

Intelligibility is Necessary for Scientific Explanation, but Accuracy May Not Be

Mike Braverman (braverm2@illinois.edu)

Department of Psychology, 603 E. Daniel Street
Champaign, IL 61820 USA

John Clevenger (jcleven2@illinois.edu)

Department of Philosophy, 810 S. Wright Street
Urbana, IL 61801 USA

Ian Harmon (iharmon2@illinois.edu)

Department of Philosophy, 810 S. Wright Street
Urbana, IL 61801 USA

Andrew Higgins (higgins9@illinois.edu)

Department of Philosophy, 810 S. Wright Street
Urbana, IL 61801 USA

Zachary S. Horne (horne2@illinois.edu)

Department of Philosophy, 810 S. Wright Street
Urbana, IL 61801 USA

Joseph Spino (spino2@illinois.edu)

Department of Philosophy, 810 S. Wright Street
Urbana, IL 61801 USA

Jonathan Waskan (waskan@illinois.edu)

Department of Philosophy, 810 S. Wright Street
Urbana, IL 61801 USA

Abstract

Many philosophers of science believe that empirical psychology can contribute little to the philosophical investigation of explanations. They take this to be shown by the fact that certain explanations fail to elicit any relevant psychological events (e.g., familiarity, insight, intelligibility, etc.). We report results from a study suggesting that, at least among those with extensive science training, a capacity to render an event intelligible is considered a requirement for explanation. We also investigate for whom explanations must be capable of rendering events intelligible and whether or not accuracy is also viewed as a requirement.

Keywords: science; explanation; psychologism; intelligibility.

Introduction

The nature of explanation has been a major topic of investigation in the philosophy of science for at least sixty years. While much is still disputed, a consensus has emerged that scientific explanations are not constituted by psychological events. Philosophers of science first made arguments to this effect in the middle of the 20th century in response to the charge that explanations play no useful role in science because they are constituted by subjectively variable psychological states (viz., familiarity and empathy).

Hospers (1946), Miller (1947), and Hempel (1965), for instance, argued that many legitimate explanations appeal to principles (e.g., Newton's law of gravitation) that were utterly unfamiliar when they were first introduced. Hempel (1942) also noted that historical explanations sometimes refer to individuals (e.g., paranoiacs) with whom most are incapable of empathizing. On such grounds Hempel famously rejects psychologistic theories of explanation in favor of a more "objective" account (1965, 426).

More recently, philosophers of science have viewed psychologistic theories as equating explanations with other feelings, such as insight, satisfaction, or "aha" feelings (Craver 2007; Salmon 1984; Trout 2007). Like Miller and Hempel, they rebut such psychologistic proposals by pointing to cases of explanation where the relevant feelings are absent. One common strategy is to point to a putative explanation, whether it be a passage of text describing a process of speciation (Trout 2007) or a computer simulation of the human nervous system (Craver 2007), whose complexity so outstrips the limits of human memory and attention that humans find it incomprehensible. Humans thus fail to derive from these putative explanations any feelings of insight, satisfaction, etc. It is frequently concluded on the basis of such examples that explanations are non-psychological - we term this the *objectivity*

hypothesis - and hence that psychological research will contribute little to our understanding of explanations.¹

We believe that this anti-psychologistic attitude is wrongheaded for a few reasons. One is that the very process of justifying theories of explanation on the basis of how well they track philosophers' classifications (i.e., as explanations or non-explanations) of various representations are laden with psychological presuppositions. Philosophers appear to assume that they have, through their exposure to science and scientists, come to possess tacit knowledge of the norms regarding the proper use of 'explanation.' Whether or not philosophical judgments mirror scientific ones is, however, an empirical matter that is best resolved through psychological investigation. We suspect, moreover, that scientists do regard a certain kind of conscious psychological event as necessary for explanation. The type of event we have in mind is not an affect-laden *feeling*, but rather the more intellectual process of understanding how or why, at least possibly, an event came about. This state is sometimes known as finding a happening intelligible or making sense of it (cf. Machamer & Woody 1994). Familiar objections have been raised against this proposal as well - namely, that there are (e.g., hyper-complex) explanations that never render anything intelligible to anyone. Our specific hypothesis, however, is that scientists will not consider a representation an explanation unless it has the *capacity* to render intelligible, which we term *intelligibility*. It would thus not undercut our position if there were explanations that never actually render the event intelligible to anyone. Consider, by comparison, that a liquid may be a solvent of salt even if at a given time it happens not to be dissolving any salt at all or even if it never does so (viz., because the opportunity never arises). It must merely have the capacity to do so, even if it is not exercised. Likewise, we suspect that scientists require not that a representation actually renders a happening intelligible to someone, but merely that it has the capacity to do so.

Notice that even if a liquid that has the capacity to dissolve salt never exercises that capacity, it would still tell us a lot about what makes that liquid a solvent of salt if we had information about the process by which it would dissolve salt were the opportunity to arise. This is best accomplished by studying cases in which solvents of salt actually do exercise their capacity to dissolve salt. Matters are somewhat more complicated with regard to explanations, but we believe that even if a certain representation that has the capacity to make things intelligible to people never actually exercises that capacity, it may well tell us a lot about what makes that representation an explanation if we had information about the process by which it would make sense of things were the opportunity to

arise. This would be best accomplished by studying cases in which explanations actually do exercise their intelligibility.

Notice also that if the reason why a liquid has failed to dissolve any salt is that there is something about the nature of the liquid that precludes it from ever dissolving any salt, the liquid then lacks even the capacity to dissolve salt and is thus no solvent of it. Likewise, if there is something about a representation that precludes it from ever rendering an event intelligible - for instance, it is too complex to ever be comprehended by anyone at all or it refers to things like extra dimensions that are utterly incomprehensible to anyone - then it lacks even the capacity to render intelligible and scientists will thus not regard it as an explanation. We suspect that laypeople also view intelligibility as a requirement and will classify such cases in a similar way.

If a representation must have the capacity to render intelligible in order for it to be considered an explanation, a further question naturally arises: Intelligible to whom? Scientists generally interact with at least some colleagues who exceed their own intelligence, so they are almost certainly cognizant of the fact that some representations are too complex to make events intelligible to everyone. Thus, they likely do not require that a representation must have the capacity to render things intelligible to just anyone in order to be considered an explanation. However, since scientists presumably do not interact with beings that possess utterly different perceptual and cognitive abilities, they might think that a representation must have the capacity to make things intelligible to beings basically like themselves in order to be considered an explanation. Then again, they may turn out to be even more liberal, merely requiring that a representation makes sense of things for sentient beings of some sort, even ones with completely alien thought processes. There are numerous possibilities here, with regard to both scientists and laypeople, that are worth investigating.

Another widely accepted view among philosophers of science is that a high degree of accuracy is essential in order for a representation to be an explanation (Craver 2007; C.G. Hempel 1965; Humphreys 1989; Salmon 1998; Trout, 2007). We call this the *accuracy* hypothesis. We think a representation must merely specify a possible way, even if it is not the actual way, in which an event occurred for the representation to be considered an explanation. We call this the *plausibility* hypothesis. An implication of the accuracy hypothesis is that scientists will be less likely to judge that a representation is an explanation if they are told that it is merely possibly accurate (which is to say that it might be inaccurate) or that it once seemed possibly accurate but was later falsified. Cases of the latter sort mirror many cases from this history of science where a theory seemed to make sense of things (e.g., the Ptolemaic theory of planetary retrograde motion) but was ultimately proven false. If scientists and laypeople do regard such cases as explanations, telling them that a representation is merely possibly accurate or conceivably accurate but factually

¹ Trout, admittedly, seems in places to be opposed to the idea that explanations are constituted by *conscious* psychological states in that he allows that explanations sometimes involve implicit learning.

inaccurate should not undermine their tendency to regard a representation as an explanation.

We tested whether people treated either intelligibility or accuracy as a necessary condition for something to be an explanation. We tested this with regard to two groups of participants: those with and without extensive science training.

Methods

Participants

The participants in this study were 297 workers recruited using Amazon Mechanical Turk with varying levels of science training.

Materials

We used stories about the origins of either life, color experiences, or gamma ray bursts, and manipulated the characteristics of these stories. Vignettes were grouped into following three sets:

Set 1: A potential explanation (viz., a passage of text) is described that has various theoretical virtues and that also supplies some level of understanding. Specifically, the representation either supplies understanding of how the target happening actually occurred (A), of a possible way in which it occurred (PA), or of a possible way in which it occurred that is eventually shown to be false (PAF).

Set 2: A potential explanation is described which is said to be incapable of rendering a happening intelligible to humans because of our cognitive limitations - that is, either the representation defies the limits of normal human working memory and attention or it refers to highly exotic, hyper-dimensional properties. The possibility is then introduced that the representation would render the happening intelligible to beings with cognitive capacities that are, roughly speaking, quantitatively better (i.e., involving augmented memory and attention span) or qualitatively different (i.e., involving the ability to think in extra dimensions). Cases of the former sort involve hyper-complex passages of text (AM). Cases of the latter sort (HD) refer to passages of text that describe hyper-dimensional properties.

Set 3: This set closely mirrors Set 2 except that the complexity and exotic nature of the putative explanation is such that it precludes the representation from rendering the target intelligible to anyone at all. This set includes passages of text (IQuant) or computer simulations (Sim) that are hyper-complex, and passages of text that refer to exotic properties (IQual).

There were three storylines and nine vignette types, yielding twenty-seven possible vignettes. Each participant read and responded to three vignettes. One vignette came from each set, and each followed a different storyline.

Collection

Participants were recruited through the Amazon Mechanical Turk (MTurk) work-distribution website. To be

eligible, workers had to be in the U.S. and have at least a 75% approval rate. Eligible workers were redirected to SurveyMonkey, where they completed the study. Afterwards, workers were directed back to MTurk, where they were compensated with \$.50.

Procedure

Participants were presented with a vignette (story) on a computer screen. Each vignette referred either to a typewritten description of a physical process or to a computer simulation thereof. For example, some participants saw the following vignette²:

Dr. Nikro is a little-known, very-gifted scientist investigating the manner in which color experiences arise in the brain. He spends years tinkering with the complex equations of neurochemistry and subatomic physics and considering the different possible locations in which color experiences might originate. He eventually chanced upon a remarkable series of calculations, which he posts to his rarely visited public webpage. They indicate that conditions like those found in the pyramidal cells of the cortex would, over the course of hundreds of milliseconds, reliably undergo a series of changes resulting in the creation of color experiences. The calculations fit quite nicely with the most widely accepted theories and observations from a variety of related fields. Yet they merely refer to ordinary things like membranes and neural firings, of which anyone could easily conceive. As a result, Dr. Nikro comes to understand the actual manner in which color experiences are generated in the brain, as would any specialist from his field who were to study the details of the material he posted.

After reading the vignette, participants were asked to rate the extent to which they agree (-3 strongly disagree, 0 neutral, +3 strongly agree) that the representation described in the vignette constitutes an explanation. There was some concern that participants might confuse this question with the question of whether or not the representation constitutes a *good* or a *satisfying* explanation. Participants were thus informed at the outset that, insofar as they do agree that the representation constitutes an explanation, they should also specify the extent to which they agree that the representation constitutes a *good* and a *satisfying* explanation (cf. Lombrozo and Carey 2004).

Mishra & Brewer (2003) found that participants paid closer attention to the contents of their study materials (i.e., real-world explanations) if they were forewarned that they would be asked a series of simple comprehension questions. We likewise informed participants in advance that they would be asked a series of simple questions. Some of these

² A demonstration version of the study can be found at www.surveymonkey.com/s/PPG6DW6.

concerned aspects of the vignettes (e.g., accuracy or intelligibility) that were particularly salient to the experiment.

This sequence of instructions, vignettes, and questions repeated two more times, each time with vignettes from a different set and on a different topic (i.e., life, color experiences, or gamma-ray bursts). Participants completed a distractor task in between each vignette. At the end of the study, participants were asked some follow-up questions, including questions about their level of science training.

Six separate studies were constructed to balance the order in which participants saw the various storylines, with one study for each permutation (i.e., Life-Color Experiences-Gamma Rays, Life-Gamma Rays-Color Experiences, etc.).

Results

Of the 297 completed studies, data from 38 of those studies were excluded because the participants involved had already completed the study at least once. After this exclusion there were 259 completed studies or 777 sets of responses to particular vignettes. If a participant failed to answer at least two of three comprehension questions correctly for a particular vignette, responses to that vignette were excluded. Using this criterion, 61 individuals had responses to one vignette eliminated, 9 individuals had responses to two vignettes eliminated, and 5 individuals had responses to all three vignettes eliminated. After eliminating these problematic responses, responses to 683 vignettes were recorded and analyzed.

We divided our sample into high-science (5 or more college-level science courses) and low-science (less than 5 college-level science courses) groups to investigate the ways in which laypeople and scientists conceive of explanations. After the division, the high-science group contained the following number of responses to the ‘constitutes an explanation’ question for each of the following conditions: 26 accurate (A), 32 possibly accurate (PA), 22 possibly accurate but false (PAF), 28 augmented memory (AM), 28 hyper-dimensional thinking (HD), 22 unintelligible-quantitative (IQuant), 30 unintelligible-qualitative (IQual), and 27 unintelligible-quantitative simulation (SIM).

The low-science group contained the following number of responses to the ‘constitutes an explanation’ question for each of the following conditions: 37 accurate (A), 58 possibly accurate (PA), 51 possibly accurate but false (PAF), 57 augmented memory (AM), 62 hyper-dimensional thinking (HD), 21 unintelligible-quantitative (IQuant), 71 unintelligible-qualitative (IQual), and 59 unintelligible-quantitative simulation (SIM).

The mean responses and standard deviations to the explanation rating task are located in Tables 1 and 2.

Table 1: Mean Responses and Standard Deviations for High-Science Participants

	A	PA	PAF	AM	HD	IQuant	IQual	SIM
Avg.	.769	.968	.090	.285	-.357	-.727	-1.1	-.370
SD	1.77	1.80	2.02	1.95	1.66	1.72	1.78	1.94

Table 2: Mean Responses and Standard Deviations for Low-Science Participants

	A	PA	PAF	AM	HD	IQuant	IQual	SIM
Avg.	1.29	1.31	-.137	.350	-.241	-.714	-.859	-.610
SD	1.66	1.42	1.78	1.67	1.64	1.87	1.87	1.65

After the division of the sample, mean ratings were measured for differences using 18 two-tailed independent samples *t*-tests. To guard against increased chance of Type I errors, we performed a Bonferroni correction to adjust α (significant results were $p < .0027$).

In the low-science group, descriptions – be they written text or simulations – that were unintelligible to anyone, resulted in lower ratings than descriptions that were intelligible to the scientists in the vignettes (A and IQuant ($t(56) = 4.22, p < .0027$), A and IQual ($t(106) = 5.87, p < .0027$) and A and SIM ($t(94) = 5.48, p < .0027$)). Low-science participants also gave lower ratings when *only* cognitively advanced beings – whether they be humans with augmented memory or beings capable of perceiving extra dimensions – found the descriptions intelligible (PA and AM ($t(113) = 3.30, p < .0027$), PA and HD ($t(118) = 5.49, p < .0027$)). Significance obtained when comparing low-science participants’ ratings of descriptions described as false to descriptions described as accurate or possibly accurate (A and PAF ($t(86) = 3.82, p < .0027$), PA and PAF ($t(107) = 4.68, p < .0027$)). Whether a description was accurate or possibly accurate, however, made no difference to low-science participants’ judgments ($t(93) = 41, p = .96$).

For the high-science group, descriptions incapable of rendering the event intelligible to anyone resulted in lower ratings than those that made the event intelligible to the scientists in the vignettes (A and IQual ($t(54) = 3.911, p < .0027$)). However, these differences did not hold across all Set 3 vignettes (A and IQuant ($t(46) = 2.95, p = .004$), A and Sim ($t(51) = 2.25, p = .028$)). For high-science participants, ratings were not significantly different for descriptions described as false, or only possibly accurate, as compared to those described as accurate (A and PAF ($t(46) = 1.24, p = .222$), A and PA ($t(56) = .422, p = .675$), PA and PAF ($t(52) = 1.67, p = .10$)). Likewise, these participants’ ratings were not significantly affected when a description was described as intelligible only to beings with quantitatively or qualitatively different cognitive capacities (PA and AM ($t(58) = 1.40, p = .16$), PA and HD ($t(58) = 2.94, p = .004$)).

Discussion of Results

This study sought to answer the following questions:

1. Does intellig-ability matter? If the objectivity hypothesis captures how laypeople and practicing scientists conceive of explanations, then altering whether a representation is described as capable of rendering a happening intelligible should not affect the judgments of either low-science or high-science participants. If the

intellig-ability hypothesis is correct, participants should be less likely to regard a representation as an explanation when told that it is, whether due to sheer complexity or to exotic constructs, incapable of rendering a happening intelligible to anyone. Low-science participants were less likely to judge a representation to be an explanation when told that the representation lacks intellig-ability. The judgments of low-science participants thus indicate that laypeople conceive of explanations in a way consistent with the intellig-ability hypothesis and inconsistent with the objectivity hypothesis.

High-science participants were less likely to regard a passage of text as an explanation when told that the representation, because of qualitative barriers (IQual), lacks intellig-ability. They were, however, not significantly less likely to regard a passage of text (IQual) or a simulation (SIM) as an explanation when told that the representation, because of quantitative barriers, lacks intellig-ability. One possibility is that quantitative barriers to intelligibility seem far less daunting than qualitative ones. Indeed, those with extensive science training should be well aware that myriad techniques have been developed for analyzing and visualizing information regarding the behaviors of complex systems with the precise point of rendering those systems intelligible to humans. This may have had enough of an impact on judgments to weaken the intellig-ability effect. Regardless, the differences in high-science ratings between (PA), on the one hand, and (IQuant) and (SIM) on the other were large enough to warrant further investigation.

The difference between (A) and (IQual) is, by itself, clearly inconsistent with the objectivity hypothesis. Admittedly, we cannot be sure that the judgments of our high-science participants mirror those of professional scientists, so further study of practicing scientists is also needed to better support our proposal that intellig-ability is a requirement for explanation. However, taken as a whole, this set of results strongly suggests that manipulating intellig-ability alters explanation judgments, and we take this to be inconsistent with the objectivity hypothesis.

2. Intellig-ability to Whom? The results above indicate that people treat intellig-ability as in some way necessary for explanation, but what is not clear is for whom a representation must be intellig-able (i.e., able to make sense of things). To address this question, participants were asked to consider cases in which potential explanations are unintellig-able to humans because of cognitive limitations. In the (AM) case, the representation exceeds the limits of normal human working memory and attention. In another case (HD), the representation refers to exotic properties such as hyper-dimensionality (see Set 2 in Materials). The possibility is then introduced that the representation may be intellig-able to beings with, roughly speaking, quantitatively augmented memory and attention or qualitatively different cognitive capacities. We take such beings to lie along a similarity continuum such that creatures who merely have augmented memory and attention are more similar to present-day humans than creatures with qualitatively different cognitive capacities.

As it turns out, relative to judgments regarding (PA) vignettes, low-science participants are less inclined to regard a description as an explanation in the (AM) and (HD) conditions. Relative to that same baseline (PA), there was no indication that high-science participants are less inclined to regard a description as an explanation in the (AM) condition. Nor were high-science participants significantly less likely to regard a description as an explanation in the (HD) condition than in the (PA) condition. Although the difference in judgments is not significant ($p = .004$), it is suggestive enough to warrant further investigation.

There are a number of ways of interpreting these findings. Starting with low-science individuals, one possibility is that they believe representations must be intellig-able to present-day humans. This condition was satisfied in the (PA) vignettes but not in the Set 2 (i.e., (AM) or (HD)) vignettes. Another possibility is that these individuals' expectations are somewhat more flexible, requiring merely that representations be intellig-able to beings fairly similar to present-day humans. It may just be that they consider the differences between present-day humans and the beings described in the Set 2 vignettes to be too stark to meet this condition. It is also possible that the reason why they were less inclined to regard the representations in the Set 2 vignettes as explanations had nothing to do with intellig-ability. It may have been, rather, that the representations did not *actually* render the target happening intelligible to anyone. On this view, it is still an open question as to what sorts of beings a representation must actually make sense of things.

Within the high-science group, there were not significant differences between ratings of (PA) and any of the Set 2 vignettes (though, as mentioned, the comparison to (HD) ($p = .004$) is suggestive). Thus, it appears that high-science individuals may be quite flexible about to whom a representation must be intellig-able in order for it to count as an explanation. They do not expect that representations must be intellig-able to present-day humans, but it is unclear whether there is some upper bound on their expectations. They may think a representation that is intellig-able to any kind of sentient being at all counts as an explanation, or they might require that it be intellig-able to beings relevantly similar to present-day humans. In the latter case, high-science individuals may have such a liberal understanding about what counts as relevantly similar that the beings described in the Set 2 vignettes still meet this condition.

Clearly, further studies are also warranted here in order to determine for whom, precisely, a representation must be capable of making sense of things in order to be considered an explanation and whether or not low-science individuals require that this capacity be exercised.

3. Does accuracy matter? The accuracy hypothesis suggests that high-science participants should be more likely to judge something an explanation when it is accurate than when its accuracy is in question or when it is false. This pattern of results did not obtain. There was no difference in judgments about descriptions depicted as accurate (A) and

those depicted as possibly accurate (PA). As mentioned above, the possibly accurate but false (PAF) vignettes mirror many historical examples in which a representation of a way that things could have occurred is eventually discredited. We hypothesized that there would be no difference between judgments regarding the (A) and (PAF) vignettes. Results for high-science participants were consistent with this prediction in that there was no significant difference in their judgments about the two types of vignettes. However, we did find that low-science participants were significantly less likely to judge that the descriptions in the (PAF) vignettes are explanations compared with the descriptions in the (A) vignettes. The fact that the (A) versus (PAF) manipulation produced an effect in the low-science group alleviates some of our concerns about the relevance of the null result found in the high-science group. We thus once again take our findings to supply tentative evidence that the accuracy hypothesis does not reflect the views of practicing scientists, though we acknowledge that additional research on the matter is desirable.

Anti-Psychologism Revisited

We were motivated to undertake this project in large part because we reject the anti-psychologistic stance regarding the study of explanation that still pervades much of the philosophy of science. Philosophers often take this stance to be justified by appeal to cases which they regard as explanations but which seem not to elicit any relevant psychological events (e.g., familiarity, insight, intelligibility, etc.). Our data suggest, however, that scientists and laypeople have different intuitions from philosophers regarding some of these cases. Insofar as there is something about a representation, whether sheer complexity or exotic constructs, that positively precludes it from rendering a certain happening intelligible, high-science and low-science participants are significantly less inclined to regard it as an explanation. Thus, the anti-psychologistic position appears to rest upon intuitions that are, without discernible justification, idiosyncratic; psychological investigations are the proper methodology for determining how scientists conceive of explanations. In addition, given that the capacity to produce a certain kind of mental state (i.e., finding intelligible) seems to be the crucial factor, empirical investigation into the actual exercise of this capacity will surely be a part of any complete portrait of why certain representations are regarded as explanations and others are not. Indeed, insofar as scientists and laypeople have, with their concepts of explanation, correctly demarcated the boundaries of a real natural or sociocultural kind, then psychological research into what it means to find a happening intelligible will probably contribute much to our understanding of what explanations are, in and of themselves. Thus, for philosophers of science wishing to know what explanations are and what role they play in our lay and scholarly lives, it would seem inadvisable to turn their backs on empirical psychology and retreat to

evaluating theories based upon how well they track their own classifications of supposedly clear-cut cases of explanation and non-explanation.

References

- Craver, C. F. (2007). *Explaining the brain*. New York: Oxford University Press.
- Hempel, C. G. (1942). The function of general laws in history. *The Journal of Philosophy*, 39(2), 35-48.
- Hempel, C. G. (1965). *Aspects of scientific explanation and other essays in the philosophy of science*. New York: The Free Press.
- Hospers, J. (1946). On explanation. *The Journal of Philosophy*, 43(13), 337-356.
- Humphreys, P. (1989). Scientific explanation: The causes, some of the causes, and nothing but the causes'. *Minnesota studies in the philosophy of science*, 13, 283-306.
- Lombrozo, T., & Carey, S. 2004. Functional explanation and the function of explanation. *Cognition*, 99, 167-204.
- Machamer, P., & Woody, A. (1994). A model of intelligibility in science: Using Galileo's balance as a model for understanding the motion of bodies. *Science & Education*, 3(3), 215-244.
- Mishra, P., & Brewer, W. F. (2003). Theories as a form of mental representation and their role in the recall of text information. *Contemporary Educational Psychology*, 28, 277-303.
- Miller, D. L. (1947). Explanation versus description. *The Philosophical Review*, 56(3), 306-312.
- Salmon, W. (1984). *Scientific explanation and the causal structure of the world*. Princeton, NJ: Princeton University Press.
- Salmon, W. (1998). *Causality and explanation*. New York: Oxford University Press.
- Trout, J. (2007). The psychology of scientific explanation. *Philosophy Compass*, 2(3), 564-591.