
Process-Tracing Methods in Social Science FREE

Derek Beach, Department of Political Science, Aarhus University

<https://doi.org/10.1093/acrefore/9780190228637.013.176>

Published online: 25 January 2017

Summary

Process tracing is a research method for tracing causal mechanisms using detailed, within-case empirical analysis of how a causal process plays out in an actual case. Process tracing can be used both for case studies that aim to gain a greater understanding of the causal dynamics that produced the outcome of a particular historical case and to shed light on generalizable causal mechanisms linking causes and outcomes within a population of causally similar cases. This article breaks down process tracing as a method into its three core components: theorization about causal mechanisms linking causes and outcomes; the analysis of the observable empirical manifestations of the operation of theorized mechanisms; and the complementary use of comparative methods to enable generalizations of findings from single case studies to other causally similar cases. Three distinct variants of process tracing are developed, illustrated by examples from the literature.

Keywords: process tracing, case studies, causal inference, causal mechanisms, theory building, theory testing

The Analytical Core of Process Tracing

Process tracing is a research method for tracing causal mechanisms using detailed, within-case empirical analysis of how a causal process plays out in an actual case. Process tracing can be used both for case studies that aim to gain a greater understanding of the causal dynamics that produced the outcome of a particular historical case and to shed light on generalizable causal mechanisms linking causes and outcomes within a population of causally similar cases.

The analytical added value of process tracing is that it enables strong causal inferences to be made about how causal processes work in real-world cases based on studying within-case mechanistic evidence. But process tracing is a single-case method, meaning that only inferences about the operation of the mechanism *within* the studied case are possible because this is the evidence gathered through tracing the process in the case. Therefore, to generalize beyond the studied case, we need to couple process-tracing case studies with comparative methods to enable us to generalize about causal processes. Comparisons across cases make generalization possible because we can then claim that, as a set of other cases are causally similar to the studied one, we should expect similar mechanisms to also be operative in these cases.

Process tracing can be used for either theory-building or theory-testing purposes. In the first, the researcher engages in both a thorough “soaking and probing” of the empirics of the case and a far-reaching search in the theoretical literature to gain clues about potential mechanisms that could link a cause and outcome together, whereas in the latter, hypotheses about the observable manifestations that a theorized mechanism might leave are tested empirically in a case.

Process tracing as a method can be broken down into its three core components: theorization about causal mechanisms linking causes and outcomes; the analysis of the observable empirical manifestations of theorized mechanisms; and the complementary use of comparative methods to enable generalizations of findings from single case studies to other causally similar cases.

What We Are Tracing—Causal Mechanisms

In theory-guided social science research, the ambition is to use causal theories to explain why something occurs either in a particular case or more broadly across a population of causally similar cases. This focus on causal explanations means that process tracing involves more than the production of detailed, descriptive narratives of the events between the occurrence of a purported cause and an outcome. Instead, process-tracing research probes the theoretical causal mechanisms *linking* causes and outcomes together.

Yet causal mechanisms are one of the most widely used but least understood types of causal claims in the social sciences (e.g., Brady, 2008; Gerring, 2010; Hedström & Ylikoski, 2010; Waldner, 2012). The essence of making a mechanism-based claim is that we shift the analytical focus from causes and outcomes to the hypothesized causal process *in-between* them. That is, mechanisms are *not* causes, but are causal processes that are triggered by causes and that link them with outcomes in a productive relationship. However, beyond this core point, there is disagreement about the nature of mechanisms. There are (at least) three distinct takes on the nature of mechanisms, each of which imply different research designs. The result of this ambiguity about the nature of mechanisms is that there is considerable confusion in the methodological literature about what process-tracing methods actually are tracing, and how we know good process tracing when we see it in practice.

Some scholars view mechanisms as a form of *intervening variable* (Gerring, 2007; Weller & Barnes, 2015; King, Keohane, & Verba, 1994). However, to make causal inferences about the effects of intervening variables (IV) requires empirical evidence in the form of variation *across* cases, measuring the difference that changes in the value of the IV have for values of an outcome, with all other things held equal (Runhardt, 2015; Woodward, 2003). In this understanding of mechanisms, there is no such thing as within-case analysis. The difference that an IV makes can either be assessed through some form of experimental design (either an actual manipulated experiment or a natural or logical counterfactual experiment) (Runhardt, 2015), or by creating “variation” by disaggregating single cases into multiple “cases” (either spatially, temporally, or substantively) to assess whether there is evidence of difference-making by the IV (King, Keohane, & Verba, 1994, pp. 219–228).

There are several challenges related to the experimental route. While actual experiments are possible, they are difficult (if not impossible) for many of the research questions in which case-study scholars are interested. Natural experiments are difficult to utilize because it is almost impossible to find two cases that are completely similar in all aspects except for the presence of the IV. Regarding logical counterfactuals, they face the critical challenge of being speculative “what ifs” without any actual empirical evidence backing them. Despite many attempts to build a methodology for logical counterfactuals (Goertz & Levy, 2007; Tetlock & Belkin, 1996; Lebow, 2000; Levy, 2015; Fearon, 1991), there are no objective empirical truth conditions for assessing a nonexistent “possible” alternative world (Beach & Pedersen, 2016a). And even more fundamentally, any form of experiment still does not tell us how the IV produces an effect, only that it does (Illari, 2011; Dowe, 2011; Bogen, 2005, p. 415). Yet understanding how a causal process works is the very reason we decide to study a mechanism in the first place.

The “one-into-many” strategy suggested by King, Keohane and Verba for producing evidence of the difference that the IV makes transforms the within-case tracing of a causal process into a co-variational analysis of patterns across subunits of the original case. Two assumptions have to hold to be able to make inferences about the IV based on evidence of difference-making within these subunits: (1) All of the subunits are causally homogeneous, and (2) the subunits are independent of each other. However, when we disaggregate single cases into many cases, these two assumptions typically never hold. For example, if we are studying a negotiation and disaggregate it into issue areas, there will be some issues that are more important for actors than others, meaning that there will be very different causal relationships playing out in these “high politics” areas than in other “low politics” issues. Additionally, assuming independence of these subunits is also highly problematic. For example, treating a before/after comparison as two independent cases is highly problematic because what happens previously in a case (e.g., failure to reach agreement) naturally impacts what happens later in the case (e.g., parties reach an agreement). An even more fundamental problem is that the goal of process tracing is to trace the workings of causal mechanisms as they operate *within* a case; shifting the analysis to another level (subunits) basically means that one is studying something different than what was intended. Indeed, one can argue that assessing mean causal effects of the IV transforms the analysis into a form of variance-based *comparative* case study. Given all of these problems, the rest of the article does not include variance-based understandings of process tracing that discuss assessing the difference that values of an IV make for an outcome.

When using process tracing, many scholars claim that we do not need to utilize evidence of difference-making across cases (or subunits of cases) in order to make inferences (George & Bennett, 2005; Collier, Brady, & Seawright, 2010; Mahoney, 2012). Instead, we can use observational within-case empirical material left by the workings of a causal mechanism within an actual case to make inferences about the existence of a mechanism in a case, a form of evidence that is termed “mechanistic evidence” in recent work in the philosophy of science (Russo & Williamson, 2007; Illari, 2011). The term “mechanistic evidence” is a more precise formulation of the type of evidence used in process tracing than the widely used term “causal process observation,” which is a broader term that refers to information about *both* context and mechanisms, defined as “... an insight or piece of data that provides information about the

context or mechanism and contributes a different kind of leverage in causal inference. It does not necessarily do so as part of a larger, systematized array of observations ..." (Collier, Brady, & Seawright, 2010, pp. 184–185).

Among case-study scholars who attempt to trace *within-case* causal processes using mechanistic evidence, two distinct takes on mechanisms can be identified in the literature: minimalist and systems understandings of mechanisms. In *minimalist understandings*, the causal arrow between a cause and outcome is not unpacked in any detail, either empirically or theoretically. Instead, within-case evidence, also sometimes called "diagnostic evidence" (Bennett & Checkel, 2014), is produced by asking "if causal mechanism M exists, what observables would it leave in a case?" However, the mechanism is not unpacked theoretically in any detail, meaning that our mechanistic evidence is somewhat superficial because we have not empirically traced the workings of each part of the mechanism. Theories of mechanisms in the minimalist understandings are therefore typically depicted merely as *Cause -> M -> Outcome*, meaning that the actual causal links in-between remain implicit (examples of this understanding include Elster, 1998; George & Bennett, 2005, p. 6; Bennett & Checkel, 2014; Falletti & Lynch, 2009, p. 1146; Mahoney, 2015).

An example of a theorized "minimalist" mechanism can be found in Nina Tannenwald's 1999 article on the impact of norms on U.S. decision-making. She theorizes that the cause, norms against the nonuse of atomic weapons (a nuclear "taboo"), contributed to U.S. decision-makers avoidance of using them (outcome), but the mechanism remains firmly within a theoretical black box because no causal mechanism is detailed whereby the cause is linked to the outcome. The closest she gets to unwrapping causal mechanisms is in the conclusion of the article where she mentions three plausible links between norms and nonuse: constraints imposed by individual decision-maker's personal moral convictions, domestic opinion, and world opinion (Tannenwald, 1999, p. 462; 2007, pp. 47–51). Yet these brief descriptions do not describe the causal process that links norms with nonuse (i.e., how does the existence of the taboo actually produce behavioral changes?). For example, how are individual moral convictions against nonuse linked to a decision to not use nukes in a situation where they could have been used? Do these individuals have to deploy normative speech acts to shame other actors? Given that we cannot answer these questions, we can conclude that the actual causal process remains "minimalist." However, in the research situation faced by Tannenwald, not unpacking mechanisms in any detail was warranted because there was a low prior confidence in the existence of a causal relationship between norms and nonuse (1999, p. 438). But after she found within-case evidence of a relationship, the natural follow-up would have been to probe mechanisms in more detail.

In a *systems understanding* of mechanisms, the ambition is to unpack *explicitly* the causal process that occurs in-between a cause (or set of causes) and an outcome and trace each of its constituent parts empirically. Here the goal is to dig deeper into how things works, but by tracing each part of the mechanism empirically using mechanistic evidence and in particular observing the empirical fingerprints left by the activities of entities in each part of the process, we are arguably able to make stronger causal inferences about how causal processes actually worked in real-world cases (Illari, 2011; Russo & Williamson, 2007). In comparison, in the "minimalist" understanding, we have less direct mechanistic evidence because we have not made the process that the evidence is explicit, resulting in weaker inferences about the operation of a causal process.

In the systems understanding of mechanisms, a causal mechanism is unpacked into its constituent parts. Mechanisms are theorized as systems of interlocking parts that transmit causal powers or forces between a cause (or a set of causes) to an outcome (examples of this understanding include Glennan, 1996, 2002; Bunge, 1997, 2004; Bhaskar, 1978; Machamer, Darden, & Craver, 2000; Machamer, 2004; Mayntz, 2004; Waldner, 2012; Rohlfing, 2012; Beach & Pedersen, 2013, 2016a).

Each of the parts of the mechanism can be described in terms of entities that engage in activities (Machamer, Darden, & Craver, 2000; Machamer, 2004). Entities are the factors (actors, organizations, or structures) engaging in activities, whereas the activities are the producers of change or what transmits causal forces or powers through a mechanism. What the entities and activities more precisely are in conceptual terms depends on the type of causal explanation, along with the level at which the mechanism works and the time span of its operation. The activities that entities engage in move the mechanism from an initial or start condition through different parts to an outcome.

When theorizing the parts of a causal mechanism, the parts should exhibit some form of *productive continuity*, meaning that each of the parts logically leads to the next part, with no large logical holes in the causal story linking a cause (or set of causes) and an outcome together (Machamer, Darden, & Craver, 2000, p. 3). The overall mechanism can be depicted (as in Figure 1), where each part of the mechanism in-between a cause and outcome is detailed in terms of entities engaging in activities. The entities can be defined as nouns, whereas the activities can be depicted as verbs.

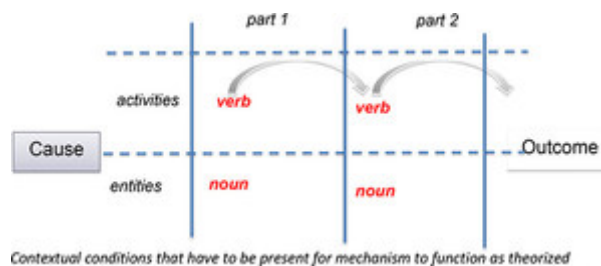


Figure 1: A simple template for a two-part causal mechanism.

The analytical value added of unpacking causal mechanisms in detail is twofold. First, unpacking mechanisms exposes the causal claim to more logical scrutiny because one cannot just postulate that a cause like norms is linked to behavioral change through “mechanisms” described with terms like “personal moral convictions” or “world opinion” (Tannenwald, 2007, pp. 47–51). By unpacking a causal process, we are better able to identify logical shortcomings in our theories and also critical links in causal stories that are particularly interesting to elaborate on. More logical scrutiny about causal logics results in better causal theories, other things being equal.

Second, by explicitly theorizing the activities that are expected to leave empirical fingerprints for each part of the mechanism, the subsequent analysis should also *study* the workings of each part empirically. If evidence is found that each part worked as theorized, then a strong

causal inference about the relationship is made possible. If evidence for one or more parts is not found, this should result in a theoretical revision of the mechanism, thereby producing more accurate theories of causal processes.

Returning to the Tannenwald example, if we were to unpack a causal mechanism linking individual moral convictions and behavior, we would first have to develop what type of causal logic we are drawing on. This could be work on the constraining impact of norm-based speech acts (e.g., Krebs & Jackson, 2007; Schimmelfennig, 2001). Using this logic, the mechanism could then be depicted as in Figure 2. Here we see two parts: (1) A believer (entity) in the taboo uses a speech act (activity) to attempt to shame proponents of use, and (2) the proponent (entity) is silenced (activity) because they are unable to deploy counterarguments because of their normative costs (clash with taboo).




Figure 2: A two-part “taboo talk” causal mechanism.

Source: Builds on Tannenwald (1999).

The context for the operation of the “taboo talk” mechanism is one where decision-makers are debating the possible use of nuclear weapons and where a medium-strength taboo against their use exists as a form of a shared norm among a number of participants in the group.

This example shows the methodological value added of both making contextual conditions explicit and unpacking the mechanism into constituent parts composed of entities engaging in activities. First, trying to make the contextual conditions explicit tells us that this initial theorization is underspecified because we cannot answer questions like: How many members of the group have to believe in the taboo? Are all participants equal or does their relative standing in the group matter for how strong the taboo acts as a rhetorical constraint in the group?

Second, by making the activities of entities explicit, this focuses our attention on the causally productive parts of the process, resulting in better theories and evidence of processes. On the theoretical side, unpacking the norms->nonuse mechanism into parts composed of entities engaging in activities forces us to make the causal logic of the process explicit. On the empirical side, while Tannenwald uses taboo talk as evidence, by making the activities explicit, we would simply need to get more direct evidence because we would have to investigate the empirical fingerprints left by the interaction process whereby making taboo-based arguments are deployed and proponents of use are silenced afterward because they are deprived of rhetorical material for rebuttals. Merely finding taboo talk only tells us that norm-based arguments were deployed, but this does not shed any light on whether they actually had an effect, and if so, how.

The Analytical Uses of Minimalist and Systems Understandings of Process Tracing

Despite the confusion that the two different understandings of mechanisms create in the literature on process tracing, it is actually helpful to view them as two distinct variants of process tracing because they are applicable in different research situations. The minimalist understanding can be used when engaging in process tracing relatively early in mechanism-focused research, when we are still unsure about what mechanisms link causes and outcomes together. For most causal theories, there are multiple plausible mechanisms that could link a given cause and outcomes together, depending on context, a phenomenon that can be termed equifinality at the level of mechanisms (Gerring, 2010; Beach & Rohlfing, 2015). Yet when we have little knowledge of which type of mechanism (or mechanisms) links a given cause and outcome, and under which conditions one or the other mechanism provides the link, it makes sense first to engage in a form of a process-tracing *plausibility probe* where mechanisms are not unpacked in any detail. In this situation, we first want to know *which* mechanism links a cause (or set of causes) and an outcome in a given context *before* we get to the question of learning about the inner workings of a particular mechanism. And after we have engaged in more intensive process tracing (systems understanding) of one or more cases, we can use a minimalist understanding to determine whether what we found in the studied case(s) also holds in other cases within the population of causally similar cases (e.g., they share similar contextual conditions).

In contrast, the systems-understanding variant of process tracing can be deployed after we have a good idea about a plausible causal process. It is when we first have a reasoned belief that there might actually be something to trace that it makes sense to engage in the intensive theoretical and empirical work of unpacking each part of the mechanism and empirically tracing the observable manifestations left by the activities of entities. By tracing each part of the mechanism, the result is the production of a richer body of mechanistic evidence, thereby enabling stronger causal inferences to be made.

Both understandings of mechanisms used in process-tracing methods share certain foundational assumptions relating to the types of causal claims being made. When working with mechanisms, asymmetric claims are being made about the links between a given cause (or set of causes) and an outcome (Beach & Rohlfing, 2015; Beach & Pedersen, 2016a). Asymmetry here means that no claims are made about what happens when a cause, outcome, or the mechanism are *not* present (for more on asymmetric claims, see Ragin, 2000, 2008; Goertz & Mahoney, 2012). This has implications for concept definition because it means that we only need to define what can be termed the “positive pole” of causal concepts (i.e., causes and outcomes) and the qualitative threshold at which the cause and outcome have causal properties in relation to a given mechanism (Beach & Pedersen, 2016a). Here one would define a cause like democracy in terms of the attributes required for the concept to have causal properties (i.e., it can trigger a causal process linked to the outcome), but one would not need to define the attributes of nondemocratic systems because they would be analytically irrelevant in relation to tracing mechanisms linking democracy to an outcome. They would just be treated as “anything but democracy.”

A more contested point in debates about process-tracing methods is whether studying mechanisms in case studies requires deterministic theorization. There is often a large degree of misunderstanding in methodological debates about determinism because many conflate ontological and epistemological determinism. Epistemological determinism would mean that empirical evidence would allow us to be 100% confident about why something happened, which of course is impossible given the inherent messiness and ambiguity of the types of empirical evidence we use in the social sciences. However, just because we cannot be empirically 100% certain does not mean that we should theoretically always hedge our bets in the form of probabilistic ontological claims about causation. Probabilistic causal claims are about trends across populations of cases, whereas deterministic causal statements claim that under specified conditions, a given cause (or set of causes) will produce an outcome through a specified mechanism (or mechanisms), but only within a small, bounded population of causally similar cases. Indeed, it can be argued that the intensive empirical tracing of mechanisms in a particular case is simply not very useful if we want to examine probabilistic causal relationships that will only manifest themselves as trends across a set of cases. If we theorized about causal mechanisms in probabilistic terms, studying a single case would tell us little about whether we should revise our theory when we have negative findings or we should just select another case because we were within the 30% of instances where the relationship did not occur for unknown reasons (Mahoney, 2008, pp. 415–416).

In contrast, when making ontological determinate claims, we are forced to tackle head-on any incongruences and anomalies that we find when engaging in process tracing, instead of discounting them as being exceptions from an overall trend, as we would if we are studying probabilistic theoretical claims. If we do not find confirming evidence of any mechanism in a process-tracing case study where our theory told us that there should be one present, we do not just discount this as an exception to an otherwise strong trend across cases. Instead, we are forced to reappraise our theory, attempting to figure out why what we expected did not occur in the case (Mahoney, 2008; Adcock, 2007). These failures of our theories are intensely interesting for case-based research and enable us to build better theories of causal processes, thereby learning more about how things work (Andersen, 2012). The result of grounding process tracing on the assumption of ontological determinism is that our causal claims become progressively refined in an iterated process of empirical research, making our knowledge less and less wrong as we better understand how causal mechanisms work and the contextual bounds in which these relationships hold.

Finally, a common assumption used when making mechanism-based causal claims is that they are sensitive to contextual conditions (also termed scope conditions) (Falleti & Lynch, 2009, p. 1152; Kurki, 2008, p. 231). Contextual conditions can be defined as the “... relevant aspects of a setting (analytical, temporal, spatial, or institutional) in which a set of initial conditions leads ... to an outcome of a defined scope and meaning via a specified causal mechanism or set of causal mechanisms” (Falleti & Lynch, 2009, p. 1152). They are not causally productive, but are merely conditions that have to be present for a relationship to work in a particular manner. Context is important because “[f]ormally similar inputs, mediated by the same mechanisms, can lead to different outcomes if the contexts are not analytically equivalent” (Falleti & Lynch, 2009, p. 1160). This means that generalization of our claims about causal mechanisms from one case to other cases can only be done after it is demonstrated that the studied case is contextually similar to other positive cases where the relationship might be present.

Using Within-Case Evidence to Make Causal Inferences about Mechanisms

How can we make causal inferences about mechanisms when we only possess within-case mechanistic evidence provided by tracing causal processes in a case? In process tracing, we are not assessing the difference that changes in values of X make for values of Y. Instead, inferences are made using the correspondence between hypothetical and actual observable manifestations of the operation of mechanisms within a selected case, what can be termed mechanistic, within-case evidence.

Increasingly, scholars are arguing that Bayesian logic provides a good logical language that enables them to ask the right questions when evaluating what collected empirical material can act as evidence of (Bennett, 2008, 2014; Beach & Rohlfing, 2015; Humphreys & Jacobs, 2015; Beach & Pedersen, 2013, 2016a; Charman & Fairfield, 2015). At the core of the Bayesian approach is the idea that science is about using new evidence to update our confidence in causal theories. Bayesian empirical updating goes in the direction of both confirmation and disconfirmation, but never reaches 100% or 0% confidence in a theory due to the inherent uncertainty of real-world empirical evidence. Yet adopting an epistemologically probabilistic approach does not mean we have to work with ontologically probabilistic causal claims; determinate claims can also be evaluated using Bayesian logic as the inferential underpinning (Howson & Urbach, 2006).

At its simplest, Bayes's theorem states that our confidence in a theory after we have evaluated new evidence (posterior) is a function of our *prior* confidence times the *evidential weight of new empirical material* collected. The amount of updating that new evidence enables is determined by *both* our prior confidence in a theory and the evidential weight of new evidence. For a longer, more technical introduction to Bayesian logic, see Bennett (2014).

Our prior confidence in a causal hypothesis matters because, if we already have a large amount of theoretical and empirical knowledge suggesting that a theory is valid, only very strong new empirical evidence can further increase our confidence in the theory. In contrast, and more typical for the situation in which we employ process-tracing case studies, when we know relatively little about a causal mechanism that potentially links a cause and outcome, even relatively weak confirming evidence can increase our confidence in a theory. As applied to process tracing, setting prior confidence requires an assessment based on the existing literature of how confident we can be in a hypothesized mechanism existing *in a given case*. This means that existing research is not necessarily very relevant prior knowledge for the selected case because there can be many contextual factors that might make the population-level trend not hold in the individual case. Therefore, unless we have good knowledge of these contextual factors (and we usually do not) that would enable us to have a qualified guess as to whether the population-level trends should hold in the selected case, Bayesian logic would suggest that we should proceed in a cautious fashion by setting our prior confidence for the selected case as relatively low. This means that most process-tracing case studies start out as a form of plausibility probe of mechanisms in the selected case, making tracing "minimalist" mechanisms the most appropriate analytical first step in many research situations.

Central to Bayesian logic is the intense evaluation of what diverse types of empirical material can potentially tell us about the veracity of causal theories. In the application of Bayesian logic to process tracing, the literature tells us that we should focus on evaluating two questions in particular: whether we have to find a given piece of empirical material (certainty of evidence), and if found, whether there are any plausible alternative explanations for finding the empirical material (uniqueness of evidence) (Van Evera, 1997; Bennett, 2014; Rohlfing, 2012; Beach & Pedersen, 2013, 2016a). While certainty is relatively straightforward, there is disagreement in the literature about what “alternative explanations” are as regards evaluating uniqueness of within-case, mechanistic evidence. Many believe that these alternatives come from competing causal hypotheses (e.g., material constraints versus normative constraints), but this type of competing theoretical explanation at the level of causes is usually only relevant when assessing evidence of difference-making, where we want to isolate the impact of one cause by controlling for all other potential causes. In contrast, uniqueness of mechanistic evidence relates to whether there are any plausible alternatives for finding the particular empirical material in the case, not whether there are other causes present or not in the case.

In a minimalist understanding, we ask ourselves about what observables the operation of a mechanism has to leave in a case, and if found, whether there are any alternative explanations for finding them. In contrast, in the systems understanding, we ask ourselves about what observables would be left in a case by the data-generating process for the activities of entities for each part of the mechanism.

The relationship among certainty and uniqueness of evidence and confirmation and disconfirmation in Bayesian logic can be illustrated using a flowchart, as in Figure 3. The arrows depict “paths” through which we might reach the found evidence or its absence.




Figure 3: Certainty and uniqueness of evidence in relation to confidence on part of mechanism being present or not.

Source: Builds on Friedman (1986a).

Our prior confidence determines the degree to which we are confident about the existence of a part of a causal mechanism before research. If we have very high prior confidence, we would be relatively confident *before* our research that we are at the “part of causal mechanism present” node, and vice versa. Note that in a minimalist understanding, we are speaking about the presence of the overall mechanism, but the logic is the same.

What finding the predicted evidence tells us about the presence of a part of a mechanism is then a function of the certainty and/or uniqueness of the evidence. If we find evidence and it is very implausible that we would find it with alternative explanations (very unique evidence), the newly found evidence makes us more confident that we have “traveled” to the found evidence through the solid-line pathway instead of along the dotted line through the lower node (“part of mechanism not present” node), meaning that the found evidence *strengthens* our confidence in the part being present. In contrast, if found evidence is not very unique, we have not updated our knowledge because we are just as uncertain about whether we reached found evidence through the “part present” and “not present” nodes as we were before we started our research.

If we do not find the evidence and it was highly certain that we had to find it, this would make us more confident that we reached “evidence absent” through the “part of mechanism not present” node, whereas if the evidence has low certainty and it is not found, we are no more confident in having reached “evidence absent” through either “part of mechanism present” or “part of mechanism not present” than we were before research started.

Assessing the certainty and uniqueness for each piece of evidence involves providing reasoned theoretical arguments in the form of propositions for why the data-generating processes of the activities associated with each part of the mechanism (or overall mechanism if minimalist) being studied would leave particular empirical fingerprints in the case (i.e., mechanistic evidence). These justifications need to include theoretical arguments *and* case-specific arguments about what the postulated evidence means in a given context. It is important to be clear about where in the empirical record one would expect the evidence to be found and the argument for why the data-generating process of the part of the causal mechanism would leave these specific fingerprints.

When assessing uniqueness of mechanistic evidence, it is also important to situate a particular piece of evidence within the full body of potential evidence in a given case. This is important in order to avoid “cherry-picking” particular pieces of evidence that do not represent the broader pattern of what happened in a case. We might find evidence in the minutes of a meeting that it was chaotic and think this is evidence of a suboptimal decision-making process. However, unless we assess whether there were other similarly chaotic meetings, one could argue that an alternative explanation for finding the evidence of a chaotic meeting is that it was just the one meeting where the air was cleared, and that the rest of the meetings were orderly. This danger of overinterpretation based on not properly assessing the representativeness of found evidence is particularly acute when we do convenience sampling based on what types of sources are most accessible. It is therefore important to evaluate the *representativeness of individual pieces of evidence* in the accessible empirical record explicitly when discussing uniqueness.

The probative value of individual pieces of mechanistic evidence is typically quite low. However, if the multiple pieces of mechanistic evidence are *independent* of each other and they point in the same direction, their probative value can be summed together in terms of the amount of updating they enable (Good, 1991, pp. 89–90). This means that adding more pieces of “weak” evidence together provides stronger evidence confirming/disconfirming a given part of a mechanism, although there are natural limits to this because Bayesian logic also tells us that it becomes increasingly difficult to update as we become ever more confident (or less

confident) in a hypothesis (Bennett, 2014). In contrast, if two pieces of evidence are completely *dependent* on each other, finding both does not tell us anything more than finding either of the single pieces by themselves. It is therefore vital to establish whether pieces of evidence are actually independent of each other before we sum evidence together. Establishing independence is often easier when we are dealing with very different types of evidence from very different sources; this is one of the reasons that *evidential diversity* is favored by Bayesians (Howson & Urbach, 2006; Bennett, 2014).

What does mechanistic empirical evidence look like in process-tracing research? Basically, relevant mechanistic evidence is any empirical material potentially left by the operation of a causal process that increases or decreases our confidence in the existence of an underlying causal mechanism or mechanisms. There are four distinguishable types of evidence that can be relevant when studying causal mechanisms: patterns, sequences, traces, and accounts. *Patterns* relate to predictions of statistical patterns in the empirical record. For example, if we are testing a part of a mechanism dealing with racial discrimination in a case dealing with employment, statistical differences in patterns of employment could be relevant evidence on which we could make inferences. *Sequences* deal with the temporal and spatial chronology of events that are predicted by a hypothesized causal mechanism. For example, if we are testing a causal mechanism about rational decision-making in a given case, certain evidence would be that decision-makers would first collect all available information, followed by an assessment of the information, and finally they would make a decision that maximizes their utility based upon this assessment. If we found in the case that decision-makers first made a decision and *then* collected information, this would be quite strong disconfirming evidence in relation to the hypothesized rational decision-making process being valid in the case.

Traces are pieces of evidence the mere existence of which provides proof. For example, if we were testing a mechanism about lobbying processes, the existence of some record of a meeting being held between a decision-maker and a lobbyist would be proof that they had met. If this predicted evidence had to be found (high certainty), and we did not find it despite having full access to the empirical record, we would then downgrade our confidence in the lobbying mechanism being present in the case. Finally, *accounts* deal with the content of empirical material, be it meeting minutes that detail what was discussed in a meeting or an oral account of what took place in a meeting.

What empirical material can be evidence of is often case-specific in process-tracing research. The basic point is that the workings of mechanisms often leave different empirical observables in different cases, despite being the same theorized mechanism. Given the very case-specific nature of the evidence implied by the mechanism in different cases, what empirical material counts as evidence in one case is not necessarily what counts as evidence in another. To develop empirical fingerprints that are sensitive to the particulars of individual cases, however, requires considerable case-specific knowledge and expertise.

Finally, there is debate about whether priors, certainty, and uniqueness should be quantified or not in process-tracing research. Some scholars believe that the values of the prior and certainty/uniqueness of each piece of evidence should be set explicitly using numbers; for instance, by stating that based on existing research, there is a 45% likelihood that the hypothesized part of a mechanism actually exists in a given case and has values for the probative value of evidence (Bennett, 2014; Humphreys & Jacobs, 2015). The argument behind this is that it makes the evaluation as transparent as possible and allows other

scholars to question the values. Other scholars contend that, given the nature of the empirical material we are working with, it is impractical or even meaningless (Beach & Pedersen, 2013, 2016a; Charman & Fairfield, 2015). In the latter, the argument is that in process tracing, we typically have heavily contextualized empirical and theoretical knowledge that enables us only to make qualified guesses about ranges of values of the probative value of evidence, at best. Therefore, assigning numbers to the priors and theoretical certainty and uniqueness would be very arbitrary at best and misleading at worst. Even more damning is the argument that quantification leads to excessive simplification, given that the probative value of individual pieces of within-case evidence derives from *complex interpretations* of what evidence means in context. Charman and Fairfield (2015) offer an extended example of quantification of the probative value of evidence in an article, concluding that "... the most probative pieces of evidence are precisely those for which quantification is least likely to provide added value. The author can explain why the evidence is highly decisive without the need to invent numbers ..." (2015, pp. 31-32).

Case Selection and Nesting Process-Tracing Case Studies

This section introduces principles of case selection in relation to process tracing. At its core, process tracing involves the detailed empirical tracing of the operation of mechanisms within an individual case. Mechanistic evidence either confirms or disconfirms our theories about the operation of a causal mechanism in the studied case (Illari, 2011).

Therefore, to generalize beyond the studied case requires that process-tracing studies are combined with comparative methods. Here the logic is that we need to demonstrate that the studied case is causally similar to a set of other positive cases, making us more confident that the causal processes found in the studied case should also be present in the other cases. Naturally, if the process-tracing research is focused on providing a comprehensive explanation of a particularly interesting historical outcome, case-selection principles are not relevant because the case is not a "case of" a narrower theoretical phenomenon, but a proper noun in the form of the Cuban Missile Crisis or the French Revolution.

Selecting appropriate cases for process tracing requires some form of prior cross-case knowledge of the population of the given theoretical relationship. Mapping the population can take the form of a very rough preliminary comparative mapping of membership of cases of a cause, outcome, and contextual conditions, but it can also be more systematic in the form of a qualitative comparative analysis (QCA) study (Beach & Pedersen, 2016b; Schneider & Rohlfing, 2013; Beach & Rohlfing, 2015).

Mapping cases using comparative methods results in differentiating a population into four different types of cases. Positive (or typical) cases are those where the cause (or set of causes), outcome, and potentially causally relevant conditions for the operation of a given causal mechanism are *all* present. In Figure 4, these are cases 4, 5, 6, and 7 in quadrant I, all of which possess the cause, outcome, and the requisite contextual conditions for the mechanism to be present. In quadrant II, there are deviant cases where either the cause and/or the requisite contextual conditions are not present, whereas in quadrant III, there are deviant cases where the outcome should occur but it does not. These are analytically

irrelevant cases when working with asymmetric causal claims as in process tracing because they tell us nothing about the causal process in-between a cause and an outcome (Goertz & Mahoney, 2004; Rohlfing, 2012).

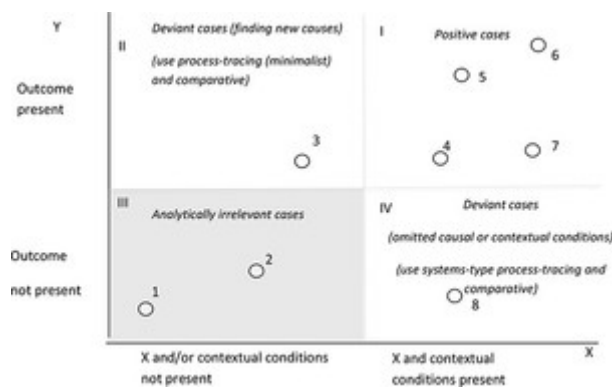


Figure 4: Mapping the population of cases and different analytical purposes of process-tracing case studies.

Which case or cases are relevant for process tracing depends upon the research situation. Positive cases in quadrant I are the only cases that should be selected for either testing or building theories of causal mechanisms based on the logic “why should we trace a mechanism in a case where we know a priori that it is not there”?

When we are in a situation in which there is uncertainty about what contextual conditions matter, the analyst can start by selecting a case where as many possible contextual conditions are present (Rohlfing, 2012, pp. 200–211). If evidence of a mechanism is found in the “maximum-context” case, we cannot automatically infer that it is also present in other cases that do not share all of the contextual conditions. Here we can engage in an iterated strategy of gradually whittling down the number of contextual conditions by assessing new cases with fewer shared conditions until we find the lower bounds of required contextual conditions, admitting more cases into the bounds of quadrant I (Rohlfing, 2012, pp. 200–211).

Note that if one accepts the assumption of ontological determinism as the basis for case-based research like process tracing, it makes no sense to discuss most- or least-likely cases within the set of positive cases. A causal relationship is either possible or not in any given case. Where likelihoods make sense relates to our *empirical knowledge* of whether the relationship actually exists or not in a studied case. Therefore, in process tracing, we select positive cases in which a causal mechanism is theoretically possible because the cause (or set of causes), outcome, and the requisite contextual conditions are present.

Deviant cases in quadrant IV are useful for detecting omitted causal and/or contextual conditions when the cause (or set of causes) should be sufficient to produce an outcome. This type of process tracing would typically only be applied after we have a good working knowledge of the mechanism(s) linking a cause and outcome together from studying multiple positive cases. Case 8 could then be selected for process tracing, attempting to trace the mechanism until it breaks down in order to gain information on when and why the mechanism failed (Andersen, 2012, pp. 421–422). This information can then inform a pair-wise comparative analysis of the deviant case with a positive case as similar as possible to

determine what conditions the deviant case lacks. After we have found an omitted contextual or causal condition, we would then use our updated theoretical knowledge to reclassify cases, with case 8 shifting to quadrant III.

Deviant cases in quadrant II can be used to search for new causes, although here it only makes sense to use process tracing in a minimalist sense. If one has no idea about causes, why should one spend so much analytical effort in trying to trace each part of an unknown mechanism to an unknown cause? But by engaging in a looser empirical soaking-and-probing, it is possible to detect potential new causes that can then be tested more systematically (Lieberman, 2005; Rohlfing & Schneider, 2013).

There is some debate in the literature about whether additional criteria should be included when selecting positive cases. Some scholars argue that one should select positive cases that are only members of one cause (or a conjunction of causes), enabling “control” for other potential causes. Gerring and Seawright write that “... researchers are well advised to focus on a case where the causal effect of one factor can be isolated from other potentially confounding factors.” (2007, pp. 122). They term this type of case a “pathway case” (for more, see also Weller & Barnes, 2015). Schneider and Rohlfing draw on this guidance in their discussion of case selection for process tracing when they state that we should choose “unique set” cases, where we “... focus on one term ... to unravel the mechanism through which it contributes to the outcome in the case under study” (2013, pp. 566–567). Goertz writes that one should avoid cases that exemplify multiple causal mechanisms (Goertz, 2012, p. 18). The logic behind these claims is based on the importance of controlling for other factors when working with difference-making evidence across cases. For instance, for an experiment to produce valid evidence of difference-making, it is imperative that only one cause “varies,” with all potential other causes of the outcome held equal.

Other scholars argue that controlling for other causes when selecting cases is not important because control can be achieved at the evidential level when each piece of evidence is evaluated for its uniqueness (Beach & Pedersen, 2016b). The logic here is that different theoretical mechanisms can be predicted to leave very different empirical observables, especially if one has broken a mechanism down theoretically into its constituent parts (systems understanding). For instance, the causal mechanism linking norms with behavior would leave very different empirical observables than if we were testing a mechanism linking rational, consequentialist thinking with behavior. In other words, the claim here is that we do not need to select cases based on a variance-based logic of controlling for other causes because we can distinguish empirically between causes when using within-case evidence of mechanisms. Additionally, if we restrict our analysis to those where only one cause is present, we cannot generalize to other positive cases where other causes are also present, because the mechanism(s) might be different in the two sets of cases. By selecting both pathway and non-pathway cases, we can determine whether the presence of other causes impacts the operation of mechanisms in the population.

Finally, there is some debate about whether it is possible to generalize to all other positive cases based on only one process-tracing case study. Many scholars claim that one theory test can be enough. Lieberman writes that “... if *one* or more intensive case studies can demonstrate the validity of the theoretical model—which had already passed muster in the LNA—by plausibly linking cause to effect in the expected manner, then the nested analysis

provides *ringing support* for the model (End analysis I ...)” (Lieberman, 2005, p. 448, emphasis added). In the flowchart illustrating his case-selection strategy, if one finds support in the single case study, he suggests we can end our analysis.

Other scholars have claimed that, given the sensitivity of mechanisms to context, it is vital to demonstrate both that the population of other positive cases is causally homogeneous with the selected case, and that mechanisms can be traced in at least one more positive case to enable the generalization of the findings about the mechanism (or mechanisms) in the studied cases to other positive cases (Beach & Pedersen, 2016b).

Variants of Process Tracing

This section develops three variants of process-tracing methods: theory-testing, theory-building, and case-centric process tracing. Process tracing can also be used for revising theories (searching for new causes or omitted causal/contextual conditions), but as this was discussed in the previous section, it is not developed further here. Guidelines are put forward for each variant, followed by an example from published research.

Theory-testing process tracing is a theory-first research method, testing whether a hypothesized causal mechanism exists in a positive case or set of positive cases by exploring whether the predicted evidence of a hypothesized causal mechanism exists in reality. By providing evidence of a mechanism linking X and Y, stronger claims of causation can be made within the studied case. At the same time, by tracing mechanisms, we gain a greater understanding of *how* X causes Y.

The first step is to theorize a plausible causal mechanism based on existing literature and logical reasoning, along with giving some thought to the contextual conditions required to be present for the mechanism to work properly. As discussed in the section on causal mechanisms, the theorized mechanism can either be in a minimalist form of Cause->M->Outcome or in a more unpacked form by theorizing each of the constituent parts of the mechanism in terms of entities and the causally productive activities that provide the causal link to the next part.

The theorized causal mechanism then needs to be operationalized in terms of developing predicted empirical observables for the mechanism or its constituent parts in a specific case, focusing on developing what the data-generating process of the overall mechanism (minimalist) is or what the activities of entities for each part are (systems). The predictions about mechanistic evidence should be as clear as possible, making it easier to determine whether they are then actually found in the subsequent case study or not.

Each predicted observable has to be evaluated using existing knowledge about whether they have to be found in the given case (certainty), and if found, whether there are any plausible alternative explanations for finding them (uniqueness). Alternative explanations can come from competing theories, but it is more typical when they are more ad hoc, case-specific explanations for finding the evidence (Rohlfing, 2014). Unless it is possible to develop highly unique observables, it is useful to process through as many possible observables that the operation of a mechanism (or the parts of a mechanism) might have left.

At the core of theory-testing process tracing is a structured empirical assessment of whether a hypothesized causal mechanism is actually present in the evidence of a given case. Empirical material is gathered to see whether the predicted evidence was present or not for each part of the mechanism, or the overall mechanism in the minimalist understanding. If this evidence is found, we can then infer that the hypothesized causal mechanism is present *in the case* and worked as we theorized. If we also want to know whether the same mechanism links a cause (or set of causes) and an outcome in the other cases, we would engage in a second theory-testing process-tracing case study in another positive case. If we then find that the mechanism is also present there, we can make the cautious inference that the causal mechanism is probably also present in the rest of the population of positive cases (quadrant I). To be even more confident of our generalization, we could follow up with a minimalist mechanism case study of a third (or more) positive case, focusing on assessing whether a key observable of a key part of the causal mechanism was present.

If evidence is not found for a given part (or for the overall mechanism if the minimalist understanding is used), the researcher should engage in a round of theory-building, using the insights gained from the empirical analysis as inspiration for building theories of new parts of the mechanism. However, if it is impossible to detect a causal mechanism between X and Y in a positive case after numerous repeated attempts, there can be two reasons for this: (1) either the case is idiosyncratic, or (2) there is no causal relationship. To determine which of the two is correct requires *comparing* the chosen case with what we know about other positive cases, and assessing whether there are any important differences between them that could potentially prevent the mechanism from working. In particular, this comparison can shed light on the contextual conditions for the proper functioning of the mechanism, enabling us to assess whether the bounds of the population have been set properly.

When presenting the empirical evidence in a theory-testing process-tracing case study, given that there are many different types of evidence, a theory test typically does *not* read like an analytical narrative. Instead, in the minimalist understanding, it involves a rigorous discussion of whether the predicted observables were found and what their probative values are. For systems understanding, presentation of our research involves a part-by-part discussion of what empirical observables we expected to find if the part of the mechanism existed and what we would learn if found/not found (certainty and uniqueness), followed by a presentation of whether the predicted evidence was actually present in the case or not. There can be situations, however, where we are unable to assess whether the predicted evidence was present or not because of access issues, especially relevant when we are researching politically sensitive topics. If this is the case, no updating takes place.

Returning to Tannenwald's 1999 article as an example, she first theorizes a minimalist mechanism linking norms with nonuse of nuclear weapons (1999, pp. 437–438). This is followed by the development of one core empirical observable, described as “taboo talk.” She first describes clearly how we would know the predicted evidence when we see it, writing that taboo talk is evidence of “... non-cost-benefit-type reasoning along the lines of ‘this is simply wrong’ in and of itself (because of who we are, what our values are, ‘we just don’t do things like this,’ ‘because it isn’t done by anyone,’ and so on)” (1999, p. 440). She then puts forth the argument that the evidence is not certain, arguing that “[a]t an even greater level of ‘taken-for-grantedness,’ the taboo might become a shared but ‘unspoken’ assumption of decision-makers” (1999, p. 440). She does claim that the observable is not very plausible with the main

alternative explanation drawn from realist theorization, writing that “[t]aboo talk is not just ‘cheap talk,’ as realists might imagine ... in this case decision makers themselves believed they were constrained by a taboo ...” (1999, p. 440). While the claim of the relative uniqueness of taboo talk could have been further elaborated on to strengthen the argument, most examples of theory-testing process tracing just proclaim uniqueness or certainty of observables without any reasoning behind it at all. If we want to claim we have mechanistic evidence of the operation of a process, we logically have to provide some justification for why empirical material can actually enable inferences to be made.

Theory-building process tracing is a more inductive form of research that in its purest form starts with empirical material and uses a structured analysis of this material to build a plausible hypothetical causal mechanism whereby a cause (or set of causes) is linked with an outcome that can be present in multiple cases, meaning that it can be generalized beyond the single case. In effect, it involves using empirical material to answer the question “how did we get here?” (Frieden, 1986, p. 582; Swedberg, 2012, pp. 6–7). Theory-building process tracing is utilized primarily when we know that there might be a relationship between a cause and an outcome, but we are in the dark regarding potential mechanisms linking the two.

After the key theoretical concepts (causes and outcome) are defined and operationalized, theory-building proceeds to investigate the empirical material in the case, using empirical material as clues about the possible empirical manifestations of an underlying causal mechanism. This involves an intensive and wide-ranging search of the empirical record, with material collected without knowing what it is evidence of. Here it can also be helpful to develop a descriptive narrative of what happened in the case to shed light on potential mechanisms.

The next step involves inferring from the observable empirical material that actual evidence reflecting an underlying plausible causal mechanism was present or not in the case. Tentative hunches about potential mechanisms (and their parts in the systems understanding) are made based on the first round of empirical probing, after which the researcher proceeds to evaluate whether any of the collected material is actually evidence of the tentative mechanism (or parts of the mechanism). This evaluation of evidence proceeds in a slightly different fashion than in theory-testing, given that it is not relevant to discuss the certainty of evidence because it has *already been found*; instead, one only evaluates the uniqueness in relation to the tentative hypothesized mechanism or its parts.

Evidence does not speak for itself. Often theory-building has a strong testing element, in that scholars seek inspiration from existing theoretical work and previous observations about what to look for. Here existing theory can be thought of as a form of grid to detect systematic patterns in empirical material, enabling inferences about predicted evidence to be made. In addition, one can look to research on mechanisms on similar research topics for inspiration about what parts of the mechanism might look like. In other situations, the search for mechanisms is based on hunches drawn from puzzles that are unaccountable for in existing work.

The final step of theory-building involves inferring that the found empirical material is actually evidence of an underlying causal mechanism, either in the form of a minimalist mechanism or unpacked into its constituent parts. In reality, theory-building process tracing is usually an iterative and creative process.

A good example of a theory-building process tracing work is Janis's book on "groupthink" (Janis, 1982). In the book, he attempts to build a causal mechanism that details how conformity pressures in small groups (cause) can have an adverse impact upon foreign policy decision-making processes (outcome), using a selection of case studies of policy fiascoes that were the result of poor decision-making practices by a small group of policymakers who constituted a cohesive group. He uses the term groupthink to describe the causal mechanism whereby conformity pressures in small groups actually produce poor decisions.

The first exploratory case that he uses in the book is an analysis of the Bay of Pigs fiasco. He notes first that groupthink was by no means the sole cause of the fiasco (Janis, 1982, p. 32), but at the same time, he notes a puzzle that existing explanations are unable to explain; why did the "best and the brightest" policymaking group in the Kennedy administration not pick out the faulty assumptions underlying the decision to support the intervention. The starting point of his case study was using a combination of existing empirical accounts, existing psychological theories of group dynamics, and other relevant political science theories as an inspiration for his search through the empirical record for systematic factors that formed parts of a possible groupthink causal mechanism. For example, he writes that "... when I reread Schlesinger's account, I was struck by some observations that earlier had escaped my notice. These observations began to fit a specific pattern of concurrence-seeking behavior that had impressed me time and again in my research on other kinds of face-to-face groups.... Additional accounts of the Bay of Pigs yielded more such observations, leading me to conclude that group processes had been subtly at work" (Janis, 1982, p. vii). Here we see the importance that imagination and intuition also play in devising a theory from empirical evidence, while at the same time, he is informed by theoretical research from other fields.

Step one then involved collecting empirical material in order to detect potential evidence of underlying causal mechanisms. Inferences were then made from empirical material to conclude that evidence of the part of the mechanism existed (step two), resulting in the secondary inference that an underlying mechanism was present in step three. He writes that "[f]or purposes of *hypothesis construction*—which is the stage of inquiry with which this book is concerned—we must be willing to make some inferential leaps from whatever historical clues we can pick up ..." (Janis, 1982, p. ix). Further, "What I try to do is to show how the evidence at hand can be viewed as forming a consistent psychological pattern, in the light of what is known about group dynamics" (Janis, 1982, p. viii).

Explaining-outcome process tracing is an iterative research strategy that aims to trace causal mechanisms in order to produce a comprehensive explanation of a particular historical outcome. In many respects, research in explaining outcome process tracing resembles abduction, where there is a continual and creative juxtaposition between empirical material and theories (for more, see Peirce, 1955; Timmermans & Tavory, 2012). The types of theoretical explanations constructed are often eclectic combinations (for more on this type of explanation, see Sil & Katzenstein, 2010).

There are two different starting points for explaining-outcome process tracing: either theory or empirics. The theory-first path follows the steps described above under theory-testing, where an existing cause (or set of causes) and the associated mechanisms are tested to see whether they can account for the outcome. In most research situations, a single existing cause

and mechanism cannot provide a sufficient explanation for an outcome, resulting in a second stage of research where either a testing or building path can be chosen, informed by the results of the first empirical analysis. If the testing path is chosen again, this would involve testing another theorized cause and associated mechanism as a supplemental explanation to see whether together they can account for the “big and important” things going on in the case. Alternatively, the theory-building path can be chosen in the second iteration, using empirical evidence to build a new mechanism that can account for the elements of the outcome that were unaccounted for using the first mechanism, following the steps discussed under theory-building. In both paths, theorized mechanisms and empirical tests are treated more pragmatically as heuristic devices for understanding important events.

A good example of explaining-outcome process tracing can be seen Schimmelfennig’s article on the Eastern enlargement of the EU (2001). The article attempts to explain a particular empirical puzzle, which is why countries like France that were opposed to Eastern enlargement of the EU ended up not opposing it (2001, p. 49). The case study proceeds using three iterations of the testing path. He takes as his point of departure two competing theorized causes and associated mechanisms from “rationalist” and “sociological” theories of international cooperation to explain the existing EU member states’ positions toward Eastern enlargement. He first tests a rationalist mechanism and finds it can account for national preferences, but not for the final decision to enlarge. Informed by the findings of his first empirical analysis, he undertakes a second round of tests to see whether a sociological mechanism can account for the outcome. He finds that it can account for the final decision of France to accept enlargement, but it cannot account for the negotiating process.

Not surprisingly, however, Schimmelfennig finds that neither mechanism can fully explain the outcome (neither is sufficient), finding them both “wanting in the ‘pure’ form” (2001, p. 76). In response, he uses the empirical results of the first two iterations to formulate an eclectic combination of the two causes and the mechanisms they produce, attempting to “... provide the missing link between egoistic preferences and a norm-conforming outcome” by developing the idea of “rhetorical action” (the strategic use of norm-based arguments). In the third iteration of the case study, he tests this eclectic conglomerate cause and associated mechanism, finding that they provide a sufficient explanation for the historical outcome. He provides relatively strong evidence suggesting that the more complex set of causes and associated mechanism are actually present in the case and that they are *sufficient* to account for the outcome. Sufficiency is confirmed when it can be substantiated that there are no important aspects of the outcome that are unaccounted for by the explanation.

Conclusions

Process tracing involves tracing causal mechanisms using in-depth case studies that provide within-case, mechanistic evidence of causal processes. How then do we know a good process-tracing case study when we see it in practice? This article has argued that “good” process tracing always involves two components: mechanisms that are clearly theorized (either in a minimalist fashion or unpacked as theoretical systems composed of parts with entities and causally productive activities clearly described), and the transparent evaluation of what causal inferences empirical material enable by assessing the certainty and uniqueness of each observable. If the study aims to generalize beyond the single case, an additional component of

process tracing is that cross-case comparative evidence has to be mustered to demonstrate that the population of other positive cases is relatively causally similar to the selected cases, thereby enabling generalizations to be made.

Process tracing has already emerged as one of the most valuable methodological tools for main causal inferences in the social sciences. But process tracing is by no means a panacea, and there are numerous tradeoffs associated with using it. These include the low degree of external validity of findings from individual process-tracing case studies, making it vital to combine them with comparative methods that map cases onto causally similar populations to enable generalizations to be made. Further, the more mechanisms are unpacked theoretically and then traced empirically, the stronger the causal inferences we are able to make, other things being equal. However, tracing each part of a complex mechanism requires deploying enormous analytical resources, meaning that in many research situations, it makes more sense to engage in a form of process tracing “lite,” where mechanisms are treated in a minimalist fashion. Yet this weakens our ability to make strong causal inferences because we have more superficial evidence of causal processes in operation.

Finally, it is important to note that mechanistic evidence does not enable us to make causal inferences about the overall causal effect of a given cause (or set of causes), because there can be situations where a cause might trigger two or more different mechanisms that have different effects on the outcome (Illari, 2011, p. 150). For example, exercise (cause) triggers both a mechanism that produces weight gain through muscle growth, but also triggers another mechanism that results in weight loss because of an increase in calorie usage (Steel, 2008, p. 68).

There is still considerable methodological work that needs to be done, in particular in developing the Bayesian inferential underpinnings of process tracing and in developing more guidelines for how theories of mechanisms can be built using empirical case studies.

References

- Adcock, R. (2007). Who’s afraid of determinism? The ambivalence of macro-historical inquiry. *Journal of the Philosophy of History*, 1, 346–364.
- Andersen, H. (2012). The case for regularity in mechanistic causal explanation. *Synthese*, 189, 415–432.
- Beach, D., & Pedersen, R. B. (2013). *Process-tracing methods: Foundations and guidelines*. Ann Arbor: University of Michigan Press.
- Beach, D., & Pedersen, R. B. (2016a). *Causal case studies: Comparing, matching and tracing*. Ann Arbor: University of Michigan Press.
- Beach, D., & Pedersen, R. B. (2016b). Case selection techniques when studying causal mechanisms as systems. *Sociological Methods and Research*. January 13, 2016. Online First.
- Beach, D., & Rohlfing, I. (2015). Integrating cross-case analyses and process tracing in set-theoretic research: Strategies and parameters of debate. *Sociological Methods and Research*. December 13, 2015. Online First.

-
- Bennett, A. (2008). Process-tracing: A Bayesian perspective. In J. M. Box-Steffensmeier, H. E. Brady, & D. Collier (Eds.), *The Oxford handbook of political methodology* (pp. 702–721). Oxford: Oxford University Press.
- Bennett, A. (2014). Appendix. In A. Bennett & J. Checkel (Eds.), *Process tracing: From metaphor to analytic tool*. Cambridge, U.K.: Cambridge University Press.
- Bennett, A., & Checkel, J. (2014). *Process tracing: From metaphor to analytic tool*. Cambridge, U.K.: Cambridge University Press.
- Bhaskar, R. (1978). *A realist theory of science*. Brighton, U.K.: Harvester.
- Bogen, J. (2005). Regularities and causality; generalizations and causal explanations. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 36, 397–420.
- Brady, H. (2008). Causation and explanation in social science, chapter 10. In J. M. Box-Steffensmeier, H. E. Brady, & D. Collier (Eds.), *The Oxford handbook of political methodology*. Oxford: Oxford University Press.
- Bunge, M. (1997). Mechanism and Explanation. *Philosophy of the Social Sciences*, 27(4), 410–465.
- Bunge, M. (2004). How does it work? The search for explanatory mechanisms. *Philosophy of the Social Sciences*, 34(2), 182–210.
- Charman, A., & Fairfield, T. (2015). Applying formal Bayesian analysis to qualitative case research: An empirical example, implications, and caveats. Unpublished paper.
- Collier, D., Brady, H. E., & Seawright, J. (2010). Sources of leverage in causal inference: Toward an alternative view of methodology. In H. E. Brady & D. Collier (Eds.), *Rethinking social inquiry: Diverse tools shared standards* (2d ed., pp. 161–200). Lanham, MD: Rowman & Littlefield.
- Dowe, P. (2011). The causal-process-model theory of mechanisms. In P. M. Illari, F. Russo & J. Williamson (Eds.), *Causality in the Sciences* (pp. 865–879). Oxford: Oxford University Press.
- Elster, J. (1998). A plea for mechanisms. In P. Hedström & R. Swedberg (Eds.), *Social mechanisms* (pp. 45–73). Cambridge, U.K.: Cambridge University Press.
- Falleti, T. G., & Lynch, J. F. (2009). Context and causal mechanisms in political analysis. *Comparative Political Studies*, 42, 1143–1166.
- Fearon, J. (1991). Counterfactuals and hypothesis testing in political science. *World Politics*, 43(2), 169–195.
- Frieden, R. D. (1986). A diagrammatic approach to evidence. *Boston University Law Review*, 66(4), 571–620.
- George, A. L., & Andrew, B. (2005). *Case studies and theory development in the social sciences*. Cambridge, MA: MIT Press.
- Gerring, J. (2007). *Case study research*. Cambridge, U.K.: Cambridge University Press.

-
- Gerring, J. (2010). Causal mechanisms: Yes but.... *Comparative Political Studies*, 43(11), 1499–1526.
- Gerring, J., & Seawright, J. (2007). Techniques for choosing cases. In *Case study research* (pp. 86–150). Cambridge, U.K.: Cambridge University Press.
- Glennan, S. S. (1996). Mechanisms and the nature of causation. *Erkenntnis*, 44(1), 49–71.
- Glennan, S. S. (2002). Rethinking mechanistic explanation. *Philosophy of Science*, 69, 342–53.
- Goertz, G. (2012). Case studies, causal mechanisms, and selecting cases. Unpublished paper.
- Goertz, G., & Levy, J. S. (Eds.). (2007). *Explaining war and peace: Case studies and necessary condition counterfactuals*. London: Routledge.
- Goertz, G., & Mahoney, J. (2004). The possibility principle: Choosing negative cases in comparative research. *American Political Science Review*, 98(4), 653–669.
- Goertz, G., & Mahoney, J. (2012). *A tale of two cultures—Qualitative and quantitative research in the social sciences*. Princeton, NJ: Princeton University Press.
- Good, L. J. (1991). Weight of evidence and the Bayesian likelihood ratio. In C. G. Aitken & D. A. Stoney (Eds.), *Use of statistics in forensic science* (pp. 85–106). London: CRC.
- Hedström, P., & Ylikoski, P. (2010). Causal mechanisms in the social sciences. *Annual Review of Sociology*, 36, 49–67.
- Howson, C., & Urbach, P. (2006). *Scientific reasoning: The Bayesian approach* (3d ed.). La Salle, IL: Open Court.
- Humphreys, M., & Jacobs, A. (2015). *Mixing methods: A Bayesian unification of qualitative and quantitative approaches*, in review.
- Illari, P. M. (2011). Mechanistic evidence: Disambiguating the Russo-Williamson thesis. *International Studies in the Philosophy of Science*, 25(2), 139–157.
- Janis, I. L. (1982). *Groupthink: Psychological studies of policy decisions and fiascoes*. Boston: Houghton Mifflin.
- King, G., Keohane, R. O., & Verba, S. (1994). *Designing social inquiry: Scientific inference in qualitative research*. Princeton, NJ: Princeton University Press.
- Krebs, R. R., & Patrick, T. J. (2007). Twisting tongues and twisting arms: The power of political rhetoric. *European Journal of International Relations*, 13(1), 35–66
- Kurki, M. (2008). *Causation in international relations: reclaiming causal analysis*. Cambridge, U.K.: Cambridge University Press.
- Lebow, R. N. (2000). What's so different about a counterfactual? *World Politics*, 52(4), 550–585.
- Levy, J. S. (2015). Counterfactuals, causal inference, and historical analysis [<http://dx.doi.org/10.1080/09636412.2015.1070602>](http://dx.doi.org/10.1080/09636412.2015.1070602). *Security Studies*, 24(3), 378–402.

-
- Lieberman, E. S. (2005). Nested analysis as a mixed-method strategy for comparative research. *American Political Science Review*, 99(3), 435-451.
- Machamer, P. (2004). Activities and causation: The metaphysics and epistemology of mechanisms. *International Studies in the Philosophy of Science*, 18(1), 27-39.
- Machamer, P., Darden, L., & Craver, C. F. (2000). Thinking about mechanisms. *Philosophy of Science*, 67(1), 1-25.
- Mahoney, J. (2008). Toward a unified theory of causality. *Comparative Political Studies*, 41(4/5), 412-436.
- Mahoney, J. (2012). The logic of process tracing tests in the social sciences. *Sociological Methods and Research*, 41(4), 570-597.
- Mahoney, J. (2015). Process tracing and historical explanation. *Security Studies*, 24(2), 200-218.
- Mayntz, R. (2004). Mechanisms in the analysis of social macro-phenomena. *Philosophy of the Social Sciences*, 34(2), 237-259.
- Owen, J. M. (1994). How liberalism produces democratic Peace. *International Security*, 19(2), 87-125.
- Peirce, C. S. (1955). *Philosophical writings of Peirce*. In J. Buchler (Ed.), *Philosophical Writings*. New York: Dover.
- Ragin, C. C. (2000). *Fuzzy-set social science*. Chicago: University of Chicago Press.
- Ragin, C. C. (2008). *Redesigning social inquiry: Fuzzy sets and beyond*. Chicago: University of Chicago Press.
- Rohlfing, I. (2012). *Case studies and causal inference*. Houndmills, U.K.: Palgrave Macmillan.
- Rohlfing, I. (2014). Comparative hypothesis testing via process tracing. *Sociological Methods and Research*, 43(4), 606-642.
- Runhardt, R. W. (2015). Evidence for causal mechanisms in social science: Recommendations from Woodward's manipulability theory of causation. *Philosophy of Science*, 82(5), 1296-1307.
- Russo, F., & Williamson, J. (2007). Interpreting causality in the health sciences. *International Studies in the Philosophy of Science*, 21(2), 157-170.
- Russo, F., & Williamson, J. (2011). Generic versus single-case causality: The case of autopsy. *European Journal of the Philosophy of Sciences*, 1(1), 47-69.
- Schimmelfennig, F. (2001). The community trap: Liberal norms, rhetorical action, and the Eastern enlargement of the European Union. *International Organization*, 55(1), 47-80.
- Schneider, C., & Rohlfing, I. (2013). Combining QCA and process tracing in set-theoretical multi-method research. *Sociological Methods and Research*, 42(4), 559-597.
- Sil, R., & Katzenstein, P. J. (2010). *Beyond paradigms: Analytical eclecticism in the study of world politics*. Basingstoke, U.K.: Palgrave Macmillan.

-
- Steel, D. 2008. *Across the Boundaries: Extrapolation in biology and social science*. Oxford: Oxford University Press.
- Van Evera, S. (1997). *Guide to methods for students of political science*. Ithaca, NY: Cornell University Press.
- Swedberg, R. (2012). Theorizing in sociology and social science: Turning to the context of discovery. *Theoretical Sociology*, 41(1), 1-40.
- Tannenwald, N. (1999). The nuclear taboo: The United States and the normative basis of nuclear non-use. *International Organization*, 53(3), 433-468.
- Tannenwald, N. (2007). *The nuclear taboo*. Cambridge, U.K.: Cambridge University Press.
- Tetlock, P. E., & Belkin, A. (Eds.). (1996). *Counterfactual thought experiments in world politics: Logical, methodological, and psychological perspectives*. Princeton, NJ: Princeton University Press.
- Timmermans, S., & Tavory, I. (2012). Theory construction in qualitative research: From grounded theory to abductive analysis. *Sociological Theory*, 30(3), 167-186.
- Waldner, D. (2012). Process tracing and causal mechanisms. In H. Kincaid (Ed.), *The Oxford handbook of the philosophy of social science* (pp. 65-84). Oxford: Oxford University Press.
- Weller, N., & Barnes, J. (2015). *Finding pathways: Mixed-method research for studying causal mechanisms*. Cambridge, U.K.: Cambridge University Press.
- Woodward, J. (2003). *Making things happen: A theory of causal explanation*. Oxford: Oxford University Press.

Related Articles

Comparative Public Policy

Qualitative Comparative Analysis (QCA) and Set Theory

Reconceptualizing Field Research in Political Science