriately can be found in philosophy (see Arabatzis, 2017).

For philosophers, this means attending to the evidential role case studies play, rather than using them only for illustrative purposes. This requires engaging with the historical details of cases, rather than taking their simplest form at face value. Doing so can both enrich and sharpen the philosophical position in question. Incorporating judgments about the aspects of canonization—salience, representativeness, scope, *et cetera*—into normal modes of critique can both further hone the philosophical position and address the tragedy of the canon, even if only a little bit each time. This does not mean that every new case study we introduce must make the canon more representative. Rather, in analogy to voluntary engagement in affirmative action, it means that the state of the canon should form part of the background against which case studies are selected.

Much of the debate over case studies concerns how we should conceive of HPS in relation to the sciences—whether treating case studies as evidence implies an analogy with scientific *methods*, particularly methods of data handling and theory construction (e.g. McAllister, 2018). We suggest instead that historians and philosophers can learn from the *attitudes* prevalent in the sciences. Historians and philosophers traditionally work alone, and we are less accustomed to thinking how our work fits within a collective project (see Ashrafi, 2007). It is this very feature that introduces so many systematic biases into the global canon. A lesson we can learn from the social and natural sciences is how to square our individual aims with communitarianism attitudes, in which individually fallible work builds toward the good of the community by keeping its interests in mind. Only so, can we put an end to this aspect of the tragedy of the canon.

¹⁵ For a thorough discussion from the perspective of the history of science, see the June 2016 “Viewpoint” section of *Isis* dedicated to *The History Manifesto* (Cohen, 2016).

References

- Allamel-Raffin, Catherine, & Gangloff, Jean-Luc (2015). Some remarks about the definitions of contingentism and inevitabilism. In Léna Soler, Emiliano Trizio & Andrew Pickering (Eds.), *Science as it could have been: Discussing the contingency/inevitability problem* (pp. 99–113). Pittsburgh: University of Pittsburgh Press.
- Anderson, Philip W. (1987). Testimony on the Superconducting Super Collider. In *Superconducting Super Collider: Hearing before the Subcommittee on Energy Research and Development of the Committee on Energy and Natural Resources, 100th Cong., 1st sess.* (pp. 63–70) (April 7).
- Ankeny, Rachel A. (2010). Historiographic reflections on model organisms: Or, how the mureaucracy may be limiting our understanding of contemporary genetics and genomics. *History & Philosophy of the Life Sciences*, 32(1), 91–104.
- Ankeny, Rachel A., & Leonelli, Sabina (2011). What’s so special about model organisms? *Studies in History and Philosophy of Science*, 42, 313–323.
- Arabatzis, Theodore (2017). What’s in it for the historian of science?: Reflections on the value of philosophy of science for history of science. *International Studies in the Philosophy of Science*, 31(1), 69–82.
- Ashrafi, Babak (2007). Big history? In Kostas Gavroglu & Jürgen Renn (Eds.), *Positioning the history of science* (pp. 7–11). Dordrecht: Springer.
- Barnes, Eric Christian (2018). Prediction versus accommodation. In Edward N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2018 Edition). <https://plato.stanford.edu/archives/fall2018/entries/prediction-accommodation/>.
- Barseghyan, Hakob (2015). *The laws of scientific change*. Cham: Springer.
- Biagioli, Mario (1996). From relativism to contingentism. In Peter Galison & David J. Stump (Eds.), *The disunity of science: Boundaries, contexts, and power* (pp. 189–206). Stanford, CA: Stanford University Press.
- Bickle, John (Ed.). (2009). *The Oxford Handbook of Philosophy and Neuroscience*. New York: Oxford University Press.
- Bolinska, Agnes, & Martin, Joseph D. (2020). Negotiating history: Contingency, canonicity, and case studies. *Studies in History and Philosophy of Science Part A*, 80, 37–46.
- Brush, Stephen, with Ariel Segal. (2015). *Making twentieth century science: How theories became knowledge*. Oxford: Oxford University Press.
- Bueno, Otávio, French, Steven, & Ladyman, James (2012). Empirical factors and structure transference: Returning to the London account. *Studies In History and Philosophy of Science Part B: Studies In History and Philosophy of Modern Physics*, 43(2), 95–104.
- Burian, Richard M. (2001). The dilemma of case studies resolved: The virtues of using case studies in the history and philosophy of science. *Perspectives on Science*, 9(4), 383–404.
- Cartwright, Nancy, Shomar, Towfic, & Suárez, Mauricio (1995). The tool box of science: Tools for the building of models with a superconductivity example. In William E. Herfel, Władysław Krajewski, Ilkka Niiniluoto & Wójcicki Ryszard (Eds.), *Theories and Models in Scientific Processes, Poznań studies in the philosophy of the sciences and the humanities* 44 (pp. 137–149). Amsterdam: Rodopi.
- Chakravartty, Anjan (2017). *Scientific ontology*. Oxford: Oxford University Press.
- Chang, Hasok (2003). Preservative realism and its discontents: Revisiting caloric. *Philosophy of Science*, 70(5), 902–912.
- Chang, Hasok (2012). Beyond case studies: History as philosophy. In Seymour Mauskopf & Tad Schmaltz (Eds.), *Integrating history and philosophy of science* (pp. 109–124). Dordrecht: Springer.
- Cohen, H. Floris (Ed.). (2016). Viewpoint: The History Manifesto and the History of Science. *Isis*, 107(2), 309–357.
- Collins, Harry M., & Pinch, Trevor (1993). *The golem: What everyone should know about science*. Cambridge: Cambridge University Press.
- Cravens, Hamilton (2004). American social science and the invention of affirmative action, 1920s–1970s. In Hamilton Cravens (Ed.), *The social sciences go to Washington: The politics of knowledge in the postmodern age* (pp. 9–40). New Brunswick, NJ: Rutgers University Press.
- Cushing, James (1994). *Quantum mechanics: Historical contingency and the Copenhagen hegemony*. Chicago: University of Chicago Press.
- Darrigol, Olivier (2012). *A history of optics: From Greek antiquity to the nineteenth century*. Oxford: Oxford University Press.
- Darrigol, Olivier (2013). For a philosophy of hydrodynamics. In Robert Batterman (Ed.), *The Oxford Handbook of Philosophy of Physics* (pp. 12–42). Oxford: Oxford University Press.
- Douglas, Heather, & Magnus, P. D. (2013). State of the field: Why novel prediction matters. *Studies in History and Philosophy of Science Part A*, 44, 580–589.
- Dresow, Max (2020). History and philosophy of science after the practice-turn: From inherent tension to local integration. *Studies in History and Philosophy of Science Part A*, 82, 57–65.
- Earman, John, & Clark, Glymour (1980). Relativity and eclipses: The British eclipse expeditions of 1919 and their predecessors. *Historical Studies in the Physical Sciences*, 11(1), 49–85.
- Galison, Peter (1987). *How experiments end*. Chicago: University of Chicago Press.
- Gardner, Michael R. (1982). Predicting novel facts. *The British Journal for the Philosophy of Science*, 33(1), 1–15.
- Giere, Ronald N. (1973). History and philosophy of science: Intimate relationship or marriage of convenience? *The British Journal for the Philosophy of Science*, 24(3), 282–297.
- Giere, Ronald N. (1984). Testing theoretical hypotheses. In John Earman (Ed.), *Testing scientific theories*, Minnesota Studies in Philosophy of Science, Vol. 10 (pp. 269–298). Minneapolis: University of Minnesota Press.
- Giere, Ronald N. (1988). *Explaining science: A cognitive approach*. Chicago: University of Chicago Press.
- Godfrey-Smith, Peter (2014). *Philosophy of biology*. Princeton: Princeton University Press.

- Griffiths, Paul (2008). Philosophy of biology. In Edward N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. Spring 2018 Edition. <https://plato.stanford.edu/archives/spr2018/entries/biology-philosophy/>.
- Guldi, Jo, & Armitage, David (2014). *The history manifesto*. Cambridge: Cambridge University Press.
- Hacking, Ian (1999). *The social construction of what?* Cambridge, MA: Harvard University Press.
- Hardin, Garrett (1968). The tragedy of the commons. *Science*, 162(3859), 1243–1248.
- Harker, David (2008). On the predilection for predictions. *The British Journal for the Philosophy of Science*, 59(3), 429–453.
- Henrich, Joseph, Heine, Stephen J., & Norenzayan, Ara (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, 33(2–3), 61–83.
- Hesse, Mary (1966). *Models and analogies in science*. Notre Dame, IN: University of Notre Dame Press.
- Hitchcock, Christopher, & Sober, Elliott (2004). Prediction versus accommodation and the risk of overfitting. *The British Journal for the Philosophy of Science*, 55(1), 1–34.
- Inkpen, S. Andrew (2017). Are humans disturbing conditions in ecology? *Biology and Philosophy*, 32(1), 51–71.
- Jamieson, Annie, & Radick, Gregory (2013). Putting mendel in his place: How curriculum reform in genetics and counterfactual history of science can work together. In Kostas Kampourakis (Ed.), *The Philosophy of Biology* (pp. 577–595). Dordrecht: Springer.
- Kennefick, Daniel (2019). *No shadow of a doubt: The 1919 eclipse that confirmed Einstein's theory of relativity*. Princeton: Princeton University Press.
- Kinzel, Katherina (2015a). Narrative and evidence: How can case studies from the history of science support claims in the philosophy of science? *Studies in History and Philosophy of Science Part A*, 49, 48–57.
- Kinzel, Katherina (2015b). State of the field: Are the results of science contingent or inevitable? *Studies in History and Philosophy of Science Part A*, 52, 55–66.
- Kohlstedt, Sally Gregory (1995). Women in the history of science: An ambiguous place. *Osiris*, 10, 39–58.
- Kuhn, Thomas S. (1962). *The structure of scientific revolutions*. Chicago: University of Chicago Press.
- Kuukkanen, Jouni-Matti (2018). Can history be used to test philosophy? *Journal of the Philosophy of History*, 12(2), 183–190.
- Lakatos, Imre (1978). *The methodology of scientific research programmes*. Cambridge: Cambridge University Press.
- Laubichler, Manfred D., & Maienschein, Jane (2007). *From embryology to evo-devo: A history of developmental evolution*. Cambridge, MA: MIT Press.
- Laudan, Larry (1989). Thoughts on HPS: 20 years later. *Studies in History and Philosophy of Science Part A*, 20(1), 9–13.
- Martin, Allison E. (2016). Outward bound: Women translators and scientific travel writing, 1780–1800. *Annals of Science*, 73(2), 157–169.
- Martin, Joseph D. (2013). Is the contingentist/inevitabilist debate a matter of degrees? *Philosophy of Science*, 80(5), 919–930.
- Martin, Joseph D. (2017). Prestige asymmetry in American physics: Aspirations, applications, and the purloined letter effect. *Science in Context*, 30(4), 475–506.
- Martin, Joseph D. (2018). *Solid state insurrection: How the science of substance made American physics matter*. Pittsburgh: University of Pittsburgh Press.
- Martin, Joseph D., & Janssen, Michel (2015). Beyond the crystal maze: Twentieth-century physics from the vantage point of solid state physics. *Historical Studies in the Natural Sciences*, 45(5), 631–640.
- McAllister, James W. (2018). Using history as evidence in philosophy of science: A methodological critique. *Journal of the Philosophy of History*, 12(2), 239–258.
- Milam, Erika Lorraine (2010). The equally wonderful field: Ernst Mayr and organismic biology. *Historical Studies in the Natural Sciences*, 40(3), 279–317.
- Mill, John Stuart (1855). *System of logic, ratiocinative and inductive; being a connected view of the principles of evidence and the methods of scientific investigation*, 1855 [1846]. New York: Harper & Brothers.
- Mizrahi, Moti (2015). Historical inductions: New cherries, same old cherry-picking. *International Studies in the Philosophy of Science*, 29(2), 129–148.
- Mizrahi, Moti (2020). The case study method in philosophy of science: An empirical study. *Perspectives on Science*, 28(1), 63–88.
- Morgan, Mary S., & Morrison, Margaret (Eds.). (1999). *Models as mediators: Perspectives on natural and social science*. Cambridge: Cambridge University Press.
- Morus, Iwan Rhys (2016). Invisible technicians, instrument makers, and artisans. In Bernard Lightman (Ed.), *A Companion to the history of science* (pp. 97–110). New York: Wiley.
- Oakes, Jason (2016). Garrett Hardin's tragic sense of life. *Endeavour*, 40(4), 238–247.
- O'Malley, Maureen A. (2016). The ecological virus. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 59, 71–79.
- Ostrom, Elinor (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge: Cambridge University Press.
- Pickering, Andrew (1984). *Constructing quarks: A sociological history of particle physics*. Chicago: University of Chicago Press.
- Pitt, Joseph (2001). The dilemma of case studies: Toward a heraclitian philosophy of science. *Perspectives on Science*, 9(4), 373–382.
- Popper, Karl (1963). *Conjectures and refutations: The growth of scientific knowledge*. London: Routledge & Kegan Paul.
- Potters, Jan (2019). Stabilization of phenomenon and meaning: On the London & London episode as a historical case in philosophy of science. *European Journal for Philosophy of Science*, 9(2), 1–30.
- Psillos, Stathis (1994). A philosophical study of the transition from the caloric theory of heat to thermodynamics: Resisting the pessimistic meta-induction. *Studies in History and Philosophy of Science Part A*, 25(2), 159–190.
- Radick, Gregory (2005). Other histories, other biologies. In Anthony O'Hear (Ed.), *Philosophy, biology and life* (pp. 21–47). Cambridge: Cambridge University Press.
- Ratti, Emanuele (2018). 'Models of' and 'models for': On the relation between mechanistic models and experimental strategies in molecular biology. *The British Journal for the Philosophy of Science*, 2, 773–797.
- Renn, Jürgen (2020). *The evolution of knowledge: Rethinking science for the anthropocene*. Princeton, NJ: Princeton University Press.
- Sagoff, Mark (2017). On the definition of ecology. *Biological Theory*, 12(2), 85–98.
- Sarkar, Sahotra (2004). *Molecular models of life: Philosophical papers on molecular biology*. Cambridge, MA: MIT Press.
- The philosophy of historical case studies. In Tilman Sauer & Raphael Scholl (Eds.). *Boston studies in the philosophy and history of science 319*, (2016). Cham: Springer.
- Schaffer, Simon (1996). Contextualizing the canon. In Peter Galison & David J. Stump (Eds.), *The disunity of science: Boundaries, contexts, and power* (pp. 207–230). Stanford, CA: Stanford University Press.
- Schaffer, Simon (2011). Easily cracked: Scientific instruments in states of disrepair. *Isis*, 102(4), 706–717.
- Schickore, Jutta (2011). More thoughts on HPS: Another 20 years later. *Perspectives on Science*, 19(4), 453–481.
- Schickore, Jutta (2018). Explication work for science and philosophy. *Journal of the Philosophy of History*, 12(2), 191–211.
- Schindler, Samuel (2008a). Model, theory, and evidence in the discovery of the DNA structure. *The British Journal for the Philosophy of Science*, 59(4), 619–658.
- Schindler, Samuel (2008b). Use-novel predictions and Mendeleev's periodic table: Response to Scerri and Worrall (2001). *Studies in History and Philosophy of Science Part A*, 39(2), 265–269.
- Schindler, Samuel (2013). Theory-laden experimentation. *Studies in History and Philosophy of Science Part A*, 44(1), 89–101.
- Schindler, Samuel, & Scholl, Raphael (2020). Historical case studies: The 'model organisms' of philosophy of science. *Erkenntnis*. <https://doi.org/10.1007/s10670-020-00224-5>
- Scholl, Raphael (2018). Scenes from a marriage: On the confrontation model of history and philosophy of science. *Journal of the Philosophy of History*, 12(2), 212–238.
- Scholl, Raphael, & Räs, Tim (2013). Modeling causal structures. *European Journal for Philosophy of Science*, 3(1), 115–132.
- Scholl, Raphael, & Räs, Tim (2016). Towards a methodology for integrated history and philosophy of science. In Tilman Sauer & Raphael Scholl (Eds.), *The philosophy of historical case studies* (pp. 69–91). Cham: Springer.
- Shapin, Steven (1989). The invisible technician. *American Scientist*, 77(6), 554–563.
- Showalter, Elaine (1971). Women and the literary curriculum. *College English*, 32(8), 855–882.
- Sober, Elliott (1984). *Philosophy of biology*. Oxford: Oxford University Press.
- Soler, Léna (Ed.). (2008). The contingentism versus inevitabilism issue. Special section, *Studies in History and Philosophy of Science Part A*, 39(2), 221–64.
- Soler, Léna, Trizio, Emiliano, & Pickering, Andrew (Eds.). (2015). *Science as it could have been: Discussing the contingency/inevitability problem*. Pittsburgh: University of Pittsburgh Press.
- Soler, Léna, Zwart, Sjoerd, Lynch, Michael, & Israel-Jost, Vincent (Eds.). (2014). *Science after the practice turn in the philosophy, history, and social studies of science*. New York: Routledge.
- Stanford, P. Kyle (2003). No refuge for realism: Selective confirmation and the history of science. *Philosophy of Science*, 70(5), 913–925.
- Tambolo, Luca (2018). The problem of rule-choice redux. *Journal of the Philosophy of History*, 12(2), 284–302.
- Vickers, Peter (2013). A confrontation of convergent realism. *Philosophy of Science*, 80(2), 189–211.
- Virmajoki, Veli (2018). Could science be interestingly different? *Journal of the Philosophy of History*, 12(2), 303–334.
- Waller, John (2002). *Fabulous science: Fact and fiction in the history of scientific discovery*. Oxford: Oxford University Press.
- Weisberg, Michael (2007). Who is a modeler? *The British Journal for the Philosophy of Science*, 58(2), 207–233.
- Weisel, Gary J. (2017). The plasma archipelago: Plasma physics in the 1960s. *Physics in Perspective*, 19(3), 183–226.
- Whewell, William (1849). Of induction, with especial reference to Mr. In J. Stuart Mill's *System of Logic*. London: John W. Parker.
- Whitehead, Alfred North (1926). *Science and the modern world*. Cambridge: Cambridge University Press.
- Whittaker, Edmund T. (1951). *A history of the theories of aether and electricity*. Edinburgh: Thomas Nelson and Sons.
- Worrall, John (1989). Fresnel, Poisson and the white spot: The role of successful predictions in the acceptance of scientific theories. In David Gooding, Trevor Pinch, & Simon Schaffer (Eds.), *The uses of experiment* (pp. 135–157). Cambridge: Cambridge University Press.
- Zahar, Elie (1973). Why did Einstein's programme supersede Lorentz's? (II). *The British Journal for the Philosophy of Science*, 24(3), 223–262.