

GRAVITY AND GRACE

Gordon Belot

University of Michigan

*This work is licensed under a Creative Commons
Attribution-NonCommercial-NoDerivatives 3.0 License
<www.philosophersimprint.org/022003/>*

BUT WHAT ABOUT such a proposition as 'I know I have a brain'? Can I doubt it? Grounds for doubt are lacking! Everything speaks in its favour, nothing against it. Nevertheless it is imaginable that my skull should turn out empty when it was operated on.

—Wittgenstein, *On Certainty* §4

1. Introduction

i. This is a paper about the bearing of the underdetermination of theory by data on the question of scientific realism. But the approach taken will be oblique. Initially, the focus will be on the thesis that anyone following the methods of science will be led closer and closer (without bound) to the truth about any given question within the purview of those methods, as more relevant data are considered. I will call this the thesis that science is *alethotropically objective* (i.e., objective in the sense that it has a tendency to turn towards the truth).

The centrality of the thesis of the alethotropic objectivity to philosophy of science has fluctuated over the years. But it seems fair to say that its popularity among philosophers is currently at a low point. My first goal below will be to kick it while it is down. I will provide some reason to think that science cannot be alethotropically objective in the actual world and some reason to think that science couldn't be objective in that sense in any world as complex as ours. And I will argue that two popular manoeuvres that promise to save at least the spirit of alethotropic objectivity can do no such thing.

Why bother with all this if, as I say, few philosophers today endorse this sort of strong objectivity thesis? Because many philosophers today espouse scientific realism and dismiss worries about underdetermination. My final task below will be to make a case that it is not so easy to dismiss underdetermination arguments once one has faced up to what scientific realism looks like in the absence of alethotropic objectivity.

ii. Some preliminaries before proceeding. First, note that even if science is alethotropically objective, there may be many questions whose answers remain forever unknown to us. There may be questions outside

the purview of scientific methods. There may be questions within the purview of scientific methods that are not pursued indefinitely. There may be questions within the purview of scientific methods that are pursued indefinitely but for which new evidence eventually ceases to come to light. And there may be questions whose answers we approach asymptotically but never reach.

iii. Next, since alethotropic objectivity is by no means the most common notion of scientific objectivity a few words may be order about its career (and about its credentials to be considered a species of objectivity).¹

According to one important strand of usage, reality is objective to the extent that it is mind-independent and methods of inquiry are objective to the extent that they allow us to discover the nature of this reality.²

When scientists are driven to discuss objectivity, it is often in reaction to some view on which the human or the mental is at least partially constitutive of the apparently non-human and non-mental aspects of the world described by physics. Thus, some early structural realists, writing in opposition to various forms of positivism, argued that scientific progress consisted in: (i) the gradual increase in objectivity via the elimination from the scientific world picture of anthropocentric elements; and (ii) the production of world pictures that more and more accurately represent the objective structural relations that are in principle accessible to all intelligent investigators, human and non-human.³ In more recent times, it has been the spectre of social constructivism, rather than positivism, that has driven Nobel Prize winners to pro-

nounce their faith that extra-terrestrial scientists should discover the same laws that we do.⁴ So it is natural to interpret them as committed at least to the view that the methods of science are alethotropic when it comes to laws (and it is tempting to think that the limitation to laws is an inessential aspect of the view).

Peirce was uninterested in the label ‘objective,’ but he took science to be alethotropic:

all the followers of science are fully persuaded that the processes of investigation, if only pushed far enough, will give one certain solution to every question to which they can be applied.⁵

And on this he based his accounts of the pragmatic content of the notions of truth and reality.

The opinion which is fated to be ultimately agreed by all who investigate, is what we mean by the truth, and the object represented in this opinion is the real. That is the way I would explain reality.⁶

Alethotropic objectivity has also figured in recent philosophy of science (even among authors who don’t share Peirce’s views about truth and reality). Earman (1992, 138) takes it to be part of the “popular image of science” and urges Bayesian philosophers to have it on their wish lists. Lloyd (1995, §2.3) considers it one of the principal senses of the polysemous notion of objectivity. Railton (2000, 186) identifies it as a component or consequence of the logical empiricist account of objectivity.

1. For surveys of notions of scientific objectivity, see Lloyd (1995) and Reiss and Sprenger (2020).

2. See, e.g., Lloyd (1995, §2) and Reiss and Sprenger (2020, §2).

3. See Planck (1909/1970; 1910/1915, Lecture 1) and Poincaré (1905/1907, Introduction & Part III). Planck was explicitly reacting to (his reading of) the phenomenalist positivism of Mach (1886/1897), Poincaré to the conventionalist-pragmatist-Bergsonian positivism of Le Roy (1901). Schlick (1936/1979) defends a view related to those of Planck and Poincaré. On the context of these works, see Daston and Galison (2007, Chapter V), Heilbron (1986, 47–60), and the editorial apparatus of Friedl and Rutte (2013, 436–464).

4. See Glashow (1992) and Weinberg (1996; 1999), who both explicitly link this faith to the objectivity of science. See also Gell-Mann (2007).

5. Peirce (1878, 299). Peirce later revised this passage, replacing “fully persuaded” by “animated by a cheerful hope,” “every” by “each,” and “can be applied” by “apply it.” See Houser and Kloesel (1992, 378 n. 18).

6. Peirce (1878, 299). For a discussion of Peirce’s views about convergence, truth, and reality (and of the evolution of these views over time), see Hookway (2004).

iv. A final preliminary point. For present purposes, I will identify scientific realism with the thesis that: (i) the sentences expressing a scientific theory are true or false; (ii) special cases aside, the relevant truth conditions are mind-independent; and (iii) the empirical success of our theories gives us (defeasible) reason to think them true.⁷ The early years of the twentieth century saw a flourishing debate over scientific realism, in which many anti-realist participants denied clauses (i) or (ii).⁸ But, for better or worse, in recent decades almost all scientific anti-realists have endorsed (i) and (ii)—and the scientific realism debate has largely been concerned with the epistemic virtues (or lack thereof) of our most successful scientific theories. That is the point at which we will be joining the dialectic below.

2. Direct Problems & Inverse Problems

v. It will be helpful for what follows to recall the distinction between direct problems and inverse problems. In fields such as medical imaging, geophysics, and astrophysics, we are often interested in the set of measurements that can be made in the region exterior to an object with a given internal structure.

THE DIRECT PROBLEM: Given the laws and the internal structure X of object S , to determine the set $O = K(X)$ of outcomes of measurements that can be made exterior to S .

THE INVERSE PROBLEM: Given the laws and the outcomes O of the measurements that can be made exterior to a body S , to determine its internal structure $X = K^{-1}(O)$.

So long as the laws of physics take their usual form, the direct problem will be solvable in principle—if one specifies the set of possible internal

structures the object of interest might have, there will be a mathematically well-defined map K that determines for each possible internal structure X , the set $K(X)$ of outcomes of ideal measurements in the region exterior to the body of interest, should it have internal structure X . But the inverse problem need not be solvable in this sense. Even if K exists as a well-defined map, it need not have a well-defined inverse—if multiple possible internal structures of our body give rise to the same set of exterior measurements, then the inverse mapping is undefined. In this case we say that the inverse problem is *underdetermined* or *unsolvable*. Among applied mathematicians, it is generally thought that solvable inverse problems are the exception rather than the rule—and that the more complex and non-linear the direct problem, the more likely the inverse problem is to be unsolvable.⁹

vi. Many of the inverse problems that arise in medical imaging are solvable. Suppose that you want to reconstruct the density of matter within my head, having measured the degree to which x-rays are attenuated as they are shot through it from various directions. This inverse problem, known as the problem of *x-ray tomography*, is solvable: if you know for each of infinitely many directions in space the degree to which x-rays travelling in that direction are attenuated as they travel through my head, then you can uniquely reconstruct the density of matter within it.¹⁰ The situation is quite different if you are restricted to work with a *finite* set of directions: in that case, for any region inside my head, any set of measurements is consistent with that region being void of matter.¹¹

vii. An iconic unsolvable inverse problem is the geophysical problem

7. This characterization is intended to be non-tendentious. It is, for instance, somewhat weaker than the characterization of Psillos (1999, *xix*), the “naïve statement of the position” offered by van Fraassen (1980, 6 f.), and the “general recipe for realism” of Chakravartty (2017, §1.2).

8. See, e.g., the references of fn. 3 above.

9. On this point, see, e.g., Snieder and Trampert (1999, 135). Sometimes, however, non-linearity constitutes a resource rather than an obstruction—see, e.g., Kurylev *et al.* (2018).

10. Theorem 4.1 of Smith *et al.* (1977): “An object is determined by any infinite set of radiographs.”

11. Theorem 4.2 of Smith *et al.* (1977): “A finite set of radiographs tells us nothing at all.” Their commentary on this result: “For some reason this theorem provokes merriment.”

of *gravimetry*: to determine the pattern of mass density within the Earth from knowledge of the gravitational field exterior to the planet. Newton (1729, Book 1, Section XII) already observed that this problem is unsolvable in the special case of a spherically symmetric distribution of matter (a uniformly dense sphere determines the same exterior gravitational field as does an equally massive hollow shell of constant thickness and uniform density). Stokes (1867, 482 f.) observed that this result can be generalized—given any distribution of mass within the Earth and any region in the interior of the Earth, there is a second distribution of mass within the Earth that gives rise to the same exterior field and in which the given interior region is void of matter.¹²

3. Two Types of Underdetermination

viii. We are interested here in the phenomenon of underdetermination of theory by evidence. It is important to distinguish two species of underdetermination.

- i) We say that a problem exhibits *regular-strength underdetermination* when any finite data set is consistent with multiple hypotheses under consideration but there exist methods of inference whose outputs are (probabilistically) guaranteed to converge to the truth in the limit of infinitely large data sets.
- ii) We say that a problem exhibits *extra-strength underdetermination* when any finite data set is consistent with multiple hypotheses under consideration and there exist no methods of inference whose outputs are (probabilistically) guaranteed to converge to the truth in the limit of infinitely large data sets.

Paradigm examples of regular-strength underdetermination: the problem of determining the bias of a coin from knowledge of the outcomes of a sequence of tosses; the problem of determining the identity of a continuous function $f : \mathbb{R} \rightarrow \mathbb{R}$ from knowledge of its values $f(x_k)$ for

a suitably distributed family of real numbers x_1, x_2, \dots . Paradigm cases of extra-strength underdetermination: the problem of determining both the bias and the nationality of a coin from knowledge of the outcomes of a sequence of tosses; the problem of determining the identity of a continuous function $f : \mathbb{R} \rightarrow \mathbb{R}$ from knowledge of its values $f(x_k)$ for a family of positive real numbers x_1, x_2, \dots .

ix. In one sense, underdetermination is endemic in science: whenever the space of hypotheses under consideration is sufficiently rich, we do not expect that collecting a finite amount of data will allow us to conclusively determine the true hypothesis. Our primary interest below will be in the distinctive difficulties associated with extra-strength underdetermination. These go beyond those raised by Hume's problem of induction.

A standard presentation of Hume's problem runs along the following lines: (i) consider the inference from the past track record of bread (nourishing every day!) to the expectation that bread will be nourishing tomorrow; (ii) no matter how large our data set, we may be disappointed tomorrow—the inference is ampliative; and (iii) the obvious ways of trying to justify this inference are unsatisfactory because they turn out to be rule-circular. At the second stage, we appeal to the fact that knowing that a binary sequence begins with a certain number of consecutive ones (encoding that fact that bread has always nourished thus far) does not guarantee that it will not eventually contain a zero. So we have a form of underdetermination (a data set consistent with multiple answers to a question of interest).

But we do not here have the feature characteristic of extra-strength underdetermination, since the strategy of guessing that bread will always nourish unless and until a day comes on which it fails to nourish is a strategy that is guaranteed to output a sequence of conjectures that eventually settle permanently on the truth concerning whether or not bread always nourishes. Extra-strength underdetermination brings with it further difficulties beyond those attendant upon Hume's problem.

12. For a modern treatment of the unsolvability of the problem of gravimetry, see, e.g., Leweke *et al.* (2018).

4. Examples of Extra-Strength Underdetermination

x. If we were to imagine for a second that we had no way of discovering what was inside of human heads other than shooting x-rays through them, then the problem of x-ray tomography would constitute an example of regular-strength underdetermination.¹³ And if we were to imagine for a second that we had no way of discovering what was inside the Earth other than measuring the gravitational field exterior to it, then the problem of gravimetry would constitute an example of extra-strength underdetermination.

xi. Here is a simple recipe for constructing examples of extra-strength underdetermination: find a physical system such that (i) our only way finding out about its internal structure is by making measurements exterior to it and (ii) the relevant inverse problem is unsolvable. This recipe is not altogether easy to follow. It is easy enough to find physical systems with unsolvable inverse problems, harder to convince yourself that the evidence *in principle* available about them is restricted to measurements made exterior to them. There are, after all, ways of finding out what is inside a planet (or a head) other than making measurements exterior to it. We want examples of systems for which it is in principle impossible to directly inspect their interiors.

xii. Stars are natural candidates to be such systems, since it is, presumably, impossible even in principle to examine their internal structure directly—our evidence is of necessity restricted to what we can measure exterior to them. But there is a hitch: not just any star will do. The Sun is hot enough to be opaque to photons—so we cannot investigate its interior using giant versions of our medical imaging machines. But ordinary stars like the Sun are transparent to neutrinos—and so in principle

we could investigate them via neutrino tomography.¹⁴ The relevant inverse problem is unsolvable in general but solvable in some special cases.¹⁵ It is, I believe, an open question whether this map is invertible when restricted to the stellar regime. If it is not, then such stars are good candidates to be obstructions to the alethotropic objectivity of science. But in any case, there exist objects—such as very newly formed neutron stars—so hot that they are opaque even to neutrinos.¹⁶ It is not implausible that the inverse problems relevant to investigating such extreme objects are unsolvable.¹⁷

xiii. Another sort of candidate is the global structure of our spacetime. Our universe (i) started out very hot and dense, (ii) will undergo permanent exponential expansion towards the future, (iii) looks approximately the same in every spatial direction at large scales, and (iv) appears to be very large.

There are many general relativistic worlds with these features. These differ from one another about various important questions, including the topology of space. But in every general relativistic universe satisfying the conditions above, there is an early time t_0 prior to which the universe is so hot as to be opaque to radiation. Let us agree that observers can exist in such universes only to the future of t_0 and that ideal observers in such universes exist eternally to the future of t_0 . Relative to any ideal observer I in such a universe, we can divide the post- t_0 universe into two regions: the region \mathcal{O} observable by I (consisting of events that can send signals to I 's worldline) and the region \mathcal{U} unobservable to I (consisting of events that cannot send signals to I 's worldline).

13. So long as we are willing to impose some weak *a priori* conditions on the relevant matter distribution ρ and to assume (perhaps unrealistically) that the noise in the data becomes arbitrarily small as the number of data points goes to infinity, then there exist algorithms that take as input suitable larger and larger data sets and give as output a sequence of conjectures that converge to ρ —see, e.g., Louis and Natterer (1983, §VI).

14. On neutrino tomography, see, e.g., Winter (2006).

15. Given the nontrivial curvature of spatial geometry within a star, the map whose invertibility is in question is not the ordinary x-ray transform of medical imaging, but the so-called *geodesic x-ray transform*. For a survey of what is known about the invertibility of this map, see Uhlmann and Zhou (2016).

16. See, e.g., Lattimer and Prakash (2004).

17. It is an open question whether the structure of a newly formed neutron star can be reconstructed from knowledge of the gravitational waves that it emits. On this question, see, e.g., Völkel and Kokkotas (2019).

Our observer I can make measurements only within \mathcal{O} —which is to say that I can investigate \mathcal{U} only by making observations in the region external to \mathcal{U} . Under reasonable assumptions we find that so long as I is able to make measurements of only finite precision, knowledge of the physical state within \mathcal{O} fails to determine the topology of the universe (even though in imposing (i)–(iii) above we have imposed very strong global constraints).¹⁸

So the inverse problem of determining cosmic topology (which we can think of as being information about \mathcal{U}) from measurements made within the observable universe (i.e., exterior to \mathcal{U}) is unsolvable. And here it is very natural to maintain that information about the unobservable portion of the universe is in principle inaccessible to direct measurement. So in universes like our own, our inability to determine the topology of space is an obstruction to the alethotropic objectivity of science.¹⁹

xiv. It would appear, then, that for some objects in our world, even the in-principle evidence must come from the region exterior to them—and in some cases the relevant inverse problems are unsolvable. So we have extra-strength underdetermination and hence a limit on the alethotropic objectivity of science: some scientific questions are intractable, and convergence to the truth cannot be guaranteed. The alethotropic objectivity of science fails! But for all that has been said so far, it fails only because it frays a little bit around the edges.

18. See Ringström (2013; 2014).

19. There is a well-established philosophical literature on indistinguishable spacetimes—see Manchak (2009) for the state of the art—that does not rely on assumptions like (i)–(iv) above and which establishes that there is a sense in which cosmic topology is underdetermined by observation at almost any general relativistic world. But this sense is dialectically precarious, as this approach (unlike that discussed in the main text above) requires one to treat as genuine possibilities spacetimes that would be dismissed as skeptical nightmares by most working scientists. For further discussion, see Belot (forthcoming, Chapter V).

5. Amazing GRACE

xv. It is natural to worry that we have gone astray. After all, people do build reliable medical imaging devices and do make reliable geophysical inferences from gravimetric data. So there must be practical ways around the fact that a finite data set hardly ever logically determines a unique model.²⁰

xvi. Already in the eighteenth century, precision pendulum measurements of gravity were used to constrain models of the shape and internal structure of the Earth—and these methods were elaborated greatly in the nineteenth century as theory and instruments improved and the project of making widely distributed precision measurements was taken up by an international collaboration.²¹ In the present century, the pair of satellites that made up the Gravity Recovery and Climate Experiment (GRACE) made it possible, for fifteen years, to make continual measurements of the Earth's gravitational field, accurate enough to allow month-by-month determination of changes in glaciers and much else besides.²²

How is all of this possible? Roughly speaking: each month, data from GRACE gave us, more or less directly, values for the strength of the gravitational field at a finite number of locations 500 km above the surface of the Earth. This information, supplemented by auxiliary hypotheses, allowed the fitting of a smoothed-out model of the Earth's gravitational field for that month.

The default assumption is that month-by-month changes in the

20. For practical purposes, the distinction between regular-strength and extra-strength underdetermination is often of limited interest. Scales and Snieder (2000, 1708): "Although the uniqueness question is a hotly debated issue in the mathematical literature on inverse problems, it is largely irrelevant for practical inverse problems."

21. For a survey of these developments, with special focus on Peirce's contributions, see Lenzen and Multhauf (1966). On the relation between Peirce's scientific and philosophical work, see Hacking (1990, Chapter 23).

22. For an overview of GRACE, its successor mission GRACE-FO, and their results see Tapley *et al.* (2019).

Earth's gravitational field are driven by processes localized near the surface of the Earth. If we idealize such processes as occurring exactly at the Earth's surface, we can translate month-by-month changes in the models of the gravitational field into models of the month-by-month redistribution of mass at the Earth's surface.²³

Given that we have excellent reason to think that large-scale changes of mass distribution at the surface of the Earth are primarily due to changes in the distribution of frozen and liquid water, this procedure allows us to give reasonable estimates of the rates at which glaciers are melting, sea levels are rising, aquifers are being depleted, and water is being lost to drought. In effect, we impose very strong auxiliary constraints in order to select among the vast family of models consistent with our data.

Sometimes this procedure leads to anomalous results. Over the southern portion of Hudson Bay, for instance, only a small fraction of the rate of increase of the gravitational field can be traced to increased water storage. There is, however, strong independent evidence that in this region the Earth's crust is still undergoing rebound from the withdrawal of the ice sheet at the end of the most recent glacial period. There are a number of competing models of this process—differing, e.g., as to the viscosity in various parts of the Earth's mantle. These lead to distinct predictions for the dynamics of the exterior gravitational field, so that GRACE data can be used to put constraints on these models.²⁴

xvii. We see something similar in the case of medical imaging. Notwithstanding the Wittgensteinian result mentioned above (see fn. 11), actual medical imaging machines get by with making only finitely many measurements—and, indeed, reliably identify pneumocephaly. How do they do it? In effect, by imposing assumptions that rule out large variations of mass density on short length-scales within human bodies.²⁵

23. See, e.g., Wahr *et al.* (1998).

24. See, e.g., Paulson *et al.* (2007).

25. See, e.g., Natterer (2001, §III.1).

xviii. In practice, underdetermination of models by data is resolved by appeal to independently well-supported empirical results from other parts of science. In medical imaging, these will include facts about fine structure of the tissues that compose the human body. In geophysics, these will include, e.g., facts about the behaviour of materials and facts concerning the structure and dynamics of the Earth.²⁶

In the best case, these assumptions will rule out all but one of the models consistent with the data. This would appear to leave room for hope that even extra-strength underdetermination can be tamed. Some scientific problems, considered in isolation, may suffer from extra-strength underdetermination. But what matters is how such problems look when we take into account the full range of our empirical knowledge. It seems at least possible that when we do so, all underdetermination is resolved in the infinite-data limit.

Or rather, this seems possible for *some* putative examples of extra-strength underdetermination—perhaps the problem of determining the structure of a proto-neutron star falls in this category. But for a problem like the determination of the topology of space, it is hard to see how appeal to far-flung empirical results could help at all: we are already modelling observers as having access to all evidence available in the part of their past lightcones lying to the future of the time at which the cosmos becomes transparent to radiation.

6. A Global Worry

xix. Now I would like to raise a more global worry.²⁷ What is science except the attempt to determine (aspects of) the unobserved structure \mathcal{X} of a single system \mathcal{S} (the world) from the class \mathcal{O} of all measurements that will ever be made? Presumably there is a map \mathcal{K} that takes as

26. An early survey of evidence that the land around Hudson Bay is rising drew on a wide variety of types of evidence, including naval history, the condition and disposition of driftwood on the raised beaches on the slopes above the bay, the derivations of Cree place-names, and changes in which routes were passable to dog teams in winter. See Bell (1896).

27. For further discussion, see Belot (2015).

input possible global ways \mathcal{X} that the world could be and gives as output the set of observations $\mathcal{O} = \mathcal{K}(\mathcal{X})$ that would be made were the world that way. Science is the inverse problem: from the set \mathcal{O} of all observations that will be made, to determine the relevant unobserved structure $\mathcal{X} = \mathcal{K}^{-1}(\mathcal{O})$ of our world. Many ordinary inverse problems are unsolvable: the salient hidden aspects of the target system cannot be reconstructed from the accessible evidence, even in the infinite-data limit. Surely it is natural, then, to suspect that this mother of all inverse problems is likewise unsolvable—and hence that the threat posed by underdetermination to the alethotropic objectivity of science is a deep and global one. At any rate, the burden of proof would appear to rest with advocates of alethotropy: here, as in ordinary cases, it appears reasonable to presume that the more complex the direct problem, the less plausible it is that the inverse problem should be solvable. And it should be clear that there is no prospect in this global case of underdetermination being resolved via appeal to independently well-supported empirical findings—there *are* no empirical questions beyond the global one we are asked to solve.

7. An Evasive Manoeuvre

xx. Before pressing on to discuss realism, it will be profitable to discuss a couple of possible responses that aim to save the spirit, if not the letter, of the thesis of the alethotropic objectivity of science.

xxi. The first response takes as its point of departure the thesis that we should believe claims produced by methods guaranteed to eventually lead us to the truth and that we should not believe claims generated by methods that lack this guarantee. Without any pretence to complete historical accuracy, I will call this *Reichenbach's thesis*.²⁸ Adopting Reich-

28. Something like it can be found in, e.g., Anderson (2004), Earman (1993), and Reichenbach (1938, Chapter V). In place of belief and lack of belief, one might consider instead other first-rate attitudes towards beliefs (such as being willing to act upon) and second-rate attitudes towards beliefs (such as alienation from).

enbach's thesis would serve to protect the alethotropic objectivity of science from the threat of extra-strength underdetermination—the idea being that, properly understood, science simply does not pronounce on questions involving extra-strength underdetermination.

xxii. Sadly, this response is untenable. The problem is that Reichenbach's thesis implies that scientific warrant fails to be closed under known logical implication in an unacceptably strong way. The problem can be seen as follows.²⁹ Suppose that Nature is revealing an infinite binary sequence, one bit at a time. Suppose also that these bits are generated by tossing a coin with some fixed bias in favour of heads, given by a number between zero and one. Before each new bit is revealed, we are required to conjecture answers to two questions:

Q1. What is the bias r of the coin?

Q2. Is r a rational number or an irrational number?

Let M be a method for answering one of these questions (i.e., a map that takes as input finite data sets and gives conjectured answers as output). For any bias r that the coin might have, we say that M *guarantees* success (failure) if there is probability one that a coin of bias r will generate a sequence that will lead M to output conjectures that (fail to) converge to the truth. Reichenbach's thesis tells us we should believe deliverances of a method if it guarantees success for each possible bias r that the coin might have; otherwise we should not believe its deliverances.

Suppose that I adopt the *straight rule* as my means of generating guesses about the bias of the coin: if the coin comes up heads k times in the first n tosses, I conjecture that the coin has bias k/n . Now, the law of large numbers tells us that the straight rule guarantees success for each possible bias that the coin might have. So Reichenbach's thesis advises me to believe its output. Whatever data I see, I will know that the conjecture I make by following this rule is a rational number. So if scientific warrant is closed under known logical implication, I should

29. For variants on the problem raised here, see Belot (2017, §3 and Appendix).

believe, no matter what data I see, that the bias of the coin is rational. But the policy of guessing the bias of the coin is rational, no matter what data I see, is not a method that is guaranteed to succeed, no matter what the bias of the coin (since it is guaranteed to fail for each irrational bias that the coin might have). So when combined with the straight rule, Reichenbach's thesis gives me advice inconsistent with the closure of scientific warrant under known logical implication: upon seeing k heads in n tosses, believe that the bias of the coin is k/n but do not believe that the bias of the coin is a rational number.

Of course, the straight rule is just one way of handling the first question. But any method M for handling that question induces a method M' for handling the second question: guess that the bias of the coin is rational if your current best estimate of its bias is rational; otherwise guess that it is irrational. Reichenbach's thesis and the closure of scientific warrant under known logical implication jointly imply that if you are committed to a method M for handling the first problem that guarantees success for each possible bias of the coin, then the method M' that M induces for handling the second problem must also guarantee success for each possible bias the coin might have. But no method for handling the second problem has this feature: for every method of handling the second problem, there are infinitely many possible values for r for which that method does not guarantee success; furthermore, if a method guarantees success for each rational number (or even for some dense set of rational numbers), then there is a dense and uncountable (indeed, co-meagre) set of irrational numbers for which it guarantees failure.³⁰ So Reichenbach's thesis is inconsistent with closure in the following strong sense: if you have a method that underwrites belief in your best estimate of the coin's bias, then you

30. See Koplowitz *et al.* (1995). The essential difficulty here is an instance of a well-known phenomenon: there are many natural problems of parametric statistical inference for which there exist statistically consistent estimators of the parameter but for which there exists no statistically consistent test of whether the parameter takes a rational or an irrational value. For history and references, see Le Cam and Yang (2000, §7.7).

must not believe any conjecture you make concerning whether the coin is rational or irrational.

Closure principles are by no means sacrosanct. Consideration of skeptical scenarios have driven some philosophers to deny closure.³¹ And certain closure principles are inconsistent with standards of statistical inference that are common in the sciences.³² Indeed, in the case at hand, it is far from unintuitive that there should be *some* failure of closure: if I am forced to estimate the bias r of the coin and to guess whether that bias is rational or irrational, knowing that there is a sense in which the second task is more difficult than the first, it seems only reasonable that in some situations I will take a more positive attitude towards my estimate \hat{r} of the bias r than I take towards my guess as to whether r is rational or not (even while knowing whether \hat{r} is rational). Reichenbach's thesis demands something much stronger: that no matter what data you see, you must not believe your conjecture as to whether the bias of the coin is rational. This is a very strange restriction. For many people, at least, there are data sets that they could see that would render them practically certain that the coin being tossed is fair.³³ Upon seeing such data, it would be very strange if we were not willing to affirm the weaker proposition that the bias r of the coin was a rational number. It would be analogous to Newton saying: I have discovered that gravity varies inversely as the square of distance—but don't ask me whether it varies inversely as some polynomial or other.

31. See, e.g., Dretske (1970) and Nozick (1981, Chapter 3).

32. See Mayo-Wilson (2018).

33. Suppose, for instance, that you are a Bayesian. You might have a prior (such as the Laplace-Bayes indifference prior) that assigns no weight to any rational number—in which case you will be certain that the bias of the coin is irrational no matter what data you see. But many Bayesians will prefer a prior that puts some weight on at least some rational numbers. If your prior puts any weight at all on the hypothesis that the coin is fair, and you are shown a data stream in which heads and tails have the same limiting relative frequency, then as the size of your data set goes to infinity your credence in the fairness of the coin will approach one.

8. Another

xxiii. The response we have just considered involved restraining the ambitions of science in order to protect the alethotropic objectivity of science from the threat posed by extra-strength underdetermination. A second response proceeds rather by restraining the ambitions of objectivity, maintaining that, properly understood, scientific objectivity requires not convergence to the truth, but mere convergence of opinion of scientists. As Hempel (1983, 75) characterizes a common view, the practice and products of science are objective to the extent that they are “independent of idiosyncratic beliefs and attitudes on the part of the scientific investigators.” Consider, then, the thesis that either there is just one scientific method of inference for handling any given problem, or there are multiple such methods, but any disagreement between them evaporates in the infinite-data limit. I will call this the thesis that science is *symphonotropically objective* (i.e., objective in the sense that it has an inherent tendency to turn towards harmony). Extra-strength underdetermination, as such, presents no threat to the symphonotropic objectivity of science.

xxiv. Sadly, this second response doesn’t work either. It requires that for any question, there be some learning method M with which all scientifically permitted methods agree in the infinite-data limit. Here the stumbling block is that contexts exist in which there is no optimal method of learning.³⁴ Suppose that Nature is revealing an infinite binary sequence to an agent one bit at a time and that immediately before each bit is revealed, the agent is required to guess whether it will be a zero or a one. An *extrapolator* is a map that takes as input a finite binary string (the bits seen so far) and gives as output a single bit (the guess as to the identity of the next bit to be revealed). Let us say that extrapolator M *learns* binary sequence σ in the infinite long run, when shown σ , M correctly predicts the next bit at least two-thirds of the time. This

34. The essential point is due to Putnam (1963). See Belot (2020) for various elaborations.

learning problem is intractable: each extrapolator learns uncountably many binary sequences and fails to learn uncountably many binary sequences—but there is sense in which the sequences learnable by M (forming a meagre subset of the space of binary sequences) are incomparably more rare than those unlearnable by M . Furthermore, there is no best extrapolator: for any extrapolator and any countable set of sequences, there is another extrapolator that learns each of the given sequences while also learning every sequence learned by the first extrapolator. So each extrapolator M is dominated by some extrapolator M^* , which is in turn dominated by some extrapolator M^{**} , and so on without end.³⁵

Let us call two extrapolators *asymptotically equivalent* if for each binary sequence σ , if fed sufficiently long initial segments of σ , the two extrapolators always agree about what they expect the next bit to be. For the problem at hand (learning patterns in infinite binary sequences), symphonotropic objectivity requires that all scientifically permitted extrapolators be asymptotically equivalent. Suppose that M is a scientifically permitted extrapolator. Then we know that there is an extrapolator M^* that learns a strictly larger set of sequences than M does. M and M^* will not be asymptotically equivalent: if σ is a sequence that M^* learns but M does not, then there will be arbitrarily long initial segments of σ that lead M and M^* to make different predictions as to the next bit. So, according to the proposal at hand, M^* cannot be scientifically permitted. At this point, one can only wonder why we should care more about the proffered variety of scientific permissibility than we care about the ability of our methods to arrive at the truth.

35. Interestingly, if we were to move to a more general setting in which we allowed merely finitely additive mixed strategies, then for any such chain of extrapolators, there would be a learning strategy that is at least as good as every extrapolator in the chain—see Schervish *et al.* (2020, Theorem 1). I will say here only that I think that a theory of human rationality can safely neglect merely finitely additive mixed strategies, since there seem to be excellent reason to think that they cannot be physically implemented at worlds like ours—see Easwaran (2014) and Earman (2020).

9. Scientific Realism

xxv. I think, then, that we should accept at face value examples of extra-strength underdetermination as constituting failures of the alethotropic objectivity of science. Let me turn, finally, to the argument from underdetermination against scientific realism.

- a) Underdetermination of theory by evidence is endemic—many (most?) theories have empirically equivalent rivals.
- b) Evidence can give us no reason to believe one rather than another of a pair of empirically equivalent theories.
- c) So we should have little confidence in the truth of even our best scientific theories.

The argument is often rehearsed but seldom endorsed—as a rule, it is wheeled on stage only so that authors can make one or more of the following points in reply.³⁶

- i) Underdetermination is an anti-realist fantasy, not a scientific reality.
- ii) The argument relies on an epistemologically inert distinction between the observable and the unobservable.
- iii) Just because two theories are empirically equivalent doesn't mean that the evidence gives us equal reason to believe them.

I think that the underdetermination argument has more going for it than is generally acknowledged and that these replies are less decisive than is generally thought. As should be clear from the foregoing, I consider (i) to be at best an exaggeration.³⁷ And I think that (ii) is a distraction: anti-realists can afford to work, if need be, with the distinction between the observed and the unobserved rather than the distinction between

36. For exceptions to the rule, see, e.g., Earman (1993) and Kukla (1996). An underdetermination argument is often attributed to van Fraassen (1980), but it is not in fact easy to find one there—on this point, see van Fraassen (2007, 347).

37. On this point see also Belot (2015).

the observable and the unobservable.³⁸ And part of the point of shifting the discussion from realism to alethotropic objectivity above was to bring out the sense in which (iii) is far from satisfactory.

xxvi. How should realists react to instances of extra-strength underdetermination?³⁹ Advocates of (iii) above tell us that we should be untroubled, because these are cases in which evidence gives us more reason to believe one of the hypotheses at hand.

But I think that we should be troubled. Consider, for instance, how things look from a Bayesian perspective. Consider a Bayesian agent with prior probability distribution Pr who faces a case of extra-strength underdetermination. Divide the hypotheses under consideration into classes of empirically equivalent hypotheses E_1, E_2, \dots .⁴⁰ For each k , let A_k be the hypothesis in E_k to which Pr assigns maximum probability, and let B_k be some other hypothesis in E_k to which Pr assigns positive probability. Define A to be the disjunction of the A_k 's and B to be the disjunction of the B_k 's. A and B are contingent propositions, and our agent assigns each of them non-zero prior probability. We have set things up so that $Pr(A) > Pr(B)$ (i.e., our agent assigns greater prior probability to A than to B). But from the definitions of A and B , it follows for any data set D , $Pr(A|D) > Pr(B|D) > 0$: our agent assigns higher *posterior* probability to A than to B , *no matter what evidence turns up*. In other words, our agent, while regarding A and B as contingent propositions (either one of which might be true) also thinks that no evidence could show that B was more plausible than A .

The realists' reply (iii) above to the argument from underdetermination has it that this is a case where evidence D gives us more reason

38. Many scientific realists are liable to complain that under these new terms, the problem of scientific realism threatens to become a mere facet of the problem of induction. I think that this is a mistake—see fn. 43 below.

39. Either of the well-understood but fairly special sort—when we are interested in determining the internal structure of a proto-neutron star or when we ask about the topology of space—or of the possibly endemic but admittedly less concrete science-as-the-mother-of-all-inverse-problems sort.

40. We will assume that these hypotheses are not stochastic in character.

to believe *A* than to believe *B*. But is this really a case in which the evidence gives our agent reason to prefer *A* to *B*? It seems, rather, like a built-in initial bias in favour of *A* over *B* is doing the work (and, of course, the same sort of issue arises on non-Bayesian approaches).

I claim that realists should find something unsettling in cases like this, where an *a priori* bias cannot be washed away by evidence. In the toy case above, agents with this prior begin life with a bias that favours hypothesis *A* over hypothesis *B*. That in itself is not disturbing: it is a truism that successful inductive learning is possible only against a background of biases that favour some hypotheses over others.⁴¹

But in the presence of extra-strength underdetermination, there is no optimal method and learning requires luck: our beliefs may depend on the evidence we see, but not in a way that tracks the truth. Where we end up is determined by our starting point rather than by the world. At this point, some may be tempted to retreat to the position that the success of science is a reason for belief only in those parts of science that will ultimately face some sort of empirical test—but as we have seen above, this commits one to accepting that in science, rational belief fails to be closed under known logical implication in a bizarrely strong way. Others may be tempted to gesture towards an epistemology on which there is a unique rationally permitted initial bias (or, perhaps, a unique family of rationally permitted biases, sharing an asymptotic behaviour). But as we saw above, this too would be costly: to adopt this kind of view requires sometimes counting as irrational methods known to be more reliable than one's own.

xxvii. The foregoing is not an argument against scientific realism. Rather, it is a plea for an end to the complacency with which realists have tended to dismiss underdetermination arguments—a plea for a recognition of the real threat such arguments pose to a comfortable realist view of science. Here is how Peirce (1877, 11) saw our situation:

41. See, e.g., Jeffreys (1933, 524 f.), Kuhn (1962, 3 ff.), Chomsky (1965, §1.8), and Hempel (1966, §2.3).

Now, there are some people, among whom I must suppose that my reader is to be found, who, when they see that any belief of theirs is determined by any circumstance extraneous to the facts, will from that moment not merely admit in words that that belief is doubtful, but will experience a real doubt of it, so that it ceases to be a belief.

To satisfy our doubts, therefore, it is necessary that a method should be found by which our beliefs may be caused by nothing human, but by some external permanency—by something upon which our thinking has no effect.⁴²

I think that most of us will admire the impulse expressed in the first quoted paragraph—even if we would not subscribe to a view quite this stern.⁴³ But those of us who have given up on the full alethotropic objectivity of science (and also on the fantasy of an account of rationality on which there is only one rational response to any given body of evidence) have to think that Peirce is asking for too much in the second paragraph.

What stable resting place is there for those of us who wish that we could stay true to the spirit of Peirce's first paragraph but know that

42. Peirce (1877, 11). Peirce later revised this passage, inserting "in some degree at least" before "ceases" and replacing "caused" by "determined." See Houser and Kloesel (1992, 377 n. 24).

43. As noted above (fn. 38), some realists beat a retreat when the debate over scientific realism threatens to open up into wider worries about the problem of induction. I would expect that they will be tempted to likewise retreat from where the present discussion has led: "The considerations that motivated Peirce here are not inherently scientific, witness the current literature in epistemology on uniqueness and arbitrarily formed beliefs—see, e.g., Schoenfield (2019; forthcoming). But I was promised some philosophy of science—I'm out of here!"

My own reaction is that there is something puzzling about this realist impulse to retreat when the debate over scientific realism comes into contact with more general epistemological debates: realists who follow Planck (1947/1949, §I), Sellars (1956, §51), and Quine (1957, §§1 f.) in taking science to be methodologically continuous with common sense ought to hope and to expect that suitably abstract problems about scientific reasoning will have their correlates in questions concerning reasoning about ordinary matters of fact. Get back in there!

the goal he sets in the second is beyond all hope?⁴⁴

References

- Anderson, E. (2004) "Uses of Value Judgments in Science: A General Argument, with Lessons from a Case Study of Feminist Research on Divorce." *Hypatia* 19 1–24.
- Bell, R. (1896) "Proofs of the Rising of the Land around Hudson Bay." *American Journal of Science* 1 219–228.
- Belot, G. (2015) "Down to Earth Underdetermination." *Philosophy and Phenomenological Research* XCI 456–464.
- Belot, G. (2017) "Objectivity and Bias." *Mind* 126 655–695.
- Belot, G. (2020) "Absolutely No Free Lunches!" *Theoretical Computer Science* 845 159–180.
- Belot, G. (forthcoming) *Lambda*.
- Chakravartty, A. (2017) "Scientific Realism." In E. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/entries/scientific-realism/>
- Chomsky, N. (1965) *Aspects of the Theory of Syntax*. Cambridge MA: MIT Press.
- Daston, L., and P. Galison (2007) *Objectivity*. Brooklyn: Zone Books.
- Dretske, F. (1970) "Epistemic Operators." *The Journal of Philosophy* LXVII 1007–1023.
- Earman, J. (1992) *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory*. Cambridge MA: MIT Press.
- Earman, J. (1993) "Underdetermination, Realism, and Reason." *Midwest Studies in Philosophy* XVIII 19–38.
- Earman, J. (2020) "Quantum Sidelights on the Material Theory of Induction." *Studies in History and Philosophy of Science* 82 9–16.
- Easwaran, K. (2014) "Regularity and Hyperreal Credences." *The Philosophical Review* 123 1–41.
- Friedl, J., and H. Rutte (eds.) (2013) *Moritz Schlick: Erkenntnistheoretische Schriften 1926–1936*. Wien: Springer.
- Gell-Mann, M. (2007) "Truth, Beauty, and ... Physics?" TED Talk. https://www.ted.com/talks/murray_gell_mann_beauty_truth_and_physics
- Glashow, S. (1992) "The Death of Science!?" In R. Elvee (ed.), *Nobel Conference XXV: The End of Science?* Lanham MD: University Press of America 23–32.
- Hacking, I. (1990) *The Taming of Chance*. Cambridge: Cambridge University Press.
- Heilbron, J. L. (1986) *The Dilemmas of an Upright Man: Max Planck as Spokesman for German Science*. Berkeley: University of California Press.
- Hempel, C. G. (1966) *Philosophy of Natural Science*. Englewood Cliffs NJ: Prentice–Hall.
- Hempel, C. G. (1983) "Valuation and Objectivity in Science." In R. S. Cohen and L. Laudan (eds.), *Physics, Philosophy, and Psychoanalysis: Essays in Honor of Adolf Grünbaum*. Dordrecht: Reidel 73–100.
- Hookway, C. (2004) "Truth, Reality, and Convergence." In C. Misak (ed.), *The Cambridge Companion to Peirce*. Cambridge: Cambridge University Press 127–149.
- Houser, N., and C. Kloesel (eds.), *The Essential Peirce: Selected Philosophical Writings*, Volume 1. Bloomington: Indiana University Press.
- Jeffreys, H. (1933) "Probability, Statistics, and the Theory of Errors." *Proceedings of the Royal Society of London. Series A* 140 523–535.
- Koplowitz, J., J. E. Steif, and O. Nerman (1995) "On Cover's Consistent Estimator." *Scandinavian Journal of Statistics* 22 395–397.
- Kuhn, T. S. (1962) *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press.
- Kukla, A. (1996) "Does Every Theory Have Empirically Equivalent Rivals?" *Erkenntnis* 44 137–166.

44. Versions of this paper were presented in Santa Barbara, Bergen, Pittsburgh, Munich, and Toronto. For helpful discussion, thanks to Anonymous, Dave Baker, Jim Brown, Stephen Humphrey, Eleanor Knox, Mohan Matthen, Laura Ruetsche, Teddy Seidenfeld, and Brian Weatherson. Whether they like it or not, John Earman and Clark Glymour are ultimately to blame.

- Kurylev, Y., M. Lassas, and G. Uhlmann (2018) "Inverse Problems for Lorentzian Manifolds and Non-Linear Hyperbolic Equations." *Inventiones Mathematicae* 212 781–857.
- Lattimer, J. M., and M. Prakash (2004) "The Physics of Neutron Stars." *Science* 304 536–542.
- Le Cam, L., and G. Yang (2000) *Asymptotics in Statistics: Some Basic Concepts* (second edition). New York: Springer.
- Lenzen, V. F., and R. P. Multhauf (1966) "Development of Gravity Pendulums in the 19th Century." *United States National Museum Bulletin* 240 301–345.
- Le Roy, É. (1901) "Un positivisme nouveau." *Revue de métaphysique et de morale* 9 138–153.
- Leweke, S., V. Michel, and R. Telschow (2018) "On the Non-Uniqueness of Gravitational and Magnetic Field Data Inversion." In W. Freeden and M. Zuhair (eds.), *Handbook of Mathematical Geodesy: Functional Analytic and Potential Theoretic Methods*. Cham: Birkhäuser 883–919.
- Lloyd, E. A. (1995) "Objectivity and the Double Standard for Feminist Epistemologies." *Synthese* 104 351–381.
- Louis, A. K., and F. Natterer (1983) "Mathematical Problems of Computerized Tomography." *Proceedings of the IEEE* 71 379–389.
- Mach, E. (1897) *The Analysis of Sensations*. Chicago: Open Court. Originally published as *Beiträge zur Analyse der Empfindungen*. Jena: G. Fischer (1886).
- Manchak, J. B. (2009) "Can We Know the Global Structure of Spacetime?" *Studies in History and Philosophy of Modern Physics* 40 53–56.
- Mayo-Wilson, C. (2018) "Epistemic Closure in Science." *Philosophical Review* 127 73–114.
- Natterer, F. (2001) *The Mathematics of Computerized Tomography*. Philadelphia: Society for Industrial and Applied Mathematics.
- Newton, I. (1729) *The Mathematical Principles of Natural Philosophy*. London: Benjamin Motte.
- Nozick, R. (1981) *Philosophical Explanations*. Cambridge MA: Harvard University Press.
- Paulson, A., S. Zhong, and J. Wahr (2007) "Inference of Mantle Viscosity from GRACE and Relative Sea Level Data." *Geophysics Journal International* 171 497–508.
- Peirce, C. S. (1877) "The Fixation of Belief." *Popular Science Monthly* XII 1–15.
- Peirce, C. S. (1878) "How to Make Our Ideas Clear." *Popular Science Monthly* XII 286–302.
- Planck, M. (1915) *Eight Lectures on Theoretical Physics*. New York: Columbia University Press. Originally published as *Acht vorlesungen über theoretische physik, gehalten an der Columbia university in the city of New York im frühjahr 1909*. Leipzig: S. Hirzel (1910).
- Planck, M. (1949) "The Meaning and Limits of Exact Science." In *Scientific Autobiography and Other Papers*. New York: Philosophical Library 80–120. Originally published as *Sinn und Grenzen der exacten Wissenschaft* (second edition). Leipzig: Johann Ambrosius Barth (1947).
- Planck, M. (1970) "The Unity of the Physical World-Picture." In S. Toulmin (ed.), *Physical Reality: Philosophical Essays on Twentieth-Century Physics*. New York: Harper & Row 1–27. Originally published as "Die Einheit des physikalischen Weltbildes." *Phyikalische Zeitschrift* 10 (1909) 62–75.
- Poincaré, H. (1907) *The Value of Science*. New York: Science Press. Originally published as *La valeur de la science*. Paris: E. Flammarion (1905).
- Psillos, S. (1999) *Scientific Realism: How Science Tracks Truth*. London: Routledge.
- Putnam, H. (1963) *Probability and Confirmation*. Washington: U.S. Information Agency.
- Quine, W. V. O. (1957) "The Scope and Language of Science." *The British Journal for the Philosophy of Science* 8 1–17.
- Railton, P. (2000) "Scientific Objectivity and the Aims of Belief." In P. Engel (ed.), *Believing and Accepting*. Dordrecht: Kluwer 179–208. This is a revised version of "Truth, Reason, and the Regulation of Belief." *Philosophical Issues* 5 (1994) 71–93.

- Reichenbach, H. (1938) *Experience and Prediction*. Chicago: University of Chicago Press.
- Reiss, J., and J. Sprenger (2020) "Scientific Objectivity." In E. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/entries/scientific-objectivity/>
- Ringström, H. (2013) *On the Topology and Future Stability of the Universe*. Oxford: Oxford University Press.
- Ringström, H. (2014) "On the Future Stability of Cosmological Solutions to Einstein's Equations with Accelerated Expansion." In S.Y. Jang, Y.R. Kim, D.-W. Lee, and I. Yie (eds.), *Proceedings of the International Congress of Mathematicians: Seoul 2014*, Volume II. Seoul: Kyung Moon SA 983–999.
- Scales, J. A., and R. Snieder (2000) "The Anatomy of Inverse Problems." *Geophysics* 65 1708–1710.
- Schervish, M. J., T. Seidenfeld, R. B. Stern, and J. B. Kadane (2020) "What Finite Additivity Can Add to Decision Theory." *Statistical Methods & Applications* 29 237–263.
- Schlick, M. (1979) "The Universe and the Human Mind." In H. Mulder and B. van de Velde-Schlick (eds.), *Moritz Schlick: Philosophical Papers*, Volume 2. Dordrecht: Reidel 499–513. Translation of "Weltall und Menschengestalt" (a manuscript of 1936).
- Schoenfield, M. (2019) "Permissivism and the Value of Rationality." *Philosophy and Phenomenological Research* XCIX 286–297.
- Schoenfield, M. (forthcoming) "Meditations on Beliefs Formed Arbitrarily." To appear in *Oxford Studies in Epistemology*.
- Sellars, W. (1956) "Empiricism and the Philosophy of Mind." In H. Feigl and M. Scriven (eds.), *The Foundations of Science and the Concepts of Psychology and Psychoanalysis*. Minneapolis: University of Minnesota Press 253–329.
- Smith, K. T., D. C. Solmon, and S. L. Wagner (1977) "Practical and Mathematical Aspects of the Problem of Reconstructing Objects from Radiographs." *Bulletin of the American Mathematical Society* 83 1227–1270.
- Snieder, R., and J. Trampert (1999) "Inverse Problems in Geophysics." In A. Wirgin (ed.), *Wavefield Inversion*. Vienna: Springer-Verlag 119–190.
- Stokes, G. G., (1867) "On the Internal Distribution of Matter which Shall Produce a Given Potential at the Surface of a Gravitating Mass." *Proceedings of the Royal Society of London* 15 482–486.
- Tapley, B. D., M. M. Watkins, F. Flechtner, C. Reigber, S. Bettadpur, M. Rodell, I. Sasgen, J. S. Famiglietti, F. W. Landerer, D. P. Chambers, J. T. Reager, A. S. Gardner, H. Save, E. R. Ivins, S. C. Swenson, C. Boening, C. Dahle, D. N. Wiese, H. Dobslaw, M. E. Tamisiea, and I. Velicogna (2019) "Contributions of GRACE to Understanding Climate Change." *Nature Climate Change* 9 358–369.
- Uhlmann, G., and H. Zhou (2016) "Journey to the Center of the Earth." arXiv:1604.00630
- van Fraassen, B. (1980) *The Scientific Image*. Oxford: Oxford University Press.
- van Fraassen, B. (2007) "From a View of Science to a New Empiricism." In B. Monton (ed.), *Images of Empiricism: Essays on Science and Stances, with a Reply from Bas C. van Fraassen*. Oxford: Oxford University Press 337–384.
- Völkel, S. H., and K. D. Kokkotas (2019) "On the Inverse Spectrum Problem of Neutron Stars." *Classical and Quantum Gravity* 36 115002.
- Wahr, J., M. Molenaar, and F. Bryan (1998) "Time Variability of the Earth's Gravity Field: Hydrological and Oceanic Effects and Their Possible Detection Using GRACE." *Journal of Geophysical Research* 103 30,205–30,229.
- Weinberg, S. (1996) "Sokal's Hoax." *The New York Review of Books* August 8 11–15.
- Weinberg, S. (1999) "Illustrative Prediction." *Science* 285 1013.
- Winter, W. (2006) "Neutrino Tomography—Learning about the Earth's Interior Using the Propagation of Neutrinos." *Earth, Moon, and Planets* 99 285–307.
- Wittgenstein, L. (1969). *On Certainty*. Oxford: Basil Blackwell.