

ford Handbook of Philosophy of Social Science, edited Harold Kincaid. Oxford: Oxford University Press: 21–45.
Ylikoski, Petri. 2013. “Causal and Constitutive Explanation Compared.” *Erkenntnis* vol. 78, no. 2: 277–297.

Do New Accounts of Causal Mechanisms Offer Practical Advice for Process Tracing?

<https://doi.org/10.5281/zenodo.823311>

Andrew Bennett
Georgetown University

When Alexander George and I wrote our 2005 book, *Case Studies and Theory Development in the Social Sciences*, a central theme was that theories about causal mechanisms, and process tracing analysis of the observable implications of hypothesized causal mechanisms, were central to qualitative case study research. Our emphasis on causal mechanisms was not especially novel at the time, but it ran counter to the preeminent role that Gary King, Robert Keohane, and Sidney Verba devoted to causal effects, and in the intervening years causal mechanisms have gained added attention from social scientists.¹

At the time we wrote our book, we based our definition of causal mechanisms largely on that of Wesley Salmon, the philosopher of science whose concept resonated most closely with our own understanding of mechanisms. Being political scientists rather than philosophers of science, we were not aware, in those days before Google Scholar, that James Woodward, building on his earlier work and on work by Judea Pearl, had critiqued Salmon’s definition and pushed forward the philosophical discussion of causal mechanisms.²

The present symposium on causal mechanisms is thus an opportune time to examine whether the latest philosophical discussions of causal mechanisms offer practical advice for political scientists engaged in process tracing. In particular, two suggestions for improved process tracing practices deserve attention: 1) the construction of directed acyclic graphs³ and the use of such graphs to improve upon and assess the completeness of explanations based on process tracing;⁴ and, 2) the use of counterfactual analysis in process tracing.⁵

The present article first notes the differences between the understanding of causal mechanisms that emerges from Wood-

ward’s manipulationist account of causation and the definition George and I used in 2005. It then assesses whether and how David Waldner’s account of the potential completeness of explanations, building on Pearl’s work and also related to Woodward’s conception of “invariance,” and Rosa Runhardt’s approach to using counterfactual analysis, building on Woodward’s concepts, offer practical advice for process tracing in political science case studies.⁶ I conclude that Waldner’s account demands more than process tracing can deliver, and Runhardt’s suggestion needs additional development to systematize the kinds of counterfactual analysis that process tracers already carry out. Yet these are promising lines of development that deserve continued research.

Defining Causal Mechanisms: Wesley Salmon versus James Woodward

In the 1980s and later, Wesley Salmon developed what is known as the *causal mechanical* approach to causal explanation.⁷ In this account, a causal process is a physical process that transmits what Salmon termed a *mark* between one physical entity and another. A mark is a change in physical structure that persists beyond the interaction of the first and second entity in the absence of any further causal interactions. Building on Salmon’s discussion of causal mechanisms, Alexander George and I defined causal mechanisms as “ultimately unobservable physical, social, or psychological processes through which agents with causal capacities operate, but only in specific contexts or conditions, to transfer energy, information, or matter to other entities, [thereby changing] the affected entity’s characteristics, capacities, or propensities in ways that persist until subsequent causal mechanisms act upon it.”⁸ We noted the obvious inferential challenges of measuring the effects of particular mechanisms, isolating them from other mechanisms, and assessing the conditions under which particular mechanisms are activated, and we discussed the ways in which process tracing could address (but not in any ultimate or decisive sense resolve) these inferential challenges.

Given the lags involved in researching, writing, and publishing our book, and given that George and I were autodidacts in the philosophy of science, we were not aware that James Woodward⁹ had already convincingly critiqued Salmon’s definition. Woodward notes that Salmon’s account of mark transmission encounters problems in explaining, for example, the motion of billiard balls after their collision, which relies on knowing the masses and velocities of the balls before the collision, the assumption of a perfectly elastic collision, and the hypothesized mechanism of the conservation of linear momentum. As Woodward points out, Salmon’s account of mark transmission does not distinguish between the explanatorily relevant momentum of the balls and explanatorily irrelevant features like the transmission of a chalk mark from one ball to

Andrew Bennett is Professor of Government at Georgetown University. He can be reached at BennettA@Georgetown.edu.

¹ King, Keohane and Verba 1994.

² Woodward 2003; Pearl 2000.

³ Pearl 2000; Knight and Winship 2013.

⁴ Waldner 2012, 2015, and 2016.

⁵ Runhardt 2015. For present purposes I do not address recent contributions to more formal Bayesian process tracing (Bennett 2015, Fairfield and Charman forthcoming), which I believe are beginning to offer very useful guidelines for process tracing even if further research is needed on how much and what kind of Bayesian formalization is practically useful in carrying out and writing up process tracing research.

⁶ Runhardt 2016; Waldner 2016.

⁷ Salmon 1989.

⁸ George and Bennett 2005, 137; see also Bennett and Checkel 2015, 12.

⁹ Woodward 2003.

another.¹⁰ Capturing this distinction, Woodward argues, requires an appeal to a counterfactual like that invoked by his own manipulationist account (that is, would the post-collision course change if there were no chalk mark, or would it change if one ball's pre-collision momentum was different).¹¹

Woodward's manipulationist account, in its simplest form, is that "if it is possible to manipulate Y [or, he adds *seriatim*, the probability distribution of Y] by intervening on X, then we may conclude that X causes Y, regardless of whether the relationship between X and Y lacks various other features standardly regarded as necessary for causation: even if X is not connected to Y via a spatiotemporally continuous process, even if there is no obvious sense in which there is a transfer of energy from X to Y, and so on."¹² This avoids the philosophical problems with Salmon's approach, and it is appealing to social scientists as we hope to find variables that policy makers can reliably manipulate to create a better social world.¹³ Of course, it is a definition of causation and explana-

tion rather than a guide to causal inference: Woodward notes that "my [2003] project is semantic or interpretive, and is not intended as a contribution to practical problems of causal inference"¹⁴ or "inference to the best explanation."¹⁵

Woodward's account relies on three key features: 1) an intervention on a causal variable X (whether the intervention is in fact possible or not),¹⁶ 2) a counterfactual assertion that Y or its probability distribution would have been different if the intervention on X were carried out, and 3) a type of invariance relationship between X and Y that justifies the counterfactual assertion. In a lengthy chapter,¹⁷ Woodward carefully defines the kind of invariance relations he has in mind. For present purposes, I use his general definition that "a generalization is invariant or stable across certain changes if it holds up to some appropriate level of approximation across those changes."¹⁸ Variance, in this view, "comes in gradations or degrees";¹⁹ more invariant relationships hold across or are robust to many changes in background conditions, and less invariant relationships break down under such changes. Invariance relations, "at least within a certain range, are immutable or unconditional in the sense that there is nothing we or nature can do to turn them off."²⁰ In particular, it is critical that a generalization, if it is to be explanatory, should "continue to hold under an intervention that changes its independent variables sufficiently (or in such a way) that the value of its dependent variable is predicted by the generalization to change under the intervention."²¹ This is consistent with Waldner's view herein, and as Waldner notes, this conception of invariance does not imply (or reject) determinism, as it allows probabilistic relationships as long as they are stable in some conditions.

Waldner argues that the kind of *invariance relations* that Woodward discusses represent the ontological structures that Waldner defines as *mechanisms*. In this view, mechanisms are structures in the world that, unlike variables, we cannot turn on and off by interventions. I take this view to be superior to my own 2005 definition²² as it usefully builds on Woodward's critique of Salmon as well as Pearl's work. Waldner further argues that the mechanistic account of explanation is prior to other approaches because it implies the manipulationist, regularity, and counterfactual accounts of causal explanation, while the reverse is not true. I find Waldner's argument on this intriguing and appealing, but I do not pursue further here the question of the priority of mechanisms among different approaches to causal explanation. Rather, with regard to my own work, Waldner's argument raises two questions. The first is whether the process tracing practices Alexander George, Jeff-

than stopping smoking (such as changing the characteristics of cigarettes) can have beneficial effects.

¹⁴ Woodward 2003, 38.

¹⁵ Woodward 2003, 374.

¹⁶ Woodward 2003, 250.

¹⁷ Woodward 2003, 239–314.

¹⁸ Woodward 2003, 239.

¹⁹ Woodward 2003, 240.

²⁰ Woodward 2003, 242.

²¹ Woodward 2003, 250.

²² George and Bennett 2005, 137; or see page 2 of this essay.

¹⁰ Woodward 2003, 351–352.

¹¹ Woodward (2003, 353) adds a second critique, which is that causal mechanical accounts cannot address "action at a distance" processes that do not allow tracing of spatiotemporal processes (like gravity); a possible rejoinder is that the mechanisms behind gravity and other action at a distance phenomena are not sufficiently understood to allow tracing of their processes (and perhaps never will be sufficiently understood for this—the related physics theories collectively known as "string theory" do not as yet have any testable observable implications). Woodward notes further that "most explanations in disciplines like psychology and economics" do not involve transfer of energy or momentum, but Alexander George's and my definition is not as vulnerable to this critique as Salmon's, as we added to Salmon's account the transfer of information between entities as a kind of mechanism. Woodward's final critique is that the causal mechanical model faces difficulties when interactions are too complex or numerous to trace, such as the interactions of the numerous molecules that make up a mole of gas. A possible response to this critique is that the mechanistic approach posits a model of how molecules interact and tests the model's observable implications at the aggregate level rather than the level of individual molecules; analogously, George and I noted that it is often more efficient to do process tracing at the structural level rather than that of individual agents, but we added that a commitment to mechanistic explanation entails that if it could be shown that individual or underlying units do not behave as a structural theory posits, the structural theory is weakened even if it remains a simplification that is useful for some purposes (Woodward 2003, 142–145).

¹² Woodward 2003, 49.

¹³ To be clear, to the extent that they are policy-oriented, social scientists seek to understand mechanisms so that they can better predict the consequences of alternative manipulations or policies. Some very robust correlational findings, such as the relationship between smoking and cancer, may be sufficiently strong evidence of an average causal effect to properly become the basis for policy interventions even if the mechanisms underlying the correlations are not understood. At the same time, a better understanding of underlying mechanisms can improve our understanding of the conditions under which, and the cases for which, a generalization will hold. Improvements in our understanding of the mechanisms that relate smoking to cancer, for example, may help identify which individuals are at especially high (or low) risk if they smoke, and which interventions other

rey Checkel, and I have advocated are compatible with the concept of causation and causal mechanisms that Woodward, Waldner, and Runhardt advance.²³ I believe they are, and as I have not read anything by these authors that indicates that they disagree, I do not pursue this question further here. The second question, which I address below, is whether there are additional or different practices of process tracing that emerge from the concept of invariant mechanisms that improve upon the process tracing practices that I and other methodologists have advocated. Here, I think there are real improvements, but they are not yet as substantial or unambiguous as Waldner and Runhardt have suggested.

David Waldner: Directed Acyclic Graphs (DAGs) and Process Tracing

On the issue of best practices in process tracing, David Waldner made an important contribution in the book I co-edited with Jeffrey Checkel by calling attention to the “completeness standard” of assessing process tracing. There, he argued that (*italics are the author’s*):²⁴

Process tracing yields causal and explanatory adequacy insofar as: (1) it is based on a causal graph whose individual nodes are connected in such a way that they are jointly sufficient for the outcome; (2) it is also based on an *event history map* that establishes valid correspondence between the events in each particular case study and the nodes in the causal graph; (3) theoretical statements about causal mechanisms link the nodes in the causal graph to their descendants and the empirics of the case studies allows us to infer that the events were in actuality generated by the relevant mechanisms; and (4) rival explanations have been credibly eliminated, by direct hypothesis testing or by demonstrating that they cannot satisfy the first three criteria listed above.

I will not explicate here the procedures of outlining causal arguments as directed acyclic graphs (DAGs),²⁵ although I find graphical representations of arguments to be an extremely useful operational procedure for clarifying my own thinking and understanding other authors’ explanatory claims. I further agree that Waldner’s four-step procedure is extremely useful as a practice in process tracing, and that completeness is an important criterion for judging process tracing. Waldner used the completeness standard in his book chapter to critique very

effectively several prominent examples of process tracing in comparative politics for having failed to assess evidence on major and obvious steps in their hypothesized processes (in contrast, Waldner judges Elisabeth Wood’s work in *Forging Democracy from Below*²⁶ to have achieved a high level of completeness). I concur that more complete process tracing accounts are, other things equal, preferable to less complete ones.²⁷

Yet I am skeptical that Waldner’s completeness standard provides much clarity, as he claims, on a “much-needed *stopping rule*” regarding when to stop collecting evidence.²⁸ How do we know, from the passage quoted above, all of the following: 1) that the nodes in the graph are “jointly sufficient,” 2) that the correspondence between the event history map and the events in the case is “valid,” 3) that the events “were in actuality generated by the relevant mechanisms,” and 4) that rival explanations have been “credibly” eliminated? Waldner’s four-step procedure and Bayesian process tracing can provide answers to these questions, but those answers are inherently provisional to some degree because they depend in part on unknown unknowns, particularly rival explanations and omitted variables we haven’t thought of and evidence we didn’t think to look for or were not able to access. So true completeness is not possible. In addition, Waldner’s second step appears to be in some tension with Bayesianism: it is not clear that completeness is necessary. What matters is the relative likelihoods of the evidence under the alternative hypotheses, and priors on the latter, and these may lead to very high or low degrees of belief even when evidence and event maps are incomplete. Waldner’s four steps thus do not solve the fundamental problem of causal inference. Of course, *no* approach to causal inference can overcome the problem of unknown unknowns, which is why the problem of causal inference is fundamental.

It is worth noting here that in his discussion of DAGs and the equations associated with them, Woodward includes an “error term” to represent the “combined influence of all the other unknown direct causes of Y that are not explicitly represented in the equation.”²⁹ Similarly, Pearl includes an error term in structural equations to “represent errors (or ‘disturbances’) due to omitted factors.”³⁰ Woodward also notes the importance of specifying the changes under which a relation is and is not invariant, or what political scientists commonly call scope conditions.³¹ Yet as Woodward does not address causal inference, he does not offer any advice on how to assess scope conditions.

Consider Waldner’s example of combustion. He notes that his illustrative account of combustion is incomplete without knowing more about inter- and intra-molecular processes and

²³ Interestingly, Judea Pearl’s account and his explication of Directed Acyclic Graphs (DAGs) begins with “special emphasis on Bayesian inference and its connection to the psychology of human reasoning under uncertainty” and adheres to “the Bayesian interpretation of probability, according to which probabilities encode degrees of belief about events in the world and data are used to strengthen, update, or weaken those degrees of belief” (Pearl 2000, 2). I take this to suggest that Pearl’s account is compatible with my own view on the value of Bayesian inference in process tracing, although Pearl does not give much attention to singular causation or inferences from individual cases and does not mention process tracing.

²⁴ Waldner 2015, 128.

²⁵ See Pearl 2000.

²⁶ Wood 2000.

²⁷ At the same time, there is a tradeoff between completeness and parsimony in explanatory accounts.

²⁸ Waldner 2015, 129.

²⁹ Woodward 2003, 43.

³⁰ Pearl 2000, 27.

³¹ Woodward 2003, 243.

sub-atomic particles. But it is incomplete in another way as well: it says nothing about the kind and concentration of the fuel, the concentration of the oxygen, the different flash point temperatures of different materials.³² This is not a terrible shortcoming for the example of fire, since there is no alternative explanation for combustion, but in more complex cases involving social behavior it is harder to tell the scope conditions of hypothesized processes and to assert with confidence that the levels of the variables in a theory were sufficient to generate the outcome observed in a case and that alternative variables or theories do not account for the outcome.

Waldner's example of combustion is also incomplete in that it does not address implicit assumptions about background conditions: in this simple example, perhaps the only consequential background assumption is the presence of gravity. While the absence of gravity does not prevent combustion when fuel, oxygen, and heat are present, combustion in zero gravity produces a very different kind of "fire" (actually it is cooler than fire on earth, both literally and figuratively: see youtube video on combustion in zero gravity³³). In social science, we often have many implicit background assumptions and many alternative explanations that complicate our task of inferring causation and specifying the scope conditions of hypothesized mechanisms.

Thus, I stand by Jeffrey Checkel's and my critique of Waldner's completeness standard, in the closing chapter of our book, namely, that "it is very ambitious to expect a theory or explanation to be fully complete, as there will always be steps in an explanation that involve variables exogenous to a theory, steps for which strong empirical evidence is not available, and steps that are at a more micro level than a researcher chooses to explore. Thus, not every step in a theoretical explanation of a process will fully determine the next step in it."³⁴

I also think Waldner's reading of Pearl³⁵ overlooks one of the most important claims that can be made for methods that focus on developing mechanistic explanations of individual cases. One of the most common and misplaced critiques of case studies is that their results do not "generalize." My standard response is that we do not have a good understanding of whether the findings of a case study will or will not generalize, or the scope conditions to which it will generalize, until we are satisfied that we understand the mechanisms that account for the outcome in the case study. This understanding often emerges from a combination of inductive and deductive process tracing on the case. Pearl makes a related point in arguing for the superiority of causal over associational knowledge:³⁶

causal models (assuming they are valid) are much more informative than probability models. A joint distribution

tells us how probable events are and how probabilities would change with subsequent observations, but a causal model also tells us how these probabilities would change as a result of external interventions—such as those encountered in policy analysis, treatment management or planning everyday activity. Such changes cannot be deduced from a joint distribution, even if fully specified.

He later adds that "true understanding enables predictions in ...novel situations."³⁷ Thus, by improving our understanding of how mechanisms operate, process tracing can help us anticipate how changes in the institutional or political context upon which past correlations were based might lead to different distributions in the future, as well as helping us anticipate or narrow the range of possible outcomes in future cases.

Waldner has done an important service in providing an account that shows how in principle we can work to explain individual cases. I share his view—and Woodward's³⁸ and Pearl's³⁹—that a philosophical account of causal explanation should encompass both type causal claims (or claims that X is causally related to Y where X and Y are variables) and token causal claims (claims that a particular X caused a particular Y in a particular case, sometimes called singular causal claims). I also agree that Woodward's and Pearl's works make important strides in this direction (even though Pearl focuses mostly on experimental and correlational methods rather than the token or singular causation of individual cases). In practice, however, explanations of cases and populations will always be potentially fallible.

Rosa Runhardt: Using Counterfactuals in Process Tracing

Rosa Runhardt, building explicitly on Woodward's work, similarly addresses the question of whether Woodward's approach offers potential improvements in process tracing.⁴⁰ Rather than focusing on the type of invariance relations envisioned by Woodward, as Waldner does, Runhardt focuses on the counterfactuals embedded in Woodward's conception. This is an interesting and promising approach, and others have also begun to explore whether and how process tracing relates to or can benefit from a potential outcomes framework.⁴¹ Yet more work needs to be done to clarify the uses of counterfactuals in process tracing, the kinds of counterfactuals that are useful, and the ways in which we can use counterfactuals to identify and address cognitive biases in our own thinking.

Runhardt critiques as an example Kristin Bakke's excellent use of process tracing to assess whether international jihadists caused the shift among Chechen rebels toward a globalized vision of Islamic revival, and toward the use of suicide tactics, through the provision of resources, training, and propaganda material framing Chechen aspirations for independence as part of a global struggle between Islam and other religions. While Bakke provides evidence on the key steps in her hypothesized

³² This information might be encoded into structural equations representing a causal graph; I am merely pointing out that a substantial amount of knowledge on scope conditions has to go into these equations even in the relatively simple case of combustion.

³³ DNews 2013.

³⁴ Bennett and Checkel 2015, 265.

³⁵ Pearl 2000.

³⁶ Pearl 2000, 22.

³⁷ Pearl 2000, 26.

³⁸ Woodward 2003, 74–79 and 209–220.

³⁹ Pearl 2000, 310–311.

⁴⁰ Runhardt 2015.

⁴¹ Mildemberger 2015; see related discussion in Psillos 2004.

mechanisms and on the timing of changes in the Chechen rebels' behavior and doctrine, Runhardt critiques Bakke for not adding sufficient detail to her assertion that Chechen terrorists did not use suicide tactics prior to their interaction with global jihadists. This assertion bolsters the counterfactual implication of Bakke's argument: had there been no contact with global jihadists, the Chechens would not have framed their struggle as global or used suicide tactics.

It is not at all clear, however, what kind of details we should expect Bakke to add, or whether any additional detail is in fact necessary or useful. The fact that Chechen rebels had attacked civilians but not resorted to suicide attacks prior to their meetings with global jihadists is common background knowledge to those who have studied Chechnya. Runhardt also does not cite the practical standards that political scientists have developed for building and assessing counterfactuals, such as projectability—accurate predictions on hitherto unobserved cases⁴²—and cotenability—"connecting principles that link the antecedent [the premise of the counterfactual] with the consequent [the posited counterfactual outcome]...that are cotenable [consistent] with each other and with the antecedent."⁴³ Runhardt also does not take note of the alternative explanations that Bakke considers or suggest counterfactuals that might be assessed for these alternatives, though she provides some additional detail on this in her contribution to the present symposium.⁴⁴ For example, a key part of Bakke's argument is that Chechen rebels adopted the rhetoric and tactics of international jihadists not so much out of the inherent ideological appeal of these jihadists' global framing of the struggle, but because the global jihadists were able to offer material incentives—funding, expertise, etc.—for adopting their approach. This suggests a projection/prediction that jihadists will be more successful in winning converts when the jihadists have resources to offer and when the potential converts are resource poor, which has in fact proved generally true (after Bakke's writing) in the recruiting efforts of ISIS and in the radicalization of Islam in places like Kosovo,⁴⁵ although there are no doubt some cases in which materially endowed individuals joined ISIS and imams in well-to-do mosques became radicalized.

There are several additional fruitful uses for counterfactual reasoning in case studies that Runhardt does not address. Counterfactuals can help expose biases in our own thinking about the mechanisms that we hypothesize: if a scholar has confidence in a causal claim, but not an equal level of confidence in the logically equivalent counterfactual claim that the causal claim entails, then that scholar needs to identify and address the source of their cognitive bias. Counterfactuals can also be assessed to the extent that they entail factual ob-

servable implications. For example, if individual A discussed the possibility that individual or entity B might behave in a certain way, and made decisions and contingency plans for what to do if B took the anticipated action, but B did not take the anticipated action, it is plausible that A would have acted on his or her plan if B had behaved as anticipated. An alternative course of action is also a plausible counterfactual if individuals in positions of power argued for that course of action and appear to have nearly won the argument over what to do, and not plausible if it lacked any powerful advocate(s).

We can also use counterfactuals to clarify researchers' *thinking* about the relative "causal strength" of variables (which is analogous to the regression coefficients and their substantive significance in a statistical study, but is not the same as estimates of those coefficients because here we are talking about subjective degrees of belief about such coefficients). Any purchase on this notion of causal weight is otherwise very difficult to assess in qualitative research. For example, I have a long-running debate with William Wohlforth of Dartmouth College's Department of Government about why Soviet leaders did not resort to force to keep the Warsaw Pact countries from breaking away from Soviet control in 1989. We disagree on whether this was primarily due to: A) Soviet leaders' desire to gain access to world trade and technology in order to boost their sagging economy, or B) the lessons Soviet leaders learned from Afghanistan. I support the latter explanation and I have asked Bill directly whether he agrees to the following counterfactual: if the Soviet economy had started growing sharply under Gorbachev's perestroika policies (restructuring of the economy and society), the Soviet Union would have used force to prevent the independence of their Eastern European allies. He has declined to endorse this counterfactual, without specifying some alternative.⁴⁶ I am willing to endorse what I take to be the logically equivalent counterfactual of my argument: if Soviet leaders in 1989 had considered their use of force in Afghanistan a success, they would have used force in Eastern Europe in 1989. We can bring some evidence to bear on these counterfactuals as noted above, but the counterfactuals are of course not fully testable in that we are not able to re-run history. Thus my greater willingness to specify and endorse a counterfactual for my argument does not necessarily mean I am correct, and I do not discount entirely that economic incentives were a factor. Yet Wohlforth's and my different commitments regarding these respective counterfactuals convey useful information about our views on the strength of our arguments.

Runhardt mentions most-similar cases as a relevant kind of counterfactual comparison. Ideally, most-similar cases are similar in all but one independent variable of interest, and different in their outcomes. Runhardt notes that similarity claims are difficult to defend, but she does not note that case study methodologists have addressed most similar case comparisons in great detail. I have argued, for example, for using process tracing in most similar case comparisons to: 1) instantiate

⁴² Tetlock and Belkin 1996, 30–31.

⁴³ Tetlock and Belkin 1996; see also Levy 2015.

⁴⁴ Woodward (2003, 216) notes the role of counterfactual analysis of alternative explanations, calling it a "characteristic pattern of eliminative argument that plays an important role in establishing singular-causal claims and the counterfactuals associated with them."

⁴⁵ Gall 2016.

⁴⁶ I recount this in part from conversations with Wohlforth; for the relevant published articles see Wohlforth 2005 and Bennett 2005.

that there is a plausible process connecting the values of the independent variable of interest in the two cases to their respective outcomes, and, 2) assess whether other potential independent variables whose values differed in the two cases might be responsible for their outcomes.⁴⁷ To be fair, Runhardt's 2015 article is focused on the inverse question: whether most similar case comparisons can be used to corroborate a counterfactual that, in turn, supports a process claim. Either way, we must remember that both case comparisons and process tracing, whether framed in factual or counterfactual terms, are potentially incomplete and fallible.

In sum, I agree with Runhardt on the importance of thinking about counterfactuals in process tracing, but I think she underappreciates the ways in which counterfactual analysis is already in use among process tracers. There are also broader ways in which counterfactuals might be used, and we need more clarity on how we can use counterfactual analysis to assess and reduce our own cognitive biases.

Conclusions

David Waldner and Rosa Runhardt have done qualitative researchers a great service by exploring whether contemporary philosophical discussions of causal mechanisms have implications for process tracing practices. They have offered useful practical advice for process tracing. These are important contributions, even if they do not warrant Waldner's ambitious claim to have solved the fundamental problem of causal inference. Runhardt's analysis, by reminding us how difficult it is to make defensible counterfactual claims, reminds us also that it is difficult to make defensible causal claims. These contributions offer fruitful directions for additional research on the intersection of philosophy of science and process tracing methods. In particular, we need research, including experimental studies, on how actual researchers, with the cognitive biases to which they are vulnerable, make and use claims about evidence, causation, counterfactuals in process tracing, and whether training in Bayesian analysis, DAGs, counterfactual reasoning, and other process tracing techniques can bring actual practices more in line with philosophical ideals.

References

- Bennett, Andrew. 2005. "The Guns that Didn't Smoke: Ideas and the Non-use of Force in 1989." *Journal of Cold War Studies* vol. 7, no. 2: 81–109.
- Bennett, Andrew. 2015. "Disciplining our conjectures: systematizing process tracing with Bayesian analysis." In *Process Tracing: From Metaphor to Analytic Tool*. Edited by Andrew Bennett and Jeffrey Checkel. Cambridge: Cambridge University Press.
- Bennett, Andrew and Jeffrey Checkel, eds. 2015. *Process Tracing: From Metaphor to Analytic Tool*. Cambridge: Cambridge University Press.
- Fairfield, Tasha and Andrew Charman. 2016. "Explicit Bayesian Analysis for Process Tracing: Guidelines, Opportunities, and Caveats." Working paper available at http://media.wix.com/ugd/9ce28d_86c60b503a61423e9720ceb430b94ecb.pdf (last accessed 1 February 2017).
- DNews. November 11, 2013. "How Fire Burns in Space." Online at <https://www.youtube.com/watch?v=9zdD7lfB0Fs> (last accessed 26 December 2016).
- Gall, Carlotta. 2016. "How Kosovo was Turned into Fertile Ground for ISIS." *New York Times* (22 May).
- George, Alexander L. and Andrew Bennett. 2005. *Case Studies and Theory Development in the Social Sciences*. Boston: MIT Press.
- King, Gary, Robert Keohane, and Sidney Verba. 1994. *Designing Social Inquiry: Scientific Inference in Qualitative Research*. Princeton: Princeton University Press.
- Knight, Carly and Christopher Winship. 2013. "The causal implications of mechanistic thinking: identification using directed acyclic graphs (DAGs)." In *Handbook of Causal Analysis for Social Research*, edited by Stephen Morgan. London: Springer Press, 275–299.
- Levy, Jack. 2008. "Counterfactuals and Case Studies." In *Oxford Handbook of Political Methodology*, edited by Janet Box-Steffensmeier, Henry Brady, and David Collier. Oxford: Oxford University Press on Demand, 627–644.
- Levy, Jack. 2015. "Counterfactuals, Causal Inference, and Historical Analysis." *Security Studies* vol. 24, no. 3: 378–402.
- Mildenberger, Matto. 2015. "Process tracing, counterfactual comparisons, and causal inference: A new bridge over the qualitative-quantitative divide." Paper presented at the American Political Science Association annual conference.
- Pearl, Judea. 2000. *Causality: Models, Reasoning, and Inference*. Cambridge: Cambridge University Press.
- Psillos, Stathis. 2004. "A glimpse of the secret connexion: Harmonizing mechanisms with counterfactuals." *Perspectives on Science* vol. 12, no. 3: 288–319.
- Runhardt, Rosa W. 2015. "Evidence for causal mechanisms in social science: Recommendations from Woodward's manipulability theory of causation." *Philosophy of Science* vol. 82, no. 5: 1296–1307.
- Runhardt, Rosa W. 2016. "Tracing the Productive Continuity of Social Mechanisms." *Qualitative and Multi-Method Research: Newsletter of the American Political Science Association's QMMR Section* vol. 14, no. 1-2.
- Salmon, Wesley. 1989. *Four Decades of Scientific Explanation*. Minneapolis: University of Minnesota Press.
- Tetlock, Philipp and Aaron Belkin. 1996. *Counterfactual Thought Experiments in World Politics*. Princeton: Princeton University Press.
- Waldner, David. 2012. "Process Tracing and Causal Mechanisms." In *The Oxford Handbook of Philosophy of Social Science*, edited by Harold Kincaid.
- Waldner, David. 2015. "What makes process tracing good? Causal mechanisms, causal inference, and the completeness standard in comparative politics." In *Process Tracing: From Metaphor to Analytic Tool*, edited by Andrew Bennett and Jeffrey Checkel. Cambridge: Cambridge University Press, 126–152.
- Waldner, David. 2016. "Invariant Causal Mechanisms." *Qualitative and Multi-Method Research: Newsletter of the American Political Science Association's QMMR Section* vol. 14, no. 1-2.
- Wohlforth, William. 2005. "The End of the Cold War as a Hard Case for Ideas." *Journal of Cold War Studies* vol. 7, no. 2: 165–173.
- Wood, Elizabeth Jean. 2000. *Forging Democracy from Below: Insurgent Transitions in South Africa and El Salvador*. Cambridge: Cambridge University Press.
- Woodward, James. 2003. *Making Things Happen: A Theory of Causal Explanation*. Oxford: Oxford University Press.

⁴⁷ George and Bennett 2005, 51 and 81.