Process Tracing: From Philosophical Roots to Best Practices

Andrew Bennett and Jeffrey T. Checkel

The **Simons Papers in Security and Development** are edited and published at the School for International Studies, Simon Fraser University. The papers serve to disseminate research work in progress by the School's faculty and associated and visiting scholars. Our aim is to encourage the exchange of ideas and academic debate. Inclusion of a paper in the series should not limit subsequent publication in any other venue. All papers can be downloaded free of charge from our website, www.sfu.ca/internationalstudies.

The series is supported in part by the Simons Foundation.

Series editor: Jeffrey T. Checkel
Managing editor: Martha Snodgrass

Process Tracing: From Philosophical Roots to Best Practices

Simons Papers in Security and Development
No. 21/2012 | June 2012

Abstract:
This paper has two overarching goals – to summarize recent developments on the philosophical and practical dimensions of process tracing, and to identify features common to best practices of it on different kinds of arguments, with different kinds of available evidence, in different substantive research domains. First, we define process tracing and discuss its foundations in the philosophy of social science. Next, we address its techniques and evidentiary sources, and advance ten criteria for judging its quality in particular pieces of research. Finally, we analyze the methodological issues specific to process tracing on general categories of theories, including structural-institutional, cognitive-psychological, and sociological. (NB: this paper is forthcoming as Chapter 1 in A. Bennett and J.T. Checkel, eds., *Process Tracing in the Social Sciences: From Metaphor to Analytic Tool*.)

About the authors:
Andrew Bennett is a professor in the Department of Government at Georgetown University. He has written on the US foreign policy process, research methods, alliance burden-sharing, and regional conflicts and peacekeeping. He is the co-author, with Alexander L. George, of *Case Studies and Theory Development in the Social Sciences* (MIT Press) and author of *Condemned to Repetition? The Rise, Fall, and Reprise of Soviet–Russian Military Interventionism 1973–1996* (MIT Press). Bennett is President and co-founder of the Consortium on Qualitative Research Methods.

Jeffrey T. Checkel holds the Simons Chair in International Law and Human Security at Simon Fraser University and is an Adjunct Research Professor at the Peace Research Institute Oslo. He is the author of *Ideas and International Political Change: Soviet/Russian Behavior and the End of the Cold War* (Yale University Press), editor of *International Institutions and Socialization in Europe* (Cambridge University Press), co-editor (with Peter J. Katzenstein) of *European Identity* (Cambridge University Press), and editor of *Transnational Dynamics of Civil War* (Cambridge).

Process Tracing: From Philosophical Roots to Best Practices

I. Introduction

Why did the Cold War end peacefully, without a shot being fired? In the post-Cold War world, civil conflicts have replaced interstate war as the dominant form of organized political violence, with rebel groups – instead of ICBMs – as a key focus of both policy and scholarship. Yet what makes such groups tick? Why do some engage in wanton killing and sexual violence while others do not? The European Union is a unique experiment in governance 'beyond the nation state,' but how are its supra-national governance structures being crafted and with what effect on the ordinary citizens of Europe?

Contemporary political scientists and sociologists have converged on the view that these puzzles, and many more on the current scholarly and policy agendas, demand answers that combine social and institutional structure and context with individual agency and decision-making. This view, together with recent developments in the philosophy of science, has led to an increasing emphasis on causal explanation via reference to hypothesized causal mechanisms. Yet this development begs the questions of how to define such mechanisms, how to measure them in action, and how to test competing explanations that invoke different mechanisms.

This book argues that techniques falling under the label of "process tracing" are particularly well-suited for measuring and testing hypothesized causal mechanisms.

Indeed, a large and growing number of political scientists now invoke the term. Despite or perhaps because of this fact, a buzzword problem has arisen, where process tracing is mentioned, but often with little thought or explication of how it works in practice. As one sharp observer has noted, proponents of qualitative methods draw upon various debates – over mechanisms and causation, say – to argue that process tracing is necessary and good. Yet, they have done much less work to articulate the criteria for determining whether a particular piece of research counts as good process tracing (Waldner, nd: 2–3). Put differently, "there is substantial

distance between the broad claim that 'process tracing is good' and the precise claim 'this is an instance of good process tracing'" (Waldner, 2011: 7).

This volume addresses such concerns, and does so along several dimensions. Meta-theoretically, it establishes a philosophical basis for process tracing – one that captures mainstream uses while simultaneously being open to applications by interpretive scholars. Conceptually, contributors explore the relation of process tracing to mechanism-based understandings of causation. Most important, we develop evaluative standards for individual process-tracing accounts – for example, criteria for how micro to go and how to deal with the problem of equifinality (the possibility that there may be multiple paths or combinations of different variables that can produce the same kind of outcome).

Ours is an applied methods book – and not a standard methodology text – where the aim is to show how process tracing works in practice. If Van Evera (1997), George and Bennett (2005) and Gerring (2007a) set the state of the art for case studies, then our volume is a logical follow-on, establishing clear standards for what is perhaps the central within case method – process tracing.

Despite all the recent attention, process tracing – or the use of evidence from within a historical case to make inferences about causal explanations of that case – has in fact been around for thousands of years. Related forms of historical analysis date back to the Greek historian Thucydides and perhaps even to the origins of human language and society. It is nearly impossible to avoid historical explanations and causal inferences from historical cases in any purposive human discourse or activity.

Although social science methodologists have debated and detailed formal approaches to inference such as statistical analysis for over a hundred years, they have only recently coined the term 'process tracing' or attempted to explicate its procedures in a systematic way. Perhaps this is because drawing causal inferences from historical cases is a more intuitive practice than statistical analysis and one that individuals carry out in their everyday lives. Yet, the seemingly intuitive nature of process tracing obscures that its unsystematic use is fraught with potential

inferential errors; it is thus important to utilize rigorous methodological safeguards to reduce such risks.

The goal of this book is therefore to explain the philosophical foundations, specific techniques, common evidentiary sources, and standards and best practices of process tracing in order to reduce the risks of making inferential errors in the analysis of historical cases. This introductory chapter first defines process tracing and discusses its foundations in the philosophy of social science. We then address its techniques and evidentiary sources, and advance ten criteria for judging the quality of process tracing in empirical research. The chapter concludes with an analysis of the methodological issues specific to process tracing on general categories of theories, including structural-institutional, cognitive-psychological, and sociological. Subsequent chapters take up this last issue in greater detail and assess the contributions of process tracing in particular research programs or bodies of theory.

II. Defining Process Tracing

The term process tracing originated in the field of cognitive psychology in the United States in the late 1960s or early 1970s.[1] As used in psychology, process tracing refers to techniques for examining the intermediate steps in cognitive mental processes to better understand the heuristics through which humans make decisions. In 1979, the Stanford University political scientist Alexander L. George appropriated the term process tracing to describe the use of evidence from within case studies to make inferences about historical explanations (George, 1979).

Because much of George's own research was in political psychology, and because the term process tracing originated in cognitive psychology, process tracing has sometimes been viewed as applying mostly or only to analyses of individual level decision making. Although process tracing does apply well to individual level and cognitive theories (see Jacobs, this volume), George made clear in subsequent writings that it can also be used to make inferences on structural or macro-level explanations of historical cases (George and Bennett, 2005: 142,

---

[1] The very first usage of the term remains unclear; the earliest relevant citation on Google Scholar is Hobarth, 1972, a PhD thesis at the University of Chicago.

214; see also Waldner, this volume). For example, many economic theories hypothesize relationships and sequences among macroeconomic variables that can be tested through process tracing at the macro level as well as that at the micro or individual level.

Similarly, because of its origins in cognitive psychology and because many of its early practitioners in that field went on to explore the errors that individuals make and the biases they exhibit in their decision making, process tracing is sometimes viewed as incompatible with rational choice theories. We concur, however, with the many prominent rational choice theorists who argue that their hypotheses should bear some correspondence with the actual processes through which individuals make decisions, and that they should therefore be amenable to process tracing (Bates, Weingast, Grief, Levi, and Rosenthal, 1998; see also Schimmelfennig, this volume).

The essential meaning that the term process tracing has retained from its origins in cognitive psychology is that it refers to the examination of intermediate steps in a process to make inferences about hypotheses on how that process took place and whether and how it generated the outcome of interest. In previous work together with George (George and Bennett, 2005: 6), one of us defined process tracing as the use of "histories, archival documents, interview transcripts, and other sources to see whether the causal process a theory hypothesizes or implies in a case is in fact evident in the sequence and values of the intervening variables in that case." We added that "the process-tracing method attempts to identify the intervening causal process – the causal chain and causal mechanism – between an independent variable (or variables) and the outcome of the dependent variable" (George and Bennett, 2005: 206).

The authors then used a metaphor to expand on this definition. If one had a row of fifty dominoes lying on the table after they had previously been standing, how could one make inferences about whether the first domino caused the last to fall through a 'domino process,' or whether wind, a bump of the table, or some other force caused the dominoes to fall? The answer, George and Bennett argued, was to use evidence on the intervening processes posited by each of the alternative explanations. Did anyone hear a succession of dominoes? Do the positions of the fallen dominoes shed light on how they fell? And so on.

While we feel this definition is still an excellent starting point, it is necessary to point out a weakness in both it and the accompanying metaphor. The term intervening variable opens the door for potential confusion because social scientists are accustomed to thinking of variables as either causal (independent) or caused (dependent). However, both the term and the metaphor of dominoes falling suggest that an intervening variable is both fully caused by the independent variable(s) that preceded it, and that it transmits this causal force, without adding to it, subtracting from it, or altering it, to subsequent intervening variables and ultimately through them to the dependent variable.

The observable events that intercede between hypothesized causes and observed effects could indeed have this character, but they might instead be caused by forces in addition to those of the specified independent variables, and they may have amplifying or dampening effects of their own, or interactions or feedback effects. In the domino metaphor, the dominoes could be falling in sequence with increasing or decreasing force depending on the spacing between them, and the table could also be shaking and the wind blowing at the same time that the dominoes are striking one another, each with some effects on the falling dominoes.

In such instances, researchers have to make theory-building choices. Are they going to model additional variables that play a role in the process as exogenous, complementary, or endogenous? Exogenous variables are those excluded from the model because they are either not powerful enough to worry about or too complex to be theorized. Complementary variables are those that add to or subtract from the effects of the main variables of interest but do so independently, or without interaction effects related to the main variables. When such variables are sufficiently simple to be theorized about, they can be added to a model without changing the main variables or mechanisms of interest. Additional variables that interact with the independent variables of interest in more complex ways need to be either brought into the model (endogenized) or identified but set aside from the model (exogenized). Methodologically, whatever way additional variables are brought into or set aside from the theory that aims to explain the case, this modification can be tested through additional process tracing.

Alternatively, and perhaps ideally, some process tracing observations are based on diagnostic evidence on events, or evidence that indicates the kind of process taking place but

does not transmit any effects to the dependent variable. This is analogous to a diagnostic medical test, such as a dye injected into a patient to enhance a CAT scan of blood flow in a damaged heart. The dye does not transmit any effects to the heart, but it provides evidence on the processes that damaged the heart. Similarly, in social and political life, the ways in which actors privately frame or explain their actions, and the actions that they refrain from, may provide diagnostic evidence on their motives without directly and publicly affecting the outcomes of interest.

We thus drop the term intervening variable and define process tracing as the analysis of evidence on processes, sequences, and conjunctures of events within a case for the purposes of either developing or testing hypotheses about causal mechanisms that might causally explain the case. Put another way, process tracing examines the deductive observable implications of hypothesized causal mechanisms within a case to test whether these might in fact explain the case (Schimmelfennig, this volume, emphasizes such a procedure). Or, it inductively uses evidence from within a case to develop hypotheses that might explain the case; the latter hypotheses may, in turn, generate additional testable implications in the case or in other cases (Petersen, this volume, gives examples of the inductive use of process tracing to develop theories).

This definition encompasses the inductive (theory development) and deductive (theory testing) sides of process tracing, and it eschews the ambiguous term intervening variable. It also reminds us that the causal relations among the events within the case, and between these events and the specified independent and dependent variables, may involve amplifying effects, dampening effects, feedback effects, and diagnostic connections. Nonetheless, this definition aspires to causal inference.

It is important as well to define "case," "within a case" and "within case analysis" as we use them. Following George and Bennett, we define a case as "an instance of a class of events" (George and Bennett, 2005: 17). This definition recognizes that classes of events – revolutions, democracies, capitalist economies, wars, and so on – are the social constructions of both political actors and the social scientists who study and define political categories. They are not simply given to us by history, but defined by our concepts, and much contestation in interpreting the

results of case study research concerns disagreements over which cases should or should not be included in a defined population.

We define within-case evidence as evidence from within the temporal, spatial, or topical domain defined as a case. This can include a great deal of evidence on contextual or background factors that influence how we measure and interpret the variables within a case. Henry Brady and David Collier provide a useful distinction here between data set observations and causal process observations. Data set observations are "an array of scores on specific variables for a designated sample of cases," and these observations provide the basis for statistical analyses. Causal process observations are "observations on context, process, or mechanism" and are used in within-case analyses such as process tracing (Brady and Collier, 2010: 12).

Process tracing is closely related to historical explanation, as that term is used by the historian Clayton Roberts. In Roberts's view, an historical explanation is not simply a detailed description of a sequence of events; rather, it draws on theories to explain each important step that contributes to causing the outcome. Roberts distinguishes here between macro-correlation and micro-correlation, the latter of which is quite similar to process tracing. Macro-correlation involves an attempt to explain historical cases at a high level of generality through universalistic theories, similar to Hempel's notion of theories as covering laws. Roberts argues that historical events are too complex to fit easily under exception-less covering laws, and efforts to explain history via covering laws "have met with little success" (Roberts, 1996:15). He urges instead that researchers should use micro-correlation, which involves "the minute tracing of the explanatory narrative to the point where the events to be explained are microscopic and the covering laws correspondingly more certain" (Roberts, 1996: 66).

One difference between Roberts's approach to process tracing and our own is that Roberts felt that – at the micro-correlational level – the theories underlying an historical explanation would be "platitudinous," and he noted that historians rarely referenced them explicitly because to do so would "hopelessly clog the narrative" (Roberts, 1996: 66–67, 87–88). We emphasize instead the importance of making explicit the hypotheses about underlying causal mechanisms that are theorized to have caused an historical outcome so that these can be rigorously assessed, even at the expense of political science narratives that are more clogged

(and alas, less likely to become best-sellers) than those of historians (see also Evangelista, this volume, for analysis of works that focus their process tracing as much on explaining an important historical case – the end of the Cold War – as on developing and testing general theories).

Our concept of process tracing differs even more sharply with time series cross sectional analysis, which involves the correlational study of data across a variety of units (often, annual data across a range of countries). Although this form of analysis might be confused with process tracing because it involves temporal data from within cases over time, it is still a form of cross-case and correlational inference, rather than the study of hypothesized processes within individual cases, and it is thus fundamentally different from process tracing.

In sum, process tracing is a key technique for capturing causal mechanisms in action. It is not simply glorified historiography, nor does it proceed by the logic of frequentist statistics. And – as we argue below – there are metrics and standards available that allow one to distinguish good process tracing from bad. However, since standards in some important sense flow from underlying philosophical positions, it is important first to clarify the meta-theory of process tracing.

III. Philosophy of Social Science and Process Tracing

On a philosophical and epistemological level, process tracing is closely related to the turn toward social science explanations based on reference to causal mechanisms (Elster, 1998; Gerring, 2007b; Mayntz, 2004), or the underlying entities that generate observed processes and outcomes. Much of the thinking about causality and causal explanation over the last two hundred years has been strongly influenced by David Hume's argument that constant conjunction – the frequent conjoint occurrence of variables A and B – is the essence of causal inference. However, more recent work by pragmatist (Johnson, 2006) and scientific realist (Wight, 2006) philosophers of science provides a meta-theoretical foundation more amenable to thinking in terms of mechanisms. Indeed, for these scholars, a causal explanation is built around contiguity and sequencing – concepts that Hume mentioned but gave insufficient attention. These concepts open a methodological space for process tracing.

One difficulty in making use of contemporary discussions in the philosophy of science is that there are at least a half-dozen variants of scientific realism (Chernoff, 2002) and even more different definitions of causal mechanisms (Mahoney, 2001; see also Hedstroem and Ylikoski, 2010). While a full discussion of scientific realism is beyond our present purposes, we concur with the emphasis it places on causal processes and causal mechanisms as central elements of causal explanation.

More important for this volume is the task of sorting out the competing definitions of causal mechanisms. These divide along three fundamental issues. 1) Are causal mechanisms in some sense unobservable? 2) Does explanation via reference to causal mechanisms involve a commitment to methodological individualism, or beyond that, to explaining human behavior by neuroscience and ultimately by sub-atomic physics? 3) Are causal mechanisms sufficient to explain outcomes in specified circumstances or contexts, or might mechanisms be inherently probabilistic or stochastic? A brief discussion of these controversies sets the stage for own definition of causal mechanisms.[2]

On the first issue, most discussions of mechanisms place them on the ontological level. This means we make hypotheses about how such ontological entities as mechanisms might work, and we test the observable implications of these hypotheses, but we do not have unmediated access to and are not able to directly observe causal mechanisms. Some proponents of mechanisms take a different view, arguing that they are at least somewhat observable. Hedstroem and Ylikoski, for example, critique Mahoney for the view that mechanisms are unobservable, and draw an analogy to our ability to observe the inner workings of an auto engine (Hedstroem and Ylikoski, 2010: 50–51; Mahoney, 2001).

Such critiques, however, miss the more fundamental point that causal mechanisms are in some sense ultimately unobservable. We do not get to observe causality – we make inferences about it. Moreover, we cannot unproblematically observe many mechanisms at lower levels of analysis – brain waves, neurons, atoms, *etc*. Rather, we rely on potentially fallible instruments of observation (brain scans, microscopes) and theories about how they work. We may be able to

---

[2] We do not address a fourth issue – the difference between causal mechanisms and theories or hypotheses – as this has been treated elsewhere. Hedstroem and Ylikoski, 2010: 54-55, for example.

push back the borders of the unobservable world by developing instruments of observation in which we have great confidence, but there will always be some still finer level of detail that we cannot observe.

The boundary between the observable and unobservable worlds is like the horizon. We can push this border back as our instruments of observation improve, but it also recedes as we move toward it, and some part of the universe always remains beyond the horizon and unobservable. In the social world, even if brain scans are beginning to reveal some of the inner workings of the brain, we do not have these in real time for actual social actors in real world settings, we cannot scan brain activity from the past, and there will still be additional micro-level brain processes that we cannot observe.

This raises the second issue concerning methodological individualism and the degree to which mechanism-based explanations have to go to minute levels of detail, tiny increments of time, and temporally distant causes of events. In our view, explanations need not always go to the individual level of analysis (or beyond); it is possible to do process tracing on hypothesized causal mechanisms at the macro level (for example, Waldner, this volume). In principle, mechanism-based explanations have to be consistent with the finest level of detail we observe; however, in practice, this does not mean we must always go to this level to have confidence that one explanation is more likely to be true than the alternatives. Many macroeconomic theories posit processes involving several temporal steps that can be tested empirically at the macro level.

The controversy surrounding this issue has led some critics to argue that explanations built on causal mechanisms – and, thus, process tracing – involve a potentially infinite regress to look at steps between steps in a hypothesized process at ever smaller increments of time and more detailed levels of analysis (King, Keohane, and Verba, 1994: 85–87). It is true that a commitment to explanation via mechanisms means that explanations are always incomplete and provisional, and that every explanation can be called into question if it can be shown that its hypothesized processes are not evident at a lower level of analysis. It is also true that there is no infallible way of deciding how far down, or how far back, to go in explaining an event. However, as we argue below, researchers can and do make defensible decisions about when and where to

begin and stop in constructing and testing explanations (see also Hedstroem and Ylikoski, 2010: 52).

The issue of when to stop is related to the third controversy of whether causal mechanisms involve relations of sufficiency or probabilism. Mahoney (2001: 580) defines mechanisms as being sufficient in specified circumstances to generate the outcome of interest, while Hedstroem and Ylikoski (2010: 51) argue that mechanisms could be inherently stochastic. This is a thorny issue, as stochastic relations – like those posited by quantum theory – have some aspects of causal explanation but lack others (Salmon, 1990: 120).

The core problem is that even if the world is deterministic, we observe it as probabilistic because of measurement error and specification error, including the omission of important variables from our models. We cannot tell with 100% confidence whether we are witnessing a probabilistic world or a deterministic one, or whether some processes are deterministic or nearly so while others are inherently stochastic. Researchers implicitly make choices on this issue in deciding when to stop trying to reduce or explain the "error term," or unexplained variation.

In sum, on the key issues in the definitional debates about causal mechanisms we argue the following. Causal mechanisms are ontological entities and processes in the world, and theories or hypotheses are in our heads; we theorize about mechanisms. Such mechanisms are ultimately unobservable, but our hypotheses about them generate observable and testable implications. Explanation via reference to causal mechanisms, unlike that via reference to covering laws, involves a commitment in principle to being consistent with the lowest level of analysis and finest degree of detail observable. We can never know with certainty whether the world in general or a particular mechanism that we hypothesize is deterministic or nearly so under specified circumstances or whether either is stochastic. We thus define 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. In doing so, the causal agent changes the affected entities' characteristics, capacities, or propensities in ways that persist until subsequent causal mechanisms act upon them. If we are able to measure changes in the entity being acted upon after the intervention of the causal mechanism and in temporal or spatial isolation from other mechanisms, then the causal mechanism may

be said to have generated the observed change in the entity (see also George and Bennett, 2005: 137).

The methodological challenge, then, is to develop theories about causal mechanisms in which we can place some confidence and understandings of the scope conditions in which they operate. Process tracing is one powerful method for addressing these challenges. Before turning to the nuts and bolts of how to do process tracing well, however, three additional foundational issues demand attention: the relationship of process tracing to generalization, to interpretive social science, and to Bayesian inference.

Generalizability and Process Tracing. Because causal mechanisms are operationalized in specific cases, and process tracing is a within-case method of analysis, generalization can be problematic. Case study methodologists have argued that a hypothesis is strongly affirmed and might be generalizable if it explains a tough test case or a case that, *a priori*, it looked least likely to explain. Conversely, the failure of a hypothesis to explain a most likely case strongly reduces our confidence in it. It has always been rather ambiguous, however, whether these inferences should apply only to the case being studied, to cases very similar to the one studied, or to a broader range of more diverse cases.

The use of process tracing to test and refine hypotheses about causal mechanisms can clarify the scope conditions under which a hypothesis is generalizable. A researcher cannot have a very clear idea whether, how, and to which populations an explanation of a case might generalize until they have a clear theory about the workings of the mechanisms involved in the case. To some degree, this theory can evolve inductively from close study of the case itself.

Indeed, a theory or explanation derived inductively from a case does not necessarily need to be tested against a different case for us to have confidence in the theory; rather, it can be tested against different and independent evidence in the case from which it was derived (Mahoney, 2012: 18). Often, this is a kind of evidence that the researcher had not thought to look for or did not recognize as relevant prior to developing the new explanation. Detectives, medical doctors, and case study researchers in many sciences and professions frequently make this move.

For example, in a study of international socialization in Europe, Checkel and collaborators theorized three mechanisms of socialization, two of which were partly derived from their own case material. The careful application of process tracing to additional, independent evidence from the cases was then used to specify better the scope conditions of each mechanism. The result and central finding was that the theory was limited in its application to the – albeit crucially important – case of contemporary Europe (Checkel, 2007: chapters 7–8).

Conversely, a researcher focusing on one or a few cases might uncover a new hypothesis that is broadly applicable, as when Charles Darwin's study of a few species led to his theory of evolution. In short, we may uncover hypothesized mechanisms through process tracing that may be either very generalizable or unique to one or a few cases, but it is almost impossible to know prior to researching a case the degree to which any inductively-derived explanations will be one or the other.

Interpretism and Process Tracing. Another important foundational issue is the relation between process tracing and interpretism, or more specifically in our own subfield of international relations, between process tracing and constructivism. Recall our earlier discussion, where we argued that scientific realism provides a possible meta-theoretical basis for process tracing. With its stress on cause, objectivity, the consideration of alternative explanations and the like, scientific realism is closer to positivism in its various guises than to interpretism (Wight, 2002: 35–36). What (meta-theoretical) space does this then leave for interpretive process tracing?

One difficulty here is that scholars have embraced many different kinds of interpretism and constructivism.[3] Most constructivists agree that structures or institutions are social as well as material, and that agents and structures are mutually constitutive; however, they differ on important epistemological issues (Adler, 2002). One common typology that we find useful distinguishes among conventional, interpretive and radical or post-modern views of social life. In this schema, conventional constructivists still aspire to causal explanation and believe that there are standards for assessing some interpretations of social life to be superior to others. Alexander Wendt, a leading constructivist in international relations who has espoused scientific realism and

---

[3] We will use these terms interchangeably in the following.

a role for causal mechanisms, fits into this school of thought (Wendt, 1999). Not surprisingly, process tracing figures prominently in the work of many conventional constructivists (Risse, *et al*, 1999, for example).

It is more challenging to reconcile the technique with a second, interpretive view, though some scholars are attempting to do so (Autesserre, 2009; Hopf, 2002, 2007, 2012; Pouliot, 2007). Here, agents and structures are so inherently mutually constitutive that it is not possible to separate events into discrete moves in which either the agent or the structure is primarily driving the process. If indeed mutual constitution is completely continuous at all levels of analysis, then it is impossible to break out "variables" as being causes or consequences of one another. However, one can often break down events and discern steps at which an agent – for example, a norm entrepreneur – is contesting social structures, and steps at which a structure prevents agents from acting upon or even conceiving of courses of action that are taboo. In fact, a similar bracketing strategy has been endorsed by several prominent (conventional) constructivists (Wendt, 1987; Finnemore, 1996).

A third, radical or post-modern view maintains that language, arguably the most central of all social structures, is inherently ambiguous and open to many interpretations. The danger here is that all narratives are reduced to story-telling, a critique that has also been raised against process tracing (Norkus, 2005). We should note, however, that even these radical forms of constructivism have increasingly developed standards of evidence. We thus now have clear 'how to' guides for conducting systematic discourse and textual analysis (Hansen, 2006; Neumann, 2008; Hopf, 2002: chapter 1). Moreover, genealogical methods – the historical reconstruction of discourses – bear a strong family resemblance to some historical forms of process tracing (Price, 1997). Finally, in recent years, there has been a strong move to 'bring practice back in' to the study of discourse (Pouliot, 2010), which provides an interpretive nod to the central importance of process.

In sum, while there are philosophical hurdles to surmount – or perhaps better said, to be bracketed – we see intriguing possibilities for developing a richer understanding of process tracing by drawing upon these various strands of interpretive social science. This is precisely the challenge that Pouliot takes up in his contribution below (Pouliot, this volume).

Bayesianism and Process Tracing. The Bayesian approach to confirming and choosing among explanations has became an increasingly popular one in the philosophy of science, and it is closely related to process tracing in ways that illuminate the latter's strengths and limits (Bennett, 2008). Both Bayesianism and process tracing rely on using evidence to affirm some explanations and cast doubt upon others, and each approach emphasizes that the probative value of evidence relative to competing explanations is more important than the number of pieces of evidence. Both argue for the possibility that a few pieces of evidence with high probative value, and in some instances even just one piece of evidence, can allow observers who approach a case with different theoretical priors to converge in their views on the proper explanation of the case.

The Bayesian approach involves using evidence to update one's beliefs in the likelihood that alternative explanations are true. Of course not all evidence is decisive, and some may be consistent with several competing explanations. In this regard, Stephen Van Evera has developed a very useful terminology for discussing the probative value of alternative kinds of evidentiary tests of theories that tracks closely with Bayesian logic. In Van Evera's view, the probative value of evidence depends on the degree to which a hypothesis uniquely predicts that evidence, and the degree to which it is certain in doing so; another way to think of this is whether finding particular evidence is necessary and/or sufficient for a theory to be likely to be true (Van Evera, 1997: 31–32; see also Bennett, 2010; Collier, 2011).

From the four possible combinations of (non)uniqueness and (un)certainty, Van Evera derives four tests. *Hoop tests* involve evidence that is certain, but not unique; failing a hoop test disqualifies an explanation, but passing it does not greatly increase confidence in that explanation. Hoop tests are thus most useful in excluding alternative hypotheses. Van Evera's example of a hoop test is: "Was the accused in the state on the day of the murder?" Failing this hoop test falsifies the hypothesis. If the accused was not in the state, they are not the murderer. Whether passing a hoop test constitutes strong evidence of the truth of the hypothesis depends on how frequently the "pass" condition occurs naturally (Mahoney, 2012: 7). In Van Evera's example, if the accused is always in the state, there is zero probability they will fail the hoop test, so the suspect's having passed the test is not very definitive. Also, the higher the number of other

suspects that were in the state at the time of the murder, the less this hoop test increases the likelihood that any particular one of them is the murderer.

Smoking gun tests are unique, but not certain. Passing a smoking gun test is sufficient for strongly affirming an explanation, but passing such a test is not necessary to build confidence in an explanation. In Van Evera's example, a smoking gun in a suspect's hands right after a murder strongly implicates that suspect, but the absence of such a smoking gun does not exonerate this suspect. Smoking gun tests are therefore usually used to instantiate a hypothesis. Again, the strength of a smoking gun test depends on how often the condition in question occurs on its own (Mahoney, 2012: 8–9). If the murder took place at a shooting range, or especially if it happened at a crowded shooting range during a reunion of the victim's dysfunctional family, the presence of a smoking gun in the suspect's hand may not be very definitive.

Doubly decisive tests use evidence that is both unique and certain, or that is necessary and sufficient to provide great confidence in an explanation. Van Evera uses the example of a bank camera that catches the faces of all those involved in robbing the bank. To convict an individual in the robbery, it is both necessary and sufficient to show that their face matches the camera footage; to exonerate an individual, it is necessary and sufficient to show that their features do not match the bank video.

Finally, tests that are neither unique nor certain – neither necessary nor sufficient – are *straw in the wind tests.* Any one such test is not very decisive, but it is possible that a series of such tests might increase confidence in one explanation and decrease that in others, just as a suspect might be convicted on the basis of many pieces of circumstantial evidence (Mahoney, 2012: 15).

Thus, in both Bayesianism and process tracing, what matters in determining our confidence in alternative explanations of a case is not so much the number of pieces of evidence as the probative value of the evidence relative to the competing explanations. Of course, it is impossible to consider *all* the potential alternative explanations, as these would include explanations that researchers have not yet conceived. Bayesians have argued that one should thus include all alternative explanations that have actually been proposed (Earman, 1992: 84–85). For

sure, this approach is incomplete and new hypotheses may be conceived in the future, which is one reason why Bayesians never attach 100% probability to the likelihood that a theory is true, an attitude that we urge practitioners of process tracing to adopt as well.

Bayesian logic gives a justification for the common intuition among process tracers that diversity and independent evidence are useful in testing explanations (Earman, 1992: 78–79, 119–120, 132; Hellman, 1997: 202–209). When evidence on a particular step of a process becomes sufficiently repetitive that it is to be expected, it loses any ability to further discriminate among competing hypotheses. In contrast, independent evidence on a new or different dimension or stage of a hypothesized process retains the potential for surprising us and forcing us to revise our priors.

Process tracing and Bayesian analysis are also similar in that both proceed as much by eliminating alternative hypotheses as by providing support for those that fit observed processes.[4] Sir Arthur Conan Doyle gives an excellent account of this 'eliminative induction' when he has his fictional detective Sherlock Holmes argue that an investigation "starts upon the supposition that when you have eliminated all which is impossible, then whatever remains, however improbable, must be the truth. It may well be that several explanations remain, in which case one tries test after test until one or other of them has a convincing amount of support" (Doyle, 1930/2003: 528–9).

IV. Techniques and Standards of Process Tracing

Process tracing usually proceeds through a mix of induction and deduction. The particular mix in a research project depends on the prior state of knowledge and theorizing about the phenomenon and case selected for study, and on whether the case is similar to a defined population of cases or is an outlier vis-a-vis this population. For phenomena on which there is little prior knowledge and for cases that are not well-explained by extant theories, process tracing proceeds primarily through inductive study. This often involves analyzing events backward through time from the outcome of interest to potential antecedent causes, much as a homicide detective might start by trying to piece together the last few hours or days in the life of a victim.

---

[4] See also the discussion of equifinality below.

In this process, the researcher takes in a lot of information that may or may not later become part of the hypothesized explanation, a phase that some have colloquially called "soaking and poking." Here, one immerses oneself in the details of the case and tries out proto-hypotheses that may either quickly prove to be dead ends or become plausible and worthy of more rigorous testing. It is important that the investigator be open to all kinds of possible explanations and be willing to follow the evidence wherever it leads. The more promising potential explanations uncovered in this way can then be rendered more formal and deductive and tested more rigorously against evidence in the case or in other cases that is independent of the evidence that gave rise to each hypothesis.

If theories that appear to offer potential explanations of a case already exist, or after such theories have been developed inductively, process tracing can proceed more deductively. A key step here is to develop case-specific observable implications of the theories in question (Bakke, 2013, for an excellent example; see also the discussion in Jacobs, this volume), as theories are seldom specified in such precise ways that they offer tight predictions on the observable implications that should be evident in particular cases. It is also important to cast the net widely for alternative explanations, including theoretical explanations in the academic literature, the more context-specific arguments that historians or regional or functional experts have offered, the implicit theories of journalists or others following the case, and the understandings participants have about what they are doing and why they are doing it. As researchers develop observable implications of hypothesized mechanisms, they should be on the lookout for particularly valuable kinds of evidence that allow for hoop, smoking gun, and doubly decisive tests.

Iteration between the inductive and deductive sides of process tracing is very common. The most important methodological safeguard here is that researchers should seek to identify additional observable implications or what Imre Lakatos called "new facts" to test each modification to a hypothesis in order to avoid confirmation bias. Particularly valuable are new testable implications that, if found, would fit only the modified theory and not the alternative explanations, or that had not already been observed and had not been used to construct the hypothesis (Lakatos, 1970). As noted above, it is not necessary for evidence to come from a

different case than that which led to the development of a new hypothesis; it is only necessary that any new evidence from within the case be independent of the evidence that generated the hypothesis.

There is a related distinction between evidence that is unavailable and evidence that is contrary to the process tracing expectations of a hypothesis. Evidence that is unavailable at the time of the research, such as classified information, lowers the upper limit of the probability one can attach to the likely truth of an explanation. One useful technique here is to predict what the unavailable evidence will indicate once it becomes available; such predictions, if borne out, provide strong confirmatory evidence. This was precisely the strategy followed by one of us where process tracing was employed to test hypotheses on socialization mechanisms in small group settings within international organizations. On the basis of interviews and a reading of primary documentation, predictions were made about socialization dynamics; these were subsequently confirmed through the release of previously classified meeting minutes (Checkel, 2003).

Evidence that is contrary to the process tracing predictions of an explanation lowers the likelihood that the explanation is true. The explanation may need to be modified if it is to become convincing once again. This modification may be a trivial one involving a substitutable and logically equivalent step in the hypothesized process, or it could be a more fundamental change to the explanation. The bigger the modification, the more important it is to generate and test new observable implications to guard against 'just so' stories that explain away anomalies one at a time.

Inferences from process tracing also depend in part on the likelihood that one would expect to be able to find and have access to certain kinds of evidence. Debates on alternative explanations often turn on the question of whether "absence of evidence" constitutes "evidence of absence." If we expect evidence to be readily accessible and doubly decisive – as when we feel around for change in our pocket – we might infer that failure to find something constitutes evidence it does not exist (although many of us have found lost keys or other items only on the third or fourth attempt of looking in the same place). If we expect evidence to be hard to find or access, as when thinking social actors have gone to great lengths to hide evidence of their actions

or motives, then absence of evidence might not greatly lower our expectation that an entity or relationship exists. Yet even when actors have motives to hide evidence, sufficient access to sites where such evidence, if it existed, would be available, may convince researchers that the evidence (and the entity or relationship related to it) does not exist.

Consider the US experience in Iraq post-2003. After U.S. military forces spent many months searching intensively for weapons of mass destruction and interviewing captured officials who had been in positions to know about any such weapons, nearly all observers concluded that Iraq never had them. The large development programs required would have left behind documentary and physical evidence.

In addition, process tracing helps address the limits of Mill's methods of comparison. Mill himself recognized that omitted variables and the possible presence of equifinality – that is, multiple paths to the same outcome – could threaten inferences based on comparison alone. Process tracing on omitted variables and residual differences (or similarities) between two cases can address the former problem. At the same time, process tracing can show that the explanation of interest was indeed likely to have been true, even if it represents only one of several possible paths to the outcome (George and Bennett, 2005: 153–160, 254).

Process tracing can also readily be combined with quantitative techniques in a mixed-method design. Building upon Lieberman's (2005, 2009) idea of nested analysis, for example, it can be applied to a few cases selected from a statistical analysis to clarify whether the direction of causal influence is indeed from the independent variable to the dependent variable, and not the reverse, and to help assess whether any observed correlations might be spurious.[5] In these ways, process tracing on the mechanisms hypothesized in statistical models can greatly increase the confidence in the causal significance of the correlations identified in them.

In a variation on the above, (quasi-) quantitative techniques such as agent-based modeling can be used to check the plausibility of inferences about causal mechanisms derived from process tracing. Consider contemporary research on civil war, where a central finding is

---

[5] Nested analysis is just one of several mixed-method designs where process tracing can play a central role. See Dunning, this volume.

that such conflicts are anything but 'civil': mechanisms of transnational diffusion play central roles. Scholars have now utilized process tracing to document a number of these mechanisms, including framing, resource mobilization and social learning (Checkel, 2013: *passim*).

Such findings can be strengthened through the careful application of agent-based modeling, where one assesses the plausibility of the mechanisms by using computer simulation. If the results of the simulations resemble the empirical patterns of conflict diffusion uncovered through process tracing, then the validity of the posited causal relation is strengthened (Nome and Weidmann, 2013; see also Hoffmann, 2008; and Hedstroem and Ylikoski, 2010: 62–63).

Standards of Good Process Tracing. With these definitional, philosophical and operational preliminaries in hand, we now return to the challenge highlighted in the chapter's opening pages. How do we know a particular piece of process tracing research is good process tracing? More colloquially, how would we recognize good process tracing if it were to walk through the door?

We argue for a three-part standard for what counts as a good application of process tracing (see also Bennett and Elman, 2007; Bennett, 2010; Checkel, 2008; and Checkel, 2013: chapter 1). *Meta-theoretically*, it will be grounded in a philosophical base that is ontologically consistent with mechanism-based understandings of social reality and methodologically plural. While we favoured scientific realism above, there is sufficient (and inevitable) uncertainty at this philosophical level to leave the door open for related approaches such as analytic eclecticism (Katzenstein and Sil, 2010) or pragmatism (Johnson, 2006).[6] *Contextually*, it will utilize this pluralism both to reconstruct carefully causal processes and not to lose sight of broader structural-discursive context. *Methodologically*, it will take equifinality seriously, which means to consider the alternative causal pathways through which the outcome of interest might have occurred.

With these three signposts in hand, we argue that good process tracing needs then to address ten specific criteria, the importance of which vary as a function of the question asked,

---

[6] We are not arguing there should be an explicit discussion of meta-theory for each empirical application of process tracing. Rather, at a minimum, there needs to be an appreciation that positivism is inadequate for dealing with concepts such as mechanism and techniques like process tracing. See also Pouliot, this volume. For a contrasting view, see Hall, 2012.

state of disciplinary knowledge, and the theories being tested. We start with four general criteria, and then discuss six that are especially relevant for process tracing.

1) Cast the net widely for alternative explanations.

As noted above, Bayesians rightly argue that explanations are more convincing to the extent that the evidence is inconsistent with alternative explanations. Failing to consider a potentially viable explanation that readily occurs to the readers and critics of a case study can make the process tracing unconvincing. The consequences of leaving out a viable explanation, which can lead not only to poor inferences but failed academic job interviews and rejections from peer-reviewed journals, are sufficiently serious that it is important to consider a wide range of alternatives despite the effort this entails.[7]

Specifically, researchers should at a minimum assess the process tracing evidence on the explanations that political scientists, sociologists, historians, regional specialists, and functional experts have offered for the specific case under study and for the class(es) of cases or phenomena of which it is an instance. In addition, it is often useful to render in theoretical terms and do process tracing upon the understandings of actor behavior offered by participants and journalists. Often these will overlap with scholars' explanations of the case, but occasionally they point to viable explanations that have been overlooked.

An additional criterion for assessing the adequacy of potential explanations is to ask whether any major theoretical categories of social explanation have been omitted. These include explanations based on actors' material power, institutional constraints and opportunities, and social norms or legitimacy (Mahoney, 2008). In our own subfield of international relations, these correspond with realist, neoliberal institutionalist, and constructivist theories. Another taxonomic dimension to check is whether both agent-based and structural explanations have been considered. Structural constraints can be material, institutional, or normative, for example, and agents can be motivated by rational calculations of material interests, cognitive biases, emotional drives, or normative concerns.

---

[7] Schimmelfennig, this volume, notes the tradeoff here between comprehensiveness and efficiency, and - compared to the present discussion - he puts more emphasis on the latter.

As process tracing often involves exploring what individuals knew when and how they behaved, there is a risk of overlooking normative or material structural contexts (see also Pouliot, this volume). For example, in earlier work, one of us used process tracing to explore the social-psychological factors that might lead decision-makers to change their minds in light of persuasive appeals (Checkel, 2003). Yet, as critics noted, the argument overlooked structural context, simply assuming that persuasive arguments were a function of individual-level dynamics alone. It was equally plausible, however, that the persuader's arguments were legitimated by the broader social discourse in which he/she was embedded. In positivist-empiricist terms, there was a potential problem of omitted variable bias, while, for interpretivists, the issue was one of missing the broader forces that enable and make possible human agency (Neumann, 2008).

Considering normative, material, power, efficiency, legitimacy, structural, and agent-based factors makes for a long list of potential alternative explanations, but not every one will merit or require detailed process tracing. Some may be irrelevant to the case at hand, and others may be easily dismissed in the footnotes on the basis of one or a few pieces of obvious or well-known process-tracing evidence. Yet the consequences of leaving out a viable explanation generally outweigh the efforts necessary to consider many potential ones.

2) Be equally tough on the alternative explanations.

Being equally tough on alternative explanations does not require going into equal depth in process tracing on every one of them. Some explanations may be quickly undermined by the evidence, while others will require deeper investigation. Some explanations may be more counterintuitive or, put another way, have a lower prior expectation of being true, and may thus require more evidence to convince ourselves and others even if initial process tracing evidence suggests they may be true. Some explanations may be more novel than others, and there may be more value-added in exploring these. There is also a tendency, and a justifiable one in the Bayesian view, to generate more detailed evidence on the explanations that appear to be more and more likely to be true as the evidence cumulates.

Research in cognitive science, however, reminds us of a common tendency toward confirmation bias, and one goal of methodology should be to counteract such biases. In this regard, fairness to alternative explanations requires that we fully consider evidence that fails to fit the explanations that interest us most, as well as evidence that fits explanations that initially interest or convince us the least. Some case studies accord an unduly privileged status to one explanation by granting it 'first mover advantage' (Caporaso, Checkel, Jupille, 2003). That is, they perform process tracing on this explanation and turn to evidence on the alternative explanations only to address the anomalies that confront the privileged first mover. A far better procedure is to outline the process tracing predictions of a wide range of alternative explanations of a case in advance, and then to consider the actual evidence for and against each explanation.

3) Consider the potential biases of evidentiary sources.

Because scholars have developed and explicated Bayesian logic largely with reference to the physical sciences, they have devoted relatively little attention to the question of how to judge, or discount, evidence provided by agents who have instrumental motives to convince observers that some explanations are stronger than others. Yet this is a pervasive problem in the social sciences, and Bayesianism offers a useful framework for addressing it. When those providing evidence may have instrumental motives for putting forth particular explanations of their own or others' behavior, researchers should apply a two-step Bayesian analysis. First, they should attach Bayesian priors to the possible instrumental motives of those providing evidence and weigh the evidence they give in light of those priors. Then, in a second step, researchers should use the evidence that sources provide to update prior expectations on these sources' motives, and use these updated priors in assessing subsequent evidence.

This sounds complex but in fact we make such judgments every day. Given A's possible motives, how much should I trust what he/she says? Given what he/she has said, what are A's likely motives? Social psychologists have long noted that audiences find an individual more convincing when that person espouses a view that is seemingly contrary to his/her instrumental goals. When Warren Buffet argues that wealthy Americans should pay more taxes, this is more convincing than when a person of moderate income argues for raising taxes on the rich. Bayesian

logic suggests this is a sensible procedure for accrediting or discounting evidence from individuals with potential instrumental goals for providing, distorting, or hiding evidence.

For similar reasons, researchers should follow established advice on considering issues of context and authorship in assessing evidence. Spontaneous statements have a different evidentiary status from prepared remarks. Public statements have a different evidentiary status from private statements or statements that will remain classified for a period of time. Statements in front of some audiences may reflect different instrumental purposes from those in front of other audiences. In addition to weighing such factors in judging what individuals say, write, or do, researchers should also consider the instrumental motivations that can lead to selection bias in which statements, documents, and other sources participants make accessible or available. Newly empowered actors in control of the archives are likely to make available only negative information about their opponents and positive information about themselves.

It is important to consider as well any potential selection biases in secondary sources. Historians are always at risk of selectively choosing the primary and secondary sources that confirm their arguments. For this reason, it is important to consider a wide range of secondary accounts representing contending historiographical schools and explanations (Lustick, 1996).

4) Take into account whether the case is most or least likely for alternative explanations.

Prior expectations on the strength and scope conditions of a theory require the most updating when it fails to explain a case in which it is most likely to apply, or succeeds in explaining a case in which it is least likely to apply. Process tracing can play an important role in ensuring that such cases are not flukes, and that the scope conditions of prior theories need to be revised. If, for example, a theory's failure in a most-likely case is caused by a variable that occurs only rarely or even only once, the scope conditions of the prior theory may need revision for only one or a few cases. However, if process tracing demonstrates that the prior theory failed due to a variable or interaction that is common, its scope conditions will require more radical revision.

5) Make a justifiable decision on when to start.

Process tracing requires a researcher to choose and justify a starting point for investigating evidence on alternative explanations. Do we begin process tracing on the Cuban Missile Crisis, for example, at the point when President Kennedy learned of the Soviet effort to deploy missiles in Cuba, with the Russian Revolution in 1917, or with the environmental context in which humans have evolved over the centuries? There is no universal answer to such questions, as a justifiable starting point depends on how a researcher defines the puzzle or question they are trying to explain: crisis decision-making, great power ideological rivalry, or the extent to which humans have genetic predispositions regarding conflictual and cooperative behavior.

Even within one well-defined research question, the proper starting point can be subject to debate. Just as any researcher's decision on how far 'down' to go in gathering detailed evidence can be critiqued for going too far or not far enough, the selection of the point in time at which to start process tracing can be critiqued for being too far back or too proximate. Robert Putnam's account of political differences between northern and southern Italy at the end of the twentieth century, for example, has been criticized for starting the explanatory story in the eleventh century, skipping over long periods of history, and downplaying or ignoring more historically proximate events that may have had powerful effects on regional politics (Putnam, 1993; Tarrow, 1996: 393).

Yet process tracing has to begin somewhere, and there are useful rules of thumb for deciding when to start. A reasonable place may be a critical juncture at which an institution or practice was contingent or open to alternative paths, and actors or exogenous events determined which path it would take. Path dependency theories suggest that institutions, once set on a particular path, often become locked in to that path by increasing returns, externalities, or other mechanisms (Pierson, 2000). A common critique of critical junctures is that they are identifiable only in retrospect, but process tracers have the luxury of always looking at them in retrospect.

Still, in choosing a critical juncture as a starting point for process tracing, researchers have to consider whether earlier and later ones might also be relevant (hence Tarrow's critique of Putnam), and they should also consider whether it is necessary to do process tracing on other potential but unrealized critical junctures before or after their chosen starting point (see also

Capoccia and Kelemen, 2007). These are the points at which institutions could have changed, perhaps due to some exogenous shock, but did not. Such potential junctures are subject to more conceptual and interpretive debate than the junctures that in fact led to institutional change. In general, to the extent that a researcher locates the starting point for process tracing in the distant past, it is important to show how institutions or practices could have reproduced themselves for long periods of time, even if resources and word limits do not allow continuous process tracing on the long period between the starting point and the outcome.

Another kind of starting point is the time at which a key actor or agent enters the scene or gains some material, ideational, relational, or informational capacity. This can be effective when alternative explanations hinge upon or work through the motivations, knowledge, and capacities of individual agents, and when particular agents behave differently, or with different effects, from their predecessors.

6) Be relentless in gathering diverse and relevant evidence, but make a justifiable decision on when to stop.

At times, social scientists using process tracing to assess alternative explanations of a case must be even more relentless than historians in tracking down primary sources or seeking interviews with participants. A single meeting or memo may prove to be the crucial piece of evidence that instantiates one explanation or undermines another. The more probative we expect evidence to be, the more time and effort we should be willing to expend to obtain it. Here, process tracers should use the Bayesian-inspired criteria above – smoking gun, doubly decisive, straw in wind, and hoop tests – to assess the potential probative value of data not yet obtained.

Furthermore, Bayesian logic indicates that process tracers should seek diverse and independent streams of evidence. If you want to know whether an animal is a duck, for example, instead of just looking at how it walks, you should also consider how it flies, sounds, looks, and so on. This insight is consistent with arguments in the social sciences concerning triangulation among diverse data sources With triangulation, a researcher cross-checks the causal inferences derived from his/her process tracing by drawing upon distinct data streams (interviews, media reports, documents, say).

Yet triangulation is not a panacea, as its successful use requires that the error term in each stream of evidence, on average, points in such a way that it cancels those in others. If all the streams of evidence are subject to the same selection bias, however, then errors can cumulate, making researchers unaware of this problem ever more confident in a false explanation (Symposium, 2007: 10; Kuehn and Rohlfing, 2009). Seemingly "diverse" sources of evidence (documents, media coverage, interviews) could actually all originate from one or a few individuals with instrumental reasons to convince observers of a particular explanation. We return below to the importance of considering such potential biases in evidentiary sources.

As it can demand both diverse and deep evidence, and may require significant "straw in the wind" evidence when the more definitive kind is not available, process tracing can be quite time consuming. Elizabeth Wood's excellent study of the Salvadoran civil war, for example, advances a rich, process-based argument that draws on an enormous amount of information. Yet, it was also fifteen years in the making (Wood, 2003, xi–xv; see also Petersen, this volume). Carefully executed process tracing thus requires that researchers think at an early point about their own financial limits and temporal constraints.

This point highlights the necessity of deciding when to stop gathering and analyzing evidence. There is no simple algorithm for deciding when to stop, and stopping at any point makes the researcher vulnerable to the possibility that just a little more research would have turned up evidence that would have greatly revised their estimate of the likely truth of alternative explanations. However, the Bayesian perspective discussed earlier offers a sensible logic here: One stops when repetition occurs. That is, a researcher should stop pursuing any one stream of evidence when it becomes so repetitive that gathering more of that same kind of evidence has a low probability of revising their estimate of the likely accuracy of alternative explanations.[8]

For each test in determining whether an animal is a duck – walk, sounds, *etc* – a small sample is sufficient. A thousand steps or quacks provide no more convincing evidence than a few. Yet in deciding when to stop, there is no escaping the de facto tradeoff between the risk of stopping too soon and making poor inferences, and the risk of stopping too late and wasting time,

---

[8] While using different language, ethnographers advance a strikingly similar decision rule. Gusterson, 2008.

effort, and resources on evidence that proves to have no effect on one's estimates on the verisimilitude of alternative explanations.

7) Combine process tracing with case comparisons when useful for the research goal and feasible.

Although some have argued that single case or "no variance" designs are weak (King, Keohane, and Verba, 1994), process tracing within single cases can in fact lead to convincing explanations if appropriate evidence is accessible. Moreover, if the explanations of these cases disprove claims of necessity or sufficiency, or if the cases are most likely for a theory that fails to explain them or least likely for an explanation that succeeds, their explanation can have more general implications for the veracity and scope conditions of contending theories. Yet for many inferential purposes, comparative case studies can be more powerful sources of inference than single case designs.

In a most-similar case comparison, in which two cases differ on one independent variable and on the dependent variable, process tracing can help establish that the one independent variable that differs is related through a convincing hypothesized causal process to the difference in the cases' outcomes. As most-similar cases rarely control for all but one potentially causal factor, process tracing can also establish that other differences between the cases do not account for the difference in their outcomes. Similarly, process tracing can help affirm that the one independent variable that is the same between two least-similar cases accounts for the similarity in their outcomes, and that similarities in other potentially causal factors do not explain the cases' common outcome.

An additional synergy between process tracing and case comparisons is that an explanation inductively derived from process tracing might lead the researcher to reconsider their case selection. If the close study of a case leads to the discovery of an omitted variable, adding this variable into the theoretical framework can change the definition of the relevant population of cases. Adding a variable can also change which cases are most-similar, least-similar, or deviant or anomalous, hence changing which cases are most useful to study for theory testing or theory development.

8) Be open to inductive insights.

One of the great advantages of process tracing is that it puts researchers at risk of stumbling upon many potentially causal factors, evident in the details and sequences of events within a case, which they had not anticipated on the basis of their prior alternative hypotheses. Encountering such surprises provides opportunities to re-think the prior explanations of the case. It may be possible to revise these prior explanations in trivial ways to accommodate unexpected facts, or it may prove necessary to build new explanations or link surprising facts to extant theories that the researcher had not previously thought would apply to the case. In any event, it is important to pay attention to the feeling of surprise and to follow it up with efforts to explain surprising facts theoretically.

9) Use deduction to ask "if my explanation is true, what will be the specific process leading to the outcome?"

Prior to embarking on process tracing, researchers should clarify as much as possible the facts and sequences within a case that should be true if each of the alternative hypothesized explanations of the case is true. Which actors should have known, said, and did what, and when? Who should have interacted with, worried about, or allied with whom? We cannot stress enough that theories are usually stated in very general terms; they must therefore be operationalized and adapted to the specific processes predicted in particular cases (see also the discussion in Jacobs, this volume).

For new explanations inductively derived from the evidence within a case, it is doubly important to forestall any confirmation bias by considering what other observable implications must be true if the new explanation is true. As noted above, these observable implications may be in other cases, but they could also be within the case from which the new theory was derived as long as they are independent from the evidence that gave rise to it. Either way, if additional observable implications can be derived from the new explanation and tested against new evidence, this can provide a check against confirmation bias.

10) Remember that conclusive process tracing is good, but not all good process tracing is conclusive.

The more continuous a narrative explanation of a case, and the closer the evidence fits some explanations and not others, the more confidence we can have in explanatory inferences based on process tracing (but see also Schimmelfennig, this volume). There may well be temporal or spatial gaps in the evidence bearing on hypothesized processes, however, such as documents that have been destroyed or remain classified, or participants who are unwilling or unable to submit to interviews. In addition, in some case studies the available evidence may be equally consistent with two or more hypotheses that offer incompatible explanations of the case. When the evidence does not allow high levels of confidence in supporting some hypotheses and discounting others, it is important to acknowledge the level of uncertainty that remains.

As the foregoing indicates, good process tracing builds in part on standard injunctions and checks that are applicable to an array of qualitative methods (criteria 1–4). These include attention to research design and potential biases in evidentiary sources, as well as caution in the application of triangulation. At the same time, it demands adherence to additional best practice standards – criteria 5 through 10 – that address problems related to testing inductively-generated insights in ways that reduce the risks of "curve-fitting."

For sure, not all ten criteria elaborated above may be relevant for any given study. However, they *should* serve as a starting point and checklist, thus maximizing the likelihood of conducting good process tracing. So defined and operationalized, the technique is far more than a temporal sequencing of events or mere "detective work" based on hunches and intuition (Gerring 2007a, 178).

V. Process Tracing on General Categories of Theories

The kinds of process tracing evidence that are relevant and the veracity, accessibility, and biases of that evidence are often specific to the explanations a researcher is considering and the cases they have chosen to study. The extent to which alternative explanations are mutually exclusive or complementary also varies greatly depending on the explanations and cases studied. Nonetheless, useful general observations can be made about the kinds of process tracing opportunities and challenges that arise with different general modes of explanation common in

the social sciences. Here, we consider process tracing on rational choice, cognitive, material/structural, normative/structural, and institutional/functional-efficiency theories.

Rational choice theories argue that actors have complete and transitive preferences and that they choose courses of behavior that maximize the expected value of likely outcomes given the information available to them. Some early theorists made 'as-if' assumptions, or assumptions that obviated the need for process tracing by arguing that it was unnecessary to show that actors actually made rational calculations so long as outcomes arose as if actors had done so. So-called thick rational choice approaches make further assumptions by presuming that actors have certain preferences, such as for gains in material resources or power. Rational choice theorists, however, have increasingly been willing to eschew as-if assumptions and engage in process tracing. That is, they accept the challenge of making only thin assumptions to the effect that actors decide through rational processes. They then seek to discover actor's preferences by observation, demonstrating empirically that actors actually do make calculations and choices through rational processes to maximize their preferences (Bates, *et al*, 1998; Schimmelfennig, 2003; see also Checkel, this volume).

This raises several challenges for process tracing. For one, there is the revealed preference problem. How can we infer actors' real preferences, given that they are often engaged in strategic contexts that provide incentives to misrepresent those preferences? In addition, how can we avoid circularity or tautology by inferring preferences separately from the behavioral choices that these preferences are supposed to explain? There is a danger that – no matter what the outcome – the researcher can change his/her measurement of the actors' preferences so that the chosen outcome was a value-maximizing one.

In view of these challenges, the only option is to infer preferences from an actor's earlier rhetoric and actions and use these preferences to explain subsequent behavior, while also investigating the possibility that preferences may change over time through learning, changing life cycle stages/aging, or other processes. In particular, if actors engage in costly signaling – rhetoric or actions that impose high political or material costs if preferences are not consistent with these statements or acts – this may be taken as a relatively reliable indicator of preferences. A good example is David Laitin's study of how Russian speakers in the non-Russian former

Soviet Republics chose between teaching their children the titular language of the country in which they resided (such as Latvian) or their native Russian. Laitin convincingly uses statements from the individuals making these choices, as well as aggregate data, to show that they conceived of them as involving a tradeoff between passing along to the children an exclusive focus on their Russian language and heritage or maximizing the children's' employment opportunities. Those who chose to have their children invest in the newly-dominant local language revealed their preference for employability over linguistic heritage (Laitin, 1998).

Even if preferences can be reliably inferred, however, rational choice arguments face a second hurdle in demonstrating that decision processes maximized actors' expected utilities given their preferences and the information at their disposal. This makes it very important to establish the information actors had and when they had it. This stage of rational choice explanations is often tested through process tracing against the alternative explanation that actors' decisions are influenced by cognitive errors and biases. David Lake, for example, uses process tracing to compare a rational choice approach, in this case a bargaining theory model, and an "error and bias" explanation of U.S. decisionmaking on the 2003 intervention in Iraq. Lake concludes that Iraqi leaders failed to consider readily available costly signals of American resolve and American leaders ignored ample evidence on the likely costs of war, so that "misrepresentation by the other side was far less of a problem than self-delusion" (Lake, 2010/2011: 45).

If rational choice explanations face a revealed preference problem, cognitive theories face a problem of accurately inferring revealed beliefs. Actors may have instrumental reasons, such as an interest in winning political support from groups or individuals, for publicly espousing ideas that they do not actually believe. One option here is to compare actors' public statements with available private deliberations that they expected would not be revealed for some time. Yuen Foong Khong, for example, compares the analogies American leaders used in public to justify the Vietnam War with those they used in private policy discussions that were de-classified many years later, checking to see if actors chose the same analogies in both settings. He concludes that with the exception of the analogy to France's disastrous experience in Vietnam, which was used only in private, they did so (Khong, 1992: 60-61; see also Jacobs, this volume).

Actors may also make statements authored by their staffs or pushed upon them by powerful individuals or groups, so it is important to establish the provenance and authorship of public statements, and to give spontaneous and unplanned statements more weight than planned ones as indicators of genuine beliefs. Also, stated beliefs that incur substantial audience costs are more likely to reflect genuine beliefs, and recollections of beliefs held in the past that are backed up by documentary evidence are more credible than those lacking such supporting evidence. In addition, research by social psychologists shows that the recall of past beliefs is likely to be more accurate the more intense was the social context surrounding their creation (Wood, 2003: 33–34). Finally, we should expect evidence that an actor holds socially stigmatized beliefs to be harder to find than evidence that an actor shares widely accepted beliefs, so we should treat absence of evidence on the former differently from absence of evidence on the latter.[9]

Theories that emphasize material power and structure require that actors be aware of power differentials and that they circumscribe their behavior when faced with more powerful opponents. This raises several process tracing challenges. First, actors engaged in strategic interaction may have incentives to either exaggerate their capabilities (to bluff) or to understate them (to preserve the option of surprising adversaries). The same applies to actors' publicly-stated assessments of other actors' power capabilities. This can create conflicting evidence on actors' estimates of their own and others' material power.

Second, power is often strongest as an explanation when it has a taken-for-granted quality. It may successfully deter actors from publicly discussing or even contemplating possible courses of action.[10] This makes it difficult to distinguish whether an actor was deterred from doing something or never had an interest in doing it in the first place. It also means that exceptions to power explanations – cases in which actors thought their higher level of commitment would enable them to prevail over better-endowed adversaries – are more evident and easier to document because these situations become overt conflicts, such as wars, labor strikes, or attempted revolutions, rather than being non-events.

---

[9] On all these points, see Jacobs' chapter below. Jacobs, this volume.
[10] It is only a short step from this understanding to what interpretive theorists call productive power, or power that is constitutive of agent interests and identities. Barnett and Duvall, 2005; see also Pouliot, this volume.

Nonetheless, it is possible to use process tracing to assess power explanations by paying careful attention to sequencing and to what information actors had and when they had it. For example, scholars have offered two competing explanations of the 1898 Fashoda crisis between Great Britain and France over control of the headwaters region of the Nile River. Christopher Layne gives a straightforward power explanation: France backed down because Britain had superior military power (Layne, 1994: 28–33). Kenneth Schultz advances an explanation based not only on power differentials, but also on the ability of democracies to credibly commit to use military power when both ruling and opposition parties support this stance (Schultz, 2001: 175–95). Schultz's process tracing makes his explanation more convincing because it explains the puzzle of why a weaker France challenged Britain in the first place, when Britain's resolve was unclear, and it demonstrates that France backed down precisely when Britain's democratic institutions made its threat to use force credible and France's democracy laid bare the political divisions that undermined its resolve.

Like material structure arguments, theories about norms – a form of social structure – need to show that norms prevented actors from doing things they otherwise would have done. A good example is Nina Tannenwald's research on the non-use of nuclear weapons since the bombings of Hiroshima and Nagasaki. To show that normative constraints explain this outcome, Tannenwald has to demonstrate that norms against the use of nuclear weapons – rather than their limited battlefield utility – explain their non-use. Accordingly, Tannenwald provides direct, process-tracing evidence that American presidents and their advisors chafed at the perceived normative constraint of the American public's revulsion at the idea of using nuclear weapons after the effects of nuclear fallout became more widely known. She also demonstrates that these same leaders often avoided even officially considering the option of using nuclear weapons for fear that their deliberations would be leaked to the public (Tannenwald, 2007).

Finally, institutional explanations that rely on functional efficiency and transaction costs must be able to demonstrate through process tracing how such costs affect compliance. A good example is Ron Mitchell's study of international environmental cooperation over the sea. In particular, he uses process tracing to demonstrate that the international regime to prevent the dumping of oil residue from tankers failed to reach high compliance levels due to high

transactions costs. In contrast, the international regime to force oil tanker owners to install expensive anti-pollution equipment succeeded in motivating high compliance because it made non-compliance transparent and provided low-cost means of sanctioning tanker owners who failed to comply (Mitchell, 1994).

VI. Conclusion and Preview

This introduction and the chapters that follow seek to consolidate the turn to process and mechanisms in contemporary social science. We do this not by exploring new substantive empirical domains or developing novel theories; rather, we focus on the prior, operational and methodological issue of how we come to know when studying process. As has been argued elsewhere, the central challenge here is to avoid "lazy mechanism-based storytelling" (Hedstroem and Ylikoski, 2010: 58, 64). We could not agree more, and hope that our conceptual parsing and applied discussions of process tracing strike readers as anything but lazy.

Preview. The volume has three parts. Part I is comprised of this introductory essay. It historicizes the term process tracing, grounds it philosophically, and advances specific criteria for distinguishing good process tracing from bad.

The six chapters in Part II are the manuscript's core, assessing the contributions of process tracing in particular research programs or bodies of theory, including ideational theory, work on international institutions, the European Union, the comparative politics subfield, the end of the Cold War, and the literature on conflict processes. These chapters are resolutely applied – connecting method to practice – with recognized experts assessing the strengths and weaknesses of process tracing as applied to particular substantive domains. They include process tracing for inductive/theory-development (ch.7 – Petersen) and deductive/theory-testing purposes (ch.4 – Schimmelfennig); micro-level process tracing on cognitive theories (ch.2 – Jacobs); macro-level process tracing on structural theories (ch.5 – Waldner); process tracing on the interplay of individuals and institutions (ch.3 – Checkel); interpretive "practice tracing" (ch.9 – Pouliot); and process tracing that focuses on explaining key historical cases (ch.6 – Evangelista).

Whatever the application and type of process tracing, all essays address the standards articulated in the present chapter and work from a common template of questions – what are the best examples of process tracing in your subfield; the evidentiary and interpretive matters relevant to the topics you research; the process tracing issues specific to the kind of theories on which you have focused – to assess critically the good and bad when process tracing is applied. Collectively, the analyses highlight issues of data quality, the role of hypothesized causal mechanisms, time and resource constraints, research ethics, multi-method strategies where process tracing is one technique in play, and theory development, among others.

In addition, contributors are asked to address policy, by making a comparison between theories and arguments that are structural in nature and those stressing mechanisms, process tracing and dynamics. Many argue that the latter – the focus of this volume – are better at capturing the complex world with which policymakers must deal. Is this correct? To paraphrase the late Alexander George, is the result better 'theory for policy' (George, 1993)?

Our answer is a cautious 'yes.' Process tracing conducted on the observable implications of mechanisms gives decisionmakers new insight on a range of potentially 'manipulable' factors. Yet, this down-in-the trenches research should be done smartly, which is to say it must simultaneously be attentive to broader environmental, ethical and structural contexts (see also Chapter 10, this volume).[11]

In Part III, we step back and – in three separate essays – explore the research frontier. Chapter 8 makes explicit a theme touched upon in several earlier contributions – the relation of process tracing to quantitative methods – and does so by highlighting the key role it can and should play in multi-method research. If this analysis bridges different methodological traditions, then Chapter 9 goes a step further, examining the role of process tracing in interpretive social science. It explores the gap that separates positivist and post-positivist understandings of the technique, and argues that an engagement around the concept of practice can minimize the meta-theoretical challenges involved in bridging such a divide.

---

[11] More formally, our intent here is to push contributors to think counterfactually: Absent process tracing and mechanisms, how and in what ways would policy recommendations in their particular area of study differ?

In Chapter 10, the co-editors revisit the standards of good process tracing, synthesize and critique the volume as a whole, and outline an agenda for future development of and research on process tracing. In particular, proponents of process tracing need to remember that method is not an end in itself; rather, it is a tool helping us build and test theory. And the latter remains a central challenge for process tracers – how to combine the technique's use with the development of cumulable social science theory (see also Hedstroem and Ylikoski, 2010: 61–62). Moreover, process tracing is only one way to capture process. Future work thus needs to integrate this volume's findings with insights gleaned from statistical approaches, agent-based modeling exercises, and applications of discourse analysis, among others.

References

Adler, Emanuel. 2002. "Constructivism and International Relations." In Walter Carlsnaes, Thomas Risse, Beth Simmons, Editors. *Handbook of International Relations*. London: Sage Publications.

Autesserre, Severine. 2009. "Hobbes and the Congo: Frames, Local Violence and International Intervention." *International Organization* 63/2: 249–80.

Bakke, Kristin. 2013. "Copying and Learning from Outsiders? Assessing Diffusion from Transnational Insurgents in the Chechen Wars." In Jeffrey T. Checkel, Editor. *Transnational Dynamics of Civil War*. Cambridge: Cambridge University Press.

Barnett, Michael and Raymond Duvall. 2005. "Power in International Politics." *International Organization* 59/1: 39–75.

Bates, Robert, Barry Weingast, Avner Greif, Margaret Levi, and Jean-Laurent Rosenthal. 1998. *Analytic Narratives*. Princeton, NJ: Princeton University Press.

Bennett, Andrew and Colin Elman. 2007. "Case Study Methods in the International Relations Subfield." *Comparative Political Studies* 40/2: 170–95.

Bennett, Andrew. 2008. "Process Tracing: A Bayesian Approach." In Janet Box-Steffensmeier, Henry Brady and David Collier, Editors, *Oxford Handbook of Political Methodology*. Oxford, UK: Oxford University Press.

-------. 2010. "Process Tracing and Causal Inference." In Henry Brady and David Collier, Editors. *Rethinking Social Inquiry: Diverse Tools, Shared Standards*, 2nd Edition. Lanham, Maryland: Rowman and Littlefield Publishers.

Brady, Henry and David Collier. Editors. 2010. *Rethinking Social Inquiry: Diverse Tools, Shared Standards*, 2nd Edition. Lanham, Maryland: Rowman and Littlefield Publishers.

Capoccia, Giovanni and R. Daniel Kelemen. 2007. "The Study of Critical Junctures: Theory, Narrative, and Counterfactuals in Historical Institutionalism." *World Politics* 59/3: 341–369.

Caporaso, James, Jeffrey T. Checkel and Joseph Jupille. 2003. "Integrating Institutions: Rationalism, Constructivism and the Study of the European Union – Introduction." *Comparative Political Studies* 36/1–2: 7–41.

Checkel, Jeffrey T. 2003. "'Going Native' in Europe? Theorizing Social Interaction in European Institutions." *Comparative Political Studies* 36/1–2: 209–31.

-------. 2008. "Process Tracing." In Audie Klotz and Deepa Prakash, Editors. *Qualitative Methods in International Relations: A Pluralist Guide*. NY: Palgrave Macmillan.

Checkel, Jeffrey T. Editor. 2007. *International Institutions and Socialization in Europe*. Cambridge: Cambridge University Press.

-------. 2013. *Transnational Dynamics of Civil War*. Cambridge: Cambridge University Press.

Chernoff, Fred. 2002. "Scientific Realism as a Meta-Theory of International Relations." *International Studies Quarterly* 46/2: 189–207.

Collier, David. 2011. "Understanding Process Tracing." *PS: Political Science and Politics* 44/4: 823–30.

Doyle, Sir Arthur Conan. 1930. "The Adventure of the Blanched Soldier." In Doyle, *The Case-Book of Sherlock Holmes*. (Published in *The Complete Sherlock Holmes* Vol. 2. New York: Barnes and Noble Classics, 2003, pp. 528–29).

Earman, John. 1992. *Bayes or Bust: A Critical Examination of Bayesian Confirmation Theory*. Cambridge, MA: MIT Press.

Elster, Jon. 1998. "A Plea for Mechanisms." In Peter Hedstroem and Richard Swedberg, Editors, *Social Mechanisms: An Analytical Approach to Social Theory*. Cambridge: Cambridge University Press.

Finnemore, Martha. 1996. *National Interests in International Society*. Ithaca, NY: Cornell University Press.

George, Alexander L. 1979. "Case Studies and Theory Development: The Method of Structured, Focused Comparison." In Paul Gordon Lauren, Editor, *Diplomatic History: New* Approaches. NY: Free Press.

-------. 1993. *Bridging the Gap: Theory and Practice in Foreign Policy*. Washington, DC: USIP Press Books.

George, Alexander L. and Andrew Bennett. 2005. *Case Studies and Theory Development in the Social Sciences*. Cambridge, MA: MIT Press.

Gerring, John. 2007a. *Case Study Research: Principles and Practices.* NY: Cambridge University Press.

-------. 2007b. "Review Article: The Mechanismic Worldview – Thinking Inside the Box." *British Journal of Political Science* 38/1: 161–79.

Gusterson, Hugh. 2008. "Ethnographic Research." In Audie Klotz and Deepa Prakash, Editors. *Qualitative Methods in International Relations: A Pluralist Guide*. NY: Palgrave Macmillan.

Hall, Peter. 2012. "Symposium: Tracing the Progress of Process Tracing." *European Political Science* 11/3.

Hansen, Lene. 2006. *Security as Practice: Discourse Analysis and the Bosnian War*. London: Routledge.

Hedstroem, Peter and Petri Ylikoski. 2010. "Causal Mechanisms in the Social Sciences." *Annual Review of Sociology* 36: 49–67.

Hellman, Geoffrey. 1997. "Bayes and Beyond." *Philosophy of Science* 64/2: 191–221.

Hobarth, Robin. 1972. *Process Tracing in Clinical Judgment: An Analytical Approach*. PhD Thesis. Chicago: University of Chicago.

Hoffmann, Matthew. 2008. "Agent-Based Modeling." In Audie Klotz and Deepa Prakash, Editors. *Qualitative Methods in International Relations: A Pluralist Guide*. NY: Palgrave Macmillan.

Hopf, Ted. 2002. *Social Construction of International Politics: Identities and Foreign Policies, Moscow, 1955 and 1999*. Ithaca, NY: Cornell University Press.

-------. 2007. "The Limits of Interpreting Evidence." In Richard Ned Lebow and Mark Lichbach, Editors. *Theory and Evidence in Comparative Politics and International Relations*. NY: Palgrave Macmillan.

-------. 2012. *Reconstructing the Cold War*. NY: Oxford University Press.

Johnson, James. 2006. "Consequences of Positivism: A Pragmatist Assessment." *Comparative Political Studies* 39/2: 224–52.

Katzenstein, Peter and Rudra Sil. 2010. *Beyond Paradigms: Analytic Eclecticism in World Politics*. NY: Palgrave Macmillan.

Khong, Yuen Foong. 1992. *Analogies at War: Korea, Munich, Dien Bien Phu, and the Vietnam Decisions of 1965*. Princeton, NJ: Princeton University Press.

King, Gary, Robert Keohane and Sidney Verba. 1994. *Designing Social Inquiry: Scientific Inference in Qualitative Research*. Princeton, NJ: Princeton University Press.

Kuehn, David and Ingo Rohlfing. 2009. "Does It, Really? Measurement Error and Omitted Variables in Multi-Method Research." *Qualitative Methods: Newsletter of the American Political Science Association Organized Section for Qualitative and Multi-Method Research* 7/2: 18–22.

Laitin, David. 1998. *Identity in Formation: The Russian-speaking Populations in the Near Abroad*. Ithaca, NY: Cornell University Press.

Lakatos, Imre. 1970. "Falsification and the Methodology of Scientific Research Programs." In Imre Lakatos and A. Musgrave, Editors. *Criticism and the Growth of Knowledge*. Cambridge: Cambridge University Press.

Lake, David. 2010/2011. "Two Cheers for Bargaining Theory: Assessing Rationalist Explanations for the Iraq War." *International Security* 35/3: 7–52.

Layne, Christopher. 1994. "Kant or Cant: The Myth of the Democratic Peace." *International Security* 19/2: 5–49.

Lieberman, Evan. 2005. "Nested Analysis as a Mixed-Method Strategy for Comparative Research." *American Political Science Review* 99/3: 435–452.

-------. 2009. *Boundaries of Contagion: How Ethnic Politics Have Shaped Government Responses to AIDS*. Princeton, NJ: Princeton University Press.

Lustick, Ian. 1996. "History, Historiography, and Political Science: Multiple Historical Records and the Problem of Selection Bias." *American Political Science Review* 90/3: 605–618.

Mahoney, James. 2001. "Review – Beyond Correlational Analysis: Recent Innovations in Theory and Method." *Sociological Forum* 16/3: 575–93.

-------. 2008. "Path Dependence in Historical Sociology." *Theory and Society* 29/4: 507–548.

------ 2012. "The Logic of Process Tracing Tests in the Social Sciences." *Sociological Methods and Research* Online First Edition (published 2 March): 1–28.

Mayntz, Renate. 2004. "Mechanisms in the Analysis of Macro-Social Phenomena." *Philosophy of the Social Sciences* 34/2: 237–59.

Mitchell, Ron. 1994. "Regime Design Matters: Intentional Oil Pollution and Treaty Compliance." *International Organization* 48/3: 425–458.

Neumann, Iver. 2008. "Discourse Analysis." In Audie Klotz and Deepa Prakash, Editors. *Qualitative Methods in International Relations: A Pluralist Guide*. NY: Palgrave Macmillan.

Nome, Martin and Nils Weidmann. 2013. "Conflict Diffusion via Social Identities: Entrepreneurship and Adaptation." In Jeffrey T. Checkel, Editor. *Transnational Dynamics of Civil War*. Cambridge: Cambridge University Press.

Norkus, Zenonas. 2005. "Mechanisms as Miracle Makers? The Rise and Inconsistencies of the 'Mechanismic Approach' in Social Science and History." *History and Theory* 44/3: 348–372.

Pierson, Paul. 2000. "Increasing Returns, Path Dependence, and the Study of Politics." *American Political Science Review* 94/2: 251-267.

Pouliot, Vincent. 2007. "'Sobjectivism': Toward a Constructivist Methodology." *International Studies Quarterly* 51/2: 359–384.

-------. 2010. *International Security in Practice: The Politics of NATO-Russia Diplomacy*. NY: Cambridge University Press.

Price, Richard. 1997. *The Chemical Weapons Taboo*. Ithaca, NY: Cornell University Press.

Putnam, Robert, with Robert Leonardi and Raffaella Nanetti. 1993. *Making Democracy Work: Civic Traditions in Modern Italy*. Princeton, NJ: Princeton University Press.

Risse, Thomas, Stephen Ropp and Kathryn Sikkink. Editors. 1999. *The Power of Human Rights: International Norms and Domestic Change*. NY: Cambridge University Press.

Roberts, Clayton. 1996. *The Logic of Historical Explanation*. University Park, PA: Pennsylvania State University Press.

Salmon, Wesley. 1990. *Four Decades of Scientific Explanation*. Pittsburgh, PA: University of Pittsburgh Press.

Schimmelfennig, Frank. 2003. *The EU, NATO and the Integration of Europe: Rules and Rhetoric*. Cambridge: Cambridge University Press.

Schultz, Kenneth. 2001. *Democracy and Coercive Diplomacy*. Cambridge: Cambridge University Press.

Symposium. 2007. "Symposium: Multi-Method Work - Dispatches from the Front Lines." *Qualitative Methods: Newsletter of the American Political Science Association Organized Section on Qualitative Methods* 5/1: 9–27.

Tannenwald, Nina. 2007. *The Nuclear Taboo: The United States and the Non-use of Nuclear Weapons since 1945*. Cambridge: Cambridge University Press.

Tarrow, Sidney. 1996. "Making Social Science Work Across Space and Time: A Critical Reflection on Robert Putnam's *Making Democracy Work*." *American Political Science Review* 90/2: 389–397.

Van Evera, Stephen. 1997. *Guide to Methods for Students of Political Science*. Ithaca, NY: Cornell University Press.

Waldner, David. 2011. "Process Tracing, Its Promise, and Its Problems." Paper presented at the Research Group on Qualitative and Multi-Method Analysis, Syracuse University, Syracuse, NY (June).

-------. ND. "Process Tracing and Causal Mechanisms." In H. Kincaid, Editor, *Oxford Handbook of the Philosophy of Social Science*. Oxford, UK: Oxford University Press (forthcoming).

Wendt, Alexander. 1987. "The Agent-Structure Problem in International Relations Theory." *International Organization* 41/3: 335–70.

-------. 1999. *Social Theory of International Politics*. Cambridge: Cambridge University Press.

Wight, Colin. 2002. "Philosophy of Science and International Relations." In Walter Carlsnaes, Thomas Risse, Beth Simmons, Editors. *Handbook of International Relations*. London: Sage Publications.

-------. 2006. *Agents, Structures and International Relations: Politics as Ontology*. Cambridge: Cambridge University Press.

Wood, Elisabeth Jean. 2003. *Insurgent Collective Action and Civil War in El Salvador*. NY: Cambridge University Press.