

Wishful Intelligibility, Black Boxes, and Epidemiological Explanation

Marina DiMarco*†

Epidemiological explanation often has a “black box” character, meaning the intermediate steps between cause and effect are unknown. Filling in black boxes is thought to improve causal inferences by making them intelligible. I argue that adding information about intermediate causes to a black box explanation is an unreliable guide to pragmatic intelligibility because it may mislead us about the stability of a cause. I diagnose a problem that I call wishful intelligibility, which occurs when scientists misjudge the limitations of certain features of an explanation. Wishful intelligibility gives us a new reason to prefer black box explanations in some contexts.

1. Introduction. Epidemiological explanation often has a “black box” character, meaning the intermediate steps between some putative cause and effect of interest are unknown. The black boxes of epidemiological explanation have been variously described as mere predictive heuristics, as obstacles, or even as threats to scientific understanding. Philosophers and epidemiologists alike have argued that specifying intermediate causes to fill in black boxes improves causal explanations by making them intelligible and better targets for public health intervention (Machamer, Darden, and Craver 2000; Hiatt 2004; Russo and Williamson 2007). Specifying the links between the ends of a causal chain is supposed to confer certainty, understanding, and reasons to expect a causal relationship to be stable or invariant across populations of interest.

*To contact the author, please write to: Department of History and Philosophy of Science, University of Pittsburgh, 1101 Cathedral of Learning, 4200 Fifth Avenue, Pittsburgh, PA 15260; e-mail: marina.dimarco@pitt.edu.

†I am grateful to Jim Woodward, Anya Plutynski, Michael Dietrich, Sandra Mitchell, Jonathan Fuller, Kareem Khalifa, Kathleen Creel, Dasha Pruss, Vivian Feldblyum, the Pitt HPS Works in Progress community, and audiences at POBAMz 2020 and the PMPOS Conceptual and Methodological Aspects of Biomedical Research conference for generous and thoughtful discussions about this argument.

Philosophy of Science, 88 (December 2021) pp. 824–834. 0031-8248/2021/8805-0006\$10.00
Copyright 2021 by the Philosophy of Science Association. All rights reserved.

I argue that adding information about intermediate causes can be an unreliable guide to improving epidemiological explanation because it may mislead us about the stability of a causal relationship and may convey a false sense of understanding. I diagnose this as an instance of a more general problem that I call wishful intelligibility, which occurs when scientists misjudge the limits of the pragmatic benefit conferred by certain features of an explanation. To illustrate this, I consider an example of epidemiological explanation involving the social determinants of health. My argument offers a new reason to prefer black box explanations in some contexts, not despite, but because of, their lack of information about intermediate causes. As a result, filling in black boxes is not a necessary source of intelligibility but a contingent one.

2. Black Boxes and Intelligibility. Specifying the links between the ends of a causal chain is supposed to improve our understanding of an epidemiological cause, but attempts to account for the intelligibility produced by filling in black boxes vary considerably. For mechanists like Machamer et al. (2000, 21), filling in intermediate links of a causal chain between some cause C and effect E at the preferred level of detail for a given scientific field makes the relationship between C and E more intelligible. This sense in which filling in black boxes is supposed to confer understanding on this account is rather vague, leaving open the possibility that their “intelligibility” is akin to the phenomenological “sense of understanding” that Trout (2002) convincingly condemns as an unreliable indicator of explanatory goodness. By contrast, other accounts of epidemiological explanation tether the value of filling in black boxes to the goal, broadly speaking, of using epidemiological explanations to design effective public health interventions (Russo and Williamson 2007; Broadbent 2011; Illari 2011). Following de Regt (2017), I call this pragmatic intelligibility. I focus on the pragmatic intelligibility argument, both because it seems to evade the phenomenological critique and because it bears a more obvious relationship to the design of public health policy.

One way that filling in black boxes is supposed to confer pragmatic intelligibility is by informing our inferences about the stability of epidemiological causes, that is, our expectation that a causal relationship observed in one context will also hold in other contexts of interest. Because many epidemiological causal relationships are observed in population studies, the design of effective public health interventions must often attend to the potential efficacy of such interventions both within and beyond the original conditions in which a causal relationship is observed (Dupré 1984). Russo and Williamson (2007) famously argue that filling in black boxes supports epidemiological causal inference in just this way: that filling in a black box with evidence of a plausible mechanism tells us about the stability of a cause and supports an expectation that the causal relationship will hold in contexts of interest that differ from the one in which it was observed. This feature would, if

true, make such evidence key to the use of epidemiological causes with regard to the design of public health interventions. In section 3, I show why this is not necessarily what we ought to expect.

3. Stability and Intermediate Causes. I agree that stability is one of the most important features of epidemiological inference relative to the goal of reliable public health intervention. However, I argue that filling in black boxes can lead us to both under- and overestimate the stability of epidemiological causes. In practice, this makes filling in black boxes an unreliable guide to the pragmatic intelligibility of epidemiological explanation.

Because Russo and Williamson take the increase in stability of a cause to mean that it is expected to hold in conditions different from the ones in which the original experiment or observation took place, I take it that they have in mind something like Woodward's (2010) notion of the stability of causes. On this account, there is a causal relationship between two variables C and E just in case some intervention on the value of C produces a change in the value of E that proceeds only through the change in C (Woodward 2000). Interventionist causal relationships are stable to the extent that they hold over a more or less universal set of background conditions, making causal stability a matter of degree (Woodward 2000, 2010). Although this account is particularly amenable to epidemiological practice because a Woodward intervention need not be the result of intentional human manipulation, one does not need to be an interventionist to appreciate this notion of stability.

Filling in black boxes entails adding links in a causal chain. As Woodward (2010) points out, however, for any causal chain, the set of background conditions or domain of invariance over which the entire chain is stable is limited to the extent to which the stability of all links in the chain overlaps. This means the chain is limited not only by the domain of the least stable link but also by the extent to which that domain of invariance is shared with the other links. Because intermediate causes constrain the stability of the chain in this way, adding information about intermediate links cannot, by itself, increase our confidence in the stability of the entire chain.

3.1. Underestimating Stability. With this account of stability in mind, a focus on filling in black boxes can mislead us about the stability of a causal relationship in at least two ways. First, identifying a single causal chain from C to E may be misleading with respect to the stability of a cause that produces its effect by multiple independent pathways. As Mitchell (2002), Fehr (2004), Dupré (2013), and Howick, Glasziou, and Aronson (2013) argue, specifying a single set of intermediate steps between cause and effect restricts our assessment of the relationship between C and E to one particular causal chain, when in fact there may be several pathways from C to E. For instance, a causal variable like socioeconomic status might cause cancer by way of its effects on

stress, nutrition, access to preventive care, and so on. Multiple pathways between a single cause and effect may be stable over different background conditions. This means that filling in a black box with a single causal chain can lead us to underestimate the stability of a causal relationship by confining our expectation to a single, overly narrow domain of invariance.

3.2. Overestimating Stability. As Fehr (2004) points out, our interest in causal intermediates need not commit us to the myopia of single mechanistic explanation. However, more sophisticated efforts to fill in black boxes are still subject to a second set of concerns about stability, namely, that filling in black boxes can lead us to overestimate the stability of a causal chain when we overlook the challenges of integrating multiple indirect causes. This is because filling in a black box is often a process of what Mayo-Wilson (2014) calls “piecemeal causal inference.” The intermediate steps between a cause and effect in epidemiology are frequently inferred in independent research contexts. As Baetu (2014), Mayo-Wilson (2014), and others have pointed out, integrating causal variables that are inferred in different research contexts increases underdetermination and uncertainty about causation. This problem is particularly pernicious and complex when it comes to assessing the stability of causal chains in epidemiology.

At least three features of epidemiological practice contribute to this difficulty. First, ethical constraints make it the case that many causal inferences in epidemiology cannot be made on the basis of manipulations or interventions in human populations. This means many links in a putative causal chain are thought to be stable with respect to humans on the basis of extrapolation from animal models or retrospective analyses of so-called natural experiments. Extrapolation and inferences of external validity are inherently epistemically risky business (cf. Reiss 2019). Second, causal variables within the same chain are often measured and described with very different degrees of precision and at different spatial and temporal scales in different research contexts. Finally, scientists in different research contexts often measure causal variables with respect to different background conditions of interest. For instance, social epidemiologists (e.g., Krieger 2008) are especially concerned to include possible social determinants of health, like socioeconomic status, as variables in their analyses, but other researchers interested in intermediate causes of the same effects, like epigeneticists, may not measure the socioeconomic status of their subjects at all. Even similar variables described and measured at similar scales are not consistently accounted for across studies in the same field. For example, “neighborhood” is variously measured by zip code, census tract, or county (Shavers 2007). When researchers want to integrate causal inferences to fill in a black box between, for instance, neighborhood and cancer mortality, these differences limit the extent to which piecemeal causal inference can tell us a causal relationship is stable with regard to the

values of these differently described variables. These failures to consistently define and comeasure background condition variables are a problem because they mean that inferences about whether links are stable with regard to the same background conditions may be much less certain than they appear.

When a black box causal relationship is filled in using intermediate causal inferences assembled from diverse research contexts, it presents a unique challenge for assessing the stability of the original causal relationship of interest. Because links in the same causal chain are often inconsistently described, measured at different scales, and demonstrated with respect to different background conditions, it is often impossible (or at least intractable) to assess the extent to which multiple links in a causal chain share a domain of invariance at all. Integrative inferences about the stability of the entire chain become much more complex. At a minimum, these factors make it difficult to identify a lowest common denominator, or least stable link in a causal chain. Failure to attend to these features of piecemeal causal inference can lead us to overestimate the stability of a causal relationship or to make an inference about its stability that is not justified by the available evidence.

3.3. Wishful Intelligibility. Since filling in black boxes is at best an unreliable guide to stability, and stability is critical to the goal of designing epidemiological interventions, it follows that we should not expect filling in black boxes to confer pragmatic benefit to epidemiological causes by way of improving our inferences about stability in all contexts. Instead, this assumption may lead us to be inappropriately confident in our understanding of a cause and in our estimation of stability in particular. We should not expect filling in black boxes to be conducive to the goals of epidemiological inference in cases in which such goals depend on information about the stability of a cause; this is a contingent, rather than a necessary, source of pragmatic intelligibility.

To mischaracterize the intelligibility of a causal claim is to misjudge the extent to which we understand it. This may (mis)inform the design of interventions, with serious consequences for public health policy. Unjustified attributions of pragmatic intelligibility constitute a special kind of second-order misunderstanding, namely, a failure to correctly assess the extent of the pragmatic benefit of some particular feature of an explanation (Steel 2016; cf. Trout 2002). I call this wishful intelligibility. An unqualified preference for filling in black boxes in pursuit of stability can lead us to wishful intelligibility (for a similar argument, see Broadbent 2011).

Wishful intelligibility has an obvious affinity with the more general problem of wishful thinking in the literature on science and values. Broadly speaking, wishful thinking may occur when certain values or cognitive biases lead us to form an otherwise unjustified or ill-justified belief; these biases may include but are not limited to a desire for the belief to be true (see Anderson

2004; Steel 2018). By contrast, wishful intelligibility concerns not whether a belief or claim is justified in general but, rather, whether we have good reason to expect that some feature of an explanation is conducive to its use in a specific context. That is, it concerns a particular set of beliefs: those about the pragmatic intelligibility conferred by certain features of an explanation.

4. Multilevel Causes of Cancer. In section 3, I argued that filling in black boxes with intermediate causes can be misleading with respect to the stability of a causal relationship and that filling in black boxes can be conducive to wishful intelligibility with regard to epidemiological explanation. Gehlert and colleagues' (2008) multilevel model of the social environment as a cause of cancer, from the University of Chicago's Center for Interdisciplinary Health Disparities Research (CIHDR), shows the difficulty of estimating the stability of a piecemeal causal inference. Their work is particularly interesting because it purports to have implications for the design of future public health interventions.

4.1. The Social Environment and Breast Cancer. This multilevel model fills in links in a (putative) causal chain by assembling multiple local causal inferences from separate studies. A CIHDR study evaluates the social environment of Black women newly diagnosed with breast cancer in "predominantly black neighborhoods of Chicago," using interviews and "publicly available data geocoded to the women's addresses" (Gehlert et al. 2008, 343). Gehlert et al. argue that features of the built environment, such as dilapidated housing and crime, contribute to social isolation (344). One example cited, Sampson, Raudenbush, and Earls (1997), is an investigation of the effects of "collective efficacy" (defined as neighborhood social cohesion) on violence in Chicago in 1995.

The next steps in the chain are the sequential links between social isolation, psychological states, and stress hormone responses. Here Gehlert et al. appeal to a combination of human and rodent studies to argue that social isolation affects HPA axis regulation and glucocorticoid (stress hormone) signaling via epigenetics. Since glucocorticoid levels have been linked to suppressed immune function elsewhere in the literature, the authors conclude that stress hormone regulation constitutes a means by which social isolation can "get under the skin" to cause cancer and promote tumor cell survival (Gehlert et al. 2008, 343).

This model is a clear case in which piecemeal assembly of links between "levels" does not (by itself) support any inference about the stability of the whole chain. Importantly, each step is inferred with respect to different background variables, and few, if any, of the same background conditions are measured for any two links in the chain. Some links, like the association between social isolation and mammary gland tumors, are manipulated in a laboratory environment, using model organisms, while others are observed

in humans. Qualitatively similar links in the chain are differently described: Sampson and colleagues' (1997) study of collective efficacy is performed in the same Chicago neighborhoods as the CIHDR study of newly diagnosed cancer patients but using a notion of "neighborhood clusters" to combine 847 census tracts into 343 clusters, from more than a decade earlier. Setting aside the fact that these neighborhoods have probably changed over 10 years, the extent to which these neighborhood clusters overlap spatially with the geocoded data used to measure neighborhoods in the CIHDR project is unclear. Some links, such as epigenetic regulation of altered HPA axis-related gene expression, are apparently assumed to be stable across human populations over which they may not have been measured at all. Gehlert and colleagues borrow links from several research contexts to fill in the black box between neighborhood and tumorigenesis.

Filling in the black box between the social environment and cancer incidence does not justify an expectation that the relationship between them will be stable. It does not even tell us anything reliable about how stable we might expect it to be. Instead, it may lead us to attribute wishful intelligibility to a social epidemiological cause where it is lacking and to misrepresent the limits of (and risks entailed by) this explanation as a potential basis for public health policy.

4.2. Integration, Stability, and Translation. One might object that the issue in this and similar cases is that we really ought to have something stronger (or at least better confirmed) in mind with regard to filling in black boxes (see Illari 2011). However, attention to the specific features of integration that might justify estimations of stability in such cases is notoriously absent from the Russo-Williamson account and from other arguments for filling in black boxes (Illari 2011; Plutynski 2018). We can have good evidence for each link in a causal chain and yet have relatively little evidence that these links are stable over some shared set of relevant background conditions (see Broadbent 2011). It is much more difficult to integrate evidence from diverse contexts to show that each link in a chain is stable with regard to the same conditions of interest than it is to justify inferences about individual links in a chain. Thus, assessing the stability of a causal chain can come apart from (and be more burdensome than) merely providing evidence of a causal chain or providing evidence of individual links.

Importantly, this does not mean that we ought to be pessimistic about multilevel causal models or about the social determinants of health more broadly. Rather, it presents an opportunity to negotiate what does make for better and worse attributions of stability. At a minimum, we might expect that coordination among researchers to improve standardization and comeasurement of causal variables and background conditions, together with explicit evaluation and justification of background assumptions across a causal chain, would improve

integration and inferences about the stability of a causal relationship. Many translational epidemiologists have recently turned their attention to these and other related features of knowledge integration in epidemiology, with the goal of improving and expediting the translation of biomedical research into public health interventions (e.g., Ioannidis et al. 2013).

Following O'Malley and Stotz (2011) and others, I take it that the details of successful knowledge integration in epidemiology will be largely pragmatic and contextually specific and, most importantly, that they will admit of degrees. My main concern is that the wishful intelligibility of filling in black boxes leads us to overlook these considerations entirely. In section 5, I argue that black boxes may prevent this sort of oversight in some contexts by preserving epistemic humility about epidemiological causes.

5. Preferring Black Boxes. So far, I have focused on the extent to which filling in black boxes may confer pragmatic intelligibility, by way of stability. Now, I will consider the relationship between intelligibility and black boxes themselves.

Previous endorsements of black boxed explanations in epidemiology have focused on the utility of black boxes in designing interventions despite ignorance of intermediate links in a causal chain (Cranor 2017; Plutynski 2018). These accounts often appeal to some version of the argument from inductive risk, on which, broadly speaking, the ethical consequences of error play a normative role in determining the amount and kind of evidence necessary to justify accepting or rejecting a hypothesis and, by extension, the decision to intervene in a particular way (Douglas 2000; Steel 2016). They purport to justify a preference for (or more accurately, a tolerance of) black boxes when the costs of agnosticism or inaction outweigh the possible negative consequences of ignorance about the intermediate steps between some putative C and E or when the costs of providing evidence for a plausible set of intermediate steps are too high, given the projected benefit of additional detail.

While I am sympathetic to these accounts, I am concerned to show that, by the same token of inductive risk, there are circumstances in which we ought to prefer black boxed explanations not despite but because of their lack of information about causal intermediates. This is because black boxes can prevent wishful intelligibility, especially where stability is concerned. Imagine a case in which misjudging the stability of a cause can be expected to have serious consequences for the success of some possible intervention. Since filling in a black box can actively mislead us with respect to the stability of a cause, we may prefer a black box for purposes of designing such an intervention. Wishful intelligibility often has a price.

The above discussion of stability can help to predict the contexts in which we might be particularly susceptible to these mistakes. In cases in which we have reason to expect a complex or nonspecific, multicausal structure, we

might be misled by estimating stability from a single causal chain. Further, I have argued that piecemeal causal inference (and especially multilevel piecemeal causal inference) may lead us to overestimate the stability of a causal chain, especially when links in the chain are poorly integrated. In both cases, black boxed explanations capture and convey an appropriate sense of uncertainty about the stability of some cause. In such cases, agnosticism about the stability of a cause may be more conducive to the design of effective interventions than wishful intelligibility would be. In this sense, my argument is concordant with Trout's (2002, 212) condemnation of the risks of "counterfeit understanding." Admittedly, the value of black boxes in these cases is indexed to the costs of being wrong about stability. However, given that stability is lauded as a critical determinant of which epidemiological causes make for good interventions, I take it that these concerns are relevant to a nontrivial number of cases.

Black boxes may thus preserve a certain humility, akin to what Pickersgill (2016) calls "epistemic modesty" about the stability of a cause that is conducive to the design of good interventions and the avoidance of bad ones. This means that black boxes are not merely to be preferred despite the risk of unknown intermediates. Rather, a black box is preferable to evidence of a causal chain in cases in which the consequences of error about stability are sufficiently undesirable. When we have thorough and well-integrated knowledge of a causal structure, my concerns about under- and overestimating stability are less troubling. Given that the appropriateness of black boxes (and of filling in) is contextual and specific, these considerations put black boxes on par with evidence of intermediate linking causes as features of explanations that may contribute to their pragmatic intelligibility in a contextual and contingent manner.

Of course, this does not mean that black boxes are to be preferred to information about intermediate causes in all or even most epidemiological explanations. The problem lies not with information about intermediate causes but with the inferences we are tempted to make about stability on the basis of this information. Furthermore, I have by no means exhausted the arguments for filling in black boxes in epidemiological explanation, and I expect we may have good reasons to prefer filling-in that outweigh these concerns about stability in some contexts. Rather, my argument shows that black boxes deserve the same status of contextual relevance to pragmatic intelligibility as do other features of epidemiological explanations, not despite but because of the absence of information about intermediate causes. This means that their inclusion (or filling-in) should be a matter of transparent negotiation and justification, rather than a default preference one way or the other.

6. Conclusion. Wishful intelligibility is a helpful diagnosis for the mistaken expectation that filling in black boxes is always a good guide to causal

stability. Filling in black boxes is by no means the only possible source of wishful intelligibility in epidemiology or elsewhere, nor does it always have this effect. Rather, I have been concerned to argue about specific epidemiological cases precisely because the features of an explanation that make it conducive to some particular goal are contextual and specific. Similarly, I have shown that there is a positive role for black boxes in preventing wishful intelligibility and in preserving epistemic humility about the stability of complex causes in some cases of interest to epidemiologists.

These cases seem to have something in common, namely, incomplete, limited, or poorly integrated knowledge of a complex causal structure. We might worry that my account could mistakenly preclude a positive and epistemically responsible role(s) for various departures from the whole truth (à la Elgin 2007) or that a preference for black boxes in such contexts perpetuates a harmful role for “ideal science” that may cripple important regulation and delay helpful interventions (Cranor 2017). But such a preference does not depend on a false dichotomy between ideal science and wishful intelligibility; after all, black boxes are themselves departures from the whole truth. Incomplete or poorly integrated information about a complex causal structure need not paralyze the design of public health interventions; it merely makes our assessments of stability a matter of second-order inductive risk. My account invites transparency and justification for such decisions on a case-by-case basis and recommends specific ways in which the state of current epidemiological knowledge should inform these considerations. To the extent that these measures avoid wishful intelligibility, they make for more trustworthy reasons to intervene in a particular way, not less.

REFERENCES

- Anderson, Elizabeth. 2004. “Uses of Value Judgments in Science: A General Argument, with Lessons from a Case Study of Feminist Research on Divorce.” *Hypatia* 19 (1): 1–24.
- Baetu, Tudor M. 2014. “Models and the Mosaic of Scientific Knowledge: The Case of Immunology.” *Studies in History and Philosophy of Science C* 45 (March): 49–56.
- Broadbent, Alex. 2011. “Inferring Causation in Epidemiology: Mechanisms, Black Boxes, and Contrasts.” In *Causation in the Sciences*, ed. Phyllis McKay Illari, Federica Russo, and Jon Williamson, 45–69. Oxford: Oxford University Press.
- Cranor, Carl F. 2017. “How Demands for Ideal Science Undermine the Public’s Health.” In *Tragic Failures: How and Why We Are Harmed by Toxic Chemicals*. Oxford: Oxford University Press.
- de Regt, Henk. 2017. *Understanding Scientific Understanding*. Oxford: Oxford University Press.
- Douglas, Heather. 2000. “Inductive Risk and Values in Science.” *Philosophy of Science* 67 (4): 559–79.
- Dupré, John. 1984. “Probabilistic Causality Emancipated.” *Midwest Studies in Philosophy* 9:169–75.
- . 2013. “Living Causes.” *Aristotelian Society Supplementary Volume* 87 (1): 19–37.
- Elgin, Catherine. 2007. *True Enough*. Cambridge, MA: MIT Press.
- Fehr, Carla. 2004. “Feminism and Science: Mechanism without Reductionism.” *NWSA Journal* 16 (1): 136–56.
- Gehlert, Sarah, Dana Sohmer, Tina Sacks, Charles Mininger, Martha McClintock, and Olufunmilayo Olopade. 2008. “Targeting Health Disparities: A Model Linking Upstream Determinants to Downstream Interventions.” *Health Affairs* 27 (2): 339–49.

- Hiatt, Robert A. 2004. "The Social Determinants of Cancer." *European Journal of Epidemiology* 19 (9): 821–22.
- Howick, Jeremy, Paul Glasziou, and Jeffrey K. Aronson. 2013. "Problems with Using Mechanisms to Solve the Problem of Extrapolation." *Theoretical Medicine and Bioethics* 34 (4): 275–91.
- Illari, Phyllis McKay. 2011. "Mechanistic Evidence: Disambiguating the Russo-Williamson Thesis." *International Studies in the Philosophy of Science* 25 (2): 139–57.
- Ioannidis, J. P. A., S. D. Schully, T. K. Lam, and M. J. Khoury. 2013. "Knowledge Integration in Cancer: Current Landscape and Future Prospects." *Cancer Epidemiology Biomarkers and Prevention* 22 (1): 3–10.
- Krieger, Nancy. 2008. "Proximal, Distal, and the Politics of Causation: What's Level Got to Do with It?" *American Journal of Public Health* 98 (2): 221–30.
- Machamer, Peter, Lindley Darden, and Carl F. Craver. 2000. "Thinking about Mechanisms." *Philosophy of Science* 67 (1): 1–25.
- Mayo-Wilson, Conor. 2014. "The Limits of Piecemeal Causal Inference." *British Journal for the Philosophy of Science* 65 (2): 213–49.
- Mitchell, Sandra D. 2002. "Integrative Pluralism." *Biology and Philosophy* 17:55–70.
- O'Malley, Maureen A., and Karola Stotz. 2011. "Intervention, Integration and Translation in Obesity Research: Genetic, Developmental and Metaorganismal Approaches." *Philosophy, Ethics, and Humanities in Medicine* 6 (1): 2.
- Pickersgill, Martyn. 2016. "Epistemic Modesty, Ostentatiousness and the Uncertainties of Epigenetics: On the Knowledge Machinery of (Social) Science." *Sociological Review* 64 (Suppl.): 186–202.
- Plutynski, Anya. 2018. *Explaining Cancer*. Oxford: Oxford University Press.
- Reiss, Julian. 2019. "Against External Validity." *Synthese* 196:3103–21.
- Russo, Federica, and Jon Williamson. 2007. "Interpreting Causality in the Health Sciences." *International Studies in the Philosophy of Science* 21 (2): 157–70.
- Sampson, R. J., Stephen W. Raudenbush, and Felton Earls. 1997. "Neighborhoods and Violent Crime: A Multilevel Study of Collective Efficacy." *Science* 277 (5328): 918–24.
- Shavers, Vickie L. 2007. "Measurement of Socioeconomic Status in Health Disparities Research." *Journal of the National Medical Association* 99 (9): 1013–23.
- Steel, Daniel. 2016. "Climate Change and Second-Order Uncertainty: Defending a Generalized, Normative, and Structural Argument from Inductive Risk." *Perspectives on Science* 24 (6): 696–721.
- . 2018. "Wishful Thinking and Values in Science." *Philosophy of Science* 85 (5): 895–905.
- Trout, J. D. 2002. "Scientific Explanation and the Sense of Understanding." *Philosophy of Science* 69 (2): 212–33.
- Woodward, James. 2000. "Explanation and Invariance in the Special Sciences." *British Journal for the Philosophy of Science* 51 (2): 197–254.
- . 2010. "Causation in Biology: Stability, Specificity, and the Choice of Levels of Explanation." *Biology and Philosophy* 25:287–318.