Synthese (2013) 190:1719–1729
DOI 10.1007/s11229-011-9931-7

Success and truth in the realism/anti-realism debate

K. Brad Wray

Received: 24 May 2010 / Accepted: 4 April 2011 / Published online: 11 April 2011

Abstract I aim to clarify the relationship between the success of a theory and the truth of that theory. This has been a central issue in the debates between realists and anti-realists. Realists assume that success is a reliable indicator of truth, but the details about the respects in which success is a reliable indicator or test of truth have been largely left to our intuitions. Lewis (Synthese 129:371–380, 2001) provides a clear proposal of how success and truth might be connected, comparing a test of success of our theories to medical tests with low rates of false positives and false negatives. But, contrary to what Lewis claims, I argue that it is not enough for the realist to undercut the claim that success is not a reliable indicator of truth. Rather, the realist must show that *our current best theories are likely true*. Further, I argue that tests in science are unlike medical tests in a number of important ways.

Keywords Pessimistic induction · Success · Truth · Realism · Anti-realism · Medical tests · Prediction

1 Introduction

Realists and anti-realists disagree about the relationship between the truth of our theories and the success of our theories (see, for example, Kitcher 2001; Magnus 2003; Saatsi 2005). Realists take the success of our current best theories to be evidence that our theories are true. More precisely, realists take the success of our current best theories as evidence that the claims these theories make about unobservable entities and processes are true or approximately true. Indeed, realists often suggest that the

K. B. Wray (✉)
Department of Philosophy, State University of New York, Oswego, 212 Campus Center, Oswego, NY 13126, USA
e-mail: kwray@oswego.edu

truth or approximate truth of our current best theories is the *best* explanation of their success (see, for example, Musgrave 1988, p. 239; Lipton 2004, pp. 202–203; Devitt 2008, p. 227). Anti-realists are skeptical that the successes of our current best theories warrant our drawing such an optimistic conclusion about the truth or approximate truth of those theories (see, for example, Van Fraassen 1980; Wray 2008). In fact, Larry Laudan (1984) argues that given that many past successful theories have turned out to be false, success is not a reliable indicator of truth. Successful but false theories have even enabled scientists to generate vindicated novel predictions (see Carrier 1991).

In this paper, I aim to shed light on the nature of the relationship between the success of our theories and the truth of our theories. This is a crucial issue in the debate between realists and anti-realists. I aim to show that *even if* success is a reliable indicator of truth, it does not follow that most successful theories are true. But, it is this latter claim, that most of our current theories are true, that the realist needs to establish.

In Sect. 2, I present Peter Lewis' argument against the pessimistic induction. Lewis argues that Laudan fails to show that success is not a reliable indicator of truth. It is worth noting that Lewis' aim is not to prove that success is a reliable indicator of truth. He merely aims to show that Laudan is unsuccessful in his attempt to show that success is *not* a reliable indicator of truth. In this respect, Lewis' aim is quite modest. He does not give us reason for thinking that realism is correct. He merely aims to show that a particular apparent threat to realism need not be take seriously.

In Sect. 3, I argue that Lewis fails to show what realists need to show, that is, that most successful theories are true. I will then draw out the implications of this failure for the prospects of a defense of realism, specifically the common realist claim that we have good evidence to believe that our current best theories are true or approximately true. Finally, in Sect. 4 I argue that the notion of success that Lewis appeals to in his argument is not a suitable notion from which to build a defense of realism. It even threatens to undermine the realists' No-Miracles Argument, an argument that takes the success of our best theories as evidence that they are likely true or approximately true.

2 Defending the success/truth connection

The pessimistic induction over the history of science is widely regarded to be one of the most threatening arguments for those aiming to defend scientific realism (see Worrall 1989, p. 99; Kitcher 1993, p. 136; Papineau 1996; Stanford 2006). There is, though, some disagreement about the nature of this powerful anti-realist argument concerning both (i) the structure of the argument, and (ii) what it purports to prove (see Saatsi 2005).

A number of philosophers have reconstructed the pessimistic induction as a *reductio ad absurdum* (see, for example, Psillos 1999, pp. 102–103; Lewis 2001; Saatsi 2005). I think this is a mistake, as the argument purports to be an *inductive* argument.

Some have taken the argument to be concerned with whether or not our theoretical terms genuinely refer (see Putnam 1978; Nola 2008), whereas others take it to be concerned with whether or not our theories are true or approximately true with respect to what they say about unobservable entities and processes. Indeed, in Laudan's

presentation of the argument, he addresses a variety of issues including (i) whether or not we have good reason to believe our theoretical terms "genuinely refer" *and* (ii) whether or not we have good reason to believe our theories are approximately true (see Laudan 1984, p. 107).

Some regard the conclusion of the argument to be that the success of our current best theories is not a reliable indicator of the truth or approximate truth of those theories (see Lewis 2001; Saatsi 2005, p. 1089), whereas others claim that the conclusion of the pessimistic induction is that our current successful theories are probably false (see Saatsi 2005, p. 1092). According to this latter interpretation, the pessimistic induction is an argument to the effect that given that many past successful theories have been discarded as false, and these past successful theories had successes not unlike those enjoyed by our current best theories, it is likely that our current best theories will also be discarded in the future. Hence, we have good reason to believe that our current best theories are false (Laudan 1984, Chap. 5).

Realists seeking to address the threat posed by the pessimistic induction have employed a variety of strategies. Some have taken issue with the various examples on Laudan's list of past successful theories that are now regarded as false, either on the grounds that the theories in question were not really successful, or on the grounds that the theories were not false or at least not completely false (see, for example, Psillos 1999, Chap. 5; Hardin and Rosenberg 1981).[1] Others have claimed that the pessimistic induction is based on a fallacy, though there is some disagreement about what sort of fallacy it commits (see Magnus and Callender 2004; Lange 2002). More recently, some realists have proposed that the history of science actually supports an *optimistic* induction, not, as anti-realists claim, a pessimistic induction (see Worrall 1994, p. 341; Kitcher 2001, pp. 137, 170; Nola 2008, p. 164; Fehrbach, forthcoming).

Lewis (2001) has taken a different approach to dismantling the pessimistic induction. He has targeted Laudan's claim that success is not a reliable test or indicator of truth (Lewis 2001, p. 374).[2] Indeed, Lewis takes this claim of Laudan's to be the conclusion of the pessimistic induction (see Lewis 2001, p. 373).[3] Whatever the proper formulation of the pessimistic induction is, Lewis' argument is interesting and it can help us better understand what is at issue between realists and anti-realists. Specifically, I believe that Lewis' argument can help us understand what realists and

[1] Though a realist, Kitcher (2001) is skeptical about the effectiveness of the strategy of whittling "down [Laudan's] inflated list of allegedly successful-but-false theories" (p. 168). Kitcher notes that "there are prominent examples from the history of science in which views we now take to be false were genuinely successful by anyone's standards," most notably "Fresnel's version of the wave theory of light" (p. 168). Kitcher's response to this situation is to distinguish between the working posits of a successful theory and the idle wheels (see Kitcher 2001, p. 170). He believes that on careful scrutiny we will find that all the successes of successful-but-false theories are attributable to the working posits. Psillos (1999) offers a similar defense of realism, calling it the "divide-and-conquer" strategy. Stanford (2006), though, takes issue with the selective realist strategies of Psillos and Kitcher (see Stanford 2006, Chaps. 6, 7).

[2] Kitcher (2001) also attempts to defend the claim that there is a strong connection between success and truth that warrants our believing that our successful theories are true. Magnus (2003) has criticized Kitcher's "Galilean strategy" as it is applied to our alleged knowledge of unobservables.

[3] Incidentally, Lewis does recognize that there are alternative formulations of the pessimistic induction. According to Lewis' reconstruction of *Putnam's version* of the pessimistic induction, the argument is intended to prove that "most current scientific theories are probably false" (Lewis 2001, p. 372).

anti-realists assume about the connection between success and truth. Consequently, Lewis' argument deserves our critical scrutiny.

In his efforts to show that Laudan is mistaken in claiming that success is not a reliable indicator of truth, Lewis draws on "the standard characterization of the reliability of a test in terms of the rates of false positives and false negatives" (2001, p. 374). This is the notion of reliability that is used in medical testing. According to this conception of reliability, Lewis explains, "a reliable test is one in which the false positive rate and false negative rate are both sufficiently small" (pp. 374–375). Applied to testing in science, false positives are successful but false theories, false negatives are true but unsuccessful theories, and true positives are theories that are both true and successful. Lewis argues that "*if* … the rate of false positives and false negatives are low … then *if* most current scientific theories are successful, it follows *deductively* that most current theories are true" (p. 375; emphasis added).[4] This looks like a promising line for realists to take. Provided they can show the antecedents of these two conditionals are true, then they can *deduce* the truth of most of our current theories! What more could realists hope for?!

Lewis notes that in the debate between realists and anti-realists the success of our theories is taken as a test of their truth or approximate truth. The sorts of successes Lewis has in mind are the explanatory and predictive successes of scientific theories (Lewis 2001, p. 372).[5] He then proceeds to show that success can be a reliable test, that is, reliable in the technical sense in which he uses the term, *even* if most successful theories are false. Lewis' point is to show that even though Laudan identifies a number of successful but false theories, he fails to show that success is not a reliable indicator of truth.

In fact, Lewis notes that if false positives (successful but false theories) outnumber true positives (successful true theories), then we should *expect* that most successful theories are false. To illustrate this, he asks us to "suppose that 1 in 25 theories are true, and that the false positive rate and the false negative rate for success being a reasonable test for truth are both 1 in 5" (p. 376). Given these assumptions, Lewis notes that "the probability that a theory picked at random is true and successful is …4/125" and "the probability that it is false but successful is 24/125" (p. 376).

The choice of these numbers on Lewis' part is not arbitrary. Rather, Lewis is consciously seeking to choose numbers that yield a result that matches Laudan's remark that "for every highly successful theory in the past of science which we now believe to be a genuinely referring theory, one could find half a dozen once successful theories

[4] Saatsi (2005) criticism of Lewis is based on a misreading of this key passage (see Saatsi 2005, p. 1095). Whereas Lewis' passage is clearly a conditional claim, according to Saatsi's reconstruction, Lewis assumes that *most of our current theories are successful*.

[5] Other realists also believe that the relevant notion of success is predictive success (see Worrall 1994, p. 335). In fact, Worrall criticizes Laudan for "working with [a] very loose notion of scientific success" (p. 335). Further, Worrall thinks that some of Laudan's examples of successful but false theories do not deserve to be counted as such.

Kitcher (2001) recognizes that the plausibility of the case for realism depends, to a large extent, on how we understand the notion of success (see pages pp. 166–167). Still, he fails to be very precise about what the relevant notion involves. Ultimately, he asks readers to assume "we have sufficient grasp of the notion of success" (p. 167). In this respect, Lewis (2001) is to be commended for he brings a degree of precision to the discussion that has largely been missing.

that we now regard as substantially nonreferring" (Laudan 1984, p. 123). The four true successful theories stand in a ratio of 1:6 to the 24 false successful theories.

In summary, Lewis' aim is to show that Laudan has failed in his efforts to show that success is not a reliable indicator of truth. Lewis purports to do this by showing that the test of success can in fact be reliable and yet yield the sorts of results that Laudan claims to find in the history of science. What makes success a reliable indicator is that the test of success has low rates of false positives and false negatives.

3 Assessing the damage

In this section, I aim to show that the realist cannot claim victory yet. The realist needs to do more than show that Laudan is mistaken in claiming that *success is not a reliable indicator of truth*. The realist must show that most successful theories are true. Only then does she have grounds for believing that most of our contemporary theories are true.

The demands on the realist are far greater than Lewis recognizes. It seems that in order to have good reason to believe that most of our contemporary theories are true, the claim that many realists ultimately want to establish, one must provide compelling reasons to believe the following claims: (i) that success is a reliable indicator of the truth, (ii) that most successful theories are true, and (iii) that most of our contemporary theories are successful.[6]

As mentioned above, Lewis (2001) takes the conclusion of the pessimistic induction to be the claim that *success is not a reliable indicator of truth* (see Lewis 2002, p. 373). Here, I think he is mistaken. I believe that from a survey of the history of science, Laudan infers that most past successful theories have turned out to be false. This claim then figures as a premise in an argument whose conclusion is that *our current best theories are probably false*.[7] But, we need not dwell on this point. It is more important to pursue the insight Lewis offers us into the connection between success and truth.

Lewis concedes that "the fact that a given theory is successful need not lead one to have much confidence that it is true" (Lewis 2001, p. 377). That is, he recognizes that, given his conception of success, most successful theories might in fact be false. But, it seems that Lewis fails to see what his admission amounts to. Indeed, I aim to show that

[6] I do recognize that there are a variety of realisms, including structural realism (Worrall 1994), entity realism (Hacking 1983), and convergent realism, and that not all these forms of realism would need to prove so much.

[7] One thing that is clear is that there is great confusion, and little agreement, about what the aims of the pessimistic induction are. Unlike Lewis, Worrall (1994) claims that Laudan's intention is to show that "the empirical success of *presently accepted theories* can … hardly be used as an argument for their approximate truth" (p. 334, emphasis added). Lange (2002) thinks that the anti-realists who appeal to the pessimistic induction for support need to claim "that at most past moments, most of the theories then accepted were false" (p. 285). I think Lange is mistaken here, as does Saatsi (2005). Devitt (2008) insists that the anti-realists' "appeal to historical details has to show not only that we were nearly always wrong in our unobservable posits but that, despite methodological improvements, we have not been getting increasingly right" (p. 233). He doubts this is the case. My chief aim is to develop a better understanding of how success and truth might be connected. So we can set aside for now the issue of determining what the proper way is to reconstruct the pessimistic induction.

this admission on Lewis' part poses a serious threat to realism. Specifically, the type of success he has defended, cannot serve the purposes of the realists who ultimately aim to show that most of our contemporary theories are true or approximately true.

Let me explain my concern in terms of the numbers that Lewis presents in his paper. We are to imagine a world in which there are 125 theories. Only five of them are true, and the remaining 120 are false. Scientists cannot directly ascertain which theories are true and which are false. Rather, they must rely on a test of their success in order to determine to the best of their abilities which are likely true and which are likely false. The test of success has a false positive rate of 1/5 and a false negative rate of 1/5. Consequently, relying on the test of success, scientists will judge 24 of the 120 false theories to be true (1/5 false positives), and four of the five true theories to be true (1/5 false negatives). As far as Lewis is concerned, because the test of success has a low false positive rate and a low false negative rate, Laudan is mistaken to claim that success is not a reliable indicator of truth.

The complaint that I am raising is that too many of the theories judged to be true by the test of success will in fact be false. In fact, only four of the 28 theories judged to be true will in fact be true. That is, only one out of seven theories that passes the test of success is true. Thus, although success *may* be a reliable test in the technical sense in which Lewis uses the term, it hardly provides one with much warrant that a particular successful theory is true.

Importantly, very little hangs on these specific numbers. A similar concern to the one I raise would arise for many different values of the relevant variables. Provided false theories significantly outnumber true theories, and there are some false positives, then most successful theories will not be true. Hence, the test of success will not ensure that most successful theories are true. Indeed, at the end of this section I will demonstrate just how persistent the problem I raise is. Even if true theories are far more common than Lewis assumed, and the test of success was significantly more reliable than Lewis assumed, the test of success is far more problematic than Lewis has led us to believe.

Lewis seems content to celebrate the fact that successful theories are more likely true than unsuccessful theories. Given his numbers, successful theories are true 1/7 of the time whereas unsuccessful theories are true only 1/97 of the time. In this respect, Lewis is correct to claim that if a theory is successful then there is a good chance that it is true, given that successful theories are more likely true than unsuccessful theories. I, on the other hand, insist that even if a theory is successful, there is a good chance that it is false, given that most successful theories are in fact false. This makes me and other anti-realists skeptical about the truth of our current theories, even though we grant they are successful.

It seems that the test of success, insofar as it is *effective*, merely helps scientists identify which theories they should *discard*. After all, for every 97 theories that fail the test of success, scientists will only be led to discard one true theory. But, the test of success is not very effective at helping scientists identify which theories they should *accept*, for six out of seven of the theories that pass the test are false. The test of success is thus like natural selection. It is effective at eliminating weak variations, but not especially effective at ensuring that only *the optimal* variation persists.

Indeed, some anti-realists have suggested that the most promising explanation for the success of science is a selectionist explanation, an explanation modeled on biological explanations that appeal to the mechanism of natural selection. The reason why many of our current theories are successful is because unsuccessful theories have been discarded. There is thus no need to invoke the truth to explain the success of our current theories (see Wray 2007, 2010; Van Fraassen 1980).

Before moving on, it is worth stressing how robust the results of my argument are, and thus how vulnerable Lewis' line of argument is to this criticism I raise. Let us assume far more favorable conditions for the realist than Lewis assumed. Let us assume that 1/10 theories is true, rather than 1/25. And, let us assume that the tests are quite a bit more reliable than Lewis assumed. For example, let us assume the rate of false positives is 1/10 and the rate of false negatives is 1/10. Assuming a world of 100 theories, then 9 theories that test successful will turn out to be false ($90 \times 1/10$), and 9 theories that test successful will turn out to be true ($10 \times 9/10$). Even in this possible world, only half of the theories that test successful will turn out to be true. This hardly gives us strong warrant to believe that a successful theory is true. Clearly, there is apt to be some debate and disagreement about how many false successful theories we are willing to accept and still regard success as a reliable indicator of truth. But, it seems beyond debate that in order for us to count success as a reliable indicator of truth, conditions are going to have to be far better than Lewis assumed, and even better than I have assumed here.

4 Scientific success: a moving target

There are additional problems with the way Lewis attempts to undermine Laudan's claim that success is not a reliable indicator of truth. Lewis suggests that the appropriate notion of reliability is the standard way the term is used in medical testing. In medical testing, a particular test is a reliable indicator of a specific disease provided there are low rates of false positives and false negatives.

A specific test of the sort Lewis has in mind will have fixed rates of false positives and false negatives. For example, a pregnancy test might have a false positive rate of 1/100, and a false negative rate of 1/20. The person testing positive is thus 99 % sure she is pregnant. In *science*, though, the relevant tests change in ways they do not in medical testing. To express the point in another way, medical tests presuppose a closed universe of possibilities, whereas with tests in science we cannot make such an assumption. The notion of success operative in science lacks diachronic stability.[8]

This fact has interesting implications for the role that success plays in the realism/anti-realism debate. Consider predictive accuracy as a measure of success. During the Renaissance, European astronomers were often satisfied when their predictions were accurate within 15°. That is, they sought to locate a planet within the proper constellation (see Dear 2001, p. 171, note 2). By Copernicus' time, astronomers expected a greater degree of accuracy. Typically, the competing theories were accurate to within

[8] I thank Reviewer #2 for these alternative ways of expressing my point.

5° (Gingerich 1971, 1975; Thoren 1967). And, by the end of Kepler's life, in the 1630s, theories were expected to yield predictions with even greater accuracy.[9]

We can see from this example that in science the measures of success change. Consequently, even if we *grant* that by the 1630s predictive accuracy is a *reliable* test of the truth of planetary theories, it is far from obvious that predictive accuracy was a reliable test in Copernicus' day, or in the Renaissance. Indeed, if the rate of false positives is 1 in 5 in the 1630s, then the rate of false positives earlier, when astronomers were satisfied with accuracy within 5° or 15°, must have been much higher. Obviously, more theories could yield predictions that are *that* accurate, and those theories would be false. Given the lower expectations of accuracy of these earlier times, the test of predictive accuracy employed in astronomy *then* was less reliable, and probably even *unreliable* in the technical sense that Lewis uses the term.

This example, which is probably quite typical for the sciences, suggests that Lewis may have been mistaken in importing the notion of a reliable test from the medical context to the debate between realists and anti-realists. It seems unlikely that we can *determine* the false positive and false negative rates of our tests of predictive success of scientific theories in the way we can determine the rates for diagnostic medical tests.[10]

I recognize that realists are apt to take the ever-increasing accuracy of scientific theories that I have appealed to here as evidence that we are converging on the truth, and thus as support for realism. But, caution is in order here. We must recognize that there are many cases where scientists have refined a theory in ways that yielded greater accuracy, even though the theory was subsequently replaced by an incompatible theory sometime in the future. Hence, even false theories can be and have been refined in ways that enable scientists to make more accurate predictions. Ptolemy's theory of the planets, after all, was the culmination of a tradition of geocentric theories, and it became the dominant theory because it could generate predictions that were more accurate than the competing theories could (see Kuhn 1957, pp. 72–72; Dreyer 1953, Chap. 9). Newton's physical theory was also refined for two centuries, generating predictions of ever-increasing accuracy, and yet it too is now regarded as false. Thus, increasing accuracy by itself does not warrant rejecting the skeptical attitude about theoretical knowledge implied by the pessimistic induction.

Indeed, the increasing accuracy of our theories is just the sort of thing that some anti-realists have suggested can be explained by means of the operation of a selection mechanism. Scientists discard theories or variations of a theory that are less accurate than alternative variations (see Van Fraassen 1980; Wray 2007, 2010). Indeed, their careers are apt to be adversely affected if they fail to do so. Consequently, we should not be surprised to see that, over the course of time, the theory accepted in a field is predictively more accurate than the theory or version of the theory accepted earlier.

[9] I have discussed the relativity of predictive success elsewhere (see Wray 2010).

[10] Saatsi (2005) raises a related concern. He thinks that it is unlikely that we can delineate "a non-arbitrary, well-defined collection of *both false and unsuccessful* theories" (p. 1096). Consequently, we cannot really determine the rates of false positives, false negatives, true positives and true negatives. It is considerations such as this that have led some to think that the pessimistic induction is prone to the base rate fallacy (see Magnus and Callender 2004).

All the work of scientists cannot be for nothing![11] And, there is no need to invoke the truth to explain the increasing accuracy of a theory.

There is an additional concern with Lewis' proposed conception of scientific success. It is unclear how we are to delineate or individuate tests in science. In our example of the pregnancy test, we administer the test, and get a clear, though possibly mistaken, result. In science, though, it is not clear if one vindicated prediction constitutes *a test*, or whether *a series* of such predictions constitute *a single test*. Clearly, our answer to this question would profoundly affect the rates of false positives and false negatives we assign to a test. If a test consists of a single vindicated prediction, the rate of false positives will be much higher than if a single test consists of a series of, say ten, vindicated predictions.

There is yet another concern I want to raise against Lewis' conception of success. There is reason to believe that many *realists* would object to Lewis' conception. The realists' Ultimate Argument or No-Miracles Argument takes the success of our theories as evidence that they are likely true or approximately true. Some realists who appeal to this argument claim that the truth of our theories is the *best* explanation for their success (see Musgrave 1988, p. 239; Lipton 2004, pp. 202–203; Devitt 2008, p. 227). These realists believe that success and truth must be intimately related. If they are not then this particular strategy of defending realism must be abandoned. Given Lewis' conception of success, specifically, the fact that it is consistent with only one in seven successful theories being true, one would not be warranted in claiming that the truth of our theories is the best explanation of their success. Lewis' conception thus undermines one of the most popular realist arguments. Hence, it is likely that many realists would want to resist Lewis' application of this particular conception of success from medical testing to the realism/anti-realism debate.[12]

Recall that Lewis claims we can only reasonably infer that our best theories are likely true on the basis of their passing the test of success *if* the false positive and false negative rates for our test are low, and *if* most of our current theories are successful. It seems clear now that even more is required. We also need to know that most successful theories are true. Only then can we make the inference from the success of our theories to their likely truth.

5 Concluding remarks

My aim in this paper has been to clarify the relationship between the success of a theory and the truth of that theory. This issue has been a pivotal one in the debates between realists and anti-realists. Realists, it seems, have often assumed that success is a reliable indicator of truth, but the details about the respects in which success is a reliable indicator or test of truth have been left largely to our intuitions.

[11] The Reviewers' comments helped me see the necessity of addressing the realist response to this fact, that theories are getting increasingly more accurate.

[12] Ghins (2001) claims that "scientific realism is indefensible when it is conceived to be a scientific explanation of the success of science" (p. 121; see also Frost-Arnold 2010). According to Ghins, scientific realism is, at best, a *philosophical* position.

Lewis provides a clear proposal of how success and truth might be connected. But, Lewis' attempt to undermine the pessimistic induction fails to support the realists' case. His conception of success, imported from medical testing, is inappropriate. And, the realist needs to do more than show that success is a reliable indicator of truth. The realist needs to show that *our current best theories are likely true*. Lewis' argument, though, gives us no reason to believe this.

Acknowledgments I thank Lori Nash, Michel Ghins, Greg Frost-Arnold, and Peter Lewis for constructive feedback on earlier drafts. I also thank my audience at the annual meeting of the Creighton Club, to whom I presented an earlier version of the paper, in Geneva, New York. Finally, I thank the referees for Synthese for their thoughtful comments.

References

Carrier, M. (1991). What is wrong with the miracle argument? *Studies in History and Philosophy of Science, 22*(1), 23–36.

Dear, P. (2001). *Revolutionizing the sciences: European knowledge and its ambitions, 1500–1700*. Princeton: Princeton University Press.

Devitt, M. (2008). Realism/anti-realism. In S. Psillos & M. Curd (Eds.), *The Routledge companion to philosophy of science* (pp. 224–235). London: Routledge.

Dreyer, J. L. E. (1953). *A history of astronomy from Thales to Kepler* (2nd ed.). New York: Dover Publications, Inc.

Fehrbach, L. (Forthcoming). How the growth of science ends theory change. *Synthese.*

Frost-Arnold, G. (2010). The no-miracles argument for realism: Inference to an unacceptable explanation. *Philosophy of Science, 77*(1), 35–58.

Ghins, M. (2001). Putnam's no-miracles argument: A critique. In S. P. Clarke & T. D. Lyons (Eds.), *Recent themes in the philosophy of science* (pp. 121–138). Dordrecht: Kluwer Academic Publishers.

Gingerich, O. (1971). The mercury theory from antiquity to Kepler. *Actes*, 57–64.

Gingerich, O. (1975). Crisis versus aesthetic in the Copernican revolution. *Vistas in Astronomy, 17*, 85–95.

Hacking, I. (1983). *Representing and intervening: Introductory topics in the philosophy of natural science*. Cambridge: Cambridge University Press.

Hardin, C. L., & Rosenberg, A. (1981). In defense of convergent realism. *Philosophy of Science, 49*, 604–615.

Kitcher, P. (1993). *Advancement of science: Science without legend, objectivity without illusions*. Oxford: Oxford University Press.

Kitcher, P. (2001). Real realism: The Galilean strategy. *Philosophical Review, 110*(2), 151–197.

Kuhn, T. S. (1957). *The Copernican revolution: Planetary astronomy in the development of western thought*. Cambridge: Harvard University Press.

Lange, M. (2002). Baseball, pessimistic inductions and the turnover fallacy. *Analysis, 62*(4), 281–285.

Laudan, L. (1984). *Science and values: The aims of science and their role in scientific debate*. Los Angeles and Berkeley: University of California Press.

Lewis, P. J. (2001). Why the pessimistic induction is a fallacy. *Synthese, 129*, 371–380.

Lipton, P. (2004). *Inference to the best explanation* (2nd ed.). London, New York: Routledge.

Magnus, P. D., & Callender, C. (2004). Realist ennui and the base rate fallacy. *Philosophy of Science, 71*, 320–338.

Magnus, P. D. (2003). Success, truth and the Galilean strategy. *British Journal for the Philosophy of Science, 54*, 465–474.

Musgrave, A. (1988). The ultimate argument for scientific realism. In R. Nola (Ed.), *Relativism and realism in science* (pp. 252–2289). Dordrecht: Kluwer Academic Publishers.

Nola, R. (2008). The optimistic meta-induction and ontological continuity: The case of the electron. In L. Soler, H. Sankey, & P. Hoyningen-Huene (Eds.), *Rethinking scientific change and theory comparison: Stabilities, ruptures, incommensurabilities* (pp. 159–202). Dordrecht: Springer.

Papineau, D. (1996). Introduction. In D. Papineau (Ed.), *The philosophy of science* (pp. 1–20). Oxford: Oxford University Press.

Psillos, S. (1999). *Scientific realism: How science tracks truth*. London and New York: Routledge.
Putnam, H. (1978). *Meaning and the moral sciences*. London: Routledge.
Saatsi, J. T. (2005). On the pessimistic induction and two fallacies. *Philosophy of Science, 72*, 1088–1098.
Stanford, P. K. (2006). *Exceeding our grasp: Science, history, and the problem of unconceived alternatives*. Oxford: Oxford University Press.
Thoren, V. E. (1967). An early instance of deductive discovery: Tycho Brahe's lunar theory. *Isis, 58*(1), 19–36.
Van Fraassen, B. C. (1980). *The scientific image*. Oxford: Clarendon Press.
Worrall, J. (1989). Structural realism: The best of both worlds. *Dialectica, 43*(1-2), 99–124.
Worrall, J. (1994). How to remain (reasonably) optimistic: Scientific realism and the luminiferous ether. *PSA, 1994*(1), 334–342.
Wray, K. B. (2007). A selectionist explanation for the success and failures of science. *Erkenntnis, 67*, 81–89.
Wray, K. B. (2008). The argument from underconsideration as grounds for anti-realism: A defence. *International Studies in the Philosophy of Science, 22*(3), 317–326.
Wray, K. B. (2010). Selection and predictive success. *Erkenntnis, 72*(3), 365–377.