AARHUS UNIVERSITY

Coversheet

This is the accepted manuscript (post-print version) of the article.

Contentwise, the post-print version is identical to the final published version, but there may be differences in typography and layout.

How to cite this publication

Please cite the final published version:

Beach, D. (2016). It's all about mechanisms - what process-tracing case studies should be tracing. *New Political Economy*, 21(5), 463-472. DOI: 10.1080/13563467.2015.1134466

Publication metadata

Title:	It's all about mechanisms - what process-tracing case studies should be tracing
Author(s):	Beach, D.
Journal:	<i>New Political Economy</i> , 21(5), 463-472
DOI/Link:	https://doi.org/10.1080/13563467.2015.1134466
Document version:	Accepted manuscript (post-print)

General Rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

'It's all about mechanisms - what Process-tracing case studies should be tracing'

By Derek Beach

University of Aarhus

It is widely agreed that the core of Process-tracing (PT) as a distinct case study methodology is that it involves *tracing causal mechanisms that link causes (X) with their effects (i.e. outcomes) (Y)* (Bennett, 2008a, 2008b; Checkel, 2008; Hall, 2008; Rohlfing, 2012; Beach and Pedersen, 2013; introduction to this special issue).¹ We trace causal mechanisms whereby a cause (or set of causes) produces an outcome to both: 1) make *stronger evidence-based inferences about causal relationships* when we have within-case evidence of each step of the causal process (or absence thereof) in-between a cause and outcome, and in particular the activities that provide the causal links in the process, and 2) because tracing mechanisms gives us a better understanding of *how* a cause produces an outcome. This tracing can be used for either building or testing theories of causal mechanisms.

Yet when we look at the methodological literature on PT, there is considerable ambiguity and discord about what causal mechanisms actually are. The result of this ambiguity about what mechanisms are clearly maps onto existing applications of PT, with most PT case studies completely ignoring the underlying theoretical causal processes; i.e. it is black-boxed. In the few PT applications where mechanisms are discussed explicitly, they are typically only done so in a very cursory fashion, with the result that there is considerable ambiguity about what theoretical process the ensuing case study actually is tracing. The analytical result is that the underlying causal mechanism is relatively unstudied; it is grey-boxed.

¹ - Indeed, given that many scholars understand the term 'process' to denote non-causal, narrative analysis, it might be more appropriate to use the term 'Mechanism-Tracing' instead to accurately describe what we should be doing in PT if we have the ambition to study causal processes. To avoid more jargon, this article sticks with the well-known PT term.

This article first attempts to provide a clear definition of causal mechanisms that provides scholars using PT with a framework for theorizing mechanisms in a fashion that is amenable to in-depth empirical analysis. I contend that PT needs to adopt an understanding of causal mechanisms where they are explicitly fleshed out by unpacking causal processes linking X and Y into a series of interlocking parties composed of entities engaging in activities that transmit causal forces from cause to outcome. It is argued that explicitly theorizing mechanisms in this fashion results in both better causal theories and in PT case studies that are more clearly focused on tracing mechanisms.

The article then illustrates the analytical problems that result in practical applications of PT when the causal process in-between a cause and outcome is not developed in sufficient detail using a detailed analysis of the theoretical claims in an article by Ziblatt (2009), where he claims he is tracing mechanisms. While his empirical analysis of the 'mechanism' does shed some light on possible underlying causal processes, the evidence provided is quite anecdotal because it is impossible to develop clear empirical tests of an underlying causal mechanism when we do not really know exactly what it is we are tracing. The result is that we do not gain the ability to make stronger inferences about causal processes linking causes with outcomes through mechanisms, which was the very reason we use PT in the first place.

The article then develops a methodological framework for building or testing theories of causal mechanisms in the social sciences, discussing how we can empirically evaluate the presence/absence of each part of a mechanism using informal Bayesian logic.

Theorizing causal mechanisms as systems

How can we understand causal mechanisms in a manner that enables us to flesh out the causal process between X and Y? While the answer to this question continues to divide methodologists working on PT, I argue that if we want to reap the analytical benefits of PT in terms of better theories and stronger evidence-based claims about causal processes, we need to unpack mechanisms into their component parts in enough detail to enable us to know exactly what we are attempting to trace, focusing in particular on the activities that provide

the causal links between each part. The following briefly reviews the debate about what causal mechanisms are, followed by the development of a set of methodological guidelines for how causal mechanisms can be theorized as systems composed of interlocking parts (see also the introduction to this special issue).

The simplest understanding of mechanisms found in the literature is to see them merely as a series of events, or a narrative story leading to an outcome (e.g. Abell, 2004; Mahoney, 2012: 571; Roberts, 1996:16; Suganami, 1996: 164-168). Yet while describing a series of events can provide a plausible *descriptive* narrative about *what* happened, they do not shed light on *why* things happened in the case. Causal explanation involves more than just tracing temporal sequences (Gryzmala-Busse, 2012: 1268-1271). In other words, the causal mechanism remains completely black-boxed.

There are a number of scholars who treat mechanisms as *intervening variables* between X and Y (e.g. George and Bennett, 2005: 6; Pearl, 2000; Morgan and Winship, 2004: 224-230; Gerring, 2007; King, Keohane and Verba, 1994: 87; Falletti and Lynch, 2009: 1146). Here the mechanism is often depicted in the form of a causal graph, with several nodes between the purported cause and outcome.

Strictly speaking, the use of the term variable implies a research design where we have variation, with multiple observations of the values of X, the intervening variable (mechanism, M) and Y that enable us to assess the net effects of X and M on values of Y (Gerring, 2007). King, Keohane and Verba suggest that to do this we can disaggregate our case, for example by transforming a single case into a set of observations over time, or by comparing two similar cases.² Yet both in effect transform the within-case study of causal processes into a cross-case analysis of patterns of variation at a lower level of aggregation, where we lose focus on the *process* between the variables (Mahoney, 2001: 578; Mayntz, 2004:244-245; Hedström and Ylikoski, 2010: 51-52; Waldner, 2012: 76-77).

² - Additionally, observations within a single case over time cannot be treated as independent, given that what happens at t_0 naturally impacts on what happens at t_1 .

The result of viewing mechanisms as intervening variables is that the actual causal process itself is *grey-boxed* (Bunge, 1997: 428; Waldner, 2012). Yet tracing the steps of the causal process between X and Y was the very reason we would want to engage in the PT of causal mechanisms in the first place.

Taking mechanisms seriously

This article contends that we should take mechanisms seriously as causal processes by treating them as *theoretical systems* linking causes and outcomes (Rohlfing, 2012; Beach and Pedersen, 2013; Waldner, 2012). In this understanding, a causal mechanism is defined as a theory of a system of interlocking parts that transmits causal forces between a cause (or a set of causes) and an outcome (Machamer, Darden and Craver, 2000; Machamer, 2004; Glennan, 1996, 2002; Bunge, 1997, 2004; Bhaskar, 1978; Waskan, 2011). Hernes has helpfully defined a mechanism as ‘...a set of interacting parts – an assembly of elements producing an effect not inherent in any one of them. A mechanism is not so much about ‘nuts and bolts’ as about ‘cogs and wheels’ – the wheelwork or agency by which an effect is produced.’ (Hernes, 1998: 78).

A theorized causal mechanism describes each part of the mechanism whereby causal forces are transferred from cause to outcome. The parts of causal mechanisms can be helpfully defined in terms of entities that engage in activities that transmit causal forces from cause to outcome (Machamer, Darden and Craver, 2000; Machamer, 2004; Beach and Pedersen 2013; Rohlfing, 2012). *Entities* can be understood as the factors (actors, organizations or structures) engaging in activities, where the *activities* are the producers of change or what transmits causal forces through a mechanism. Note that it is the activities that entities engage in that move the mechanism from an initial causal condition through different parts to an outcome. This denotes a key distinction from a theory as a causal graph in which the activities are not described, but are depicted merely as causal arrows linking one entity to the next. In the following I depict entities as letters (A, B, C) and activities as arrows ($\rightarrow$). A part of the mechanism is therefore the combination of entity and activity (e.g. $A \rightarrow$).

Ideally, when theorizing the parts of a causal mechanism, the parts will exhibit *productive continuity*, meaning that each of the parts logically leads to the

next part, and there are not large logical holes in the causal story linking X and Y together (Machamer, Darden and Craver, 2000: 3; Darden, 2002: 283). If a mechanism is represented schematically as: $X \rightarrow [A \rightarrow B \rightarrow C] \rightarrow Y$, productive continuity lies in the arrows and their transferal of causal forces from one part of a mechanism to the next. A missing arrow, namely, the inability to specify an activity connecting A and B, leaves an explanatory gap in the productive continuity of the mechanism (Machamer, Darden and Craver, 2000: 3). In real-world research, we are often not able to get much beyond what Machamer terms a 'mechanism sketch' (Machamer, 2004). However, even a crude depiction of a process in terms of *both* entities *and* activities is better than no information on process.

When conceptualizing a causal mechanism we should therefore be able to identify clearly the different parts and how they are related through the nature of the activities. Explicitly theorizing causal mechanisms results both in: 1) better causal theories, and 2) actual empirical tracing of causal processes in cases, enabling stronger inferences to be made as a result. Better causal theories result from having to make explicit each part of the causal link between a cause and an outcome, enabling detailed scrutiny of the causal logic of each of the links in the causal process. Second, by actually tracing causal processes in detail in an empirical case study, evidence is provided suggesting either that the mechanism was present, the mechanism should be revised because it did not work as theorized, or there is no causal link, suggesting that in the case there is no causal relationship. In the philosophy of science, the term 'mechanistic evidence' is used for this type of empirical material, focusing our attention on the connection between the type of empirical material used as evidence and the type of causal claim being made (Russo and Williamson, 2007; Illari, 2011). Mechanistic evidence is contrasted to 'evidence of difference-making', which is cross-case evidence of variation on values of X and Y gained from either experimental manipulation of causes or (in a much weaker form) observational research (Ibid).

It is vital to note that this template for theorizing causal mechanisms as theoretical systems can be applied to many different types of theoretical explanations. Naturally, what the entities and activities more precisely are

depends on the type of causal explanation and the level at which the mechanism works and the time span of its operation. However, many existing recommendations for theorizing mechanisms have particular theoretical positions hard-wired into them. For example, Goertz and Mahoney write that when studying mechanisms, we, '...inevitably carry out an over-time, processual analysis of the case... The analyst will normally identify historical junctures where key events directed the case toward certain outcomes and not others. She or he may well pause to consider how small changes during these junctures might have led the case to follow a different path...The overall explanation likely will be rich with details about specific events, conjunctures, and contingencies.' (2012: 89). All of this has an eerie resemblance to Historical Institutional (HI) theories of the importance of temporality, critical junctures and path dependency. Other (Rational Choice-inspired (RC)) scholars have claimed that we always need to theorize multiple mechanisms when studying macro-level phenomena, describing mechanisms linking the macro to the micro-level (situational mechanisms), micro-micro (action-based mechanisms), and micro to the macro-level (transformational mechanisms) (e.g. Hedström and Swedburg, 2008). Yet in both instances, scholars have imported theories into their recommendations for what theorized mechanisms should include. However, there is no logical reason why a mechanism has to look like an integrative, macro-micro-macro RC theory, or a highly contingent, HI theory with path dependencies and critical junctures. A theory is *not* a research methodology, and vice versa. Our theory should tell us what type of causal mechanism to expect, but theory should not dictate our understanding of mechanisms themselves.

Existing applications of PT often black-box mechanisms

As a representative example of how existing case study research that aims to trace mechanisms often end up either black- or grey-box mechanisms, I utilize Ziblatt's 2009 article where he claims to assess the '...causal process through which this relatively abstract and distant concept of landholding inequality actually operated on the ground leading to electoral fraud...' (p. 14).

In the article, the actual mechanism is never fleshed out in sufficient detail to allow us to either evaluate the underlying theoretical logic or to assess

systematically whether there is empirical evidence of the 'capture' mechanism.³ The closest Ziblatt gets to a theoretical description of a causal mechanism is where he writes, '... they [landed elites] exert influence indirectly via the capture of rural public officials such as mayors, county commissioners, police officials, and election officials, who in turn are the actors that interfere with free and fair elections. In its most acute form, capture occurs as socioeconomic interests infiltrate the state by using their own personnel to staff the state...' (2009: 14).

While telling us something about the process whereby landholding inequality (X) is linked to electoral fraud (Y) - it works through local officials - he does not provide sufficient details of the causal mechanism. In particular, he tells us precious little about the actual causal process whereby landed elites are able to capture local officials.

For instance, what types of power resources do landed elites deploy to capture officials? Does capture occur through the use of material resources such as the power to control revenue or through control of appointment processes? Or perhaps by deploying more discursive resources? And do landed elites have to *actively* intervene to capture officials, or do officials *anticipate* what local officials want? In the next step of the causal process, when and why should local officials be responsive? And once captured, what is the process whereby local officials actually engage in electoral fraud? What types of actions do they use? It could range from something like removal of voters from electoral rolls to pressuring poll officers, etc.?

By not explicitly theorizing the parts of a causal mechanism in sufficient detail, we are left unable to answer basic questions about the underlying causal logic linking X and Y together, making it difficult to evaluate whether the theorized causal process is logically consistent, and even more problematic, it is very difficult to trace systematically whether there is evidence of the process when we are not told what the process is that is being traced.

³ - In assessing hundreds of articles from top journals in the past decade, the author has found precious few examples of case studies that explicitly theorize causal process in-between X and Y. Therefore, when compared to almost all of the existing article-length research, Ziblatt's study provides us with much more theoretical information on the causal process than is typical. But this does not mean that Ziblatt's theorized mechanism could not be improved further (see below).

The subsequent empirical case study does illustrate that both X and Y are present and that there is some *anecdotal* evidence suggesting that the link is through something like a 'capture' process. However, given that the mechanism is not explicitly theorized, the actual analysis ends up merely producing small empirical vignettes that *insinuate* the existence of an underlying mechanism without providing strong empirical evidence. Indeed, it is difficult to determine whether the presented empirical material actually confirms the underlying mechanism, given that we are left guessing about what the underlying mechanism is.

The anecdotal nature of the evidence is clear in the following quote that is representative of the types of empirical material presented. Ziblatt writes, 'As one Landrat from Posen reported in his memoirs in 1894, "I had to join the local branch of the Agrarian League, because everyone I interact with socially – and everyone I hunt with – is a member!" (2009: 16). It is obvious that this piece of empirical material relates in some fashion to an underlying part of mechanism whereby landed elites can pressure local officials. But by not detailing the underlying causal mechanism, we are left unsure about basic questions such as whether social pressure is the only means whereby local officials are captured. Further, by not telling us what empirical fingerprints (predicted evidence) each of the parts of the mechanism can be expected to leave in a case - a natural result of not theorizing explicitly each of the parts -, the presentation of empirical material seems very unsystematic. A scholar critical of case study methodology might go so far as to state that it is just 'cherry picking' empirical observations instead of being a systematic analysis of whether there is empirical evidence that confirms or disconfirms the theoretical mechanism being present in the case.

Unfortunately, by not unpacking the mechanism in sufficient detail, Ziblatt's case study does not enable us to examine the underlying causal logic linking X with Y, nor does it enable us to assess systematically whether there is empirical evidence that the hypothesized 'capture' mechanism actually worked in a given case. The result is that the 'capture' mechanism remains in a 'grey box'.

The theorized mechanism could be improved by fleshing out the activities that provide the causal links between the two parts of the mechanism. While the mechanism could be fleshed out even more, I recommend keeping the theorized mechanism as simple as meaningful in relation to the research question. Tracing a mechanism with thirty parts empirically would be a daunting task, and there is the serious question of whether we could formulate a generalizable mechanism that is so complex.

Ziblatt's theory deals with how landed elites might be able to produce electoral fraud through 'capture'. To transform the theory into a mechanism, we would want to disaggregate the causal process into a series of parts. A suggested mechanism that is more explicitly described is depicted in figure 1, composed of two parts linking X and Y is theorized based upon the empirical insights drawn from the case study.

The core difference between the causal theory used by Ziblatt in his article and the 'capture' mechanism as described in figure 1 is that my revision *explicitly* theorizes what is going on in each part of the causal mechanism - focusing in particular on the activities that provide the causal links - thereby telling us much more about *how* and *why* capture is theorized to happen. The first part of the mechanism details the how and why process whereby local officials are captured, followed by theorizations on the process whereby captured officials influence elections, resulting in electoral fraud.

Causal condition (X)	Capture causal mechanism		Outcome (Y)
	Part 1 'pressuring'	Part 2 'crafting proposals'	
Landed elites seeking to preserve electoral dominance in context where traditional social power eroding	<u>Landed elites</u> put pressure using control of appointment processes and local revenue in an attempt to influence local public officials (capture)	<u>Local officials</u> exploit the powers of their office to benefit landed elites, using instruments such as removing voters from electoral rolls, etc.	Electoral fraud

--	--	--	--

Figure 1 – A ‘capture’ causal mechanism linking landed elites with electoral fraud.

Note : Entities are underlined and activities are in italics.

In the subsequent empirical analysis we would systematically trace whether there is evidence suggesting that the observable implications of the two parts of this simple mechanism were present instead of the empirical vignettes used by Ziblatt; vignettes that shed some light on the workings of the mechanism but in an ad hoc and unsystematic fashion. The empirical part of the PT case study would focus solely on tracing the ‘capture’ causal mechanism, asking whether there is empirical evidence of each of the parts in a given case. For example, do we find evidence of landed elites actually putting pressure on local officials through appointments? Do we find evidence of systematic removal of voters from electoral rolls by local officials, or were other instruments used? By making the theorized mechanism more explicit in terms of a system of interlocking parts that transmits causal forces from the cause (or set of causes) to the outcome, we then know where to look for the empirical fingerprints in the subsequent PT case study. We now turn to a short example of how causal mechanisms can be studied empirically, introducing the use of informal Bayesian logic as a tool to focus our attention on what questions we should ask when developing the observable manifestations of parts of mechanisms for empirical testing.

Studying causal mechanisms in practice - developing empirical fingerprints

In PT case studies, inferences are made by assessing the match between predicted and actually found evidence for each part of the mechanism. This can take the form of empirics-first research, where we collect empirical material and then evaluate what causal process it might be evidence of, or theory-first research, where clear observables are developed for each part of the theorized mechanism that can then be assessed empirically.

Mechanistic evidence as used in PT does *not* take the form of patterns of variation between values of X and Y *across* cases (i.e. evidence of difference-making). Instead, PT uses an inferential logic similar to that widely used in observational research in fields as diverse as paleontology, geography and biology, and in practical situations such as courts of law. The logic of inference in PT is inspired by the Bayesian approach, where the core idea is that science is about using new empirical evidence to update our confidence in causal theories being valid (Beach and Pedersen, 2013, 2016; Rohlfing, 2012; Bennett, 2008a).⁴

Evaluating evidence involves three steps when testing theories using PT: 1) making predictions about what empirical fingerprints each part of a mechanism would leave in a given case, and evaluating whether the predicted evidence in theoretically certain and/or unique in relation to the theorized mechanism, 2) collecting empirical evidence and assessing whether one actually found the predicted evidence, and 3) evaluating whether we can trust the found evidence. Only after we have evaluated both what found evidence can tell us in relation to theory (certainty and uniqueness) and whether we can trust the evidence, can we make inferences about the presence/absence of hypothesized mechanisms. When building theories, the sequence is basically reversed, although here there is considerable more back-and-forth between theories and empirics (for more, see Beach and Pedersen, 2016).

What kinds of empirical material can act as evidence for the existence of parts of mechanisms? Mechanistic evidence is not just a series of events between the occurrence of the cause and outcome, as events can be the empirical fingerprint of many different things. Evidence is also not just cross-case variation in values of X and Y. Instead, mechanistic evidence is *any* observable manifestation of our theorized causal mechanism(s) that has a probative value in relation to determining whether it was present or not in the case.

We can distinguish between four particular types of evidence that can be the fingerprints of a part of a mechanism. First, pattern evidence relates to statistical patterns in the empirical record. If we are testing a theory relating to

⁴ - Note that there is considerable debate about whether this use of Bayesian logic should be formalized in terms of quantified probabilities, or used in a more informal, Folk Bayesian fashion as in law. For more on this discussion, see Beach and Pedersen (2016).

racial discrimination in employment, we would expect that we would find statistical differences in rates of employment in the empirical record. A second form of evidence relates to sequences, making predictions about temporal or spatial chronologies of events. When testing a theory about rational decision-making, relevant evidence might be whether decision-makers first collected all relevant information, evaluated it, and then took the decision that they believed best solved the policy problem they face. If the analyst then found that the decision was taken before information was collected (as often happens in the real world), we would downgrade our confidence in the theory being valid in the case. Third, trace evidence refers to material where its mere existence provides proof. If our theory states that lobbyists had to meet with decision-makers, and there were few alternative explanations in the case for them meeting, finding that they did would strengthen our confidence in the theory. Fourthly, account evidence relates to material where it is the content that matters. This can be in the form of what participants tell us in interviews, or the content of relevant documents like legislative proposals.

When testing theories of mechanisms, developing predictions of what evidence we might find involves answering the question ‘what empirical fingerprints might each part of the causal mechanism, if valid, leave in the selected case’? Here we need to think carefully but also creatively about the data-generating processes that the operation of the theorized causal process will leave. As one piece of evidence is usually not conclusive, it is important to game through as many fingerprints as possible that the data-generating processes might have can have left in the case. Predictions are also typically case-specific, as the same mechanism might leave very different empirical fingerprints in different cases.

In contrast, when we are building theories of mechanisms, we engage in an extensive soaking and probing a case, starting by reading existing accounts, and probing available empirical records in a search for patterns in the record that might be empirical fingerprints of underlying causal processes. When we uncover an empirical observation that we have a hunch might be an empirical fingerprint of a causal process, we start to engage in a form of backwards induction to uncover the underlying parts of the mechanism in operation, asking

ourselves what mechanisms the found material is potentially evidence of. Here the process of the assessment of evidence described below flows in the opposite direction, from empirics to theory.

After we have described what the predicted evidence is in a theory-testing design, we then need to elaborate what the evidence can tell us about the studied causal relationship (van Evera, 1997). Here there are two relevant questions to ask: 1) do we have to find the evidence? (theoretical certainty), and 2) if found, are there plausible alternative explanations for finding it? (theoretical uniqueness). Theoretical certainty relates to the disconfirming power of evidence. E.g. the suspect being in the town when a murder is committed is a certain prediction. If the predicted evidence is not found, then we can disconfirm the hypothesis that the suspect is guilty. Theoretical uniqueness describes the expected probability of finding the predicted evidence if the hypothesis is not true, and relates to the confirming power of evidence. E.g. a unique prediction would be expecting to find 'smoking gun' evidence linking the suspect to the crime (powder burns on the suspect's hands, suspect's fingerprints are on the gun). If we found this predicted evidence, it would be difficult to account for this unless the suspect actually used it to commit the crime (strong confirmation). However, we would not be certain to find the smoking gun, and if we did not find it we would be unable to make any inferences beyond 'we did not find the smoking gun'. Note that theoretical uniqueness is a form of empirical 'control' for other causes, but relates to other causes for finding the evidence itself and therefore is not typically drawn from rival theoretical claims about the causes of a given outcome (Rohlfing, 2014). Rival theories are only relevant when they actually provide a competing explanation for finding a particular piece of evidence.

After predictions are developed for each part of the mechanism in a theory-testing case study, empirical material is then collected and assessed before inferences are made. This empirical assessment involves three steps: 1) assessing whether we have actually found the predicted evidence (content assessment), 2) assessing what it means if we do not find the predicted evidence in the case (what is the probability that we will empirically observe the predicted evidence?), and 3) assessing whether we can actually trust the found evidence

(what is the probability that our measure is accurate)? Only after we have transparently assessed each question can we make inferences based on the actual empirical evidence collected.

As in a court of law, before we can admit observations as evidence upon which we can make inferences, we first need to assess using contextual knowledge whether the observations we have collected match the evidence we predicted we would find. Second, if the predicted evidence was not found, the research has to assess what this means. While an empirical fingerprint might be theoretically certain, if we are unable to gain access to a particular record in the archives that might enable us to assess whether the evidence is there or not, we would not be able to downgrade our confidence in the hypothesis based on not finding the evidence. Only when we have engaged in an extensive search and we are able to document that we had access to all known sources, can we claim that by not finding the predicted theoretically certain evidence we have downgraded our confidence in the hypothesis being valid (here absence of evidence would be evidence of absence). Finally, we need to engage in extensive source criticism to evaluate whether we can trust the found evidence, or whether the source might have motives for producing a biased account. Other things equal, we attribute stronger evidential weight to pieces of evidence that we can trust.

Taken together, after the empirical evidence is evaluated for what it can tell us theoretically (certainty and uniqueness) and whether we can trust it empirically, are we able to sum together the different pieces of evidence for each part of the mechanism, enabling us to infer whether there is evidence of the mechanism being present in a given case.

Conclusion

If we want to reap the methodological benefits of PT as a tool for gaining greater understanding of causal processes using in-depth case studies, we need to take the study of mechanisms seriously by explicitly theorizing them in sufficient detail to enable us to know exactly what it is we should be tracing empirically.

PT is by no means a panacea – and in many respects is a very limited methodological tool. Like an electron microscope, it has only a few different uses, but what it does it does powerfully. PT requires the deployment of extensive

analytical resources, meaning that one is only able to realistically assess a small number of cases. Additionally, PT taken on its own only enables within-case inferences about causal processes, meaning that PT case studies have to be nested in comparative designs to enable cross-case generalizations to be made. But by tracing mechanisms using in-depth case studies we can analytically pry open the causal processes that link causes and outcomes.

This article has provided a set of clear methodological framework for tracing mechanisms by suggesting they should be: 1) understood as systems linking causes and outcomes, 2) unpacked into a series of interlocking parts composed of entities engaging in activities that transfer causal forces from one part to the next, and 3) operationalized by developing predictions of what evidence we should find if each part of the mechanism is present.

Naturally there will be theoretical claims that are more difficult to translate into causal processes than others, in particular when there are feedback loops and other forms of non-linear processes. In addition, just because we have made the causal mechanism explicit does not mean we can trace each of its parts empirically, given that especially psychological processes might not leave sufficient empirical fingerprints to determine whether they were present or not in a case. However, unless we explicitly theorize mechanisms in sufficient detail, we do not know what we should be tracing in PT.

References

- Beach, Derek and Rasmus Brun Pedersen. 2013. *Process-tracing Methods: Foundations and Guidelines*. Ann Arbor: University of Michigan Press.
- Beach, Derek and Rasmus Brun Pedersen. 2016. *Causal Case Studies: Comparing, Matching and Tracing*. Ann Arbor: University of Michigan Press.
- Bennett, Andrew. 2008a. Process-Tracing: A Bayesian Perspective. In *The Oxford Handbook of Political Methodology*, ed. Janet M. Box-Steffensmeier, Henry E. Brady and David Collier, 702-721. Oxford: Oxford University Press.
- Bennett, Andrew. 2008b. The mother of all “isms”: Organizing political science around causal mechanisms. In *Revitalizing causality: realism about causality in philosophy and social science*, ed. Ruth Groff, 205-219. London: Routledge.
- Bunge, Mario. 1997. Mechanism and Explanation. *Philosophy of the Social Sciences* 27(4): 410-465.
- Bunge, Mario. 2004. How Does It Work?: The Search for Explanatory Mechanisms. *Philosophy of the Social Sciences* 34(2): 182-210.
- Checkel, Jeffrey T. 2008. Tracing Causal Mechanisms. *International Studies Review* 8(2): 362-370.
- Darden, Lindley. 2002. Strategies for Discovering Mechanisms: Schema Instantiation, Modular Subassembly, Forward/ Backward Chaining. *Philosophy of Science* (Supplement PSA 2000 Part II) 69: S354-S365.
- George, Alexander L. and Bennett Andrew. 2005 *Case studies and Theory Development in the Social Sciences*. Cambridge Massachusetts London: MIT Press.
- Gerring, John. 2007. *Case Study Research*. Cambridge: Cambridge University Press.
- Glennan, Stuart S. 1996. Mechanisms and the Nature of Causation. *Erkenntnis* 44(1): 49-71.
- Goertz, Gary and James Mahoney. 2012. *A Tale of Two Cultures*. Princeton: Princeton University Press.
- Hall, Peter A. 2008. Systematic Process Analysis: when and how to use it. *European Political Science* 7(3): 304-317.
- Hedström, Peter and Richard, Swedberg (ed). 1998. *Social Mechanisms an Analytical Approach to Social Theory*. Cambridge: Cambridge University Press.

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

Hernes, Gudmund. 1998. Real Virtuality. In *Social Mechanisms an Analytical Approach to Social Theory*, ed. Peter Hedström and Richard Swedberg, 74-101. Cambridge: Cambridge University Press.

Illari, Phyllis McKay. 2011. 'Mechanistic Evidence: Disambiguating the Russo-Williamson Thesis.', *International Studies in the Philosophy of Science*, 25(2): 139-157.

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

Machamer, Peter, Lindley Darden and Carl F. Craver. 2000. Thinking about Mechanisms. *Philosophy of Science* 67(1): 1-25.

Machamer, Peter. 2004. Activities and Causation: The Metaphysics and Epistemology of Mechanisms. *International Studies in the Philosophy of Science* 18(1): 27-39.

Mahoney, James. 2001. Beyond Correlational Analysis: Recent Innovations in Theory and Method. *Sociological Forum* 16(3): 575-593.

Mahoney, James 2012. The Logic of Process Tracing Tests in the Social Sciences., *Sociological Methods and Research* 41(4): 570-597.

Rohlfing, Ingo. 2012. *Case Studies and Causal Inference*. Houndmills: Palgrave Macmillan.

Rohlfing, Ingo. 2014. 'Comparative hypothesis testing via process tracing.', *Sociological Methods and Research* 43(4): 606:642.

Russo, Federica and Jon Williamson. 2007. Interpreting Causality in the Health Science. *International Studies in the Philosophy of Science*, 21(2): 157-170.

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

Waldner, David. 2012. Process Tracing and Causal Mechanisms. In *Oxford Handbook of the Philosophy of Social Science*, ed. H. Kincaid, 65-84. Oxford: Oxford University Press.

Waskan, Jonathan. 2011. 'Mechanistic explanation at the limit.', *Synthese*, (2011) 193: 389-408.