UC Riverside

UC Riverside Previously Published Works

Title

Pathway Analysis and the Search for Causal Mechanisms

Permalink

<https://escholarship.org/uc/item/1423q1z4>

Journal

Sociological Methods & Research, 45(3)

ISSN

0049-1241

Authors

Weller, Nicholas

Barnes, Jeb

Publication Date

2016-08-01

DOI

10.1177/0049124114544420

Peer reviewed

Sociological Methods & Research

<http://smr.sagepub.com/>

Pathway Analysis and the Search for Causal Mechanisms

Nicholas Weller and Jeb Barnes

Sociological Methods & Research published online 12 August 2014

DOI: 10.1177/0049124114544420

The online version of this article can be found at:

<http://smr.sagepub.com/content/early/2014/08/12/0049124114544420>

Published by:

<http://www.sagepublications.com>

Additional services and information for *Sociological Methods & Research* can be found at:

Email Alerts: <http://smr.sagepub.com/cgi/alerts>

Subscriptions: <http://smr.sagepub.com/subscriptions>

Reprints: <http://www.sagepub.com/journalsReprints.nav>

Permissions: <http://www.sagepub.com/journalsPermissions.nav>

>> [OnlineFirst Version of Record](#) - Aug 12, 2014

[What is This?](#)

Pathway Analysis and the Search for Causal Mechanisms

Nicholas Weller¹ and Jeb Barnes¹

Abstract

The study of causal mechanisms interests scholars across the social sciences. Case studies can be a valuable tool in developing knowledge and hypotheses about how causal mechanisms function. The usefulness of case studies in the search for causal mechanisms depends on effective case selection, and there are few existing guidelines for selecting cases to study causal mechanisms. We outline a general approach for selecting cases for pathway analysis: a mode of qualitative research that is part of a mixed-method research agenda, which seeks to (1) understand the mechanisms or links underlying an association between some explanatory variable, XI , and an outcome, Y , in particular cases and (2) generate insights from these cases about mechanisms in the unstudied population of cases featuring the XI/Y relationship. The gist of our approach is that researchers should choose cases for comparison in light of two criteria. The first criterion is the expected relationship between XI/Y , which is the degree to which cases are expected to feature the relationship of interest between XI and Y . The second criterion is variation in case characteristics or the extent to which the cases are likely to feature differences in characteristics that can facilitate hypothesis generation. We demonstrate how to apply our approach and compare it to a leading example of pathway analysis in the so-called resource curse literature, a prominent example of a correlation featuring a nonlinear relationship and multiple causal mechanisms.

¹ University of Southern California, Los Angeles, CA, USA

Corresponding Author:

Nicholas Weller, Department of Political Science, University of Southern California, 327 Von KleinSmid Center, Los Angeles, CA 90089, USA.

Email: nweller@usc.edu

Keywords

causal mechanisms, mixed-method research, case selection, hypotheses generation, qualitative methods

Introduction

Political scientists, sociologists, and economists, who agree on little else, have embraced the search for causal mechanisms, the links or pathways between explanatory variables and outcomes (e.g., Elster, 1998; Gerring 2010; Heckman and Smith 1995; Hedstrom and Ylikoski 2010; Imai et al. 2011; Kiser and Hechter 1991; Malhotra and Krosnick 2007; Mayntz 2004; Waldner 2007). Indeed, it is often expected that researchers include some account of how one variable generates another in making causal claims. In the words of David Waldner (2007:146), “Explanatory propositions are distinguished from nonexplanatory propositions by the inclusion of causal mechanisms” (see also Kiser and Hechter 1991:5; Mayntz 2004:14).

While there are many possible ways to search for mechanisms, an increasingly common approach is to employ mixed-method research. This turn to mixed methods reflects a conviction that quantitative and qualitative studies have complementary strengths, which can be leveraged in exploring the underlying relationship between some explanatory variable, X_1 , and an outcome, Y , controlling for other factors (X_2). On one hand, quantitative approaches are well suited for identifying patterns of association between X_1 and Y within large data sets and estimating the relationship between an explanatory variable and an outcome, controlling for X_2 . However, as every graduate student knows, “correlation is not causation” and so researchers often need to supplement what is learned from standard regression techniques applied to observational data with alternative approaches to better understand *how* X_1 generates Y (see, e.g., Achen 1986; Chatfield 1995; Freedman 1991; Gerber et al. 2004; Kittel and Winner 2005; Winship and Sobel 2004).

On the other hand, scholars have long recognized the usefulness of case studies for generating hypotheses and developing theory (e.g., Eckstein 1975; George and Bennett 2005; Lijphart 1971; see also Munck 2004),¹ and case studies seem particularly apt for peering inside the “black box” of correlation and exploring causal pathways between X_1 and Y . Their richness of detail enables researchers to plot the sequence of variables, detect interactions among them, and assess the direction of causality (Collier, 2011;

George and Bennett 2005). Recent advances in qualitative methods have reinforced these strengths by helping scholars devise strategies for generating “causal process observations” in the course of conducting case study research, including research aimed at plotting the causal pathways between variables (Collier, Brady, and Seawright 2004; see also Bennett 2010; Freedman 2008; but see Beck 2006, 2010).

Given these strengths, case studies have the potential to pick up where quantitative studies leave off. The usefulness of case studies in mixed-method research in the search for causal mechanisms, however, depends on effective case selection. The current case selection guidelines for mixed-method research on causal mechanisms focus on simple linear relationship between $X1/Y$ and between the mechanisms and therefore may mislead scholars when the underlying relationship between a causal variable, $X1$, and an outcome, Y , is nonlinear and/or there are multiple causal pathways between them (equifinality).² Under these circumstances, the existing guidelines may result in poor case selection that can produce false negatives if scholars fail to observe a relevant mechanism because the relationship with the key explanatory variable is small. Conversely, existing guidelines may produce false positives if researchers pick a case that involves a large but atypical effect of $X1$ on Y or anomalous mechanisms. Either scenario could lead to inaccurate conclusions about the unobserved cases and erroneous theoretical claims.

In this article, we outline a general approach for selecting cases for pathway analysis: a mode of qualitative research that is part of a mixed-method research agenda, which seeks to (1) understand the mechanisms or links underlying an association between some explanatory variable, $X1$, and an outcome, Y , in particular cases and (2) generate insights from these cases about mechanisms in the unstudied population of cases featuring the $X1/Y$ relationship. The gist of our approach is that researchers should choose cases for comparison in light of two criteria. The first is the expected relationship between $X1/Y$, which is the degree to which cases are expected to feature the relationship of interest between $X1$ and Y in light of existing theory, empirical studies, and large- N data. The second is variation in case characteristics or the extent to which the cases are likely to feature differences in criteria that can facilitate hypothesis generation. Our comparative approach stands in stark contrast to common advice in the field, which stresses the selection of single cases based on extreme values.³

Our discussion is roughly divided into three parts. We begin by providing some background, defining what we mean by “mechanism” and clarifying why we use “pathway analysis” instead of the more common term “process

tracing.” We then provide an overview of our approach. We conclude by applying it and comparing it to a leading example of pathway analysis in the “resource curse” literature, a prominent example of an $X1/Y$ correlation featuring a nonlinear relationship and multiple causal mechanisms.

Before turning to the argument, several caveats are in order. This article is about case selection, not how researchers should proceed once they have chosen a case for pathway analysis. As such, our work is distinct from (and complementary to) the growing number of texts that describe process-tracing methods. In our view, it is telling that these works often analogize social science researchers to detectives trying to resolve a particular crime (as opposed to a crime spree; (see Collier 2011 for a recent example of this analogy). Although this analogy may be useful for thinking about how we can reach causal inferences (whodunit) from a small N , it is not particularly useful to thinking about multimethod research, because the hypothetical “detective” on a specific case does not have to select which case to investigate from among a large number of possible cases. Equally important, the detective does not have to consider how findings from one investigation generalize to other crimes that have not been studied. In conducting pathway analysis, however, scholars have to choose which cases to investigate and often want to infer something about the other, unobserved, unstudied cases that feature the relationship of interest.

Moreover, this article sets forth a method of case selection, not empirical findings. Our goal is to discuss research strategies and not critique or revise substantive findings. Nor would that be appropriate. In the context of pathway analysis, case selection does not guarantee particular findings; it provides a rationale for selecting cases from a population and a basis for developing general hypotheses that may apply to the unstudied cases. From this vantage, a “better” approach to case selection implies an improved research design that is more likely to lead to useful knowledge, given a set of analytical goals.

Definitions

Terminology in this area is tricky, as many of the key concepts are contested and some have been arguably stretched to include multiple and sometimes incompatible meanings (Elster 2007; Gerring 2008, 2010; Hedstrom 2005; Mayntz 2004; Norkus 2004). As a result, it is important to clarify at least some of our terms, beginning with the nettlesome concept of “mechanism.” In a thoughtful review, John Gerring (2010) finds that the literature on mechanisms features many definitions, including the following:

(a) the pathway or process by which an effect is produced, (b) a micro-level (microfoundational) explanation for a causal phenomenon, (c) a difficult-to-observe causal factor, (d) an easy-to-observe causal factor, (e) a context dependent (tightly bounded or middle-range) explanation, (f) a universal (i.e., highly general) explanation, (g) an explanation that presumes probabilistic, and perhaps highly contingent, causal relations, (h) an explanation built on phenomenon that exhibit law-like regularities, (i) a technique of analysis based on quantitative or case study evidence, and/or (j) a theory couched in formal mathematical models. (p. 1500-01)

We have no desire to parse these conflicting definitions or make strong claims about which definition “best” captures the essence of the concept. Nevertheless, it is important to locate our definition within the array of competing alternatives in the literature. For starters, we agree that the meaning of mechanisms is highly context-specific, dependent on both the underlying type of research being conducted and the state of technology. A cognitive scientist might think of a mechanism differently than a political scientist, even though both may be interested in studying decision making. Similarly, technological changes can make today’s unobserved mechanisms into tomorrow’s well-measured variables.

Because the concept of a mechanism depends on the nature and state of the relevant research agenda, it is important to relate our working definition to pathway analysis. By its nature, pathway analysis explores the underlying links between some explanatory variable ($X1$) and some outcome (Y), controlling for other factors ($X2$). Accordingly, for purposes of pathway analysis, mechanisms are unobserved. (These unobserved factors may or may not be at a lower level of analysis than $X1$, $X2$, or Y .) This does not imply that mechanisms are *unobservable*, only that they are currently *unmeasured* in the large- N data. In addition, the definition of pathway analysis implies that mechanisms lie between $X1$ and Y in a causal chain, so that $X1$ is a cause of the mechanism and the mechanism is a cause of the outcome. Finally, variables in this chain can be seen as mechanisms for some research questions or as explanatory variables for other questions, depending on what part of the causal chain remains unexplored and is of interest.

Because pathway analysis is part of a mixed-method research agenda, it also implies certain epistemological commitments. First, the mixed-method nature of pathway analysis means that quantitative and qualitative studies must share a common set of concepts, so any definition of mechanism must be compatible with the relevant estimation techniques. This means mechanisms, in the context of pathway analysis, are treated as conceptually

analogous to mediating or intervening variables in standard regression analyses (Baron and Kenny 1986; Gerring 2012; Imai, Tingley, and Yamamoto 2010), which simply implies that the mechanism is caused by $X1$ and occurs on a causal pathway between $X1$ and Y .

Second, our working definition of mechanisms implies that they can, at least in principle, be manipulated (Gerring 2012). We recognize that this assumption is also a matter of debate among philosophers of science, but we believe that it is reasonable. As Gerring (2012) notes, claims about mechanisms imply that a mechanism's absence will have some effect on the underlying correlation between two variables. This suggests the ability, at least in theory, to remove or alter the mechanism; or, in other words, to manipulate it. As a substantive matter, this assumption is consistent with why at least some social scientists search for mechanisms. In seeking to develop policy prescriptions, for example, it may not be enough to know that there is a robust relationship between some explanatory variable ($X1$) and an outcome. Policy interventions may focus on manipulating a mechanism rather than the explanatory variable for which there is an established association, because it may be very difficult and/or politically infeasible to modify the key explanatory variable. Given these considerations, it would be odd to define mechanisms in a way that precluded manipulation. In sum, for our purposes of pathway analysis, mechanisms are unobserved factors that lie between an explanatory variable and an outcome in a causal chain. They are analogous to mediating or intervening variables that can, at least in theory, be manipulated.

Another definitional issue concerns our use of the term pathway analysis instead of the more familiar term process tracing. Process tracing has been used in diverse types of research and in connection with a wide variety of claims (see generally, Bennett 2010; Collier 2011; George and Bennett 2005). It can be used to identify intervening variables and mechanisms as well as plotting the sequence of variables, probing the direction of causality, exploring the analytic boundaries of theories, examining the causes of specific events, and much else. Process tracing might be part of a mixed-method research agenda or not. Pathway analysis, by contrast, is much more targeted. It is part of a mixed-method research agenda, which seeks to elucidate causal mechanisms underlying related variables in a population of cases (Gerring 2007).

For these reasons, although process-tracing techniques are obviously relevant in the search for causal mechanisms and we are reluctant to add yet another term to the somewhat confusing lexicon of multimethod research, we believe that pathway analysis better captures the role of case studies in

mixed-methods, mechanism-centered research while avoiding some of the confusion that might arise by using the broader (but perhaps stretched) term process tracing. Note also that our term of pathway analysis differs from Gerring's narrower concept of "pathway cases" (2007:124-27). For Gerring, pathway cases must feature the $X1/Y$ relationship. As discussed later, pathway analysis often leads to studying cases that feature varying relationships between $X1$ and Y to improve hypothesis generation in the face of uncertainty about how mechanisms function. As such, pathway analysis always encompasses pathway cases; it often includes other types of cases as well.

Overview of Approach

As noted at the outset, pathway analysis aims to (1) gain insight into the mechanisms that connect some explanatory variable ($X1$) to some outcome (Y) in specific cases and (2) use the insights from these cases to generate hypotheses about mechanisms in the unstudied population of cases that feature the $X1/Y$ relationship. These two goals, in turn, imply several principles for case selection. The first goal of pathway analysis suggests the *expected relationship* criteria, which means the degree to which individual cases are expected to feature the relationship of interest between $X1$ and Y , given existing theory, empirical knowledge, and large- N studies. It is perhaps obvious, but studying mechanisms that underlie the $X1/Y$ relationship requires identifying cases where the $X1$ variable is related to the Y , controlling for possible confounds ($X2$; Gerring 2007). If the relationship between $X1$ and Y differs based on the values of $X1$, then a researcher also needs to understand how the relationship depends on the value of $X1$. The second goal of pathway analysis implies the need to consider *variation in case characteristics*, meaning the extent to which the cases selected vary in terms of the $X1/Y$ relationship, the X values, and the Y values. If it is not known how $X1$ generates Y , then comparing cases that feature different likely "levels" of the $X1/Y$ relationship can help gain perspective on the findings and generate hypotheses. Indeed, it is only through multiple case studies that we can begin to map the underlying pathways between $X1$ and Y and gain confidence about the underlying structure of the $X1/Y$ relationship.

Readers will note that choosing cases based on the expected $X1/Y$ relationship does not guarantee that the selected cases will, in fact, feature the $X1/Y$ relationship; *expected* relationships are not the same as *observed* relationships. Even if researchers have strong theoretical or empirical reasons to believe that specific cases will feature particular $X1/Y$ relationships or believe they know how the large- N data relate to the presence of mechanisms

in specific cases, there is no guarantee that these expectations will be met. Theories or empirical knowledge could be wrong, there might be some mechanisms that block the $X1/Y$ relationship, or researchers might just be unlucky and happen to pick anomalous cases.

The need to select cases using observed characteristics when the actual interest is in yet-unobserved characteristics (i.e., causal mechanisms) is a key problem for researchers interested in using mixed methods to study causal mechanisms. Although we argue it is useful to employ large- N data to guide case selection—as well as existing theoretical and empirical knowledge—the use of quantitative data to understand the expected $X1/Y$ relationships at the level of individual cases does not eliminate the possibility that a researcher's expectations will fail to matchup with actuality.

The potential gap between the expected $X1/Y$ relationship and the actual relationship underscores the question of how to best use large- N data to select individual cases. There is no foolproof way to do this. All things being equal, we argue it is useful to select cases using what is known about the values of the explanatory variable ($X1$), the controls ($X2$), the outcome (Y), and an estimate of the expected relationship between the $X1$ and Y variables. So, for example, suppose researchers run an experiment on getting out the vote and find that a phone call from a neighbor 24 hours before an election has a significant effect on turnout. They then might want to understand *how* the phone call affected turnout; that is, what are the unobserved mechanisms that connect the call to turnout? Is it the tone of the call or the caller's voice? The timing of the call? The specific content of the conversation? The degree to which a caller has a personal connection to the voter? Or some combination of these factors or other ones?

To understand these questions, it is critical to understand the potential gap between the expected $X1$ relationship and the observed $X1/Y$ relationship (regardless of the exact method for assessing the expected relationship). This is one reason why we urge scholars to adopt a comparative approach: Selecting multiple cases based on the available information—theoretical, substantive, and large- N data—serves as a hedge against the possibility that some of the cases selected might fail to feature the $X1/Y$ relationship as expected. However, even cases in which prior expectations about the $X1/Y$ relationship are not met can be useful, provided that a researcher asks the right types of questions in the case study. Specifically, if a researcher finds that the expected $X1/Y$ relationship does not materialize in the case, the questions to ask include the following: Why was the expectation about the $X1/Y$ relationship incorrect? Was there a mechanism that blocked the expected effect of $X1$ on Y ? Does the case simply fail to feature the $X1/Y$ relationship? If so,

what does this failure suggest about the prior understanding of the $X1/Y$ relationship? Is the case an anomaly or does it require rethinking the association between $X1$ and Y ? Tracing the emergence of $X1$ over time within the case study would be particularly useful in considering these questions, all of which are relevant to the search for mechanisms and the quest to better understand the $X1/Y$ relationship.

Assessing the trade-offs between expected relationships and variation in case characteristics can be difficult. If it is known that there is only one mechanism linking $X1$ and Y , and that it functions consistently across both the values of $X1$ and the expected $X1/Y$ relationship, then it is not necessary to be concerned with variations in case-level attributes. Often, however, a researcher will not know the number of mechanisms, how they interact, or whether they function consistently. Given this uncertainty, it is important to select more than one case for pathway analysis, and these cases should vary across a variety of relevant dimensions, so a set of cases *taken together* provide opportunities for comparisons that offer some leverage on both criteria.

Implementing Our Approach

While there is no mechanical formula for selecting cases for pathway analysis, four steps are useful. First, researchers must assess whether the analytic requisites of pathway analysis are met, including the existence of a robust $X1/Y$ relationship and relevant data. Second, researchers should review the literature on the $X1/Y$ relationship to ascertain what is already known about its underlying structure. Third, researchers should visualize variation among the key variables by, for example, creating a histogram of the values of $X1$ and a scatterplot of $X1$ values, values of what we call the “expected relationship” between $X1$ and Y , and Y values. Finally, researchers should select cases using this information that feature interesting variation and, where appropriate, using case control strategies that target cases with relevant combinations of similarities and differences. Each step is discussed in turn.

Step 1: Assess Whether the Requisites of Pathway Analysis Are Met

Given the goals of pathway analysis, the literature on the underlying $X1/Y$ relationship must meet two minimal requirements. First, it must establish a robust relationship between $X1$ and Y that is likely to represent a causal relationship. We are well aware that confidence in whether an $X1/Y$ relationship represents a causal effect will vary across settings. Scholars will need to make an assessment whether sufficient reason exists to believe that

the relationships in our examples are causal, so that there is reason to investigate causal mechanisms. In practice, this assumption should not be made lightly and should reflect a variety of factors including the quality of the underlying model/empirical results and scholarly agreement about the $X1/Y$ relationship.

Second, large- N data sets are needed that help understand (a) the functional form of the relationship between $X1$ and Y , (b) the values of $X1$ and Y (and relevant controls) in specific cases, and/or (c) the expected magnitude and direction of the relationship between $X1$ and Y in individual cases. If the existing literature does not establish a robust relationship and provide useful data, then pathway analysis is not appropriate.

Step 2: Review the Literature on the $X1/Y$ Relationship

Pathway analysis is primarily directed at understanding an $X1/Y$ relationship, so reviewing the literature to identify gaps in existing knowledge about that relationship and to define the types of questions to be explored in the field is important. The difficulty is that what is known about an $X1/Y$ relationship can be hard to ascertain for a variety of reasons, including that relevant insights might be spread over a number of studies, which may or may not engage one another; researchers often bury assumptions about the relationship in their models; and they often do not use the same terms when describing underlying processes or mechanisms. Under these circumstances, it is useful to develop some heuristics for organizing existing studies and identifying (and aggregating) their core findings as well as providing a tool to help reveal a researcher's own assumptions about the $X1/Y$ relationship and how these assumptions might affect their research.

The key to this process is recognizing that there are multiple possible relationships among $X1$, Y , and the related mechanisms, and these relationships imply distinct types of questions related to the case studies and interpretation of the findings. One way to conceptualize these differences is to consider the four scenarios in Table 1, which represent paradigmatic examples of relationships between a key explanatory variable, the intervening mechanisms, and the outcome. In Table 1, $X1$ represents a key explanatory variable, M represents a mechanism, Y represents the outcome, and the arrows capture the direction of the relationship between these three. The scenarios are a simple way to represent our understanding of the relevant relationships. If $X1$ is directly connected to Y , it implies a direct effect of $X1$ on Y . If $X1$ is connected to Y through a mechanism (M), then that is considered an indirect effect of $X1$ on Y . These scenarios describe the *overall* or aggregate pattern

Table 1. Common Scenarios for Underlying Relationships Among Causal Variables, Outcomes, and Mechanisms.

Scenario (Description)	Graphical Representation of the Overall Structure of the X1/Y Relationship	Interpretation of Graph
Scenario 1 (Single Pathway Scenario)	<pre> graph LR X1 --> M1 M1 --> Y </pre>	There is only one mechanism; X affects Y only through M1
Scenario 2 (Direct and Indirect Pathways Scenario)	<pre> graph LR X1 --> Y X1 --> M1 M1 --> Y </pre>	There is only one mechanism; Mechanism represents one of the ways that X affects Y, but X may also directly affect Y regardless of M1
Scenario 3 (Multiple, Exclusive Pathways Scenario)	<pre> graph LR X1 --> M1 M1 --> Y X1 --> M2 M2 --> Y </pre>	Mechanisms are mutually exclusive; Mechanisms M1 and M2 represent possible ways that X affects Y, but X may also directly affect Y regardless of M1 and M2
Scenario 4 (Multiple, Nonexclusive Pathways Scenario)	<pre> graph LR X1 --> M1 M1 --> Y X1 --> M2 M2 --> Y X1 --> Y M1 <-.-> M2 </pre>	Mechanisms can occur independently or simultaneously; Observing M2 may be more likely if M1 is present than if M1 is absent or vice versa; The two mechanisms may interact to lead to a larger indirect effect on Y than the addition of each mechanism's indirect effect would suggest; The mechanisms may cancel out each other's effects

of relationships between $X1$, Y , and the M s, and not whether the indirect, direct effects, or both occur in a *specific* case. Indeed, mapping the various pathways will typically take multiple rounds of pathway analysis, which is another reason why it is critical for scholars to be transparent in their case selection, so that the findings of different studies can be aggregated. It is also important to note that all of these scenarios are simplifications in that they feature only a very small number of mechanisms and the relationships are relatively straightforward. There are likely to be many more mechanisms and the relationships between them are likely to be quite complicated, but these simple scenarios encapsulate key analytic differences among distinct types of $X1/Y$ relationships.

In the simplest case, scenario 1, the Single Pathway Scenario, the literature posits a single mechanism ($M1$) and a single pathway between $X1$ and Y , so that the entire $X1/Y$ relationship occurs through the mechanism and the relationship between $X1$, Y , and $M1$ is either positive or negative. When there is only a single path and mechanism between $X1$ and Y , the identification of a causal effect of $X1$ on Y implies the existence of the mechanism $M1$, because $X1$ can only affect Y via the $M1$ pathway. In this scenario, researchers might choose to focus their attention on two different issues. First, it may be useful to further unpack the relationship between $X1$ and Y by developing either a deeper understanding of how the $X1$ variable triggers the mechanism or how the mechanism affects the outcome (i.e., understanding how the arrows in the diagram function in practice). Second, it may be useful to know whether $M1$ functions similarly across all values of $X1$. Because scenario 1 seems unlikely for the complex phenomena studied by social scientists, it would also be useful to look for evidence that would suggest that this is the inappropriate scenario, such as whether there are other variables that cause both $X1$ and Y , whether $X1$ has direct effects on Y , or whether there are other variables that might cause both $X1$ and $M1$.

Scenario 2, the Direct and Indirect Pathways Scenario, is more complex. There is a single mechanism but more than one path between $X1$ and Y : a direct effect of $X1$ on Y and an indirect one caused by $M1$. This scenario is equivalent to partial mediation in which even after accounting for the mechanism/mediator, there is still a direct relationship between X and Y . If the causal relationship involves complete mediation, it means that the mechanism captures the entire effect of X on Y , and therefore the direct relationship is absent, which essentially reduces it to scenario 1, the Single Pathway Scenario.

Scenario 2 suggests a variety of related questions for pathway analysis. For example: Does the direct effect of $X1$ on $M1$ only exist for certain values of $X1$?

Where $M1$ is present, is the relationship with the outcome entirely mediated by $M1$ (as in scenario 1)? If this is possible, then a reasonable purpose of pathway analysis is to shed light on what values of $X1$ are associated with the direct and/or indirect effect. Researchers must be careful in how any single case under this scenario is interpreted. Unlike the Single Pathway Scenario, the failure to observe $M1$ in scenario 2 does not necessarily raise questions about the $X1/Y$ relationship as posited, because $X1$ can directly affect Y even without $M1$. Moreover, even if $M1$ is observed in a single case, that case cannot offer an answer as to whether that mechanism functions similarly across different effect sizes of $X1$ or whether particular effect sizes of $X1$ are associated with the presence or absence of $M1$, which are both important questions for understanding how $M1$ links $X1$ and Y . In addition, it is important to try to understand whether $M1$ occurs without $X1$ and whether $M1$ co-varies with other X s. This information is useful in trying to understand more fully the ways that the key explanatory variable affects an outcome.

Scenario 3, the Multiple, Exclusive Pathways Scenario, features multiple mechanisms and paths (so that there are direct and indirect effects), and the mechanisms are mutually exclusive. In general, case studies in research settings such as this one should focus on developing a better understanding of the mechanisms. Why are the mechanisms mutually exclusive? If one mechanism occurs does it block the other from occurring? Or, are certain values of $X1$ associated with $M1$ rather than $M2$? Are both mechanisms positively associated with the outcome or do the mechanisms differ in their relationship with the outcome? Answering all of these questions would be useful to fully understand the $X1/Y$ relationship.

Scenario 4, the Multiple, Nonexclusive Pathways Scenario, presents the most complicated situation: multiple pathways and multiple, possibly interactive mechanisms, which adds the possibility that $M1$ and $M2$ are related. In pursuing case studies, researchers must ask the following questions: Are the mechanisms related? Under what conditions does one, the other, or both occur? Do values of $X1$ associate with any pattern of the mechanisms? The possible interactions between $X1$ and $M1$ and $M2$ in this scenario limit the general conclusions reached from case studies, because it is difficult to tell whether a given case is typical of all of the possible relationships. In this situation, researchers would want to eventually understand all of the various relationships in the diagram, which is a huge challenge that also presents tremendous opportunities for case studies to add to the substantive understanding of the $X1/Y$ relationship.

To this point, we have assumed that the literature allows a researcher to determine whether an $X1/Y$ relationship resembles one of the scenarios in

Table 1. However, what is known about the causal mechanisms will often be quite limited, thus researchers will simply not know which scenario best captures the underlying relationship. There are, for example, many studies that establish robust positive or negative relationships but provide limited insight into the possible mechanisms underlying them. In these cases, the association among variables is largely a black box and the underlying relationship might fall under any of the four scenarios. In myriad academic areas, the relevant literature may suggest that multiple pathways link a variable and an outcome but fail to fully specify the paths. Other studies may specify multiple likely mechanisms but do not address how they function across different case characteristics, how they might interact with each other, or how they relate to other variables. Under either of these conditions, a researcher can have confidence that neither scenario 1 nor 2 applies (because there are multiple mechanisms) but cannot be sure whether the relationship more closely resembles scenario 3 or 4 (because there is uncertainty about the relationships between the multiple mechanisms).

If there is a basic lack of knowledge about the structure of the relationship between $X1$, Y , and Ms , then the primary task of pathway analysis is hypothesis generation about which scenario best captures the underlying relationship. To do this, it is necessary to adopt a comparative research strategy, and the comparisons are primarily useful for contributing to our understanding of four issues: (1) the number of likely mechanisms underlying the $X1/Y$ relationship (or the lower bound estimate of them); (2) the relationships among different mechanisms, including whether they interact and their potential observational equivalence; (3) the extent to which mechanisms (and how they function) are related to values of other key explanatory variables and possible confounds; and/or (4) the observable implications or measures of mechanisms.⁴

Step 3: Visualize Variation in the Data

To zero in on promising cases for analysis, we suggest that researchers create a visualization of the available information. The specific form of this visualization will depend on the context, the nature of the data, and the underlying models in the relevant literature. As a practical matter, it will often be useful to begin by examining the $X1$ values in a histogram (or other distributional plot) to understand the distribution of this variable. A histogram helps to identify common values of $X1$, which is useful to both identify interesting variation in case characteristics and understand how the cases we select compare to other cases in the sample.

Once there is a sense of the distribution of the $X1$ variable, the next step is to assess the expected relationship for each case. This differs from trying to identify the average treatment effect of the $X1$ variable, because rather than identifying the average relationship across the cases, we want to estimate the *expected* $X1/Y$ relationship for each case in the data set. There are a variety of approaches to how researchers might assess the expected relationship, in addition to using existing theory and empirical studies, which might be scant. (We present two related approaches in this article, but for more approaches, see Weller and Barnes [2014] including approaches that use matching methods in lieu of regression-based case selection.)

One way to assess the expected relationship is to compare the residuals or predicted probabilities from a regression model that includes $X1$ to the residuals/predicted probabilities from a model that excludes $X1$.⁵ In either case, the expected relationship is highest for cases where the inclusion of $X1$ is associated with a large difference in the predictions from the regression.⁶ It is important to remember that this is *not* sampling on an unobserved mechanism; instead, it is using the data to help locate cases that feature the expected $X1/Y$ relationship, which will be unpacked in the qualitative analysis of the case.

Another way to assess the expected relationship is to estimate the relationship between $X1$ and Y value at multiple values of $X1$.⁷ This information can then be used to select cases in which the value of $X1$ is associated with a large estimated change in Y . This approach, called the marginal effect approach, focuses on the expected relationship at different values of $X1$ rather than on the expected relationship for specific observations/cases. This approach can be used to examine how the expected relationship between $X1$ and Y changes when we modify the value of one of the other variables ($X2$) in the statistical model. For instance, in studying war, a researcher might have reason to believe that the mechanisms that link natural resources (the $X1$) to civil conflict (Y) depend upon a country's level of democracy ($X2$).⁸ In such a case, a researcher needs to examine the estimated relationship between natural resources and conflict at different levels of democracy in an attempt to identify which cases to select, so as to understand how a mechanism might depend on the level of democracy. In general, this approach is most useful when researchers want to investigate particular values of $X1$ or $X2$ and need a way to estimate the expected relationship, given these values. This approach can also aid in identifying the likely expected relationship for cases where there are some missing data, and therefore the approach of comparing residuals or predicted probabilities cannot work.

These preceding steps, along with an existing large- N data, offer several pieces of information for locating interesting patterns of variation in case characteristics. The challenge becomes how to organize these bits of information into a useful format. Again, there is no mechanical way of doing this, but one useful strategy is to make a scatterplot with $X1$ values on the x-axis, the expected $X1/Y$ relationship on the y-axis (how this measure is generated will be discussed later), and each point on the scatterplot labeled with the relevant information.

In selecting cases, the goal is to keep in mind the state of the existing literature and the ultimate goals of the multimethod research agenda. Specifically, we want to stress that using comparisons can help gain a better sense of the underlying structure of the $X1/Y$ relationship in a number of ways. It can shed light on the number of mechanisms connecting $X1$ and Y (or at least lower bound estimates of them). It can help probe the relationship among multiple mechanisms, including (a) whether all mechanisms are positive or negative, (b) whether some mechanisms cancel each other out, and (c) whether the mechanisms occur in the absence of expected effects or outcomes. It can also provide insight into the relationship between the mechanisms and other variables, such as the possibility of observing the key explanatory variable without a mechanism appearing and the relationship between mechanisms that link $X1$ and Y and other explanatory variables.

Step 4: Select Cases

The preceding step is likely to identify a large number of potentially promising cases. Given limited time and resources, researchers should seek to maximize the analytic leverage of their cases by using case-control strategies in case selection. In case-control research designs used in epidemiology or medicine, a researcher studies subjects who contracted a disease and subjects who did not have the disease. Ideally, the two groups of subjects are quite similar except for the hypothesized cause of the disease. The researchers then study the subjects' behavior and backgrounds prior to contracting the disease to see whether those with the disease differ systematically from those without the disease in a way that we suspect would be related to contracting the disease. Although this is a very weak research design for causal inference, it can help to generate hypotheses about the differences between the two groups that would warrant investigation via another research design. The same basic approach can be used for pathway analysis. Researchers could select cases that differ in the outcome and the value of the key explanatory variable but that appear to be similar on the other dimensions (i.e., the predicted

probability or outcome from a regression, the multivariate distance in a matching approach, or perhaps a single attribute like such as geographic region or time period). More generally, the idea of case control means that researchers intentionally choose cases that are similar on known dimensions but that differ in ways that allow generation of knowledge about a particular research question. Again, there is no cut-and-dry rule for how to do this, but throughout this article, we provide guidance about common strategies of case control that build on familiar comparative research methods.

Applying Our Approach

To illustrate our approach, we focus on the relationship between natural resources and civil conflict because it is important substantively and because there are examples of both large- N (Collier and Hoeffler 2004) and small- N (Ross 2004) research related to how a country's primary exports (as a percentage of gross domestic product [GDP]) affect civil war. Collier and Hoeffler (2004) use large- N methods to explore the relationship between primary exports and civil war. Their article has been cited more than 3,000 times on Google Scholar, suggesting that theirs is an influential argument. The presence of a relationship between primary exports and conflict has led several researchers to propose different causal mechanisms that might underlie the primary export–civil war relationship, which Ross (2004) investigates via a number of case studies that he selected based on the secondary literature without using Collier and Hoeffler's large- N data. As such, this literature offers a good opportunity to both apply our method and compare it to a leading example in the field.

Step 1: Assess Whether the Requisites of Pathway Analysis Are Met

The first step is to assess whether the literature establishes a robust relationship between levels of primary exports ($X1$) and civil war (Y) and provides the data needed to apply our method. We will assume that this relationship meets our threshold criteria, but several points are worth noting about the $X1/Y$ relationship at this stage of the process. First, in the original Collier and Hoeffler's article, the outcome variable is dichotomous and the authors model it using a logit equation. This implies that the relationship between any given independent variable and the outcome is not constant, and the relationship between the key explanatory variable and the outcome will depend on the values of the other variables in the model. Therefore, we would not want to simply pick the largest values of the key predictor variable and

assume that the effect of the predictor is greatest in those countries. Second, the Collier and Hoeffler's model includes primary exports and the square of primary exports, and both terms are significantly related to conflict. As we see in greater detail later, the negative coefficient on the squared level of primary commodities is associated with an inverted U-shaped relationship between primary exports and probability of civil war.

Step 2: Review the Literature on the $X1/Y$ Relationship

For the purposes of this analysis, we focus on Ross's characterization of the literature because we eventually compare our approach to his case selection. Based on the Ross's article, knowledge about the underlying mechanisms is fairly limited. It appears that there are at least four potential mechanisms that might connect primary exports and the onset of civil war: looting, grievances, incentives for separatism, and state weakness. However, we cannot be sure whether this is an exhaustive list of mechanisms or how these mechanisms might interact; indeed, Ross (appropriately) seeks to use his case studies to search for the possibility of new mechanisms. To use the language discussed earlier, we know that scenarios 1 and 2, the Single Pathway Scenario and Direct and Indirect Pathway Scenarios, do not apply because there are multiple mechanisms, but we cannot be sure how many mechanisms there are or if these mechanisms are exclusive (as in scenario 3) or interactive (as in scenario 4). The primary task of pathway analysis is to gain insights into the underlying structure of the $X1/Y$ relationship.

Step 3: Visualize Variation in the Data

With this background in place, we can begin to select cases using the data from Collier and Hoeffler (2004). In the Collier and Hoeffler's data set, each observation is a country-time period, and the time period is five years. The data range from 1960 to 1995, and there are at most eight observations for each country in the data set. There are a total of 46 cases in which there was a civil conflict in a country during a five-year time period.

Using these data, we begin by examining the distribution of the key explanatory variable: primary exports as a percentage of a country's GDP. This can be done most easily via a histogram, such as the one in Figure 1. At this stage, we are assessing the overall distribution of $X1$. The histogram makes it clear that the distribution is skewed with the vast bulk of countries having relatively low levels of primary exports. Therefore, if it is important to pick cases that have common values of

Figure 1. Distribution of primary exports.

the primary exports variable, we might want to focus on cases with values below 0.2 of primary exports because over 70 percent of our cases have levels of primary exports/GDP that are less than 0.2. Among countries in the Collier and Hoeffler's data set that experienced a civil war, the average level of primary exports/GDP was 0.15, the 25th percentile was 0.07, and the 75th percentile was 0.20. This is useful to know, because it indicates that many of the cases that experience civil conflict have common values of the key explanatory variable.

Once we have a sense of the distribution of primary exports (X_1), the next step is to examine the expected relationship between primary exports and civil war (Y). Because the outcome in this case is a dichotomous variable, not a continuous one, comparing residuals based on ordinary least squares estimates as suggested by some is inappropriate and does not match the regression approach used by Collier and Hoeffler in their 2004 article. Instead, we compute the predicted probability of a civil war (Y) with and without the primary exports variable (X_1) using a logit regression. In Table 2, we present the exact regression from Collier and Hoeffler in column 1 and a reduced form of the same equation that excludes the primary exports variables in column 2. Using the two regressions in Table 2, we then can compute the predicted probability of a war including the primary exports variables and the predicted probability of a war without these variables.

Table 2. Relationship Between Primary Exports and Civil Conflict.

	Full Model	Reduced Model
Primary exports	18.1 (6.00)**	—
Primary exports squared	−27.44 (11.22)*	—
Post–Cold War period	−0.32 (0.47)	−0.33 (0.45)
Male enrollment, secondary education	−0.025 (0.01)*	−0.019 (0.009)**
GDP growth	−0.12 (0.04)**	−0.13 (0.04)**
Peace duration	−0.0025 (0.0016)	−0.002 (0.01)
Previous war	−.46 (0.54)	0.69 (0.52)
Mountainous terrain	0.013 (0.009)	0.006 (0.008)
Geographic dispersion	−2.21 (1.03)*	−1.40 (0.98)
Social fractionalization	−0.0001 (0.00009)	−0.00006 (0.00009)
Log of population	0.67 (0.16)**	0.37 (0.12)
Constant	−12.33 (2.71)**	−6.56 (1.94)**
N	688	688
Number of wars	46	46

Note: GDP = gross domestic product. Results in Reduced model column (column 2) restricted to the same set of observations used in Full model column (column 1); eight observations are excluded from column 2 by restricting the observations to be consistent with column 1.

*Significant at .05 level. **Significant at .01 level.

The next step is to combine this measure of the expected relationship with information about the values of X_1 and Y to get a better sense of how specific cases relate to the other cases in the sample. Here, a scatterplot with this information is useful. In Figure 2, the y-axis is our measure of the expected relationship: the difference in the predicted probability between the model with primary exports and the model without primary exports. A positive number indicates that the predicted probability of conflict is higher in the model that includes the primary exports variable than in the model that excludes that variable. The x-axis is the value of the key independent variable: primary exports/GDP. The figure also includes information about Y : where war occurred and where it did not occur. The “war” cases are marked by X s whereas the “nonwar” cases are marked by 0s. Each point contains three pieces of information: expected value, primary resource level (X_1), and the presence/absence of war (Y).

We will return to Figure 2, which we believe provides the most straightforward presentation of the data, but note that there are other ways to visualize the expected relationship. For example, we could plot the estimated change in Y values (probability of civil conflict) at various values of the independent variable (primary exports). In a logit regression, this relationship is not constant, as

Figure 2. Expected relationship and levels of primary exports.

it would be in a linear relationship; therefore, we must evaluate the relationship at different values of primary exports and at different values of the other variables in the regression. The use of the primary exports variable and its squared value also means that we must account for both of these factors in estimating the relationship between changes in the value of primary exports and civil conflict. One might use a figure such as this to estimate the expected relationship for cases, where we know the value of primary exports, but other data for the case are missing and therefore the case is not used in the regression, which precludes directly comparing predicted probabilities.

In Figure 3, we present the estimated relationship between primary exports and the probability of civil conflict, holding the value of other variables at their means.⁹ The figure demonstrates the expected inverted-U relationship between the level and effect of primary exports. One important piece of information from this curve is that beyond a primary exports value of about 0.4 higher values of primary resource dependence are associated with a lower probability of conflict. The figure also shows that at the most extreme values of primary exports, there is little relationship between them and the likelihood of conflict. In addition, this plot indicates that a one-unit change in primary exports has its biggest relationship with the probability of conflict in cases with intermediate values of primary exports—between 0.2 and 0.4, with the peak being just below 0.4.

Figure 3. Marginal effect of primary exports on probability of civil war.

Step 4: Select Cases

The final step is to look for patterns in the data in light of the distribution of primary export values, the relationship between primary exports and the probability of civil war, and whether war occurred. It is useful to return to the scatterplot in Figure 2 to look for patterns of variation. A couple of things emerge from this figure. When we look at the expected relationship along the y-axis, it appears that cases fall in to three general groups: those with a negative change in the predicted probability, those with a small positive increase in predicted probability (<0.1), and those with a larger positive increase in predicted probability (>0.1). When we consider the level of primary exports along the x-axis, it appears that there are two general groups of cases: those with primary exports below 0.2 and a small number of cases with values above that level.

With knowledge of the overall patterns in the data, we can target specific cases. At first glance, Iran provides a promising set of overtime comparisons, which will be useful to help us observe how changes in the key explanatory variable relate to causal mechanisms.¹⁰ At the very beginning of the data set, in 1960, Iran had high levels of primary exports (0.47) but no civil conflict.¹¹ However, in both 1970 and 1975, with no change in the level of primary exports, there is civil conflict. Why? Were new mechanisms (associated with primary exports) present in 1970 and 1975 that account for the conflict? Did

several mechanisms serve to cancel each other out in the earlier period? In addition, the expected relationship is higher in 1975 than 1970 but the outcomes were the same. Were the same mechanisms present? Did they function in the same way despite different levels in the expected relationship?

By 1980, the level of primary exports in Iran and the expected relationship have sharply decreased, and yet civil conflict continues. Again, did the same mechanisms exist and function in 1970, 1975, and 1980? If so, did they function in the same way or do the links between primary exports and civil conflict change over time and do the mechanisms differ as the expected relationship differs? Alternatively, is the decline in primary exports associated with different mechanisms? From 1980 to 1985, the levels of primary exports slightly decreased from 0.14 to 0.07, and the expected relationship shifts from positive to negative and there is no longer civil conflict in Iran. Did the mechanisms disappear? Were there new mechanisms that cancelled the effect of primary export on civil conflict in 1985? Looking across all of the time periods in Iran, how many mechanisms were present? Did some mechanisms always appear together? Were some more prominent at different levels of primary exports or levels of the expected relationship? Angola between 1975 and 1990 and Nigeria from 1980 to 1995 also feature potentially interesting within-case variation, giving rise to similar questions and offering opportunities to compare pathways across contexts and eventually map the underlying structure of the relationship.

From the perspective of expected relationship, cases like Iran 1975 and Nigeria 1995 are appealing; however, these cases are quite different from the bulk of the data—in terms of both their high measures of the expected relationship and the high levels of primary exports—which raises concerns about whether hypotheses generated from these cases would apply to other cases with more typical primary export levels or cases with more common scores of the expected relationship. This concern would lead us to also select cases that have more typical values on these two dimensions. In this vein, a promising set of comparisons would be to examine a number of Central American countries in 1975 that have much more typical primary export values.¹² One promising comparison might be Nicaragua and El Salvador in 1975, because both have similar levels of reliance on primary exports but different changes in the probability of conflict. Because the primary export scores are high, near the .20 level or 75th percentile, we would ideally add Guatemala in 1975, because this case has a more typical primary resource score of .13, which is close to the average score of .15. An additional feature is that these cases allow us to make comparisons at a similar point in time and the three countries are geographically proximate.¹³ One downside of the Central

American cases is that the mechanisms that cause civil conflict may be difficult to observe, because the expected relationship is smaller than it is in the Iran, Angola, and Nigeria cases.

A complete understanding of mechanisms would also entail selecting cases in which the outcome does not occur. If scholars want to select cases of that sort, it seems useful to examine India during the 1960s and 1970s, because it features very low primary exports, and in 1965 there is conflict, whereas in 1970 there is not a civil conflict. Studying a case such as this, with both small $X1$ value and negative expected relationship, allows researchers to consider whether similar mechanisms appear in cases that are very different on the two dimensions we use to identify interesting case comparisons. Researchers should keep in mind that the primary goal in conducting pathway analysis is to develop hypotheses about how mechanisms function, not to make causal inferences about either the specific effect of mechanisms in a case or their general effect across cases.

This exercise underscores several general points. It may be obvious, but our ability to make claims about unobserved cases depends on whether we know enough to argue that the handful of cases we have observed represent the important aspects of the unobserved cases. If we lack sufficient knowledge to make confident generalizations, we must develop second best strategies for hypothesis generation using the data at hand. This requires us to study mechanisms in cases that feature common values of $X1$ and a variety of relationships between $X1$ and Y . To the extent that we can combine these criteria with standard case-control strategies, such as comparing countries within the same region at the same time or making within-country comparisons over time, so much the better. At a minimum, by considering the distribution of $X1$ values along with the expected relationship and variation in case characteristics, researchers will be able to clarify the inevitable trade-offs that underlie their choice of cases.

Comparing Approaches

In an ambitious and well-executed study, Michael Ross (2004) conducts 13 case studies to investigate the causal mechanisms purported to underlie the link between primary exports and civil conflict. Ross does not call his research pathway analysis, but he was clearly engaged in what we consider pathway analysis because he focused on exploring causal mechanisms across a number of cases in order to shed light on this important substantive research question. Ross was appropriately aware of the limits of his case studies. He recognized that he could not identify average effects or definitively demonstrate or rule

out particular mechanisms using his research design, but his goal was to determine whether there was evidence for the existence of a purported causal mechanism and endeavored to identify new mechanisms if they existed. As noted earlier, Ross was uncertain about which scenario (from Table 1) applied and was using pathway analysis to probe the basic structure of the $X1/Y$ relationship, especially the number and identity of the key mechanisms linking natural resource wealth and civil conflict.

Ross considered a variety of different mechanisms that relate to the onset, duration, or intensity of civil war. We focus on the mechanisms he identified as related to the onset, because this is the relationship explored by Collier and Hoeffler (2004) in their large- N study of civil conflict, which we use in the next section to help contextualize Ross's cases. As noted earlier, Ross identifies four mechanisms from the existing literature that purportedly connect primary export levels to civil conflict: looting by potential rebels, grievances among locals, incentives for separatism, and state weakness due to reliance on revenue from natural resource. To see whether the hypothesized mechanisms actually existed, Ross selected 13 "most likely" cases, which he defined as cases in which a civil war occurred, and his reading of secondary source material suggested that primary exports played a role in the origin of the conflict. It is unclear if most likely means that observing the relationship is most likely, if observing a particular mechanism is most likely, or if it means something different all together. Regardless, Ross's case selection process was (like ours) clearly designed to use some observable indicators to pick cases that he examined to study whether as-yet-unobserved causal mechanisms were present. In Table 3, we list the 13 cases Ross studied as well as his conclusions regarding the presence of a particular mechanism that connects primary resources to civil conflict.

In general, Ross found little support for the purported mechanisms but did identify two unanticipated mechanisms: "sale of future rights to war booty" and intervention from a third party (Ross 2004:50). Overall, he argued that there was "no evidence in the sample of the looting mechanism, and little if any evidence of the grievance mechanism" (p. 50). This only pertained to the 13 cases in Ross's sample, and therefore any general conclusions a researcher might draw depend crucially on how these cases compare to the many unstudied cases. Put differently, it is necessary to gain some perspective on Ross's cases.

Ross' Case Selection and the Expected Relationship Criterion

We can evaluate Ross's case selection in light of our criteria using the results from Collier and Hoeffler's large- N analysis. At the outset, it should be

Table 3. Cases Chosen to Study Primary Exports and Civil Conflict.

Country	Year in Collier and Hoeffler Data	Primary Export/GDP	Expected Relationship	Marginal Effect of Primary Exports
Congo Republic	1995	0.505	−0.003	0.04
Angola	1975	0.476	0.14	0.059
Liberia	1990	0.393	0.	0.076
Indonesia	1975	0.219	0.17	0.055
Congo, Democratic Republic	1995	0.141	—	0.037
Peru	1980	0.130	0.008	0.037
Sierra Leone	1990	0.120	−0.007	0.021
Colombia	1980	0.094	−0.013	0.021
Sudan	1980	0.086	−0.07	0.021
Myanmar/Burma	1980	0.078	−0.017	0.01
Cambodia	1970	0.052	—	0.01
Afghanistan	1990	0.033	—	0.01

Note: GDP = gross domestic product. Cases based on Ross (2004). Marginal effects were estimated from the regression in column 1, table 7.2 at every 0.05, from 0.00 to 0.50, and the listed effect corresponds to the estimated effect at the nearest level of primary exports with other variables held at their means.

stressed that there are reasons to be concerned about a “most likely”/variable-based approach if Collier and Hoeffler are correct about the functional form of the relationship between primary exports and civil war. If the underlying relationship is nonlinear, then we may not be able to observe a relationship between primary exports and conflict across the entire range of primary exports. Recall that the prior figures showed the expected relationship between primary exports and civil conflict. For both low and high values of conflict, there is essentially no expected effect of resources. Therefore, it would not seem surprising if causal mechanisms are not apparent in these cases, and it may not be informative about whether the mechanisms appear at other values of primary exports. Put differently, if we choose cases based on extreme value of primary exports, then we run the risk of false negatives given the functional form of the underlying relationship.

Table 3 lists the cases Ross chose in his case studies, as well as each case’s value on the key explanatory variable, each case’s expected relationship from the logit regression, and the estimated marginal relationship based on each case’s value of the primary export variable. A few things jump out regarding the context of these cases. First, in most of these cases, the inclusion of the

primary exports variable barely increases the predicted probability of a civil conflict. In fact, for a number of the cases, the inclusion of primary exports appears to have *lowered* the predicted probability of a war. Consistent with our estimates, Ross finds very limited support for the existing purported causal mechanisms in his analysis of these 13 cases. However, this may not be a surprise, given that these cases have relatively small expected relationships. Second, in the few cases that we could not directly compute the expected relationship, the marginal effects approach suggests a small relationship between primary exports and civil conflict. Again, this implies that it might have been difficult to actually observe a causal mechanism in these cases.

Ross' Case Selection and Variation in Case Characteristics

In addition to the expected relationship, it is important to consider how the chosen cases support hypothesis generation to other, unobserved cases. Much to his credit, Ross is keenly aware of problems with making claims about other cases, explicitly stating: “[T]he findings cannot be generalized to some larger set of unexamined cases” (2004:37). Although we agree that generalizations can be difficult, a researcher’s ability to make general claims about mechanisms will improve if cases are chosen that vary on important dimensions. It will also improve if a researcher adopts an explicitly comparative research strategy that allows observation of how mechanisms function across a wide range of relationships between $X1$ and Y , while also considering the distribution of $X1$ values. Put differently, we may not be able to generalize, but we may be able to generate plausible hypotheses that we expect to apply to as-yet unstudied cases. In his case studies, Ross does not adopt an explicitly comparative research strategy that involves picking particular cases with a goal of comparing them to other selected cases. Comparative designs allow a researcher to directly ask questions about how aspects that are different (or similar) between cases are associated with different (or similar) causal mechanisms. Comparisons are crucial for building knowledge about the substantive $X1/Y$ relationship. In this particular example, one might use the large- N results from Collier and Hoeffler to extend Ross’s research by picking additional cases that are good comparisons to the ones already studied. For example, as a comparison to Indonesia in 1975, one might choose to study Nigeria in 1975, which features a similar $X1$ value and similar expected relationship, yet Nigeria did not experience a civil conflict. In a comparison such as this one, a researcher might examine whether there are inhibitory mechanisms in the Nigeria case or study if the mechanism requires some other factor to lead to war that is missing in Nigeria but present in Indonesia.

As discussed in the previous section, we might also want to focus on other comparisons like Iran or Angola that feature interesting within-case variation or cases with more typical levels of exports such as Nicaragua, El Salvador, and Guatemala in 1975.

The broader point is that our ability to make claims (or generate relevant hypotheses) about unobserved cases depends on the state of existing empirical and theoretical knowledge. The ability to make claims about unobserved cases depends on whether we know enough to argue that the handful of cases we have observed are representative of the unobserved cases. Lacking this knowledge, we must use what we know to develop second best strategies for hypothesis generation. This requires us to think about the relationship between the distribution of the key explanatory variable ($X1$) and the variation of its effects on the outcome (Y) and look for cases that will allow us to observe mechanisms in cases that feature common values of $X1$ and a variety of relationships between $X1$ and Y . To the extent that we can combine these criteria with standard comparative case selection strategies, such as comparing countries within the same region at the same time, making within-country comparisons over time or adopting “most different” designs, so much the better. At a minimum, by considering the distribution of $X1$ values along with the expected relationship and variation in other case characteristics, researchers will be able to clarify the inevitable trade-offs that underlie their choice of cases.

We should add that even if the contributions to the search for causal mechanisms are limited by the state of the literature, applying our approach yields substantively interesting empirical puzzles about the differential effects of oil wealth and civil conflict in light of the existing large- N data, which may be of interest to scholars and policy makers. A critical case approach fails to provide that benefit.

Conclusion

The search for causal mechanisms has produced an increased interest in the use of mixed-method research. Relatively little, however, has been written about practical strategies for how to select cases using the existing quantitative literature, especially when the underlying relationship is nonlinear and there are potentially multiple pathways underlying the relationship. This article puts forth a method of case selection for pathway analysis and demonstrates how our approach suggests selecting different cases than other approaches. The most important difference is our emphasis on *comparison*. Specifically, to explore causal mechanisms systematically requires selecting

cases based on both the expected relationships and variations in case characteristics using information about the values of key variables *and* the relationship in individual cases.

A critical point is that applying these principles depends on knowledge of the expected relationship between $X1$ and Y , the nature of the outcome, and the state of knowledge about causal pathways, which is often limited and uncertain. This means that researchers cannot take a mechanical approach to selecting cases. Instead, they must build a flexible tool kit for assessing the expected $X1/Y$ relationship and case variation and be aware of the limits of using different tools when interpreting the results of their case studies.

Finally, selecting cases based on the expected $X1/Y$ relationship and variation in case characteristics does not guarantee that a researcher will observe “typical” mechanisms, or even any mechanisms at all, though the quantitative data can provide useful information about the distribution of the $X1$ variable and the expected $X1/Y$ relationship. These bits of information can reveal interesting puzzles—such as why some cases have similar $X1$ values and expected relationships but different outcomes—that will lead to probing of the $X1/Y$ relationship across settings. Ultimately, it is the work of case studies to fill in the unobserved links between the variables; effective case selection can only identify promising puzzles to study. Exploring these puzzles will provide a reasonable basis for generating hypotheses about issues relating to the basic structure of the $X1/Y$ relationship, such as the number of mechanisms, whether mechanisms interact and covary with variables and outcomes, and whether mechanisms function similarly across different levels of effect. Insights on these issues can significantly contribute to a researcher’s understanding of the $X1/Y$ relationship, which is the primary goal of pathway analysis, and advance the elusive search for causal mechanisms. Over time, as cases accumulate and findings converge, we will be able to map the underlying structure of the $X1/Y$ relationship. Before we can do that, however, we need to be more systematic in how we select our cases for pathway analysis.

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

Notes

1. For purposes of this article, we define case studies as intensive analyses of single units observed at a specific time or over a specific period of time, with the goal of offering insights into a population of cases (Gerring 2007).
2. We do not argue that scholars have failed to consider how to select cases for other types of research, given causal complexity. There is a vast and very useful literature on these topics (e.g., Brady and Collier 2004; George and Bennett 2005; Gerring 2007; King, Keohane, and Verba 1994; Ragin 2000). The argument is that these issues remain undertheorized in the context of pathway analysis, which has very distinct analytic goals and thus requires distinct approaches to case selection.
3. A longer, more extensive treatment of the issues raised in this article and related issues around case selection and causal mechanisms can be found in Weller and Barnes (2014).
4. Imai and Yamamoto (2013) present a similar list of goals for building knowledge about causal mechanisms, but present no guidelines for how such knowledge should be developed.
5. Where Y is a continuous variable, it is possible to compare residuals between regression models that include and exclude $X1$; where Y is not continuous, it is possible to compare the predicted probability of a given outcome for each observation with and without $X1$.
6. This approach builds on Gerring (2007) and Seawright and Gerring (2008) and is akin to asking: How does knowledge of the key explanatory variable improve our ability to predict a case's observed outcome? The expected relationship does not have a causal interpretation but provides one tractable way for scholars to assess the likely relationship between $X1$ and Y for individual cases.
7. This could be a useful strategy even for linear models like ordinary least squares, if the estimated model contains higher-order terms or interactions between variables.
8. This is a more complicated scenario because in this scenario, another variable, $X2$, works in conjunction with $X1$ to determine which mechanisms are present.
9. The estimated relationship and the shape of the curve are substantively similar if we hold other variables at their medians, 25th or 75th percentile. This is a good thing for choosing cases because it suggests that our results are not highly sensitive to the value of other variables.
10. For purposes of discussion, we focus on case selection based solely on an analysis of the large- N data only. Of course, researchers might have good empirical, theoretical, or even practical reasons for focusing on other cases, such as access to data, language skills, or existing contacts in the field. Under these circumstances,

researchers can use the tools we describe to contextualize their cases or to identify promising cases for secondary analysis.

11. For reasons of missing data, we cannot estimate the expected relationship for Iran 1960.
12. If a researcher is constrained to a single-case strategy, and a comparative strategy is impractical for reasons of time, resources, or access to sources, there are ways to use the large- N data to identify single cases for hypothesis generation, although these are far less powerful than the comparative strategies. Nevertheless, if a researcher has no choice, given the extant state of knowledge about mechanisms for the resource curse literature, there seem to be two reasonable strategies for selecting a single case for the purpose of mechanism identification. One, pick a case that represents common values of $X1$ and resist the urge to make general claims about mechanisms across values of $X1$. Two, pick cases that feature a theoretically interesting effect (perhaps the maximum marginal effect or the absence of an effect) and again place any conclusions within the context of the general absence of knowledge about how mechanisms vary across cases.
13. If geographic proximity is not important, then we would add Indonesia in 1975, a case with analogous levels of primary exports as Nicaragua and El Salvador but a much higher increase in predicted probability.

References

- Achen, C. H. 1986. *The Statistical Analysis of Quasi-experiments*. Berkeley: University of California Press.
- Baron, R. M. and D. A. Kenny. 1986. "The Moderator-mediator Variable Distinction in Social Psychological Research: Conceptual, Strategic and Statistical Considerations." *Journal of Personality and Social Psychology* 51:1173-82.
- Beck, Nathaniel. 2006. "Is Causal-process Observation an Oxymoron?" *Political Analysis* 14:347-52.
- Beck, Nathaniel. 2010. "Causal Process 'Observation': Oxymoron or (Fine) Old Wine." *Political Analysis* 18:499-505.
- Bennett, Andrew. 2010. "Process Tracing and Causal Inference." Pp. 207-19 in *Rethinking Causal Inference: Diverse Tools, Shared Standards*. 2nd ed., edited by Henry Brady and David Collier. Lanham, MD: Rowman & Littlefield.
- Brady, Henry E. and David Collier, eds. 2004. *Rethinking Social Inquiry: Diverse Tools, Shared Standards*. New York: Rowman & Littlefield.
- Chatfield, Chris. 1995. "Model Uncertainty, Data Mining and Statistical Inference." *Journal of Royal Statistical Society, Series A (Statistics in Society)* 158:419-66.
- Collier, David. 2011. "Understanding Process Tracing." *PS: Political Science & Politics* 44:823-30.

- Collier, David, Henry E. Brady, and Jason N. Seawright. 2004. "Sources of Leverage in Causal Inference: Toward an Alternative View of Methodology." Pp. 229 in *Rethinking Social Inquiry*, edited by Henry E. Brady and David Collier. Lanham, MD: Rowman & Littlefield.
- Collier, Paul and Anke Hoeffler. 2004. "Greed and Grievance in Civil War." *Oxford Economic Papers* 56:563-95.
- Eckstein, Harry. 1975. "Case Studies and Theory in Political Science." Pp. 94-137 in *Handbook of Political Science*, vol. 7 *Political Science: Scope and Theory*, edited by Fred I. Greenstein and Nelson W. Polsby. Reading, MA: Addison-Wesley.
- Elster, Jon. 1998. "A Plea for Mechanisms." Pp. 45-73 in *Social Mechanisms: An Analytical Approach to Social Theory*, edited by P. Hedstrom and R. Swedberg. New York: Cambridge University Press.
- Elster, Jon. 2007. *Explaining Social Behavior: More Nuts and Bolts for the Social Sciences*. New York: Cambridge University Press.
- Freedman, David A. 1991. "Statistical Models and Shoe Leather." *Sociological Methodology* 21: 291-313.
- Freedman, David. 2008. "On Types of Scientific Inquiry: The Role of Qualitative Reasoning." Pp. 300-18 in *Oxford Handbook of Political Methodology*, edited by J. Box-Steffensmeier, H. E. Brady, and D. Collier. Oxford, UK: Oxford University Press.
- George, Alexander L. and Andrew Bennett. 2005. *Case Studies and Theory Development in the Social Sciences*. Cambridge, MA: MIT Press.
- Gerber, Alan, Donald Green, and Edward Kaplan. 2004. "The Illusion of Learning from Observational Data." Pp. 251-73 in *Problems and Methods in the Study of Politics*, edited by Ian Shapiro, Rogers Smith, and Tarek Masoud. New York: Cambridge University Press.
- Gerring, John. 2007. *Case Study Research: Principles and Practices*. New York: Cambridge University Press.
- Gerring, John. 2008. "The Mechanismic Worldview: Thinking Inside the Box," *British Journal of Political Science* 38: 161-79.
- Gerring, John. 2010. "Causal Mechanisms: Yes, But..." *Comparative Political Studies* 43:1499-526.
- Gerring, John. 2012. *Social Science Methodology: A Unified Framework*. 2nd ed. New York: Cambridge University Press
- Heckman, James J. and Jeffrey A. Smith. 1995. "Assessing the Case for Social Experiments." *The Journal of Economic Perspectives* 9:85-110.
- Hedstrom, Peter. 2005. *Dissecting the Social: On the Principles of Analytical Sociology*. New York: Cambridge University Press.
- Hedstrom, Peter and Petri Ylikoski. 2010. "Casual Mechanisms and the Social Sciences." *Annual Review of Sociology* 36:49-67.

- Imai, Kosuke, Dustin Tingley, and Teppei Yamamoto. 2013. "Experimental Designs for Identifying Causal Mechanisms (with discussions)." *Journal of the Royal Statistical Society, Series A (Statistics in Society)* 176:5-51.
- Imai, Kosuke, Luke Keele, Dustin Tingley, and Teppei Yamamoto. 2011. "Unpacking the Black Box of Causality: Learning about Causal Mechanisms from Experimental and Observational Studies." *American Political Science Review* 105: 765-89.
- Imai, Kosuke and Teppei Yamamoto. 2013. "Identification and Sensitivity Analysis for Multiple Causal Mechanisms: Revisiting Evidence from Framing Experiments." *Political Analysis* 21:141-71.
- King, Gary, Robert O. Keohane, and Sidney Verba. 1994. *Designing Social Inquiry: Scientific Inference in Qualitative Research*. Princeton, NJ: Princeton University Press.
- Kittel, Bernhard and Hannes Winner. 2005. "How Reliable is Pooled Analysis in Political Economy? The Globalization-Welfare State Nexus Revisited." *European Journal of Political Research* 44:269-93.
- Kiser, Edgar and Michael Hechter. 1991. "The Role of General Theory in Comparative-historical Sociology." *American Journal of Sociology* 97:1-30.
- Lijphart, Arendt. 1971. "Comparative Politics and Comparative Method." *American Political Science Review* 65:682-93.
- Malhotra, Neil and Jon A. Krosnick. 2007. "Retrospective and Prospective Performance Assessments during the 2004 Election Campaign: Tests of Mediation and News Media Priming." *Political Behavior* 29:249-78.
- Mayntz, Renate. 2004. "Mechanisms in the Analysis of Social Macro-phenomena." *Philosophy of the Social Sciences* 34:237-59.
- Munck, Gerardo L. 2004. "Tools for Qualitative Research." Pp. 105-22 in *Rethinking Social Inquiry*, edited by Henry Brady and David Collier. New York: Littlefield & Rowman.
- Norkus, Z. 2004. "Mechanism as Miracle Makers? The Rise and Inconsistencies of the 'Mechanismic' Approach in Social Science and History." *History and Theory* 44:348-72.
- Ragin, Charles. 2000. *Fuzzy Set Social Science*. Chicago: University of Chicago Press.
- Ross, Michael L. 2004. "How Do Natural Resources Influence Civil War? Evidence from Thirteen Cases." *International Organization* 58:35-67.
- Seawright, Jason and John Gerring. 2008. "Case Selection Techniques in Case Study Research: A Menu of Qualitative and Quantitative Options." *Political Research Quarterly* 61:294.
- Waldner, David. 2007. "Transforming Inferences into Explanations: Lessons from the Study of Mass Extinctions." Pp. 145-76 in *Theory and Evidence in*

Comparative Politics and International Relations, edited by R. N. Lebow and M. I. Lichbach. New York: Palgrave Macmillan.

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

Winship, Christopher and Michael Sobel. 2004. "Causal Inference in Sociological Studies." Pp. 481-503 in *Handbook of Data Analysis*, edited by Melissa A. Hardy and Alan Bryman. Thousand Oaks, CA: Sage.

Author Biographies

Nicholas Weller is an assistant professor in the Department of Political Science and School of International Relations at the University of Southern California.

Jeb Barnes is an associate professor in the Department of Political Science at the University of Southern California.