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. & Pedersen, R. B. (2016). Selecting Appropriate Cases When Tracing Causal Mechanisms, *Sociological Methods & Research*. doi: 10.1177/0049124115622510

Publication metadata

Title:	Selecting Appropriate Cases When Tracing Causal Mechanisms
Author(s):	Beach, D. & Pedersen, R. B.
Journal:	<i>Sociological Methods & Research</i>
DOI/Link:	https://doi.org/10.1177/0049124115622510
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.

Selecting appropriate cases when tracing causal mechanisms

Derek Beach and Rasmus Brun Pedersen

1. Introduction

The last decade has witnessed resurgence in the interest in studying the causal mechanisms *linking* causes and outcomes in the social sciences. This article explores the overlooked implications for case selection when tracing mechanisms using in-depth case studies. Our argument is that existing case selection guidelines are appropriate for research aimed at making cross-case claims about causal relationships, where case selection is primarily used to control for other causes.

However, existing guidelines are not in alignment with case-based research that aims to trace mechanisms, where the goal is to unpack the causal mechanism *between* X and Y, enabling causal inferences to be made because empirical evidence is provided for how the mechanism actually operated in a particular case. The in-depth, within-case tracing of how mechanisms operate in particular cases produces what can be termed *mechanistic* evidence, which can be contrasted to *difference-making* evidence that are produced by variance-based, experimental designs.¹

In observational variance-based designs, case selection principles follow the principles of experimental research where the goal is to select cases where we can assess the difference that a given cause (aka treatment) has on values of the outcome, holding all other potential causes of difference constant. In contrast, the case selection principles that are appropriate when working with mechanistic evidence are quite different. Here the focus is on building as causally homogeneous a population as possible to enable us to infer from the studied cases to the rest of the small, bounded population of causally similar cases.²

There is intense debate about how mechanisms should be understood in the social sciences,³ and also about the nature of case studies and how to study mechanisms properly

¹ - Russo and Williams, 2007. This distinction overlaps to some degree with ‘data-set’ versus ‘causal process’ observations, but is even more focused on the relationship between empirical material and causal inferences.

² - Note that there are research purposes where studying only a particular historical outcome are relevant, such as the Cuban Missile Crisis (e.g. Allison and Zelikow, 1999). When research is aimed at accounting for the ‘big and important’ aspects of a particular case, the case is not a ‘case of’ a narrow theoretical phenomenon. Here cases are selected because they are the French Revolution, or the Cuban Missile Crisis, , making case selection strategies irrelevant. For a discussion of this, see AUTHOR.

³ - See Gerring, 2010 and Heström and Ylikowski, 2010 for good introductions to debates on mechanisms. For the systems understanding of mechanisms, see e.g. Heström and Ylikowski, 2010; Bunge, 1997, 2004; Machamer, Darden and Craver, 2000; Machamer, 2004; Waldner, 2012; Waskan, 2011; Hedström and Swedburg, 1998.

using case studies.⁴ While we discuss different interpretations briefly in the next section, this article is not about these debates. Instead, this article explores the implications that a mechanism-as-systems understanding has when we take individual cases as our analytical point of departure.⁵ where causation is viewed in deterministic and asymmetric terms at the case-level. In this type of research, the focus is on ensuring causal homogeneity to enable cross-case inferences to be made to small, bounded populations of similar cases. We and finally where mechanisms are understood as more than just intervening variables but instead a system of interacting parts that transfers causal forces from causes to outcomes.⁶

Yet existing guidelines for case selection⁷ are not in alignment with the underlying assumptions of case-based research and mechanisms as systems. Despite recent guidelines going beyond the ‘all cases’ set of guidelines, where any case that is ‘representative’ of a causal relationship in terms of being on or close to the regression line,⁸ to focus more particularly on ‘pathway’ cases when studying mechanism,⁹ i.e. that are positive on X and Y but where all other potential causes are not present, existing guidelines are still more in alignment with the assumptions underlying variance-based designs, where the role of case selection is to *control* for other causes, enabling the production of evidence of difference-making across a set of cases.¹⁰

However, case studies play an adjunct role in variance-based designs, acting as a double-check to ensure that there are no important omitted causal variables that could act as confounders of cross-case causal patterns,¹¹ and to make plausible using counterfactuals in a case study the claim that found cross-case correlations actually are causal relationships.¹²

If we instead adopt the underlying assumptions of case-based designs and the in-depth tracing of mechanisms as systems,¹³ existing case selection guidelines become problematic. In case-based research, the purpose of case selection is to enable cross-case inferences from the studied case to causally similar cases; case selection is *not* about controlling for other causes. This ‘control’ is achieved at the empirical level by assessing whether a particular piece of

⁴ - Here there are two positions: one that takes cases as the point of departure, with the result that case-based research is viewed as being a different ‘culture’ than variance-based designs (Brady and Collier, 2010; George and Bennett, 2005; Goertz and Mahoney, 2012; AUTHOR). This contrasts with variance-based research that takes the population and mean causal effects at this level as the point of departure, resulting in a ‘one logic’ position (Gerring, 2011; King, Keohane and Verba, 1994).

⁵ - For more on case-based research and the assumptions underlying it, see Ragin, 2000, 2008; Goertz and Mahoney, 2012; George and Bennett, 2005; AUTHOR.

⁶ - Machamer, Darden and Craver, 2000; Machamer, 2004; Waldner, 2012.

⁷ - Gerring, 2007; Schneider and Rohlfing, 2013; Nielsen, 2014; Lieberman, 2005.

⁸ - Lieberman, 2005; Nielsen, 2014.

⁹ - Gerring and Seawright, 2007; Weller and Barnes, 2014.

¹⁰ - Gerring and Seawright, 2007; Nielsen, 2014; Schneider & Rohlfing, 2013.

¹¹ - This is most explicitly seen in Nielsen’s case selection guidelines (Nielsen, 2014).

¹² - Nielsen, 2014; Lieberman, 2005.

¹³ - Goertz and Mahoney, 2012; AUTHOR; Waldner, 2012.

empirical evidence of a part of a mechanism can be accounted for by any plausible alternative explanation. And as we will develop later, if we did select a case where only one cause is present to trace a mechanism between this X1 and Y, we would not be able to infer that similar mechanisms link X1 and Y in cases where other causes are also present.

As will be developed in this article, existing guidelines that for instance suggest we should also select negative cases for studying mechanisms,¹⁴ where neither the cause nor outcome are present, do not fit with the assumptions of case-based research. Yet studying mechanisms in particular cases focuses on using in-depth empirical analysis of what *actually* happened in the case to shed light on how a cause (or set of causes) produces an outcome through the operation of a mechanism.

If we are interested in tracing a causal mechanism *linking* X and Y, we want to trace it in cases where it could have been present, at least in theory. Tracing a non-existent mechanism in a case where we a priori knew it was not present tells us nothing about how the mechanism works in cases where it is present. Using case selection guidelines that are appropriate for variance-based designs when tracing mechanisms using in-depth case studies creates the risk that we select analytically irrelevant cases, or that we make flawed inferences to other cases based on studying cases where only X or only Y are present.

Tracing causal mechanisms using in-depth case study methods like process-tracing can have two functions that result in the selection of different cases.¹⁵ In what we term *mechanism-centered* designs, mechanisms are traced to gain better understanding of the causal relationship between X and Y within a bounded population, resulting in mechanistic evidence that sheds light on causal processes. Here the theorized mechanisms are front-and-center in our analysis, with the analytical focus on assessing how the causal arrow(s) in-between X and Y actually works in particular cases. In contrast, tracing mechanisms can be an auxiliary tool in relation to the main goal of revising a causal theory, for example tracing mechanisms to uncover when a mechanism breaks down that provides information that is used in a systematic comparison with a case where the mechanism works to shed light on omitted causal or contextual conditions from the model. We term this a *condition-centered* design. We explore which cases are appropriate to select in the two different types of research.

¹⁴ - Most explicitly found in Nielsen, 2014: 3-9, where he discusses case selection using a most-similar system design, selecting both the positive-positive and negative-negative cases, with other potential explanatory factors held constant (see also Tarrow, 2010). In Lieberman (2005) and Gerring and Seawright (2007), typical cases are all cases that are regression-line on-liers, meaning that some will be what we understand as negative-negative cases (no X, no Y present).

¹⁵ - There is also a third variant that we do not explore here: idiographic, case-centered designs that attempt to develop comprehensive explanations of a particularly important historical outcome (e.g. the Cuban Missile Crisis, see Allison & Zelikow 1999). For more, see AUTHOR. Also Van Evera, 1997; Eckstein, 1975.

The article proceeds in three steps. We first develop the underlying assumptions of studying mechanisms using case-based methods, focusing on mechanisms understood as systems, ontological determinism, causal asymmetry and causal homogeneity and the importance of context. We then develop a set of case selection guidelines that are in methodological alignment with these underlying assumptions. Section 4 develops guidelines for research where the mechanism is the primary focus, contending that only typical cases where both X, Y and the requisite contextual conditions are present should be selected. We compare our guidelines with the existing, finding that practices like selecting most/least-likely cases are not compatible with the underlying assumptions of tracing mechanisms. Section 5 present guidelines for deviant cases, focusing on tracing mechanisms until they breakdown as a tool to shed light on omitted contextual and/or causal conditions.

2. The assumptions of case-based tracing of mechanisms

This section discusses the four assumptions that arguably form the ontological foundations for studying causal mechanisms using in-depth, case-based methods like process-tracing.¹⁶ Again, the purpose of this article is not to justify this set of assumptions; all of which have been discussed extensively in recent work on case-based research in general, and the study of mechanisms using process-tracing in particular. Instead, the article focuses on the methodological implications they have for case selection guidelines.

2.1. Mechanisms-as-systems

At the most fundamental level, the understanding of mechanisms-as-systems builds on a different ontological understanding of causation from a counterfactual/manipulation understanding of the relationship between causes and outcomes.¹⁷ Causation is seen as more

¹⁶ - For recent developments in process-tracing methods, see Bennett and Checkel, 2014; and AUTHOR.

¹⁷ - Note there is some debate about whether mechanisms are really a different understanding of causation, as advocated by Waskan (2011), or whether the mechanism understanding at the deepest level are just more elaborate forms of counterfactual statements about the links between the parts of the mechanism (Woodward, 2005). Woodward claims that adherents to the mechanism understanding cannot logically account for what links each part of a mechanism to each other, and therefore he contends that mechanisms bottom-out in counterfactual claims that if a given part of a mechanism did not exist, the next part of the mechanism would also not have existed. While persuasive and potentially damaging, Waskan has tried to counter this critique through relatively pragmatic claims that what we are interested in studying when tracing mechanisms is not the hypothetical, ‘what if’ (sometimes terms ‘how possible’) that links each part with each other, but how the parts of system actually conspire to produce the happening (2011: 393). We contend that this debate is irrelevant for case selection, given that Woodward is here talking about the links within mechanisms and *not* the relationship between cause and outcome.

than just X:Y patterns of constant conjunction, or manipulation of X to produce a difference in the value of Y. Given the need to ensure methodological alignment between the nature of the causal relations we want to study and the methods we use to investigate them,¹⁸ adopting a mechanism understanding of causation has significant methodological implications. Understanding causation in terms of mechanisms means that the focus of analysis is on unpacking the causal arrow(s) linking cause and outcome into a series of interlocking and interacting parts that transmit causal forces from X to Y. Only by unpacking the causal arrow and studying it empirically by tracing how the mechanism works in a particular case can we make strong inferences that a causal relationship is present in a case by using what can be termed ‘mechanistic evidence’.¹⁹

How then should we define a mechanism? Glennan defines a mechanism as a ‘complex system, which produces an outcome by the interaction of a number of parts.’ (1996: 52). Andrew Bennett has defined causal mechanisms as ‘...processes through which agents with causal capacities operate in specific contexts to transfer energy, information or matter to other entities.’ (2008b: 207). A causal mechanism is a theorized link²⁰ between a cause (or set of causes) and an outcome, where each part of the mechanism is clearly described in an ordered sequence, and in particular in terms of entities engaging in activities that transfer causal forces.²¹

Each part of the theorized mechanism can helpfully be disaggregated into entities engaging in activities, with entities being the factors engaging in activities, whereas the activities are what actually transmit causal forces from X to Y.²² In the system understanding, each of the parts of a causal mechanism has no independent existence in relation to producing Y; instead they are integral parts of a system that transmits causal forces to Y. They are therefore not intervening variables; instead, theorizing a mechanism can be thought of as developing the causal story linking X with Y in more or less abstract terms. A good theory of

¹⁸ - Hall, 2003.

¹⁹ - Russo and Williams, 2007. While we advocate a systems-understanding, the arguments in this article also are applicable when using more ‘minimalist’ understandings of mechanisms where the causal process is not unpacked in any detail. In both understandings of mechanisms, what is being produced is within-case evidence of a causal process operating in an actual case, i.e. we are studying the empirical fingerprints left by the causal arrow in-between X and Y. In both understandings, mechanistic evidence does not tell us anything about cross-case difference-making – something that requires cross-case variation.

²⁰ - Note that the mechanism itself is a theoretical abstraction. Once it has been operationalized in terms of developing the observable manifestations for each of its parts, it has a potential empirical presence.

²¹ - Machamer, Darden and Craver, 2000; Machamer, 2004; Hernes 1998; Hedström and Swedburg, 1998; AUTHOR.

²² - Machamer, Darden, and Craver, 2000.

a mechanism does not have logical holes, thereby ensuring what Machamer, Darden and Craven term the 'productive continuity' of the mechanism in-between causes and outcomes.²³

The main contending understanding of mechanisms is to see them as intervening variables.²⁴ However, strictly speaking the use of the term variable implies a research design where we have variation that provides evidence of the difference-making of independent and/or intervening variables, with multiple observations of the values of X, the intervening variable (mechanism, M) and Y that enable us to assess the mean causal effects of X and/or M on values of Y.²⁵ King, Keohane and Verba suggest that to do this we can disaggregate our case, for example by transforming a single case (national-level) into multiple cases by observing values of variables at a lower-level (e.g. the state-level), or by comparing the same case at different times (before/after). Yet this in effect transforms the within-case study of causal mechanisms into a cross-case analysis of patterns of variation where we lose focus on the *process between* the causes and outcome.²⁶ Yet studying the process between X and Y was the very reason we would want to trace causal mechanisms in the first place. The analytical result is that the actual mechanism is 'black-boxed'. Another way of getting 'variation' is, as Nielsen suggests, to use counterfactual thought experiments to assess the difference a mechanism makes when it is present and absent. However, this suggestion uses hypothetical 'thought experiments' as evidence, whereas in case-based tracing of mechanisms we investigate the actual empirical fingerprints that they leave in actual cases, assessing using mechanistic evidence the degree of match between predicted evidence and the facts of the case in hand in order to make inferences.²⁷ In other words, tracing mechanisms using case-based methods like process-tracing makes inferences *not* by assessing difference-making by varying causes and mechanisms, but instead by assessing whether there is empirical evidence in the actual case that enables us to infer the mechanism was present.²⁸

The analytical value-added of tracing mechanisms-as-systems is that we both gain greater confidence in a causal relationship actually existing between X and Y because we have found empirical evidence of it in a case study, and we gain a greater understanding of how X is causally related to Y. Reaping these gains requires that the mechanism between X and Y is theorized explicitly, typically a mid-range theory that can be present within a

²³ - Machamer, Darden and Craver, 2000: 3.

²⁴ - e.g. Gerring, 2007; King, Keohane and Verba, 1994.

²⁵ - Gerring, 2007.

²⁶ - Mahoney, 2001: 578; Mayntz, 2004:244-245; Hedström & Swedberg 1998; Hedström and Ylikoski, 2010: 51-52.

²⁷ - Bennett and Checkel, 2014; Waldner, 2012; AUTHOR.

²⁸ - Russo and Williams, 2007.

bounded population of cases. However, when we look at the literature, many scholars claim they are theorizing mechanisms but never unpack the causal process in-between X and Y.

For example, Ziblatt claims that he theorizes a 'capture' mechanism that details how landholding inequality (X) is linked to electoral fraud (Y) (2009: 12-18). However the actual causal mechanism linking with X and Y is not explicitly theorized in the article. The closest we get to the theoretical development 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 between X and Y – i.e. it goes through local officials - he does not detail a causal mechanism in enough detail to trace it empirically. In particular, he tells us precious little about the actual causal process whereby landed elites are able to capture local officials. For example, what types of power resources do landed elites deploy to capture officials?

By not theorizing the mechanism more explicitly, we cannot evaluate the underlying causal logic, nor are we able to assess systematically whether there is empirical evidence that the hypothesized 'capture' mechanism actually worked in a given case. The evidence provided for the mechanism existing is therefore anecdotal, at best. For instance, he 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.

If we were going to trace Ziblatt's proposed 'capture' mechanism, we would first want to flesh out the mechanism in more detail. This can be done by theorizing about it in terms of a series of interlocking parts linking landed elites with electoral fraud, drawing on existing theorization and empirical research to better understand each of the causal links. In figure 1 we depict a slightly revised mechanism that more explicitly describes the activities that provide the causal links in the mechanism, 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.

>> INSERT FIGURE 1 HERE <<

Before we turn to the other assumptions underlying case-based research in mechanisms, it is important to note that mechanisms linking causes and outcomes do *not* logically have to be sufficient or necessary to produce Y. While the assumption that a mechanism is sufficient to produce an outcome is found in the literature,²⁹ there are no logical reasons the mechanism itself has to be necessary or sufficient in relation to the outcome. A mechanism logically cannot be causally more or less than the cause that triggers its operation. If we theorize that a cause like economic development is 'merely' a contributing cause of democratization, the causal mechanism linking the two cannot be theorized to be sufficient to produce democratization. The mechanism is merely the link, and therefore cannot be causally more important than the cause or set of causes that triggers the mechanism. This point has implications for case selection that have been overlooked in the existing literature; we return to this in the next section.

2.2. Case-based methods and determinism

The argument that case-based research builds on deterministic ontological foundations is well-established in social science methodology³⁰, but also widely misunderstood. The case selection guidelines developed in this article build on the assumption that mechanisms are deterministic at the ontological level.

Case-based research takes as its point of departure causal relationships within individual cases, whereas variance-based research builds on population-level mean causal effects in which probabilistic assumptions make most sense.³¹ Without deterministic claims being made, the study of individual cases makes little sense because if we do not find evidence of a relationship in a single case we will not know whether that was an exception to an otherwise strong cross-case relationship or not. This is why case studies play a secondary role in probabilistic, variance-based methods, acting as a means to double-check the findings of cross-case analyses of mean causal effects.³²

When operating with deterministic assumptions, if we find strong disconfirming mechanistic evidence in a particular case, we would not just discount this as an exception to

²⁹ - Mahoney, 2001: 580; Mayntz, 2004: 241. 253; Anderson, 2012: 416; Waskan, 2011: 403.

³⁰ - E.g. Goertz and Mahoney, 2012; Schneider and Wagemann, 2012; Mahoney, 2008; AUTHOR.

³¹ - For an excellent overview of this debate, see Mahoney, 2008.

³² - This is arguably seen in Lieberman's nested approach, where the starting point is a statistical analysis at the cross-case level.

an otherwise strong cross-case trend. Instead, we are forced to reappraise our theory of the mechanism, attempting to figure out why what we expected did not occur in the studied case, e.g. because an unknown scope condition that had to be present was not there.³³ These 'failures' are intensely interesting for case-based research, enabling us to build better theories of mechanisms.

Ontological determinism does not imply that we have perfect knowledge about why things happen in the empirical world. While the social world at the ontological level can be claimed to be deterministic (i.e. things do not happen randomly), our empirics-based knowledge about why things occur will always be imperfect.³⁴ *Epistemological* probabilism means that if we find disconfirming evidence of a mechanism in a case, we would only make disconfirming inferences *after* we have assessed transparently the level of trust we can have in the measurements being accurate.³⁵ The more we can trust our empirical measures, the stronger the inferences we can make. If we cannot justify transparently that our measures are relatively accurate, we would not begin revising the theorized mechanism based on one failure.

The ontological assumption of determinism fits well with the understanding of mechanisms-as-systems, viewing mechanisms as invariant, deterministic processes that will be triggered if the cause and relevant scope conditions that allow it to function are present.³⁶ Some argue that mechanisms as systems can also be probabilistically related to outcomes,³⁷ but here we contend that what might empirically appear to be a probabilistic relationship is typically the product of an inadequate understanding of the scope conditions under which a given mechanism functions.

2.3. Causal homogeneity and the importance of context for mechanisms

Ontological determinism has the down-stream consequence that case-based research has to operate with small, bounded populations of cases that are causally homogeneous.³⁸ Naturally, if we want to make inferences to a broader, more heterogeneous population, we have to hedge our bets by only making probabilistic claims, whereas we can make deterministic causal claims about small populations - as in case-based research.³⁹ A causally homogeneous

³³ - Dion, 2003: 106-107; Adcock, 2007.

³⁴ - The combination of ontologically deterministic theories with epistemological probabilism is well-known in the Bayesian approach to science. See Howson and Urbach, 2006.

³⁵ - For more, see AUTHOR.

³⁶ - Anderson, 2012; Mayntz, 2004; Waldner, 2012.

³⁷ - E.g. Hedström and Ylikowski, 2010: 51.

³⁸ - Collier and Mahoney, 1996: 68-69; Goertz, 2006: 29; Mahoney, 2007; Skocpol and Somers, 1980.

³⁹ - Skocpol and Somers, 1980; Mahoney, 2008.

population is where the same cause (or set of causes) is linked to the same outcome through the same causal mechanism(s). If we are able to empirically demonstrate that the population of cases is causally homogeneous by using comparative, we can then make cross-case inferences from studied cases to other causally similar cases. But comparative methods by themselves do not enable causal inferences (unless they are experiments that provide difference-making evidence); here the causal inferences are enabled by the use of within-case, mechanistic evidence.

Causal homogeneity is particularly important when studying causal mechanisms between X and Y. For mechanisms to operate correctly in a case, the requisite scope conditions need to be present.⁴⁰ Contextual (or scope) conditions are 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.’⁴¹ The same cause in a different context can in principle be linked to the same outcome through different mechanisms, meaning there can exist equifinality at the level of mechanisms unless we can demonstrate that the population of cases is causally homogenous also with regards to contextual conditions present.⁴² For example, hypothetically economic development (X) and democratization (Y) can be linked through a number of different mechanisms, depending on the contextual conditions.

Given that we almost never have complete knowledge of the contextual conditions that enable or prevent mechanisms from being triggered, or that result in different mechanisms linking the same cause and outcome, there is a very real risk of making flawed cross-case inferences about mechanisms being present unless we both: 1) only operate with small, bounded population of cases to which we intend to infer to, and 2) always trace mechanisms in multiple cases to ensure that the mechanisms we found in one case also are present in other cases.

2.4. Case-based methods and asymmetric causation

The final assumption that has implications for case selection guidelines is that case-based tracing of mechanisms involves making asymmetric causal claims. While it is well-known that case-based research makes asymmetric claims,⁴³ what is less recognized is that asymmetry becomes even more important when we are studying what causally links X and Y

⁴⁰ - Falletti and Lynch, 2009; Goertz and Mahoney, 2004: 655, 660-661.

⁴¹ - Falletti and Lynch, 2009: 1152.

⁴² - Gerring, 2010; George and Bennett, 2005, Chapter 10; Falletti and Lynch, 2009.

⁴³ - Ragin, 1987; 2000, 2008; Goertz and Mahoney, 2012.

together. An asymmetric causal claim could be that economic development (X) is necessary for democratization to occur (Y). Here no claims are made about whether less-developed countries ($\sim X$) are not democratic ($\sim Y$).⁴⁴

When we then are interested in tracing what mechanism(s) link X and Y together, we are making strongly asymmetric claims. For example, making claims about the causal mechanism that links mutual democracy (X) with avoidance of war in crisis situations (Y) does not make any claims about what mechanisms link autocracies ($\sim X$) and war ($\sim Y$). In cases where $\sim Y$ is the outcome, we would expect very different mechanisms to produce this outcome. Therefore, it makes little sense to compare the mechanisms that produce an outcome with those that produce its negation, given that they would be very different at both the theoretical level and in their empirical manifestations.⁴⁵

The strong asymmetry of mechanism-based claims implies that we only operate with *difference-of-kind* distinctions when conceptualizing cases, meaning cases are either in the set of cases where a causal relationship is possible or not (we return to this below in more detail). Differences-of-degrees that are captured using ordinal or higher scales are either causally irrelevant in relation to the core categorical distinction included in a given concept (the kind-difference), or at worst conflate degree- and kind-differences, resulting in a causally heterogeneous population where different values have different causal properties instead of mere degree-differences in the magnitude of the causal relationship.⁴⁶ In the rest of this article, we only operate with categorical in/out distinctions that capture the causally relevant differences-of-kind in concepts.

3. Case selection only after a causally homogeneous population of cases is mapped

Selecting appropriate cases requires some form of prior cross-case knowledge of the population of the given theoretical relationship. As discussed above, key to our ability to infer that the causal relationships we found in the single case are present in other similar cases is that we can justify that the set of cases we are working with is causally homogeneous through the use of comparative methods.⁴⁷ Therefore, we have to produce a mapping of the population in terms of what cases are members of the set of X, Y and the theorized contextual conditions

⁴⁴ - Ragin, 2008: 15.

⁴⁵ - Rohlfing, 2014.

⁴⁶ - Michell, 2011.

⁴⁷ - Collier and Mahoney, 1996; Skocpol and Summers, 1980; Ragin, 1987; George and Bennett, 2005: 164; Tarrow, 2010.

that enable the relationship to occur, and that plausibly result in the same mechanism linking X and Y within the population.⁴⁸

This mapping is depicted in figure 2, where cases are divided into four quadrants depending on whether they are members of X and the hypothesized contextual conditions, and/or Y, resulting in four types of cases. In quadrant I are 'typical' cases, understood as the cases where a priori we can expect the theorized X->Y relationship through the theorized mechanism to be present. Quadrant II and IV are two different types of deviant cases; quadrant II where X and/or the required contextual conditions are not present but Y is, and quadrant IV where X and contextual conditions are present but Y is not.⁴⁹ Cases in quadrant II are useful for finding new causes of the outcome other than X, whereas cases in quadrant IV are useful for detecting omitted causal and/or contextual conditions. Cases in quadrant III are analytically irrelevant when analyzing asymmetric causal claims.⁵⁰

>> INSERT FIGURE 2 HERE <<

Given that there can be both theoretical and empirical uncertainty about the exact threshold of membership for cases, we recommend a degree of caution about selecting cases close to the qualitative threshold that demarcates causally relevant differences-of-kind. When demarcating the threshold defining cases in and out of sets, there can be theoretical reasons why we are not exactly sure where the appropriate threshold is.⁵¹ For example, there might be considerable disagreement about whether measures of X need to include a particular attribute or not. Further, there will always be a degree of empirical uncertainty about whether cases close to the threshold are actually in or out. For one case we might have poor quality secondary data that we estimate has a low measurement validity, meaning that there is a considerable degree of empirical ambiguity about whether the case really is in or out of the set of the concept. When there is residual uncertainty either because of theoretical or empirical reasons, we recommend only selecting cases that are widely accepted by case experts as being within the set of the concept; i.e. there are clear substantive and theoretical arguments for set membership that can be documented in a transparent fashion. In figure 2, we might want to avoid selecting case 4 because it might in reality actually be in quadrant II or IV. In this

⁴⁸ - For more on how to map populations, see Collier and Mahoney, 1996; Rohlfing, 2012 and AUTHOR.

⁴⁹ - Note here that we might find an omitted contextual condition by studying a deviant case in quadrant IV. We would then add this to the list of contextual conditions included in the figure, meaning that the deviant case would be reclassified as being in quadrant III because it lacks the newly found contextual condition that was required.

⁵⁰ - For more on this type of case and the debate about them, see Seawright, 2002; Goertz and Mahoney, 2004; Rohlfing, 2012.

⁵¹ - Ragin, 2008; Schneider and Wagemann, 2012: 33-40; AUTHOR.

instance, if case 4 was actually in quadrant IV but we believed it was in quadrant I, our case study would have led us to downgrade our confidence in their being a causal mechanism linking X and Y together, which would have been a flawed inference.

Disregarding cases in quadrant III as analytically irrelevant clashes with most existing case selection guidelines. According to Lieberman, we should select cases for in-depth study based on the size of their residuals in relation to a regression-based analysis of X:Y correlations.⁵² Gerring and Seawright make similar claims, defining typical cases of a given X:Y relationship as regression on-liers.⁵³ Yet these guidelines are only relevant when investigating the magnitude of causal effects, where variation is key to our ability to make inferences.⁵⁴ In terms of figure 2 above, we can imagine a regression line running from case 1 to 12. Following these guidelines, cases 1, 2 and 3 would also be close to the regression line, making them candidate cases. However, what would tracing non-existent mechanisms in case 1 tell us about causal process in cases where the mechanism was present? If we are studying mechanisms linking mutual democracy (X) to avoidance of war (Y), tracing a mechanism in a case where two near democracies ($\sim$ X) went to war ($\sim$ Y) would shed light on mechanisms that lead to war, but not that link X with Y. Studying 'near misses' and negative cases can provide important information about contextual conditions, but they tell us nothing about the actual mechanisms and how they work in positive, typical cases.

4. Mechanism-centered designs – select only typical cases from quadrant I

When our research goal is to make strong within-case inferences about whether X is linked with Y, and/or shedding light on how X is linked with Y, we trace mechanisms only in cases where they can in theory be present. This means in terms of figure 2 that we only select typical cases in quadrant I that are members of both X, Y and the required contextual conditions (cases 4 to 12). We argue that typical cases are the *only* type of case where it makes sense to test whether a hypothesized causal mechanism was present, or build a theory about the mechanism linking X and Y, irrespective of whether one theorizes that X is a sufficient condition or not.

When X (or a conjunction of X's) is theorized to be sufficient, we are restricted to choosing cases within quadrant I based on the argument 'why should we investigate whether a mechanism is present linking X and Y when we know a priori based on values of X and Y

⁵² - Lieberman, 2005: 444.

⁵³ - 2007: 89.

⁵⁴ - E.g. King, Keohane and Verba, 1994: 142-146; Gerring, 2007; Lieberman, 2005; Tarrow, 2010: 249.

that it is *not* present'? If X is sufficient, we would know beforehand that no mechanism linked X with Y in case 13. The situation is slightly more complex when X is not a sufficient cause of Y. Here a mechanism can in principle be present in deviant cases in quadrant IV. X and the mechanism can *in theory* be present, but given that X is not sufficient but is either necessary or merely a contributing cause, Y has not occurred due to the lack of other relevant causal conditions that together produce the outcome with X.

However, despite this complication, we contend that we should still only choose typical cases within quadrant I even when X is not theorized as a sufficient cause. The reasoning is as follows. Selecting a case within quadrant IV to trace a mechanism from X to Y is not an analytically wise strategy given that there are three logical possibilities: 1) the mechanism is not present in the case because there is no causal relationship between X and Y, 2) the mechanism is not present because of omitted contextual conditions, or 3) the mechanism is present but X is not sufficient to produce Y. If we are testing a theorized mechanism between X and Y, situations 1 and 3 would tell us something about the mechanism and the causal relationship. In contrast, situation 2 would tell us nothing about how the mechanism itself works, or whether it is actually present when it should be present in theory. The problem is that we do not know before we have studied a case whether we are in situation 1, 2 or 3 – and given that situation 2 tells us nothing about the mechanism itself as it is not present - a safer strategy when focusing on mechanisms is to only select cases where we expect that the mechanism could *in principle* be present. However, deviant cases within quadrant IV do have important uses when our research is more focused on revising theories (condition-centered). We return to this below.

While not depicted in figure 2, all cases within zone I are theorized to be in the set of relevant contextual conditions. When we are uncertain about what contextual conditions have to be present for a given mechanism, we should start by selecting a case where as many possible contextual conditions are present. If we then find the mechanism in this case, we cannot automatically infer to other cases where fewer of the contextual conditions are present.⁵⁵ We would then want to study another case with fewer contextual conditions, gradually becoming more confident about what conditions have to be present for the mechanism to function through repeated case studies.⁵⁶

Typical cases can be used to either test or build theories of causal mechanisms linking X and Y. When we are testing whether there is mechanistic evidence suggesting that a

⁵⁵ - While this guidance is similar to what Tarrow terms 'progressively testing scope conditions' (2010: 251), here we are talking about within-case analysis that trace mechanisms instead of using paired comparisons as a research strategy.

⁵⁶ - Rohlfsing, 2012: 200-211.

mechanism actually exists between X and Y, we deduce a theory from the existing literature and then test whether there is evidence that a hypothesized causal mechanism is actually present in a given typical case. An example is Owen's work on the democratic peace,⁵⁷ where the analytical ambition is to test whether the observable implications of the existence of a more general, parsimonious theorized causal mechanism linking democracy and peace are present in a set of parallel single case studies, comparing the predictions of what evidence he should find with the evidence he actually finds in the empirical record.

When we are in the dark regarding a plausible mechanism between X and Y we use a more theory-building variant of mechanism-centered design, where the ambition is to build a theoretical mechanism from the empirical evidence of a particular case, inferring that a more general causal mechanism exists that should also be found in the population of the phenomenon. In its purest form, theory-building starts with empirical material found in a typical case and uses a structured analysis of it to detect the observable manifestations of a plausible causal mechanism whereby X is linked with Y. An example is Janis' work,⁵⁸ where he utilizes a typical case to start building a theorized mechanism that links small group dynamics with premature consensus through a groupthink mechanism. The first case chosen to build the theorized mechanism is the Bay of Pigs fiasco in 1961, where the 'best and the brightest' policy-making group in the early Kennedy administration decided to support an intervention that was doomed to fail from the start due to the faulty assumptions underlying the decision.⁵⁹

4.1. Do we need to control for other causes when selecting typical cases?

Variance-based scholars question whether we can make within-case inferences about mechanisms when we do not control for other causes by selecting cases in which only one potential cause is present. 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.'⁶⁰ They term this type of case a 'pathway case'. 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

⁵⁷ - Owen, 1997.

⁵⁸ - Janis, 1983.

⁵⁹ - Janis, 1983.

⁶⁰ - 2007: 122.

mechanism through which it contributes to the outcome in the case under study.’.⁶¹ Goertz writes that one should avoid cases that exemplify multiple causal mechanisms.⁶²

Our argument is that while the logic of controlling for other potential causal *conditions* is relevant for variance-based designs aimed at assessing probabilistic claims about the mean causal effects of individual causes at the population level (i.e. evidence of difference-making), it is not relevant for the study of mechanisms in case-based research because process-tracing offers us analytical tools that can enable us to control for other causes working at the empirical level. When evaluating empirical evidence of the workings of parts of mechanisms using process-tracing, we assess whether the particular piece of evidence is theoretically unique, or whether finding it is just as plausible with other explanations (be they competing theories, or case-specific factors).⁶³ If finding the particular piece of evidence is just as plausible with alternative explanations, then no confirming causal inferences based on it are possible. Control for other causes occurs therefore at the level of evidence *within* a case, which matches with the level of our causal inferences being made (within-case).

The difference between case-based and variance-based principles of case selection are most pronounced when dealing with potential overdetermination, defined as multiple sufficient causes being present in a case. If there is overdetermination in a given case, variance-based designs are unable to disentangle which cause actually produced the outcome because there are more than one possible cause of difference, making control for other causes through case selection even more important. However, overdetermination is not a serious problem when studying mechanisms-as-systems because we can isolate the workings of individual mechanisms *empirically* from each other through our evaluation of the theoretical uniqueness of *each piece* of mechanistic evidence, asking whether finding the evidence can be accounted for with any other plausible explanation. Basically, we should expect to find when we operate at the level of within-case mechanistic evidence that different causes will be linked to outcomes through different mechanisms that would leave empirical fingerprints that can be distinguished from each other. Therefore, choosing a case where both X1 and another potentially sufficient cause X2 are both present should not matter because we can distinguish empirically the workings of the two causes and their mechanisms from each other.

For example, using the classic example from the philosophy of science, a person found dead after wandering in the desert. It is found that there was poison in the person’s canteen, but there was also a hole in the canteen. In this instance, both poison (X1) and thirst (X2) are

⁶¹ - 2013: 8.

⁶² - Goertz, 2012: 18.

⁶³ - Note that this builds on van Evera’s original distinction between certainty and uniqueness (1997), which have later been linked with Bayesian logic and updating (Bennett, 2008a; 2014; AUTHOR).

potentially sufficient causes, but only one can actually have had a causal relationship with the outcome, i.e. he either died of poisoning or dehydration. However, overdetermination at the theoretical level does not mean that we cannot figure out why the man in the desert died in practice. This is because we can distinguish at the *empirical evidence level* between the fingerprints that causes X1 and X2 would leave. For example, we would expect to find poor skin turgor, tinting of skin, sunken eyes, and/or dry galea and dry organ surfaces if the man died of dehydration (Madea and Lachenmeier, 2005), whereas poisoning by cyanide could be detected as increased levels of cyanide in the blood if the autopsy is conducted relatively quickly after death (Muschhoff et al, 2002).

A natural counterargument would be to claim that parts of mechanisms in the social sciences often do not leave theoretically unique empirical evidence. However, we claim that we should not throw in the towel, but instead we should expend more effort on creatively developing as many potential observables as possible in order to uncover some that have a degree of uniqueness. Additionally, while one piece of evidence might not be very unique, when independent pieces are combined, they can together become quite unique (Bennett, 2014). Together they can form a 'signature' that is quite unique to one particular part of a mechanism at the empirical-level.

Systematically selecting pathway cases where *only* X1 is present also has the downside that it reduces even the population of potential cases that can be inferred to afterwards. Given the sensitivity of mechanisms to context, there is a risk that X1 is linked to Y through M1 in cases where only X1 is present, whereas in cases where X1 and X2 are present, X1 is linked to Y through M2. This situation is depicted in figure 3. If we follow Gerring and Seawright's advice, after studying a pathway case (X1, Y, ~X2) we would only be able to infer to other cases where only X1 and Y are present instead of the full population. We suggest if there are significant concerns about the potential impact of X2 on which mechanism is working that we engage in a two-step process where we first select two cases for tracing mechanisms where X1 and Y are present but not X2, followed by two cases where X1 and X2 are present in order to assess whether the presence of other causes impacts on which mechanism works.

>> INSERT FIGURE 3 HERE <<

4.2. Should we select typical cases that are most- or least-likely?

Within the set of typical cases, one could argue that depending on prior confidence in a mechanism being present, one should select most- or least-likely cases. Unfortunately, the

logic of most/least-likely cases is not compatible with deterministic claims in case-based research. Selecting most/least-likely cases is a long-standing qualitative practice that goes back at least to Eckstein in the mid-1970s.⁶⁴ The basic logic is Bayesian-inspired, where the argument is that finding a relationship where one least expects it enables one to infer across cases that it should therefore be everywhere (least-likely), and vice versa. However, while intuitively pleasing, the logic both conflates theoretical and empirical likelihood, and also within-case and cross-case likelihoods in ways that are not compatible with the underlying ontological assumptions about the nature of causality in case-based research. These problems lead us to suggest dropping the distinction between ‘most’ and ‘least’ likely when selecting cases in case-based research because it is a vestige of variance-based designs. Instead, we should use a logic that is in alignment with the assumptions about deterministic and asymmetric causation underlying case-based research, focusing only on ‘typical’ cases without any causally irrelevant distinctions.

Before we proceed, it is important to note that there are a variety of different ways in which most- and least-likely cases have been defined in the literature (see table 1). Common to most definitions is the idea that a most-likely case is one where other causal conditions except the X in focus suggests that Y should occur but it does not, implying that we can disconfirm X being a cause across the population. A least-likely case is where other causal conditions except X point in the direction of Y not occurring but it does, enabling us to infer that given that it occurred where we least expected it, it should also occur in more probable places. It is vital to note that the likelihood of a causal relationship occurring in a case is based on *theoretical* reasons, i.e. contextual conditions that are more/less conducive determine the likelihood of the causal relationship occurring. As can also be seen in the table, there is also a variety of ways in which cases are scored, either by scores on causal conditions and outcomes [variables] (or changes therein, e.g. Eckstein), or based on the assumptions of theories. We have also included Schneider and Rohlfing’s definition of ‘most-typical’ case given that is similar to a most-likely case but is understood in more set-theoretic terms.

>> INSERT TABLE 1 HERE <<

The most serious problem with most/least-likely logic is that it conflates theoretical (ontological) and empirical (epistemological) likelihood. The use of the term likelihood in relation to causation requires a probabilistic *ontological* understanding of causation, where ‘a

⁶⁴ - Eckstein, 1975.

cause raises the probability of an event occurring'.⁶⁵ Causal relationships are more likely in more conducive circumstances (more contextual conditions present), and less likely in less conducive (contextual conditions mostly absent). However, this type of ontological probabilism is not compatible with case-based designs. In case-based research we want to produce explanations that can account for what happened in particular cases instead of just producing theoretical odds for whether something might occur. In a deterministic understanding of causation, if a sufficient cause is present along with the requisite contextual conditions, the outcome *will* occur. Any case that fulfils these conditions should demonstrate the relationship, meaning that logically the relationship is not more or less likely in any particular case.

Understood in deterministic and asymmetric causal terms, the relevant distinction is therefore not most/least-likely but simply possible and not possible. If an outcome occurs when we did not theoretically expect it (a least-likely case), this should result in revision of the theory about what contextual factors matter instead of us making a strong cross-case inference that the relationship is present across the population!

Second, and related, case studies produce within-case mechanistic evidence, updating our confidence in a causal relationship being present within a single case. Inferring *beyond* the studied case to the rest of the population requires that the chosen case is causally similar to the rest of the typical cases. However, given the large differences in the contextual conditions present in least-likely, most-likely cases, and more 'normal' typical cases, and given the sensitivity of causal processes to contextual conditions that are typically assumed in case-based research, we should expect the different types of cases to exhibit high degrees of causal heterogeneity. This means that we cannot just infer that because we found a relationship in a least-likely case that it should also be present in other, *dissimilar* cases throughout the population.

Schneider and Rohfling (2013) have developed a definition of most-likely cases that is more compatible with asymmetric causation. They suggest that causal relationships are easier to study when both X and Y are strongly present (high set membership scores). Yet this assumes a 'more-is-better' logic in the theory we are studying in relation to whether it is empirically observable, and assumes that degree-differences make the causal effect of causes stronger. First, on an empirical level, using an analogy, if the outcome is fire, the outcome would have higher values the more fuel is present, and given that more fire is more observable, other things equal, it would also be easier to study empirically up to the point

⁶⁵ - Gerring, 2005: 167.

where the fire gets so big that you would be scorched. And at the theoretical level, if X is theorized to be necessary, it is necessary in a case when the cause and requisite contextual conditions are present in the case; a cause cannot be 'more' or 'less' necessary.

Concluding, most/least-likely logic conflates ontological and epistemological probabilism, and within-case and cross-case evidence in ways that are impossible to untangle. We therefore recommend the use of the term 'typical' case without any qualifications about likelihood, referring to cases where the theorized causal relationship is possible because both X, Y and the requisite contextual conditions for it to occur are present.

At the same time, we do recommend thinking about whether one has access to empirical material or not when selecting typical cases, and the chosen case is one where the mechanism might actually have left empirical fingerprints that can be observed. In some cases, the mechanism might leave very distinct and unique traces that enable strong within-case inferences to be made, whereas in others the data-generating processes might leave few traces. For example, bureaucratic politic mechanisms might leave few traces in a highly oral context, whereas in a context where extensive written records are kept, it might be easier to study. Second, accessibility concerns are also important. If we do not speak Spanish, it might be difficult to interview civil servants in Costa Rica, meaning we probably would not select it as a case unless we have no alternative. Similarly, studying elite decision-making is easier when archives are declassified, other things equal. Given that tracing mechanisms is very data-demanding, selecting typical cases that are expected to be data-rich is a wise strategy.

4.3. Inferring from the studied typical case to other typical cases

What types of inferences are possible if we find confirming mechanistic evidence of a causal mechanism operating in one typical case? Existing case selection guidelines suggest that one case can be enough to make a cross-case inference to other cases. Lieberman writes that in testing a theory using a case study, '...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...)'.⁶⁶ In the flow-chart illustrating his case selection strategy, if one finds support in the single case study, he suggests we can end our analysis. Similarly, Schneider and Rohlfing in their flowchart state that if a theory is

⁶⁶ - 2005:448 – italics added.

confirmed using a single case study we can end our analysis, enabling us to infer that the found mechanism is also present in other cases that are members of X and Y.⁶⁷

In case-based tracing of mechanisms, given the sensitivity of mechanisms to context, we would be averse to making any cross-case claims about similar mechanisms being present in similar cases to the studied unless we can demonstrate transparently that the population of other typical cases is causally homogenous, meaning the same cause should be linked to the same mechanism and outcome because similar sets of contextual and causal conditions are present in the other cases. Unless we can provide strong evidence of causal homogeneity, which is very difficult given that we often have imperfect knowledge of contextual conditions, the best solution is to trace mechanisms in multiple typical cases. Finding mechanistic evidence of a causal mechanism between X and Y in case 5 in figure 2 (above) does not necessarily mean that the same causal mechanism links X and Y in case 7. If we find that economic development contributed to produce democracy in the South Korean case through a middle-class mechanism, this does not necessarily mean that economic development contributed to producing democracy in Taiwan through the same causal mechanism.

This implies that we should never be content with a one-shot case study of mechanisms, but in contrast to existing case selection guidelines, we should always select two or more typical cases within quadrant I. If we can muster strong evidence that the population of typical cases is (relatively) causally homogeneous, we can make cross-case inferences more confidently based on fewer cases, other things equal. If on the other hand we are not able to provide strong evidence of homogeneity – usually because the population is larger (i.e. more than a handful of cases) – we then need to trace mechanisms in more typical cases to be confident about the same mechanism linking X and Y within quadrant I. Note that we will never be 100% certain when we try to transfer our findings from the single case to similar cases in the population; our confidence in the external validity of our within-case findings is a function of the strength of the evidence we can transparently provide in the degree of causal homogeneity of the bounded population, in particular whether there are similar contextual conditions. Unfortunately there is always a trade-off between internal and external validity. Selecting more cases enables us to be more confident about the external validity of our

⁶⁷ - 2013: 561. While they later in the article do suggest one can engage in multiple case studies of typical cases to bolster confidence that the mechanisms operating are the same across cases, they do not require this (see their flowchart). Further, they suggest that when selecting multiple cases one selects one that has maximum scores on X and Y, and another that has minimum within the sets of X and Y. As we will explain later, selecting this type of ‘least-likely’ case with low scores on X and Y involves a high risk of selecting cases that are not actually members of one or more causal conditions.

findings because we have studied a higher ratio of cases in the population but it would result in the mechanistic evidence to be used being more superficial (lower internal validity). However, tracing mechanisms using in-depth case studies requires considerable analytical effort to make strong within-case inferences (i.e. high internal validity), and therefore in most research situations we can only employ two or three process-tracing case studies. Yet by selecting so few cases, our ability to make stronger cross-case inferences (external validity) is impaired. Therefore we are forced to choose between stronger internal and external validity; but when push-comes-to-shove case-based researchers usually opt for the former.

If we are not able to detect a causal mechanism between X and Y in a typical 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 typical cases, assessing the values of the case and other typical cases both on X, Y, contextual conditions, and a range of other potential causal conditions. In particular, this comparison can shed light on the contextual conditions for the functioning of the mechanism, enabling us to assess whether the bounds of the population have been set properly. This procedure is the same as used when comparing deviant cases in quadrant IV with typical cases in order to detect omitted conditions (we discuss this more below in section 5). However, if we cannot detect any significant differences between the chosen case and other typical cases in the population, we can make the cautious inference that there is not a causal relationship between X and Y. We would conclude that while there is a correlation between X and Y, there is no underlying causal mechanism linking the two, meaning that the postulated causal relationship is spurious.

5. Condition-centered designs

Tracing mechanisms in deviant cases can also help us revise theories about causal conditions. It is however important to note that tracing mechanisms does not stand-alone in this type of design; indeed the analytical heavy-lifting in condition-centered designs is done by systematically comparing a deviant with a typical case (see below under deviant cases in quadrant IV) in what can be thought of as a ‘paired comparison’.⁶⁸

In condition-centered designs we only select deviant cases. Deviant cases come in two forms: 1) cases within quadrant II, where Y is present but existing causes are not present

⁶⁸ - Tarrow, 2010. However, in contrast to Tarrow, we detail in the following the role that process-tracing of mechanisms plays in conjunction with a systematic paired comparison, whereas Tarrow contends that a paired comparison enables confirming inferences by itself based on it being an approximation of a randomized experiment (p. 244).

(deviant case new cause); and 2) cases within quadrant IV, where X is present but Y is not present (deviant case omitted conditions). Deviant cases within quadrant II are used to find new causal conditions that have been overlooked in the literature, whereas deviant cases within quadrant IV help us detect omitted causal and/or contextual conditions in an existing X->Y causal relationship.

5.1. Deviant cases in quadrant II – Finding new causes (or mission impossible?)

Cases within quadrant II are deviant cases where existing causal theories (X) are unable to account for the outcome (Y is present). This type of deviant case can be used to find new causes of the outcome, although we are much more skeptical about whether tracing mechanisms is the most efficient methodological tool for these purposes, in contrast to the arguments found in the existing literature.⁶⁹

The basic idea is that one traces backwards from the occurrence of the outcome to find a new cause. But it is far from straightforward what one is actually tracing under these circumstances. Existing suggestions that we engage in backwards-tracing using case studies builds on an understanding of mechanisms that sees them as mere events. Yet tracing ‘events’ is not the same thing as tracing causal mechanisms that can be understood as theoretical systems that can be found in multiple cases. And if we have no idea about the cause (X), we also have no clue about the mechanism(s) linking the mystery cause with the outcome, meaning that we are in effect blindly groping in the dark after a cause. We therefore claim that a comparative design or a congruence case study would be a much more efficient analytical first-step in detecting new causes.⁷⁰ For example, we might compare systematically two cases that are similar in all aspects except the occurrence of the outcome (Y, ~Y). We would then want to know what new, undiscovered X is different between the two cases. After having found this using a paired comparison, it would be possible to engage in a form of in-depth case study that focuses on this new candidate cause, attempting to discern whether the new condition is actually causally related to Y.

5.2. Deviant cases in quadrant IV – why does the causal mechanism between X and Y not work in the case?

Cases within zone IV have two different purposes depending on whether one is theorizing that X is a sufficient cause for Y, or that X is merely a contributing cause. The following offers the

⁶⁹ - E.g. Lieberman, 2005: 443; Rohlfing, 2008: 1510; Schneider and Rohlfing, 2013.

⁷⁰ - See AUTHOR.

first set of comprehensive guidelines that illustrate how we combine comparative and within-case methods to uncover omitted causal or contextual conditions.⁷¹

If we theorize that X is a *sufficient* cause of Y, deviant cases where X is present but where Y is not present are useful to investigate the *contextual conditions* that have to be present to trigger the mechanism that will produce Y. If X is *not* theorized as a sufficient cause, deviant cases within quadrant IV can be used to detect omitted *causal conditions* that together with X would be sufficient to produce Y.

In both instances we only employ this type of design *after* we have positive results when tracing a mechanism in one or several typical cases within quadrant I. The argument here is that there is no reason to investigate mechanism breakdown before we are more confident about the actual existence of a mechanism linking X and Y in one or more typical cases. However, once we are confident about what is going on in typical cases, investigating this type of deviant case is very important for developing better causal theories. Using an analogy, once we are certain about the mechanisms that enable airplanes to fly, we would want to investigate accidents to develop a better understanding of the contextual conditions under which planes can fly safely.

The within-case component of a condition-centered design involves tracing a mechanism using in-depth process-tracing in a deviant case until the mechanism breaks down in the case. We want to uncover when and why the mechanism failed.⁷² Existing theories of causal mechanisms that were either built or tested on typical cases within quadrant I provide the foundation for the tracing in deviant cases until mechanism breakdown. Finding out when and why a mechanism breaks down gives us clues about omitted contextual or causal conditions, although to repeat, the tracing component here is auxiliary, and the main analytical method is a systematic paired comparison of the deviant case with a typical case.

A condition-centered design therefore relies on an analytical two-step, where we use a deviant case first to trace where and why the theorized mechanism breaks down. We then use these insights to inform a pairwise comparison of the deviant and typical case to uncover omitted causal and/or contextual conditions. Using as a hypothetical example research focused on studying links between economic development (X) and democratization (Y) through some form of ‘middle-class’ mechanism that details a causal mechanism whereby the growth of the middle-class spurs democratization, we might find in a deviant case like Poland

⁷¹ - While Schneider and Rohlfing (2013) do discuss designs that compare a typical case and a deviant case in zone IV, they do not offer guidance on how the comparative and within-case methods are combined, nor do they offer any guidance when we do not theorize that X is necessary or sufficient, but is just causally related to Y. Tarrow (2010) offers no guidance on how to combine paired comparisons and within-case tracing of mechanisms.

⁷² - Anderson, 2012: 421-422.

that while a middle-class was produced by economic growth (X), the next part of the mechanism was not present, with demands for democratization instead coming from elite actors. We would then compare the single deviant case (Poland) with what we know more broadly about other typical cases in the population where we know the middle-class mechanism functioned as expected (e.g. South Korea) to uncover whether there is a missing causal or contextual condition that can explain why the mechanism did not work as predicted in the deviant case. Here we might find that when parallel state-building processes were underway as in Poland, the middle-class mechanism did not function properly, suggesting that lack of parallel state-building processes is an important contextual condition in which economic development (X) is linked to democratization (Y) through a middle-class causal mechanism. Poland would also then be reclassified as a case outside of the relevant contextual conditions, moving into quadrant III. After we have found an omitted contextual or causal condition, we would have to reclassify our cases, with previously deviant cases becoming irrelevant (they are moved to zone III), as they lack a causal condition that has to be present for Y to occur (sufficiency), or lack a contextual condition that enables the mechanism to function properly.

6. Conclusions

This article has argued that existing selection guidelines build on methodological assumptions that are not compatible with the study of mechanisms-as-systems using case-based designs. Existing guidelines therefore result in the selection of inappropriate cases, or too few cases, resulting in flawed inferences being made.

Section 2 explored the assumptions underlying in-depth tracing of mechanisms, including viewing mechanisms-as-systems, determinism, the importance of causal homogeneity, and causal asymmetry. Together, these four ontological assumptions imply very different case selection strategies than those that are appropriate in variance-based designs where control for causes at the cross-case level is the primary purpose of case selection.

Section 3 introduced the argument that we cannot select appropriate cases unless we have mapped the population of potential cases. Four types of cases were developed: 1) typical cases where both a cause (or set of causes), the requisite contextual conditions for the cause(s) to be linked to an outcome through a particular causal mechanism, 2) deviant cases where the cause is absent but the outcome is present, 3) deviant cases where the outcome is absent but the cause is present, and 4) analytically irrelevant cases. Ensuring causal homogeneity in the population of typical cases is especially important when studying mechanisms, as for instance

the same causes can be linked to the same outcome through different mechanisms depending on context.

Section 4 developed case selection guidelines for research focused on the mechanisms linking causes with outcomes. Here only typical cases are selected. Given that there often is uncertainty about what contextual conditions are required, we suggested starting with a case where as many potential contextual conditions are present as possible in order to detect whether there is a mechanism linking X and Y. If we find the mechanism in this case, we would then want to trace the mechanism again in a typical case that has fewer potential contextual conditions present. If we have found the same mechanism across two or three different typical cases, we can then infer cautiously to the rest of the population of cases in quadrant I about X and Y being causally related through the found mechanism. Note we do not want to select cases close to qualitative thresholds because there often is ambiguity about the exact placement of the threshold in practical research situations. It is also important to note that selecting typical cases for tracing mechanisms is not just cherry-picking confirming cases. Instead, it involves testing/building theories of mechanisms in cases where it is possible they can exist. One does not go moose hunting in Manhattan. If one wants to have any chance of shooting a moose, one should go hunting where they can in principle be present, such as the backwoods of Alaska or Maine.

Condition-centered designs were discussed in section 5. Here tracing mechanisms plays an adjunct role. We suggested that, in contrast to existing guidelines, there is little reason to select deviant cases in quadrant II (detecting a new cause) given the extreme difficulty in tracing a mechanism when we have no idea what we are looking for. Instead, we contend that one gets much more analytical bang-for-the-buck if one utilizes either comparative tests or weaker, congruence case studies to search for new causes.

Deviant cases in quadrant IV are intensely interesting because they shed light on omitted contextual or scope conditions that have to be present for a mechanism to work. Tracing mechanisms until they breakdown sheds light on what went wrong, giving us clues for example about potential omitted contextual conditions that have to be present for the mechanism to work properly. The results of the tracing until breakdown case study sheds light on which factors we would want to focus our attention in a pairwise, systematic comparison of the deviant case with a typical case to detect either omitted causal or contextual conditions.

References

Adcock, Robert. 2007. "Who's Afraid of Determinism? The Ambivalence of Macro-Historical Inquiry." *Journal of the Philosophy of History* 1(2007): 346-364.

Allison, Graham and Philip Zelikow 1999. *Essence of Decision Explaining the Cuban Missile Crisis*. New York: Longman.

Anderson, Holly. 2012. "The case for regularity in mechanistic causal explanation." *Synthese*, 2012(189): 415-432.

AUTHOR

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.

Bennett, Andrew. 2014. "Appendix.", in Bennett, Andrew and Jeffrey Checkel (eds.) *Process Tracing: From Metaphor to Analytic Tool*. Cambridge: Cambridge University Press.

Bennett, Andrew and Jeffrey Checkel. 2014. *Process Tracing: From Metaphor to Analytic Tool*. Cambridge: Cambridge University Press.

Brady, Henry E. 2008. "Causation and Explanation in Social Science." In *The Oxford Handbook of Political Methodology*, ed. Janet M. Box-Steffensmeier Henry E. Brady and David Collier, 217-270. Oxford: Oxford University Press.

Brady, Henry E. and David Collier. (eds). 2011. *Rethinking Social Inquiry: Diverse Tools Shared Standards*. Second Edition. Lanham MD: Rowman Littlefield.

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.

Collier, David and James Mahoney. 1996. "Research Note: Insights and Pitfalls: Selection Bias in Qualitative Research." *World Politics* 49(1): 56-91.

Dion, Douglas. 2003. "Evidence and inference in the comparative case study." in Gary Goertz and H. Starr (eds). *Necessary conditions: Theory, methodology, and applications*. Oxford: Rowman and Littlefield Publishing, pp. 95 – 112.

Eckstein, Harry. 1975. "Case Study and Theory in Political Science." In Greenstein and Polsby (eds) *Handbook of Political Science, Vol. 7*. Reading MA: Addison-Wesley, pp. 79-138.

Falleti, Tulia G. and Julia F. Lynch (2009) "Context and Causal Mechanisms in Political Analysis." *Comparative Political Studies*, 42:1143-1166.

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

Gerring, John. 2005. "Causation: A Unified Framework for the Social Sciences." *Journal of Theoretical Politics* 17(2): 163-198.

Gerring, John. 2007. *Case Study Research*. Cambridge: Cambridge University Press.

Gerring, John. 2010. "Causal Mechanisms: Yes But..." *Comparative Political Studies* 43(11): 1499-1526.

Gerring, John. 2011. *Social Science Methodology – a unified framework*. Cambridge: Cambridge University Press.

Gerring, John and Jason Seawright. 2007. "Techniques for Choosing Cases." In *Case Study Research*. Cambridge: Cambridge University Press, pp. 86-150.

Glennan, Stuart S. 1996. "Mechanisms and the Nature of Causation." *Erkenntnis* 44(1): 49-71.

Goertz, Gary. 2006. *Social Science Concepts: A User's Guide*. Princeton: Princeton University Press.

Goertz, Gary. 2012. "Case studies, causal mechanisms, and selecting cases." Unpublished manuscript. Version 5, August 18, 2012.

Goertz, Gary and James Mahoney. 2004. "The Possibility Principle: Choosing Negative Cases in Comparative Research." *American Political Science Review* 98(4): 653-669.

Goertz, Gary and James Mahoney. 2012. *A Tale of Two Cultures*. Princeton: Princeton University Press.

Hall, Peter A. 2003. "Aligning Ontology and Methodology in Comparative Politics." In *Comparative Historical Analysis in the Social Sciences*, edited by J. Mahoney and D. Rueschemeyer. New York: Cambridge UP, 373-404.

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.

Howson, Colin and Peter Urbach. 2006. *Scientific Reasoning: the Bayesian Approach*. Third Edition. La Salle IL: Open Court.

Janis, Irving L. 1983. *Groupthink: Psychological Studies of Policy Decisions and Fiascos*. Boston: Houghton Mifflin Company.

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

Levy, Jack S. 2008. "Case Studies: Types, Designs, and Logics of Inference." *Conflict Management and Peace Science* 25: 1-18.

Lieberman, Evan S. 2005. "Nested Analysis as a Mixed-Method Strategy for Comparative Research." *American Political Science Review* 99(3): 435-451.

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.

Madea, Burkhard and Dirk W. Lachenmeier. 2005. 'Postmortem diagnosis of hypertonic dehydration.', *Forensic Science International* 155 (2005): 1-6.

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

Mahoney, James. 2007. "Qualitative Methodology and Comparative Politics." *Comparative Political Studies* 40(2): 122-144.

Mahoney, James. 2008. "Toward a Unified Theory of Causality." *Comparative Political Studies* 41(4/5): 412-436.

Mayntz, Renate. 2004. "Mechanisms in the Analysis of Social Macro-Phenomena." *Philosophy of the Social Sciences* 34(2): 237-259.

Michell, Joel. 2011. "Qualitative research meets the ghost of Pythagoras." *Theory & Psychology* 21(2): 241-259.

Musshoff, Frank Ph.D.; Schmidt, Peter M.D.; Daldrup, Thomas Ph.D.; Madea, Burkhard M.D.. 2002. 'Cyanide Fatalities: Case Studies of Four Suicides and One Homicide.', *American Journal of Forensic Medicine & Pathology* 23(4):315-320.

Nielsen, Richard A. 2015. "Case Selection via Matching." *Sociological Methods & Research* first published on October 27, 2014.

Owen, John M. 1997. *Liberal Peace Liberal War: American Politics and International Security*. Ithica: Cornell University Press.

Ragin, Charles C. 1987. *The Comparative Method: Moving beyond Qualitative and Quantitative Strategies*. Berkeley: University of California Press.

Ragin, Charles C. 2000. *Fuzzy-Set Social Science*. Chicago: University of Chicago Press.

Ragin, Charles C. 2008. *Redesigning Social Inquiry. Fuzzy Sets and Beyond*. Chicago: University of Chicago Press.

Ragin, Charles and Garrett Andrew Schneider (2011) "Case-Oriented Theory Building and Theory Testing.", in Malcolm Williams and W. Paul Vogt (eds) *Sage Handbook of Innovation in Social Research Methods*. London: Sage, pp. 150-166.

Rohlfing, Ingo. 2008. "What You See and What You Get: Pitfalls and Principles of Nested Analysis in Comparative Research." *Comparative Political Studies* 41(11): 1492-1514.

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 Sciences." *International Studies in the Philosophy of Science*, 21(2): 157–170.

Schneider, Carsten and Ingo Rohlfing. 2013. "Combining QCA and Process Tracing in Set-Theoretical Multi-Method Research." *Sociological Methods and Research*, 42(4): 559-597.

Schneider, Carsten and Claudius Wagemann. 2012. *Set-theoretic Methods for the Social Sciences*. Cambridge: Cambridge University Press.

Seawright, Jason. 2002. "Testing for Necessary and/or Sufficient Causation: Which Cases Are Relevant?." *Political Analysis* 10(2): 178-193.

Skocpol, Theda and Margaret Somers. 1980. "The Uses of Comparative History in Macrosocial Inquiry." *Comparative Studies in Society and History* 22(2): 174-197.

Tarrow, Sidney. 2010. "The Strategy of Paired Comparisons: Toward a Theory of Practice." *Comparative Political Studies* 43(2): 230-259.

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 (183): 389-408.

Weller, Nicholas and Jeb Barnes. 2014. *Finding Pathways: Mixed-Method Research for Studying Causal Mechanisms*. Cambridge: Cambridge University Press.

Woodward, James. 2005. *Making Things Happen: A theory of Causal Explanation*. Oxford: Oxford University Press.

Ziblatt, Daniel. 2009. "Shaping Democratic Practice and the Causes of Electoral Fraud: The Case of Nineteenth-Century Germany." *American Political Science Review* 103(1): 1-21.