Part 1: Millian Methods of Comparison and Case Selection
Tommaso Pavone – tpavone@princeton.edu
November 19th, 2015

Introduction

Suppose you wish to explore the conditions under which a military coup will occur (or any discrete outcome/event). And suppose you are particularly interested in the causal story underlying this outcome. In this case, a comparative case study is a particularly useful way to identify the source of a particular outcome (i.e. of probing the "causes" of an observed "effect"). But how to start?

To conduct a comparative case study, the first step is to identify your dependent variable/outcome of interest, along with one/a few possible explanatory variables. Most of the time, you achieve this by reading the scholarly literature or exploring the empirical record. But the second step is to select which cases would be appropriate for comparison. "Case selection" is an incredibly important matter for comparative case study research, and if you select the wrong cases you might be unable to make any causal inferences.

To facilitate case selection, comparative researchers often rely on "Millian methods" – so named because they were developed by English philosopher John Stuart Mill - to help them identify appropriate cases for comparison. A brief overview of Millian methods is provided below.

Method 1: Mill's Method of Difference

Mill's *method of difference* (also known as the *most similar case design*) leverages variation in your dependent variable/outcome. It seeks to select two or more cases where, despite being very similar, the outcome nonetheless differs across the cases. The logic is that because the cases are so similar, we can isolate the independent variable that actually does vary across the cases as a possible cause. For example, consider the table below:

	DV	IV$_1$	IV$_2$	IV$_3$	IV$_4$	IV$_5$
Case A	1	1	0	1	0	1
Case B	0	1	0	1	0	0

In this research design, Case A and Case B are quite similar: they have similar values for a series of independent variables (IV1, IV2, IV3, and IV4). Nonetheless, the dependent variable/outcome (DV) varies across the cases. Therefore, we can reject IV1 and IV3 as *insufficient* for the outcome (notice that for Case B, IV1 and IV3 are both present, even though the outcome is absent. If IV1 and IV3 were sufficient, then whenever they are present the outcome must always occur). We can also reject IV2 and IV4 as *unnecessary* for the outcome (notice that IV2 and IV4 are absent for Case A, even

though the outcome is present. Consequently, IV2 and IV4 cannot be necessary for the outcome, because necessity implies that the outcome cannot occur unless IV2 and IV4 are present).

But look at IV5: It varies in a way that corresponds to the outcome of interest. This is purely correlative evidence (remember, correlation does not imply causation!), but it suggests that it would be worthwhile to process-trace the relationship between IV5 and the outcome of interest across both cases. In addition, we have managed to eliminate IV1, IV2, IV3, and IV4 as causally irrelevant, so we do not have to waste time process-tracing their potential relationship to the outcome. We are well on our way!

Note that in case study research, it is sometimes easier to think of our dependent/independent variables as outcomes/events that are either present or absent. Below you will find the same table as that displayed previously, but using the language of outcome/events as opposed to dependent variable/independent variables:

	outcome	event1	event2	event3	event4	event5
Case A	1	1	0	1	0	1
Case B	0	1	0	1	0	0

For example, suppose we are interested in the conditions under which military coups occur, and we identify a series of plausible explanations. Ideally, we could try to find two cases that are very similar, but where one experiences a military coup and the other one does not. For example, look at the table above and assign the following labels to the outcome/events:

Outcome = military coup   *Event3 = international sanctions placed*
Event1 = urban protest occurs   *Event4 = dictator is struck by health problems*
Event2 = peasant rebellion occurs   *Event5 = loss of support from military generals*

In this setup, we could conclude that urban protests/international sanctions are insufficient to explain when a military coup occurs; we could reject peasant rebellions/the dictator being struck by health problems as unnecessary for the occurrence of a military coup; and we would isolate the loss of support from military generals as a plausible explanatory variable. We would then try to trace the process linking the loss of military general support to a military coup in Case A and the continued support from military generals to the absence of a military coup in Case B.

Method 2: Mill's Method of Agreement

Mill's *method of agreement* (often referred to as a *most different case design*) leverages the absence of variation in our dependent variable. It picks two or more cases that are very different, but nonetheless register the same value of our dependent variable/experience the same outcome. The

logic is that this approach helps us isolate the independent variable that is similar across the otherwise very different cases as a possible cause. To visualize this, consider the following table:

	DV	IV1	IV2	IV3	IV4	IV5
Case A	1	1	1	0	1	1
Case B	1	0	0	1	0	1

In this research design, Case A and Case B both have the same value for our dependent variable/experience the same outcome, even though they are quite different: IV1, IV2, IV3, and IV4 all vary across the two cases. Consequently, we can reject IV1, IV2, IV3, and IV4 as *unnecessary* for our dependent variable/outcome of interest (notice that Case B experiences the outcome even though IV1, IV2, and IV3 register a value of 0 for that case; and notice that Case A experiences the outcome even though IV3 registers a value of 0 for that case. Consequently, neither of these variables can be necessary for the outcome).

Yet look at IV5: It is the one independent variable that does not vary across the two cases. This is correlative evidence that IV5 might be the cause of our dependent variable/outcome of interest. Once again, correlation does not imply causation, so we would still have to process-trace the link IV5 has to our DV. But at least we have managed to eliminate a series of other explanations as causally irrelevant, and we have identified an independent variable as a likely cause.

Once again, it is sometimes helpful to conceive of our dependent/independent variables as outcomes/events that are either present or absent across our cases. To return to our military coup example, assign the following labels to the below table:

Outcome = military coup *Event3 = international sanctions placed*
Event1 = urban protest occurs *Event4 = dictator is struck by health problems*
Event2 = peasant rebellion occurs *Event5 = loss of support from military generals*

	outcome	event1	event2	event3	event4	event5
Case A	1	1	1	0	1	1
Case B	1	0	0	1	0	1

In this setup, we could conclude that urban protests, peasant rebellions, international sanctions, and the dictator being struck by health problems are all unnecessary for a military coup. We would consequently move to process-tracing the possible link between the loss of support from military generals and the military coups that occurred in both cases.

Assumptions / Limitations of Millian Methods

Although Millian methods are extremely useful, they have important limitations – and make some important assumptions – of which you should be aware. These are:

1) *Millian methods assume a deterministic understanding of causation.* That is, they require causal theories that implicitly or explicitly make claims about the necessity or sufficiency of a particular explanatory variable for an outcome of interest (rather than probabilistic claims). It is only in this way that we can eliminate rival explanations as either unnecessary or insufficient for the outcome of interest. (This does not mean you have to be deterministic in your conclusions – there is always uncertainty! But it means that the types of causal claims you are trying to support/disprove are expressed deterministically).

2) *Millian methods can only eliminate independent variables as causally irrelevant – they cannot confirm the causal relevance of a particular independent variable.* In order to make causal, as opposed to correlational, claims, you would then have to trace the causal process that plausibly links your explanatory variable to your outcome of interest.

3) *Millian methods do not take account of the fact that your independent variables might be ordered in a temporal sequence.* That is, it is possible that your independent variables (in the tables above, IV1, IV2, IV3, IV4, and IV5) are not independent of one another, but are actually ordered across time. As a result, once you begin process-tracing you should be open to the possibility that an independent variable that you have provisionally excluded might nonetheless be important for a part of the causal process you are tracing.

4) *Millian methods require you to impose scope conditions on your conclusions.* That is, technically Millian methods only allow you to infer causality across the cases you are studying. In order to generalize your causal inferences to a broader set of cases, you would have to identify other cases that are similar to the ones you have analyzed, such that the causal story you are telling could travel to those cases. But you have to be humble and careful about making such claims.

5) *Millian methods assume that the causal process will be the same across both of your cases.* In other words, Millian methods assume the absence of "equifinality" –the presence of multiple causal paths to the same outcome. That is, we must assume that if an independent variable is necessary for the outcome in Case A, then it must also be necessary for the outcome in Case B. Sometimes you will find when process-tracing that the causal processes across your cases vary in very interesting and important respects. Mill's methods cannot help you explore such possibilities – but they are important to note when you write up your results!

To learn more, see:
Falleti, Tulia, and James Mahoney. 2015. "The Comparative Sequential Method." In *Advances in Comparative-Historical Analysis*, Mahoney & Thelen, eds. New York: Cambridge UP.

Part 2: Necessary/Sufficient Conditions and Causal Processes

Introduction

Process-tracing is the primary way to make causal inferences when you are conducting a single case study, and is usually a critical component of comparative case studies as well. Process-tracing opens up the "black box" of causation – it asks: "if X causes Y, what is the story, exactly, whereby X causes Y?" That is, it tries to unearth the sequence of events and causal mechanism(s) that link the presence or absence of X, your independent / explanatory variable, with the presence or absence of Y, your dependent / outcome variable.

Process-tracing is facilitated if we can model, or visualize, the causal process we are trying to describe through primary and secondary source evidence. Most of the time, the causal process is represented as a series of temporally ordered events, or conditions, that are either necessary or sufficient for the outcome. The events are linked together via causal mechanisms. In this brief memo, we review necessary / sufficient conditions and show one way to diagram causal processes.

Necessary / Sufficient Conditions

When we conceptualize causation as a series of events / conditions that are either necessary or sufficient for the outcome, we are embracing a deterministic understanding of causation. But what does it mean for something to be necessary or sufficient?

Let us begin with a necessary condition. A _necessary condition_ is a state of the world whose counterfactual absence would prevent the outcome from occurring. That is, if A is necessary for B, then A's absence will always produce B's absence. As a concrete example, suppose you are a prosecutor who is investigating an armed robbery of a bank, and are trying to figure out if the suspect that has been detained is guilty of the crime. You might begin by asking: Was the suspect in the immediate vicinity of the bank at the time the crime occurred? The logic for asking this question is that being proximate to the bank is necessary for the suspect to have actually committed the armed robbery. Having established this as a necessary condition for the suspect to be guilty, we would engage in process-tracing to uncover evidence of the subject's whereabouts at the time of the bank robbery – perhaps by consulting video surveillance footage or interviewing friends/family of the suspect.

We can visualize a necessary condition using _set theory_. In set theory, A is necessary for B if A is a superset of B (and if B is a subset of A). A visual diagram is provided below:
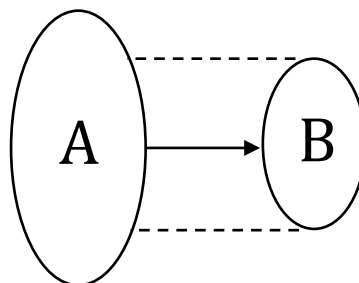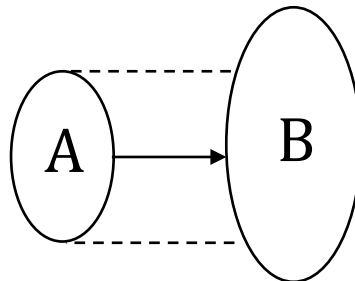
You should read the diagram from left to right, where the transition from left to right indicates the passage of time. As you move horizontally from left to right, notice that in order to end up in the set B, you must travel through set A first. The arrow represents a causal mechanism, which explains the logic that connects event/condition A with event/outcome B. To return to the bank robbery example, we could conceive B as being the set of "guilty" individuals, and A as being the set of "individuals proximate to the bank at the time of the robbery." The causal mechanism would then outline the logic that to rob a bank at gunpoint, you must be physically near/inside the bank to do so.

Importantly, notice that just because you pass through set A does not automatically mean that you end up in set B, since set A is larger than set B. To put it in probabilistic terms, if A occurs, the probability that B occurs is positive, but still less than 1. However, if B has occurred, then the probability that A occurred prior to B is 1. At least, this holds if A truly is necessary for B.

Let us transition to sufficient conditions. We say that A is *sufficient* for B if B always occurs when A occurs. That is, when event/condition A is present, the probability that outcome B occurs is 1. To use the bank robbery example, we might say that video surveillance evidence of the suspect robbing the bank is sufficient to establish his guilt. That is, if we have clear video evidence of him robbing the bank at gunpoint, that is sufficient evidence to conclude that he is guilty.

We can visualize this logic using set theory as well. In set theory, A is sufficient for B if A is a subset of B (and B is A's superset). To visualize this using the temporal logic of process tracing, consider the diagram below:


Again, read the diagram from left to right, where the transition from left to right indicates the passage of time. As we move horizontally from left to right, if we pass through set A we always end up in set B. That is, if event / condition A is present, then event / outcome B always follows. Notice, however, that it is possible to end up in set B without first passing through set A. To put it another way, even though A's presence always causes B, A is not necessary for B. In our bank robbery example, this would mean that although conclusive video evidence is enough to prove the suspect's guilt, just because we lack video evidence does not mean that we cannot establish the suspect's guilt in some other way.

In extremely rare occasions, an event / condition is both *necessary and sufficient* for an outcome. We say A is necessary and sufficient for B if A's counterfactual absence would prevent B from occurring *and* if A's occurrence always produces B. There are very few concrete examples in social science of necessary and sufficient conditions, but in our bank robbery example, we might consider

passing a DNA test to be a necessary and sufficient condition to establish a suspect's guilt beyond the shadow of a doubt. That is, if the DNA test is positive, then the suspect's guilt is clear; if the DNA test is negative, then the suspect is exonerated.

We can visualize necessary and sufficient conditions using the same set-theoretic logic as follows:
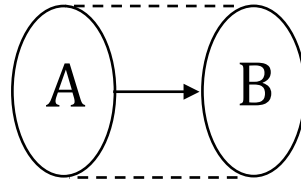
Here, as you move horizontally from left to right, the only way to end up in set B is to first pass through set A. In other words, if you are not in set A, then the outcome B never occurs (probability = 0), and if you are in set A, then the outcome B always occurs (probability = 1).

Causal Processes

We can use the logic we have developed in the previous section to illustrate causal processes as concatenations of necessary / sufficient conditions. For example, consider the below diagram:
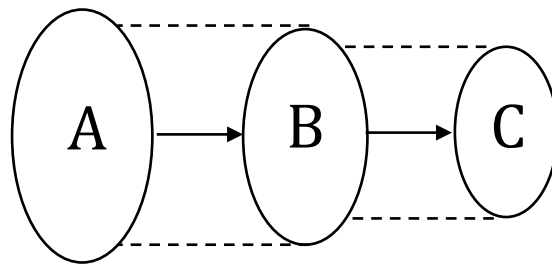

This causal process shows how event A is necessary for event B to subsequently occur, and event B is necessary for event C, or the outcome, to subsequently occur. Notice that since B is necessary for C, and in turn A is necessary for B, we can make the broader conclusion that A is necessary for C. These three events / conditions are linked together by arrows, which represent causal mechanisms.

For example, suppose A represents rapid economic growth, B represents the development of civil society, and C represents democratization. In this case, a possible causal mechanism linking A (rapid economic growth) with B (the development of civil society) is that civic associations require material resources (money) to get off the ground, which are supplied by economic growth; A possible causal mechanism linking B (the development of civil society) with C (democratization) is that civic associations can organize protests that eventually force an authoritarian regime's collapse.

Notice how the causal chain we just examined could itself be broken up into an even more fine-grained series of events and causal mechanisms. This is an inherent quality to all causal processes – events have a fractal character, meaning that they themselves are composites of smaller causal processes. In the end, how much you "zoom in," as it were, depends upon your research design: The broader the geographic and temporal scope of your project, the less fine-grained you will need to go; the more restricted your scope in time and space, the more detailed you will have to be.

When we take causal processes seriously, we cannot assume that our independent and dependent variables are always linked by the exact same causal process, even if they are very strongly correlated. That is, even if A always seems to produce B, it could be that A is producing B in different ways. This is known in qualitative research as *equifinality* – or the idea that there can be multiple causal paths to the same outcome.

For example, let us return to our case of how rapid economic growth causes democratization. Consider the following causal processes: The first is an exact replica of the one we analyzed before, but the one below it is a causal process we uncover from a comparative case:
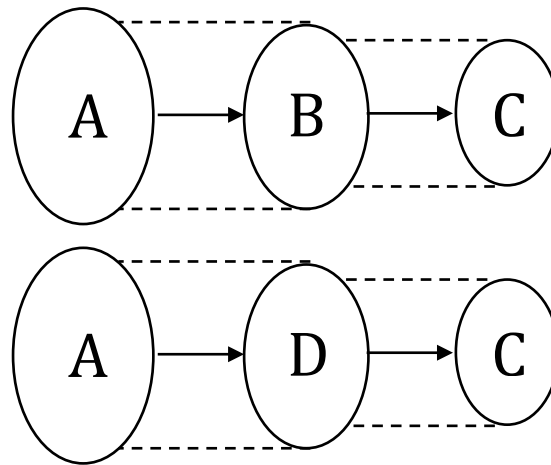
In this situation, whereas in our first case rapid economic growth (A) causes the development of civil society (B) which pressures the regime to democratize (C), in the second case rapid economic growth (A) causes democratization (C) through some other event, D. D could, for example, be the fact that as economic growth occurs and the country's residents get wealthier, they are better able to lobby foreign governments to pressure their home government to democratize. This is a clear case of equifinality: The independent variable that begins the causal process, A, is the same across both cases, and it eventually leads to the same outcome, C. However, the two causal processes are clearly distinct in interesting and important ways.

Finally, taking causal processes seriously also means that we have to take the order, or *sequencing*, of events seriously. That is, up until now we have just assumed that it is the presence or absence of a series of independent variables / events that explains whether our outcome of interest occurs or does not occur. But it could be that it is not the presence / absence, but the ordering, of our independent variables / events that matters for explaining the outcome.

For example, consider the causal processes diagrammed on the next page. In the first causal process, A being followed by B leads to outcome C. However, in the second causal process where the temporal ordering of A and B is flipped, C fails to occur.
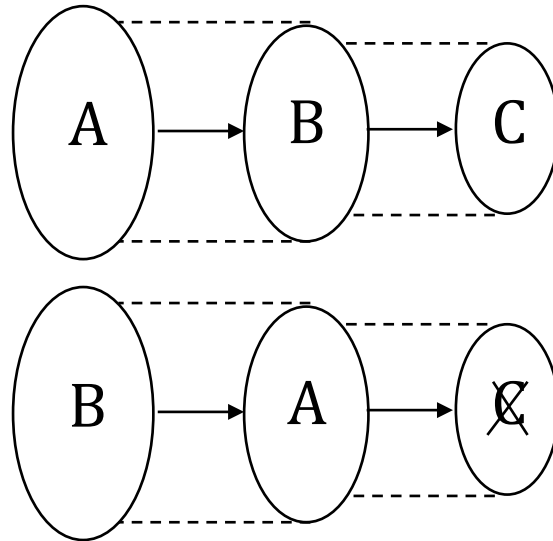
What could be a concrete example of the above scenario? Consider Tulia Falleti's argument in her book, *Decentralization and Subnational Politics in Latin America* (2010). To simplify her argument slightly, she is interested in when decentralization – the devolution of national powers to subnational administrative bodies – actually increases local political autonomy. She shows that when fiscal decentralization (A) precedes electoral decentralization (B), local autonomy is increased (C). This is because fiscal decentralization provides local districts with the monetary resources necessary to subsequently run an election effectively (elections are expensive affairs!). However, when the order is reversed, such that electoral decentralization (B) precedes fiscal decentralization (A), the local district actually becomes less autonomous. Why? Because although the district is being offered the opportunity to hold local elections, it lacks the monetary resources to administer them effectively. It consequently has to turn to the national government for money, and the national government now has added leverage over the district by controlling the power of the purse and placing conditions on how fiscal decentralization will occur.

Extensions / Limitations

Conceiving causal processes as a series of temporally ordered events that are either necessary or sufficient for the outcome of interest can be very helpful for causal inference within a single / a small number of cases. However, there are some limitations to this approach, which sometimes can be addressed via extensions utilizing alternative methods.

The most important limitation is that demonstrating necessity / sufficiency is oftentimes really difficult, particularly if the conditions we are arguing are necessary / sufficient are not trivial (that is, the presence of oxygen in the air is a necessary condition for pretty much every single causal process social scientists might be interested in, but because it is so common it is not a very good 'explanatory variable'). This is why we usually limit the number of cases for which we conduct process tracing: To show, say, that the development of civil society was necessary for an authoritarian regime's subsequent collapse, we need to dig up a lot of primary and secondary sources, which is costly. Consequently, we will usually want to qualify any claims about necessity / sufficiency with <u>scope conditions</u> – a logic about where the causal process – or a particular

component of the causal process – is unlikely to travel, and where, within those conditions, it is more likely to generalize.

For example, if the causal process you uncovered through a careful case study is *economic growth* → *development of civil society* → *democratization*, then it is unlikely that this causal process will generalize to countries that have not experienced sustained economic growth. We might also say that the causal process is unlikely to generalize to countries where civic organizations have a history of unsuccessfully mobilizing against an authoritarian regime. These are two scope conditions on the generalizability of our findings: We suggest to the reader what the limitations of our findings are likely to be across space and time.

One way to probe the generalizability of our findings is to *turn to quantitative methods* where this is possible. That is, if your independent and dependent variables are quantifiable, you might run a regression on a much larger number of cases to see if a positive correlation emerges. If it does, particularly when we control for other factors that might also be affecting the dependent variable, then we have some suggestive evidence that the causal process we have uncovered might be at work across a wider variety of cases. However, always remember that correlation does not imply causation (even if we include a series of control variables in our regression, we might still be suffering from omitted variable bias or other endogeneity concerns), and that equifinality might be at work, such that although our independent variable tends to cause our dependent variable across a wide range of cases, the specific processes through which it does so may vary from case to case.

A final way to probe the generalizability of our findings is to *assemble a constellation of cases* that are similar to the case(s) we have analyzed in detail. To return to our example where *economic growth* → *development of civil society* → *democratization*, we might identify countries where an authoritarian regime was present, rapid economic growth occurred, and some time thereafter the regime fell. Such cases display a clear family resemblance to the one(s) we have carefully process-traced, and we could plausibly suggest that these are particularly likely contexts where our results will generalize.

To read more:
Collier, David. 2011. "Understanding Process Tracing." *PS: Political Science & Politics* 44 (4): 823-830.

Falleti, Tulia. 2010. *Decentralization and Subnational Politics in Latin America*. New York, NY: Cambridge University Press.

Mahoney, James, and Rachel Sweet Vanderpoel. 2015. "Set Diagrams and Qualitative Research." *Comparative Political Studies* 48 (1): 65-100.