The Uniformity Principle vs. the Disuniformity Principle

Abstract
The pessimistic induction is built upon the uniformity principle that the future resembles the past. In daily scientific activities, however, scientists sometimes rely on what I call the disuniformity principle that the future differs from the past. They do not give up their research projects despite the repeated failures. They believe that they will succeed although they failed repeatedly, and as a result they achieve what they intended to achieve. Given that the disuniformity principle is useful in certain cases in science, we might reasonably use it to infer that present theories are true unlike past theories. Hence, pessimists have the burden to show that our prediction about the fate of present theories is more likely to be true if we use the uniformity principle than if we use the disuniformity principle.

Seungbae Park
Ulsan National Institute of Science and Technology
The Republic of Korea
nature@unist.ac.kr

1. Introduction

David Hume's (1978: 89) uniformity principle holds that unobserved instances are similar to observed instances and that the future resembles the past. In contrast, what I call the disuniformity principle holds that unobserved instances are dissimilar to observed instances and that the future differs from the past. Importantly, we use the uniformity principle in some contexts but the disuniformity principle in other ones. An interesting question, then, is whether we should use the uniformity principle or the disuniformity principle to assess the fate of present scientific theories. If we use the uniformity principle, we would infer that present theories, like past theories, will eventually be discarded. If we use the disuniformity principle, however, we would infer that, unlike past theories, present theories will not be discarded.

This paper disputes pessimists' use of the uniformity principle in assessing the status of present theories. In Section 2, I clarify how the uniformity principle is used in the pessimistic induction. In Section 3, I present two examples from science to show that the disuniformity principle is useful certain cases in science. I then argue that the burden is on pessimists to justify the choice of the uniformity principle over the disuniformity principle. In Section 4, I explicate how my criticism against the pessimistic induction differs from the criticism (Fahrbach 2011; Park, 2011a; Mizrahi, 2013a) that the pessimistic induction commits the fallacy of biased sample. In Section 5, I argue that my criticism against the pessimistic

induction is on the same boat as the argument from underconsideration (van Fraassen, 1989; Ladyman et al., 1997; Khalifa, 2010; Wray, 2008; Wray, 2012).

2. The Uniformity Principle

In this section, I show that the uniformity principle is indispensable to the pessimistic induction. But let me first unpack the pessimistic induction. Different writers formulate the pessimistic induction differently. For example, Larry Laudan (1977) and Hilary Putnam (1978) formulate it as follows:

> Most of the past theories of science are already suspected of being false; there is presumably every reason to anticipate that current theories of science will suffer a similar fate. (Laudan 1977: 126)

> ..the following meta-induction becomes overwhelmingly compelling: Just as no term used in the science of more than fifty (or whatever) years ago referred, so it will turn out that no term used now (except maybe observational terms, if there are such) refers. (Putnam, 1978: 25)

Laudan and Putnam would more or less agree with my formulation that just as successful past theories, such as the Ptolemaic theory, the phlogiston theory, the ether theory, and the caloric theory, were discarded, successful present theories, such as the Copernican theory, the oxygen theory, the special theory of relativity, and the kinetic theory, will also be discarded. Note that the pessimistic induction consists of the premise that past theories were thrown out, and the conclusion that present theories will also be thrown out.

As a preliminary discussion of the pessimistic induction, consider the following two examples of inductive inference. First, imagine that you ask me what percentage of Americans believe that God exists. I investigate some Germans as opposed to Americans, and then argue that 20% of Americans believe that God exists because 20% of Germans believe that God exists. You would be perplexed by my argument. You would think that I should have used a sample of Americans, not of Germans, because you wanted to know about Americans, not about Germans.

Second, imagine that you ask me how long people live on average today. I show you a one-hundred-year-old medical book that says that one thousand people randomly chosen from the entire population died on average at the age of forty-three, and I conclude that the average life expectancy today is forty-three years. Again, you would be perplexed by my argument. You would think that I should have taken my sample from people living today, not from people living one hundred years ago, because you wanted to know the average life expectancy of people living today, not of people from the past.

These two examples are intended to show that if we want to know whether members of a certain population have a certain property or not, we should draw a sample from that very population, not from some other population, and then we should check whether members of the sample have that property or not. To obtain a sample from a population other than the one under investigation commits what I call the fallacy of irrelevant sample. The fallacy of irrelevant sample is different from the fallacy of biased sample. While the former occurs when a sample is drawn from the wrong population, the latter occurs when a sample is drawn from the right population, but is not representative of that population. However, the fallacy of irrelevant sample and the fallacy of biased sample are similar in that both concern samples that do not adequately represent the population under investigation.

Let me turn back to the pessimistic induction. Pessimists aim to establish the conclusion that present theories are false or unwarranted. They are, however, in a dilemma between committing the fallacy of irrelevant sample and the fallacy of begging the question.

They have to draw their sample either from the set of past theories or from the set of present theories. If they draw the sample from the set of past theories, realists would readily agree that members of the sample are false. Pessimists, however, commit the fallacy of irrelevant sample. If they draw the sample from the set of present theories, they can avoid the fallacy of irrelevant sample, but they commit the fallacy of begging the question. It is under the dispute between pessimists and realists whether present theories are true or false. To claim that members of the sample from the set of present theories are false is to assume the very point under the dispute between pessimists and realists.

This is the context in which the uniformity principle rescues pessimists. The uniformity principle implies that present theories resemble past theories with respect to the issue of whether they will be discarded or retained. It bridges the gap between the sample of past theories and the population of present theories. Thus, the uniformity principle makes it legitimate for pessimists to select the members of the sample from the set of *past* theories, despite the fact that their conclusion is about the population of *present* theories. To put it differently, the uniformity principle shields the pessimistic induction from the charge that it commits the fallacy of irrelevant sample. Thus, the uniformity principle is a vital part of the pessimistic induction.

3. The Disuniformity Principle

Whether or not we ought to accept the pessimistic induction depends on the plausibility of its premise and on whether it is *better* to use the uniformity principle or the disuniformity principle to predict the future of present theories. 'Better' in this context means being more likely to lead to a true prediction about present theories. In this section, I argue that in some cases in science the use of the disuniformity principle is better than that of the uniformity principle, and hence that the burden is on pessimists to justify the choice of the uniformity principle over the disuniformity principle.

Imagine that scientists make several attempts to obtain publishable research results, but all of them fail. Should they use the uniformity principle and predict that they will continue to fail in the future as they did in the past? Or should they use the disuniformity principle and predict that they will succeed although they failed? Let me present two real examples in which scientists chose the disuniformity principle over the uniformity principle before they made significant scientific achievements.

In the 1990s, Ian Wilmut's research team made many failed attempts before they created the first cloned sheep, Dolly. Cloning a sheep involves a complicated process. In Dolly's case, the research team extracted 277 nuclei from the cells of one ewe, inserted them into the enucleated egg cells (i.e., cells whose nuclei have been removed) of another ewe, and then fused the nuclei and enucleated egg cells together with the use of electrical pulses. The resulting eggs were cultured and eventually inserted into surrogate mothers. Only one out of 277 such attempts was successful, and Dolly was the very result of that one successful attempt.

As another example, scientists have recently developed a technique to restore damaged heart muscles in monkeys using human embryonic stem cells (Chong et al., 2014 June). Their technique, if applicable to humans, will benefit those who suffer from heart diseases. The corresponding author of the article, Charles E. Murry, says that his research team made many failed attempts before developing the technique. Although he does not recall the exact number of times their experiments failed, he says, "If 25% of our experiments worked, that would be wonderful!" (Personal Communication).

These two examples indicate that scientists fail multiple times before they produce publishable research results, but that they do not give up their research projects despite the

repeated failures. Why don't they give up their research projects? It is an incomplete answer to say that they receive funding from governments, universities, and private industries. Such an answer raises a further question of the same kind. Why do governments, universities, and private industries continue to fund scientists who have failed to produce the intended results multiple times?

My answer is that scientists believe that they will succeed in the future, even though they have failed in the past. Wilmut's and Murry's teams believed that the future differs from the past, and hence that their future experiments would produce the intended results, even though their past experiments did not. The two research teams did not believe that the future resembles the past, and hence that their future experiments would fail as did their previous experiments. Scientists often use the disuniformity principle to fuel their everyday research activities. Governments, universities, and private industries also use the disuniformity principle when they make decisions to keep funding scientists.

Imagine that pessimists were in Murry's laboratory when Murry was about to conduct the first experiment to infuse human embryonic stem cells to a monkey's damaged heart muscles. Pessimists give Murry the following advice: "You should make sure that your first experiment succeeds. If it fails, that will constitute an inductive rationale for thinking that the next experiment will also fail. If the second experiment fails, the failures of the first experiment and the second experiment will jointly constitute an even stronger inductive rationale for thinking that the third experiment will also fail." Such advice is in perfect line with the pessimistic induction that the demise of past theories constitutes an inductive rationale for anticipating the demise of present theories. Murray, however, would not have taken the advice seriously.

Moreover, pessimists' advice has the following problematic implication. Imagine that you are the founder of a new research field in science and that you just have made five attempts to produce an intended experimental result. Unfortunately, however, all of them failed. An interesting question arises. Which of the five attempts had the highest probability of being successful, other things being equal? Of course, it was the first attempt because it did not have a failed forerunner. The second attempt had a chance to be successful, but its chance was lower than that of the first attempt because the pessimistic specter was conjured up by the failure of the first attempt. Indeed, the more you make mistakes, the stronger the pessimistic specter becomes, and the less likely it becomes that your next attempt will succeed. So the best way to avoid this pessimistic trap is to be successful in the first attempt.

Relatedly, imagine again that you are the founder of a new research field. Pessimists' advice to Murry above implies that if you want to see your field flourish, your first attempt should succeed. If it fails, the prospect becomes gloomy that a successful attempt will finally emerge in your field. After all, as we noted earlier, if the first attempt fails, the second attempt becomes less likely to be successful than it would in the absence of the first attempt. Thus, your first attempt might ruin all the subsequent research in your field. As the founder of a research field, you should keep in mind that your first failure has such grave repercussions for all subsequent research.

The phrase 'the achievement after trials and errors' is not in pessimists' dictionary. On their account, what comes after trials and errors are only other trials and errors. After all, the future resembles the past! So it is pointless to make repeated failures to achieve a goal. If you failed in the first attempt in a new research field, you are advised to leave the field and find a new one because you have a better chance to succeed in the new field than in the old field. If you do not leave the old field, the chances are that you will flounder in the pessimistic trap indefinitely. When you move to a new field, you should make sure that no one previously made failed attempts in the new field because if someone already made failed attempts

unbeknownst to you, you are likely to fall prey to the pessimistic specter hidden in the new field. Also, when you leave your old field, you are asked to leave a note for prospective researchers that you made failed attempts in the field. With the note, they will make an informed choice as to whether they will jump in the field or not. It is unethical for you to leave the field without such a note. Why unethical? Imagine that you created a big hole on the road and leave the place without setting up a warning sign. Such an action is immoral because innocent cars behind you will fall into the hole. There is no relevant difference between the two cases. If you leave your old research field without the note, subsequent researchers will only waste their time.

The preceding discussion indicates that there is something terribly wrong with applying the uniformity principle to daily scientific activities and that we should use instead the disuniformity principle. It is not clear how many times scientists must fail before they succeed. Presumably, different scientists make different kinds of mistakes and fail at different rates, depending on how much background knowledge they have, how lucky they are, how difficult their tasks are, and so forth. What is important for the purpose of this paper is that as scientists fail, they get closer to a point that I call the 'flipping point.' Once they go above the flipping point, their new attempt is dissimilar to their previous failed attempts, i.e., the new attempt is successful. To define 'flipping point' in general terms, an unobserved instance is similar and dissimilar to observed instances below and above a flipping point, respectively. Our prediction about an unobserved instance below and above a flipping point should be built upon the uniformity principle and the disuniformity principle, respectively. Otherwise, the prediction is false. Let me emphasize that we ought to use the disuniformity principle instead of the uniformity principle above a flipping point to gain a true belief and to avoid a false belief about an unobserved instance.

In certain cases, I admit, there is no flipping point, and hence we should give up our research project once we underwent repeated failures. Astrologers, for example, thought that there was a causal relationship between astronomical events and humans events. If we repeatedly failed to produce intended results in such projects, we should at some point think that we will fail indefinitely into the future. In such cases, repeated previous failures constitute the good reason for thinking that the uniformity principle applies indefinitely, and hence a further attempt is pointless.

It is clear, however, that the disuniformity principle is useful in certain cases in science and that we *might* reasonably use it to appraise the fate of present theories, inferring that present theories are true because past theories were false. So the burden is on pessimists to justify the choice of the uniformity principle over the disuniformity principle. In virtue of what are pessimists justified in using the uniformity principle to assess the prospect of present theories, given that scientists, governments, universities, and private industries are sometimes justified in using the disuniformity principle to assess the prospects of the various research programs? What is the reason for thinking that the use of the uniformity principle is more likely to produce a true belief about the status of present theories than the use of the disuniformity principle is? In the absence of an answer to this question, it is an open question whether the pessimistic induction is correct or incorrect, and hence whether our present theories will go the way of our past theories or not.

It is an open question, I admit, whether we should use the uniformity principle or the disuniformity principle concerning the status of present theories. The history of science indicates that scientists made many failed attempts to produce true theories about the world. If they had made enough mistakes, their last attempts would have been above flipping points, and hence we should use the disuniformity principle and infer that present theories are true unlike past theories. If they had not yet made enough mistakes, they would need to make

more mistakes before their next attempts go above flipping points, and hence we should now use the uniformity principle and infer that present theories are false like past theories. The pessimistic induction requires the choice of the uniformity principle, whereas realism requires the choice of the disuniformity principle. Since the choice of the uniformity principle is not justified yet, we have no reason for thinking that the pessimistic induction is correct.

Let me introduce two fallacies: the uniformity fallacy and the disuniformity fallacy. The uniformity fallacy occurs when an unobserved instance is above a flipping point, but we erroneously use the uniformity principle and infer that it is similar to observed instances. The disuniformity fallacy is the opposite of the uniformity fallacy. It occurs when an unobserved instance is below a flipping point, but we erroneously use the disuniformity principle and infer that it is dissimilar to observed instances. My criticism against the pessimistic induction suggests that pessimists might commit the uniformity fallacy.

A referee objects that the uniformity principle and the disuniformity principle should not be as symmetrical as I assume. Specifically, when we predict the future, we should take the uniformity principle to be the default assumption. In other words, we should use the disuniformity principle only when we have special reasons to use it; when we do not have such reasons, we should instead use the uniformity principle.

The referee has raised an interesting issue. I propose that a principle that is used more often than the other is the default assumption. Thus, if we use the uniformity principle more often than the disuniformity principle, the uniformity principle is the default assumption; if we use the disuniformity principle more often than the uniformity principle, the disuniformity principle is the default assumption. Only an empirical study can determine which principle is the default assumption. I leave this research project to psychologists.

4. Difference
Ludwig Fahrbach (2011: 148), Seungbae Park (2011a: 79), and Moti Mizrahi (2013a: 3220) observe that Laudan's (1981: 33) list of past theories is composed exclusively of *distant* past theories, although the set of *recent* past theories is much larger than the set of distant past theories. Recent past theories are more successful than distant past theories, and the former have not yet been refuted, although the latter were refuted before the 20th century. Since Laudan's list does not include recent past theories, the pessimistic induction based on it commits the fallacy of biased sample.

How does my criticism differ from Fahrbach, Park, and Mizrahi's foregoing criticism against the pessimistic induction? The two criticisms attack different aspects of the pessimistic induction. Recall that the pessimistic induction consists of the premise that past theories were rejected and the conclusion that present theories will be rejected. My charge that the pessimistic induction might commit the uniformity fallacy is directed at the inference from the premise to the conclusion. After all, my point is that the inference is questionable because we might have to use the disuniformity principle instead of the uniformity principle. In contrast, the three philosophers' charge that the pessimistic induction commits the fallacy of biased sample is leveled at the premise. After all, their point is that the premise is dubious because the sample provided to justify the premise is biased. Therefore, my criticism and their criticism have different targets.

Two conditions have to be met for the pessimistic induction to be convincing. First, the sample of *past* theories should adequately represent the population of *past* theories. Second, the use of the uniformity principle should be better than that of the disuniformity principle to appraise the fate of present theories. If at least one of these two conditions is not met, the pessimistic induction is questionable, and the pessimistic conclusion about present theories is unwarranted. Fahrbach, Park, and Mizrahi's criticism is that the first condition is not met. My

criticism is that the second condition is not met. The two criticisms, although different, can form an alliance against the pessimistic induction.

5. Similarities

My criticism against the pessimistic induction is similar to a famous antirealist critique of inference to the best explanation. The antirealist critique is known by three names: 'the threat of the unknown hypothesis' (Forrest, 1994), 'the argument from a bad lot' (Psillos, 1999: 215), and 'the argument from underconsideration' (Wray, 2008; Khalifa, 2010; Wray, 2012; Mizrahi, 2013b). Suppose that scientists thought up a few rival theories and compared them. One of them is better than the others. Realists conclude that the best available theory is true. The argument from underconsideration says that there *might* be rival theories of which scientists have not yet conceived. To conclude that the best available theory is true, realists should first show that a true theory is more likely to exist in the range of conceived rival theories than in the range of unconceived rival theories. Bas C. van Fraassen puts the argument from underconsideration as follows:

> So our selection may well be the best of a bad lot. … For me to take it that the best of set X will be more likely to be true than not, requires a prior belief that the truth is already more likely to be found in X, than not. (van Fraassen, 1989: 143)

James Ladyman, Igor Douven, Leon Horsten, and van Fraassen (1997) emphasize that van Fraassen does not have to present reasons for thinking that a true theory is more likely to be in the set of unconceived theories than in the set of conceived theories:

> So whereas Psillos challenges van Fraassen to show that it is more likely that the truth is outside the range, van Fraassen need only ask the proponent of IBE for reasons for believing that the truth is inside it. (Ladyman, Douven, Horsten, and van Fraassen, 1997: 306)

On these philosophers' account, the mere possibility that a true theory might lie in the range of unconceived alternatives gives rise to the burden for realists to justify the belief that the true theory is more likely to lie in the set of conceived alternatives than in the set of unconceived alternatives. The argument from underconsideration is embraced not only by Ladyman et al. but also by Kareem Khalifa (2010: 95-97) and Brad Wray (2008: 322; 2012: 277).

The argument from underconsideration and my criticism against the pessimistic induction are similar in the following three important respects: First, Ladyman et al. and I created the burden of justification for their opponent and my opponent by presenting alternative possibilities. Ladyman et al. created the burden for realists to justify the belief that a true theory is more likely to lie in the range of conceived alternatives than in the range of unconceived alternatives by pointing out that a true theory *might* lie in the range of unconceived alternatives. Similarly, I have created the burden for pessimists to justify the belief that it is better to use the uniformity principle than the disuniformity principle by pointing out that the disuniformity principle *might* be reasonably used.

Second, recall that van Fraassen does not have to justify the belief that a true theory is more likely to exist in the range of unconceived alternatives than in the range of conceived alternatives. The burden of justification falls on realists because it is realists, not van Fraassen, who believe that the best of conceived theories is true. Analogously, I do not have to justify the belief that the use of the disuniformity principle is better than that of the uniformity principle for predicting the fate of present theories. I need only ask pessimists for reasons for

choosing the uniformity principle over the disuniformity principle because it is pessimists, not me, who use the uniformity principle to infer that present theories will be abandoned.

Third, Ladyman et al.'s point is not that realists should provide a deductive rationale but that they should provide an inductive rationale for thinking that the best of conceived theories is true. They do not demand that realists should prove that it is 100% probable that the true theory is in the range of conceived rivals, but only that the true theory is *more* likely to exist in the set of conceived theories than in the set of unconceived theories. They would be persuaded, if realists show, for example, that it is 60% probable and 40% probable, respectively, that a true theory exists in the set of conceived rivals and in the set of unconceived rivals. Therefore, Ladyman et al. are not imposing an inordinately high standard of justification on realists.

Analogously, my point is not that pessimists should provide a deductive rationale but that they should provide an inductive rationale for their belief that present theories will be overthrown. I am not demanding that pessimists should prove that the uniformity principle should definitely be used, but only that it is *better* to use the uniformity principle than the disuniformity principle to assess the fate of present theories. In other words, pessimists have to justify the position that the use of the uniformity principle is more likely to lead to a true belief about the fate of present theories than the use of disuniformity principle is. They should show, for example, that if we use the uniformity principle, it is 60% probable that our belief about the fate of present theories is true and that if we use the disuniformity principle, it is 40% probable that our belief about the fate of present theories is true. Therefore, I am not imposing an excessively high standard of justification on pessimists.

So far I have stressed the similarities between the argument from underconsideration and my criticism against the pessimistic induction. The point of stressing the similarities is to show that they rise or fall together. If my criticism against the pessimistic induction is illegitimate, so is the argument from underconsideration. It is wrong to accept the argument from underconsideration but reject my criticism against the pessimistic induction on the grounds, for example, that I am imposing too high a standard of justification on pessimists.

6. Conclusion

Scientists use the disuniformity principle before they achieve success in their research field. Scientific progress would be impossible if scientists reasoned in accordance with the uniformity principle and believed that the future, like the past, holds nothing but failures. Given that the disuniformity principle applies to certain cases in science, we *might* reasonably use it to assess the prospects of present theories, and hence the burden is on pessimists to justify the position that our belief about the fate of present theories is more likely to be true if we use the uniformity principle than if we use the disuniformity principle.

I admit, though, that my criticism against the pessimistic induction merely conjures up a burden of justification for pessimists. It is one thing to create the burden. It is another to establish realism. Establishing realism requires a good argument for it, and I did not provide any positive argument for it in this paper, although I did it elsewhere (Park, 2011b: 23-25). This paper takes no position on whether or not current theories are true, or on whether or not we are justified in believing that they are true. The aim of this paper is not to justify realism but to defuse the pessimistic induction by taking issue with pessimists' use of the uniformity principle.

Finally, failure is a part of scientific life. Scientists are faced with the choice between the uniformity principle and the disuniformity principle all the time. So far as I know, there is no philosophical theory telling us when to use the uniformity principle and when to use the disuniformity principle. Such a theory, if developed, will have both philosophical and

practical implications. A philosophical implication is that it will tell philosophers whether to use the uniformity principle or the disuniformity principle to predict the fate of present theories. A practical implication is that it will tell funding agencies whether to keep funding a research team or not. I leave the task of developing such a theory on a future occasion.

References
Chong, James J. H., Xiulan Yang, Creighton W. Don, Elina Minami, Yen-Wen Liu, Jill J. Weyers, William M. Mahoney, Benjamin Van Biber, Savannah M. Cook, Nathan J. Palpant, Jay A. Gantz, James A. Fugate, Veronica Muskheli, G. Michael Gough, Keith W. Vogel, Cliff A. Astley, Charlotte E. Hotchkiss, Audrey Baldessari, Lil Pabon, Hans Reinecke, Edward A. Gill, Veronica Nelson, Hans-Peter Kiem, Michael A. Laflamme & Charles E. Murry (2014, June). "Human Embryonic-Stem-Cell-Derived Cardiomyocytes Regenerate non-Human Primate Hearts", *Nature* 510: 273-277.

Fahrbach, Ludwig (2011). "How the Growth of Science Ends Theory Change", *Synthese* 180 (2): 139-155.

Forrest, Peter (1994). "Why Most of Us Should Be Scientific Realists: A Reply to van Fraassen", *The Monist* 77 (1): 47-70.

Hume, David (1978). *A Treatise of Human Nature*. L. A. Selby-Bigge and P. H. Nidditch (eds.), Oxford University Press.

Khalifa, Kareem (2010). "Default Privilege and Bad Lots: Underconsideration and Explanatory Inference", *International Studies in the Philosophy of Science* 24 (1): 91-105.

Ladyman, James, Igor Douven, Leon Horsten, and Bas van Fraassen (1997). "A Defense of Van Fraassen's Critique of Abductive Inference: Reply to Psillos", *The Philosophical Quarterly* 47 (188): 305-321.

Laudan, Larry (1977). *Progress and Its Problems: Towards a Theory of Scientific Growth*. California: University of California Press.

---------- (1981). "A Confutation of Convergent Realism", *Philosophy of Science* 48 (1): 19-49.

Mizrahi, Moti (2013a). "The Pessimistic Induction: A Bad Argument Gone Too Far", *Synthese* 190 (15): 3209-3226.

---------- (2013b). "The Argument from Underconsideration and Relative Realism", *International Studies in the Philosophy of Science* 27 (4): 393-407.

Park, Seungbae (2011a). "A Confutation of the Pessimistic Induction", *Journal for General Philosophy of Science* 42 (1): 75-84.

---------- (2011b). "Coherence of Our Best Scientific Theories", *Foundations of Science* 16 (1): 21-30.

Psillos, Stathis (1999). *Scientific Realism: How Science Tracks Truth.* New York: Routledge.

Putnam, Hilary (1978). *Meaning and the Moral Sciences*. London: Routledge & K. Paul.

van Fraassen, Bas C. (1989). *Laws and Symmetry*. Oxford: Oxford University Press.

Wray, K. Brad (2008). "The Argument from Underconsideration as Grounds for Anti-Realism: A Defence", *International Studies in the Philosophy of Science* 22 (3): 317-326.

---------- (2012). "Epistemic Privilege and the Success of Science", *Noûs* 46 (3): 375-385.