Process Tracing in Case Study Research *

Andrew Bennett (Georgetown University) and
Alexander L. George (Stanford University)

MacArthur Foundation Workshop on Case Study Methods
October 17-19, 1997

Introduction

Covariations have important limitations as sources of causal inference, whether these covariations are established by statistical methods or by congruence methods in case studies. Because of these limitations, philosophers of science and social science methdologists have given increasing emphasis in the last two decades to "causal mechanisms," or the causal processes and intervening variables through which causal or explanatory variables produce causal effects, as a second basis for causal inferences, and to process tracing as a method for identifying and testing causal mechanisms. This paper first looks at some of the philosophy science issues related to process tracing, including the roles of causal effects and causal mechanisms in causal explanation and the logic of process tracing. It then looks at how process tracing relates to common misunderstandings of case study methods, such as those on the "degrees of freedom problem" and the limitations of Mills' Methods. The paper then discusses the inductive and theory-testing uses of process tracing, and it concludes with observations on the limits as well as the strengths of process tracing.

Causal Effects and Causal Mechanisms as Bases for Causal Inferences

Tests of covariations between observed outcome variables and their hypothesized causal variables, whether they take the form of case study congruence tests or statistical correlations among many cases, involve efforts to estimate the causal effects of variables. The causal effect of an explanatory variable is defined here as the change in the probability and/or value of the dependent variable that would have occurred if the explanatory variable had assumed a different value. Because this is a counterfactual definition -- regarding what would have happened if one variable had been different and all others had been held the same -- it raises what has been termed the "fundamental problem of causal inference." (KKV p. 79, following Holland, 1986) This problem derives from that fact that we cannot re-run history and change only one variable in a perfect experiment that would allow us to observe the actual causal effect of that variable. Tests of covariation are an avowedly imperfect alternative to perfect experiments. These tests try to measure causal effects by controlling for the effects of variables other than the variable of interest. In correlational studies, this control is attempted through statistical methods that use large numbers of observations to estimate the partial correlations between each independent variable and the dependent variable. In case studies, the method of congruence testing attempts to control for all but one independent variable at a time by using theories to create expectations about the dependent variable that can be compared to its actual value. (George, 1997) These methods of estimating covariations are limited in ways both similar and different.

Congruence testing and statistical correlations are useful components of broader means of

Congruence testing and statistical correlations are useful components of broader means of making causal inferences, but because covariations have important and well-known limitations as sources of causal inference, philosophers of science and social science methodologists have given increasing emphasis in the last two decades to causal mechanisms as a second basis for causal inferences. (Sayer, 1992: 104-105; Dessler, 1991: 343; Miller, 1987: 139; Yee, 1996; Salmon, 1989; Marini and Singer, 1988) Causal mechanisms are defined here as the causal processes and intervening variables through which causal or explanatory variables produce causal effects. While the notion of causal mechanisms is often explicated in terms of physical causal processes, it is applicable as well to social processes, including intentions, expectations, information, small group and bureaucratic decisionmaking dynamics, coalition dynamics, strategic interaction, and so on. (Little, 1995)

It is interesting to note that of the three sources of causal inferences recognized by the philosopher David Hume, covariance, which Hume termed "constant conjunction," was only one. The other two, temporal succession and contiguity, relate to what modern philosophers of science consider to be causal mechanisms. One school of thought in particular, self-designated among philosophers of science and methodologists as the "scientific realist" school, has defined itself in part by the view that covariation has been over-emphasized relative to causal mechanisms as a source of causal inferences. David Dessler, for example, has argued that (1991:345):

> the "why?" question is construed as meaning, "By the workings of what structures is the phenomenon produced?" Here events are distinct from the structures that produce them, so causation cannot be dismissed or reduced to behavioral regularity. . . Recall that at the correlational level we cannot differentiate between causal and accidental sequences, nor can we conceptually distinguish the correlates of an event that are its causes (say, a cold front and a thunderstorm) from those that are not (say, a falling barometer and a thunderstorm).

Similarly, the philosopher of science Andrew Sayer has stated that "what we would like . . . is a knowledge of *how* the process works. Merely knowing that 'C' has generally been followed be 'E' is not enough; we want to understand the continuous process by which 'C' produced 'E,' if it did." (Sayer, 1992: 106-107) Sayer adds that whether the causal powers inherent in a variable are activiated in a particular case depends on the presence of contingent conditions. Because such contingent conditions and counteracting forces can ennable, conceal, or override the causal powers of a variable, he argues, "the discovery of what a given mechanism can and cannot do requires considerable effort and ingenuity and . . . the search for regularities is inadequate." (Sayer, 1992: 110) In short, an adequate scientific explanation must include both arguments and measures on the causal effect of an independent variable and the hypothesized and observed causal mechanisms through which it achieves this effect.

In this view, new and more sophisticated statistical methods, and the philosophy of science notions of probabilistic causality that underly them, are not by themselves sufficient for defining causal explanations, nor can they rectify the inherent limits of covariation as the observational basis for causal inference. Wesley Salmon, a philosopher of science who has played an important role in the development of notions of probabilistic causality, came to this conclusion after surveying three prominent theories of probabilistic causality in the mid-1980s. He later noted that "the primary moral I drew was that causal concepts cannot be fully explicated in terms of statistical relationships; in addition, I concluded, we need to appeal to causal processes and causal interactions." (Salmon, 1989: 168, citing Suppes, 1970, Reichenbach 1956, and Good, 1961-2) New statistical methods have allowed more sophisticated modelling of non-linear processes and interaction effects, and some computer methods even suggest possible model specifications. However, none of these methods, which generally allow for tighter correlations between data and theoretical models, shows promise of addressing the version of the fundamental problem of inference that applies to statistical and other methods of assessing covariations: observed correlations do not provide a solid basis for inferring underlying causality.

For scientific realists, the disjuncture between correlation and causality is related to but not synonomous with that between "explanation" and "prediction." Because the ability to predict does not necessarily provide a satisfactory explanation to the causal "Why?" question, it is possible to have non-predictive explanations (as in the theory of evolution) and non-explanatory predictions (such as the use of barometer readings to predict the weather). (Sayer, 1992: 131-132; Salmon, 1989: 129-130) Causal mechanisms are also relevant to policy interventions because the ability to predict outcomes may not confer the ability to affect them through manipulable variables. For example, even if it continues to hold up to empirical scrutiny, the observed correlation that democracies do not engage in wars with one another even though they have frequently fought wars against non-democratic states is only a starting point for policy prescriptions. It would be best to understand both the causal mechanisms behind the apparent democratic peace and those behind the emergence of democracy, and to identify those mechanisms that policy interventions can affect. Indeed, some studies suggest that transitional democracies are quite war-prone and may fight other democracies (Snyder and Mansfield, 1995) Similarly, microbiologists devote great attention to understanding the causal mechanisms behind cancer, AIDs, and other diseases so that they can establish more options for appropriate interventions at various points in the causal process. In these instances, it should be noted, knowledge of causal mechanisms can be of practical use even when the entire causal process or path is not fully understood.

None of the above should be taken as suggesting in any way that causal effects are not relevant to definitions of causality, or that attempts to observe covariations are not relevant to making causal inferences. Causal effects and causal mechanisms are both essential to the definition of causality and attempts to make causal inferences. A subtle but important debate has emerged on this issue that obscures this central point. On one side, some have attempted to address causality and explanation primarily in terms of causal effects, while downplaying the status of causal mechanisms. For example, KKV argue that their definition of causal effect, similar to that adopted here,

> is logically prior to the identification of causal mechanisms . . .we can define a causal effect without understanding all of the causal mechanisms involved, but we cannot identify causal mechanisms without defining the concept of causal effect .. . Identifying causal mechanisms can sometimes give us more leverage over a theory by making observations at a different level of analysis into implications of the theory. The concept can also create new causal hypotheses to investigate. However, we should not confuse a definition of causality with the nondefinitional, albeit often useful, operational procedure of identifying causal mechanisms. (1994: 86)

While this view recognizes that identification of causal mechanisms can play a role in both

While this view recognizes that identification of causal mechanisms can play a role in both testing theories and developing hypotheses, it risks conflating the definition of "causality" with the definition of "causal effect" and relegating the identification of causal mechanisms to an inferior, "operational" status. The observation of putative causal mechanisms, through process tracing, is indeed an operational procedure, just as methods of examining covariations are operational procedures for assessing the observable analogues of causal effect. The final sentence in the passage quoted above could be easily transposed to read, with equal authority, "we should not confuse a definition of causality with the nondefinitional, albeit often useful, operational procedure of estimating observed covariations." Methodological procedures and ontological definitions should indeed not be conflated. Both causal mechanisms and causal effects are theoretical entities that are central to notions of causality and explanation. A variable cannot have a causal effect on an outcome unless there is an underlying causal mechanism, and it makes no sense to define any entity as a causal mechanism if it has no causal effect.

An opposite and equally misleading practice is to accord causal mechanisms a status that is somehow superior to that of causal affects. Albert Yee risks adopting this approach when, in citing and responding to the passage from KKV quoted above, he states that "causal mechanisms and capacities are ontologically prior to, and account for, both statistical associations and controlled experimental results." (Yee, 1996:84) This sentence, while defensible, does not improve upon the passage that it seeks to refute. The debate over whether causal effects are "definitionally" prior to causal mechanisms, as KKV argue, or causal mechanisms are "ontologically" prior to causal effects, as Yee maintains, misses the more fundamental point that both are essential to causality, explantion, and the building and testing of explanatory social theories.

This would be a harmless debate of the chicken and egg variety were it not for the fact that it entails, on each side, the inclination to priviledge the particular methodology that is best suited to addressing the element of causality that each side favors. Large N statistical studies are strongest at documenting partial correlations that attempt to measure causal effects as one basis for causal inference. Yee argues that the analysis of symbolic discourse is useful for identifying causal mechanisms related to the causal role of ideas. Similarly, in case study methods, the identification of causal mechanisms through process tracing is a stronger methodological basis for causal inferences than the estimation of covariation through congruence tests. Different methods do indeed present different strengths and weaknesses in attempting to measure either causal effects or causal mechanisms. But this does not suggest that definitions of causality should be driven by researchers' choice of methods. Rather, causality involves both causal effects and causal mechanisms and its study requires a diversity of methods, of which some are better adapted to the former and some to the latter. The search for causal explanations in any given research project can focus on either causal effects or causal mechanisms alone or on both together, but research programs as a whole, if they aspire to be progressive, should include efforts to link the two.

A related problem has arisen in the context of deductive theories, particularly rational choice theories. Some rational choice theorists have argued that their theories can be treated as causal explanations for particular cases even if the causal mechanisms that these theories posit were not observably present in these cases. Christopher Achen and Duncan Snidal, for example, have argued that "rational deterrence [theory] is agnostic about the actual calculations that decision makers undertake. It holds that they will act as if they solved certain mathematical problems, whether or not they actually solve them." (Achen and Snidal, 1989: 164) Similarly, John Ferejohn and Debra Satz have argued that "a successful social-science explanation presupposes that the behavior in question be describable as being intentionally brought about by human agents seeking goals and holding beliefs," but they add in a footnote that "this constraint does not require that the behavior in question was *actually* brought about intentionally, and certainly not that it was done consciously; it requires only that it *could* be described as having been done in that way." (Ferejohn and Satz, 1995:74, and footnote 3, p.83, their emphasis). Such formulations fail to recognize the distinction between prediction and explanation: even if rational choice theories do achieve impressive predictive performance, this does not inherently mean that they are useful causal explanations. Well-defined causal mechanisms that are consistent with the available process-tracing evidence are essential for explanations of individual behavior, just as theories of weather formation are necessary to explain why the barometer is a good predictor of weather. Rational choice theories, like other deductive theories, cannot be exempted from empirical tests against observed causal mechansims, particularly since alternative psychological theories of individual behavior make profoundly different assumptions and do hold themselves up to the empirical standard of consistency with observed decision-making processes and outcomes.

Process Tracing as a Mode of Inferences on Causal Mechanisms

Unfortunately, scientific realists have not said much on the topic of methods for generating and assessing evidence on causal mechanisms. For this, we have to turn to what is termed here as "process tracing" (following George 1979, George Causal nexus, George and McKeown), though variants of process tracing have been called "pattern-matching" (Campbell), or the "modus operandi" method (Scriven). The general method of process tracing is to generate and analyze data on the causal mechanisms, or processes, events, actions, expectations, and other intervening variables, that link putative causes to observed effects. In other words, of the two kinds of evidence on the theoretical causal notions of causal effect and causal mechanisms, tests of covariation attempt to address the former, and process tracing assesses the latter.

Within the general method of process tracing there are two very different approaches. The first, which we term "process verification," involves testing whether the observed processes among variables in a case match those predicted by previously designated theories. The second, which we term "process induction," involves the inductive observation of apparent causal mechanisms and heuristic rendering of these mechanisms as potential hypotheses for future testing. It is important that the distinction between process verification and process induction be maintained, and we have detailed each method below. Before turning to this, however, a metaphor may help illustrate the concepts of process tracing that these two approaches share.

The metaphor of dominoes helps illustrate the idea of process tracing. Suppose that Dr. Moriarity, a metaphysician, sets up fifty numbered dominoes standing in a straight line with their dots facing the same way on a table in a room, but puts a blind in front of the dominoes so that only numbers one and fifty are visible. You enter the room, and observe that dominoe number one and dominoe number fifty are lying flat with their tops pointing in the same direction.; that is, they covary. Does this mean that either domino caused the other to fall? Not necessarily -- covariation is not enough. Dr. Moriarty could have pushed over only dominoes number one and fifty, or bumped the table in a way that only these two dominoes fell, or that all the dominoes fell at once. It is essential to remove the blind and dominoes fell, or that all the dominoes fell at once. It is essential to remove the blind and look at the intervening dominoes, as they give evidence on potential processes. Are they, too lying flat? Do their positions suggest they fell in sequence rather than being bumped or shaken? Did any reliable observers hear the tell-tale sound of dominoes slapping one another in sequence? From the positions of all the dominoes, can we eliminate rival causal mechanisms, such as earthquakes and wind, as well as human intervention? If the dominoes fell in sequence, can we tell by looking at whether the dots are face up whether the direction of the sequence was from number one to number fifty or the reverse? Tracing potential processes through such inquiries helps narrow the list of potential causes. Even if the experiment is repeated, it will be difficult to resolve fully the problem of equifinality, or of similar outcomes occurring through different causal processes, and to eliminate all potential rival explanations but one, especially when human agents are involved who may be doing their best to conceal causal processes. But process tracing forces the investigator to take equifinality into account, and it offers the possibility of mapping out one or more potential causal paths that are consistent with the outcome and the process tracing evidence in a single case. With more cases, the investigator can begin to chart out the repertoire of causal paths that lead to a given outcome and the conditions under which they obtain, which relates to the topic of "typological theory" (Bennett and George, 1997b).

The domino metaphor thus illustrates not only the method of process tracing but the problem of equifinality and the potential complexity of differentiating alternative causal paths. If we extend it to many lines of dominoes, with many potential starting points, intersections, and diverting branches, and possibly several lines intersecting with the final domino in question, the potential complexity of causal mechanisms is apparent. A single domino may set off many branches, one or more of which may reach the final domino, or there may be many different dominoes that begin chains toward the final domino, perhaps with no one of these chains being necessary but each of them being sufficient to knock it over. Any proposed causal mechanism must be able to demonstrate an uninterrupted cause-effect path from the independent variable to the dependent variable.

Process-Tracing and Historical Explanation

The question is sometimes asked whether process-tracing is similar to historical explanation and, indeed, whether process-tracing is anything more than "good historical explanation." It is not unreasonable to respond to such an observation by asking what is a good historical explanation! A process-tracing explanation differs from a historical narrative, as it requires converting a purely historical account that implies or asserts a causal sequence into an analytical explanation couched in theoretical variables that have been identified in the research design. This raises the familiar objections from some historians that to convert a rich historical explanation into an analytical one may result in losing important characteristics, or the "uniqueness" of the case. Information loss does indeed occur when this is done, and the investigator should be aware of this and consider the implications for his/her study of the fact that some of the richness and uniqueness of the case is thereby lost. But, at bottom, we justify the practice of converting historical explanations into analytical theoretical ones by emphasizing that the task of the political scientist who engages in historical case studies for theory development is not the same as the task of the historian.

Nonetheless, the nature and logic of historical explanation is of direct importance for use of the process-tracing method. The requirements, standards, and indeed the "logic" of historical explanation are matters that have long been discussed and debated by philosophers of history. This discussion is studded with important disagreements and controversies of which we should be aware, since they are pertinent to process-tracing, even though we cannot and need not resolve them. We have in preparation a detailed discussion of issues associated with the logic and method of historical explanation insofar as it helps us to articulate better the requirements of process-tracing. Here a brief discussion will have to suffice.

We have found particularly useful for gaining a better appreciation of the nature of historical explanation and controversies surrounding it the book by Clayton Roberts, The Logic of Historical Explanation (University Park, Pennsylvania: Pennsylvania State University Press, 1996). Roberts offers a detailed statement of his own position that is, on the whole, remarkably consistent with our concept of process-tracing in "within-case" causal inference.

Roberts rejects as do we the view advanced in the past by some commentators that historical explanation is no more than - and requires no more than - a description of a sequence of events. In principle, he holds, each step or link a causal process should be supported by an appropriate "law," defined for historical explanation by Carl Hempel as a statement of a regularity between a set of events. A distinction is made, however, between universalistic and probabilistic laws. While the Hempelian "covering law" model is deductive in form, it is clear that no explanation using probabilistic laws can be strictly deductive. Moreover, the covering law model cannot explain, Ernest Nagel observed, "collective events that are appreciably complex." Given this problem, Roberts, observes, "historians rarely seek to explain the occurrence of a complex event by subsuming it solely under a covering law," a process that he chooses to call "macrocorrelation." Attempts to rely on macrocorrelation to explain complex events have failed: "The vast majority of historians do not use macrocorrelation to explain the occurrence of events they are studying, and those who do have met with little success*"

How, then, Roberts asks, do historians explain the occurrence of complex historical events if not by subsuming them under covering laws? The answer Roberts provides is that they do so "by tracing the sequence of events that brought them about." The similarity to what we call "process-tracing" is clear. Roberts notes that a number of earlier writers have made the same point, referring to process-tracing variously as "a genetic explanation" (Ernest Nagel), "a sequential explanation" (Louis Mink), "the model of the continuous series" (William Dray), "a chain of causal explanations" (Michael Scriven), "narrative explanations" (R. F. Atkinson), and "the structure of a narrative explanation" (Arthur Danto). Roberts chooses to call this explanatory process "colligation," drawing on earlier usages of this term and clarifying its meaning.

Roberts chooses to call this explanatory method "colligation," drawing on earlier usages of this term and clarifying its meaning. Roberts' contribution is to explicate better than earlier writers the logic of such historical explanations. Laws that embody but are no more than "regularities" and "correlations," he argues, are not adequate explanations. A mere statement of a correlation, such as that between smoking and cancer, may have some explanatory power but it is incomplete and unsatisfactory unless the causal relation or connection between the two terms is specified. He notes that historians and philosophers have given many names to such causal connections. (Later, Roberts refers approvingly to the recent philosophy of science of "scientific realism" and the emphasis it places on the need to identify "causal mechanisms.")

Given that a correlation is not a substitute for investigating causation, how then can one determine whether some correlations are causal and others not? Roberts asserts (as others, including ourselves do) that it is only through colligation (process-tracing) that this can be done. He notes that historians, as with geologists, often rely on "process" explanations to answer the question "what has happened [to bring this about]?"

Roberts regards efforts to explain complex events solely by invoking a covering law insupportable for two reasons: formulation of general covering laws for this purpose is generally not possible, and reliance solely on them foregoes the necessary process-tracing of the sequence in the causal chain. Each step in such a causal sequence, Roberts holds, should be supported with an appropriate, though necessarily circumscribed covering law. He labels this the practice of "microcorrelation" to distinguish it from efforts at "macrocorrelation" for supporting efforts to explain complex events. As Roberts puts it, microcorrelation "is the minute tracing of the explanatory narrative to the point where the events to be explained are microscopic and the covering laws correspondingly more certain."

We offer an example from our own work that illustrates the difference between "macrocorrelation" and "microcorrelation" and depicts reliance on microcorrelation for explaining a complex phenomenon. In *States and Social Revolutions* Theda Skocpol was interested to provide a causal explanation for three social revolutions (the French, Russian and Chinese revolutions). She identified and worked with two independent variables: international pressures on the state and peasant rebellion. To show how these two variables were causally related to the revolutionary social transformation in each of these countries, Skocpol did not rely on macrocorrelation but instead employed a complex form of microcorrelation. Thus, she used the process-tracing procedure to identify a complex sequence of events to depict how each of the two independent variables set into motion a complex causal chain. She also showed how the two causal sequences came together to trigger a revolutionary social transformation in each country. The first of the following three charts (which we have extracted from her analysis) depicts how the explanatory problem would have been formulated had she attempted a macro-type explanation. The second and third charts depict the tracing of the process that connected each of the two independent variables through a sequence of events with the dependent variable outcome. The causal logic she employed for tracing each step (link) in the causal chain was supported by combining Mill's methods with micro-process-tracing. Be it noted, therefore, Skocpol did not attempt to support the causal relationship between the two independent variables and the outcome of the dependent variable by means of macro-type covering laws but identified a sequence of several steps or links between each independent variable and the outcome, supporting each by a form of micro process-tracing.

Roberts recognizes that some explanations, particularly those supported by available probabilistic laws, will be weak, and he discusses various strategies historians employ to develop stronger explanations. One of them is of particular interest for the present study of process-tracing. "Redescription," as he labels it, describes the event to be explained in a less concrete, more abstract manner. Doing so may enable the investigator to use a credible covering law. This brings to mind a familiar practice in political science research of "moving up" the ladder of generality in formulating concepts. A similar practice is frequently employed in statistical studies, "cell reduction" being a way of obtaining enough cases in a broader cell to permit statistical analysis. The new, larger cell necessarily requires a less concrete, more abstract concept than those concepts attached to the smaller ones absorbed into the larger one.

Roberts is particularly supportive of another strategy for strengthening weak explanations. "Microcorrelation," to which he referred earlier as noted above, strengthens an explanation via "the minute tracing of the explanatory narrative to the point where the events to be explained are microscopic and the covering laws correspondingly more certain." At the same time, Roberts recognizes that "the more microscopic the event to be explained, the more likely that the covering law will be a platitude* or a truism."

As is implicit in the preceding account of Roberts views, he rejects the widespread belief that historians do not make use of covering laws. He attributes this misconception to the fact that most of the laws historians make use of are not only "parochial" but also are not generally visible in their historical narratives. Such laws are not visible because they are generally implicit in the explanatory accounts historians provide. Roberts defends this practice on the ground that many of the covering laws are "platitudinous" and, therefore it would be tedious continually to list them and to assert their validity. Besides, these covering laws are so numerous in historical narratives that to list and justify them "would hopelessly clog the narrative." Roberts recognizes that historians have an obligation to make sure that the implicit covering laws they employ are true. But he does not address the question of how this can be or is done, contenting himself with the observation that "reviewers and perceptive readers" can readily tell the difference between histories based on sound covering laws and those that are naïve and superficial." He adds that historians will occasionally make their supportive generalizations explicit, particularly when a controversy arises among historians over the truth of an explanation.

In theory-based process-tracing, it is less likely and less desirable than in the historians' practice Roberts describes to rest explanations on implicit laws. Besides, the method of structured, focused comparison and process-tracing are employed not only in studies that attempt to provide explanations for specific cases but also to test and refine available theories and hypotheses, to develop new theories and to produce generic knowledge of a given phenomenon. Roberts discussion of the use of covering laws in historical explanation does not cover the task of causal inference in these other types of investigations.

In chapter 6, "The Logic of Colligation," Roberts distinguishes eight different forms that process-tracing may take. Several of these are of interest for the present study. The simplest form of process-tracing, "linear colligation," depicts "a straightforward chain of events" which is often a naïve simplification of a complex phenomenon. "Convergent colligation" depicts the outcome to be explained as flowing from the convergence of several conditions, independent variables, or causal chains. Skocpol's study, discussed above, is an example of how two processes set into motion, one by international pressures causing state breakdown and the other by peasant rebellions, converged to cause revolutionary social movements.

Another type of process-tracing, "repetitive colligation," provides the basis for Roberts' consideration of the relation of history to theory and science. Whereas history often limits itself to searching for the cause of a single event, "the purpose of science is to discover the laws governing the behavior of a phenomenon*," although laws of a correlational nature are used in the covering-law model of explanation. "To explain why a law exists, why a correlation occurs, one needs a theory," one which contains "a model that shows how the system works, the system that gives rise to the uniformities observed." It appears, here, that Roberts, is alluding to what others have referred to as "causal mechanisms." As we shall presently note, he later refers to the importance of causal mechanisms.

Roberts notes that the corpus of historical writing contains few theories, the reason being that historians have been unable to find any general laws that stood the test of time. The failure generally of the social sciences (with economics a partial exception) to find meaningful laws, Roberts observed, has led Jon Elster to conclude that "the basic concept in the social sciences should be that of a mechanism rather than of a theory." Roberts takes Elster's observations as consistent with his own concept of historical explanation as being "a marriage of colligation [process-tracing] and correlation.

Process Tracing and the "Degrees of Freedom Problem"

One of the most widespread misinterpretation of case studies is that they inherently suffer from what statisticians call a "degrees of freedom problem." (Achen and Snidal, 1989:156-7) In statistical terms, the number of degrees of freedom is defined as "the number of quantities that are unknown minus the number of independent equations linking these unknowns." (Blalock, 1979:205) The degrees of freedom problem arises because in order to obtain a unique solution for simultaneous equations, it is necessary to have the same number of unknowns (or cases) as equations. Thus, when a researcher has many independent variables but only one or a few observations on the dependent variable, the research design is indeterminate, and there is a very limited basis for causal inferences apart from simple tests of necessity or sufficiency.

Stated in these generic terms, there is indeed a degrees of freedom problem in any case study research design in which there are few observations but many independent variables. An important misinterpretation arises on this issue, however, from using definitions of "case," "variable," and "observation" that are excessively narrow. One common but potentially misleading definition describes a case a "phenomenon for which we report and interpret only a single measure on any pertinent variable." (Eckstein, 1975:85, KKV, 1994:52) This definition would lead naturally to the conclusion that any case study research design with fewer cases than independent variables would have a degrees of freedom problem. In fact, however, each qualitative variable has many different dimensions rather than providing a "single observation." Statistical researchers tend to aggregate variables together into single indices to get fewer independent variables and more degrees of freedom, but case study researchers do the reverse: they treat variables qualitatively, in all of their relevant dimensions, and they try to distinguish qualitatively different types of each independent and dependent variable. For example, rather than constructing a single index of "Soviet military interventionism," a case study researcher might look at the number of Soviet troops deployed, the kinds of weapons used, the rules of engagement, the amount of military aid, and so on. (Bennett, 1992) An independent variable affecting the level of Soviet interventionism, rather than the "Soviet economic situation" or "Soviet GNP," might include production in specific sectors, such as production of Soviet forces for power projection, production of weapons available for export, and so on. A case study researcher seeking to explain Soviet retrenchment in Afghanistan in the late 1980s, to continue the example, could test whether the many dimensions of independent variable(s) are congruent with those of the dependent variables. In this instance, Soviet military aid to Afghanistan increased even as Soviet troops withdrew, and the Soviet economy and particularly Soviet military production did not decline sharply until after the Soviet withdrawal, casting doubt on whether economic constraints were a major factor in causing this withdrawal.

on whether economic constraints were a major factor in causing this withdrawal.

In addition, within a single case there are many possible process tracing observations along the hypothesized causal paths between independent and dependent variables. A causal path may include many necessary steps, and they may have to occur in a particular order. At each step, the researcher may measure the magnitudes and signs of intervening variables to see if they are as the hypothesis predicts. These many predicted observations may provide sufficient "degrees of freedom," or many more observations than variables, even when the researcher is studying a single case and using several independent variables. KKV, for example, note that "defining and then searching for these different causal mechanisms may lead us to find a plethora of new observable implications for a theory." (KKV, 1994: 225)

This is why the accomplished methodologist Donald Campbell reversed his criticism that an inherent degrees of freedom problem plagued case study methods. Admirably setting out to "correct some of my own prior excesses in describing the case study approach," Campbell noted that:

> I have overlooked a major source of discipline (i.e., degrees of freedom if I persist in using this statistical concept for the analogous problem in nonstatistical settings). In a case study done by an alert social scientist who has thorough local acquaintance, the theory he uses to explain the focal difference also generates predictions or expectations on dozens of other aspects of the culture, and he does not retain the theory unless most of these are also confirmed. In some sense, he has tested the theory with degrees of freedom coming from the multiple implications of any one theory. (Campbell, 1975: 179, 181-2)

As Campbell describes this process, which he terms "pattern matching," it involves elements of two modes of within-case analysis, the "congruence method" and "process tracing." The congruence method (addressed in detail in a related conference paper by Alexander L. George) involves testing for whether the outcome of a case, in its various dimensions, is congruent with the various dimensions of the independent variable(s) and the expectations of the underlying theory linking the two. Process tracing involves testing whether all of the intervening variables were consistent with the expectations of the causal theory under consideration and the causal mechanisms that it posits. Thus, as long as sufficient evidence is accessible for congruence tests and process tracing, case study researchers have the means to resolve the degrees of freedom problem. (KKV, 1994:119-120) Instead of arguing over whether there is an inherent degrees of freedom problem in case study methods, methodological critiques can focus on whether in a particular study a researcher has generated enough process-tracing and congruence predictions and tested them against sufficient data to make valid inferences.

At the same time that they have allayed the standard degrees of freedom criticism of case studies, however, KKV have suggested that case studies suffer from the exact opposite of the degrees of freedom problem. In this view, "there always exists in the social sciences an infinity of causal steps between any two links in the chain of causal mechanisms," raising the danger of an "infinite regress" in process tracing. (KKV 1994:86) Thus, ironically, case studies have been criticized for drawing on data that some see as too scarce and others wrry is potentially infinite. In fact, however, the resolution of the "infinite regress" problem is no more difficult than the problem of selecting from among the infinite possible hypotheses and observations to be used in studies of covariance. The domains of hypotheses to be traced and observations to be used are restricted by the definition of research objectives traced and observations to be used are restricted by the definition of research objectives, the selection of the levels of analysis that are of interest, the relevant body of alternative hypotheses on the relationship under study, and the distinction between enabling variables and immediate causes. Thus, it is not necessary to examine all links between links. (Sayer, 1992: 120, Yee, 1996:84)

There are two other constraints on process tracing that pose more telling limits on its actual practice. Process tracing provides a strong basis for causal inference only if it can be established whether an uninterrupted causal path existed linking the putative causes to the observed effects, at the appropriate level(s) of analysis as specified by the theory being tested. Evidence that a single necessary intervening variable along this path was contrary to expectations strongly impugns any hypothesis whose causal affects rely on that causal path alone. The inferential and explanatory value of a causal path is weakened, though not negated, if the evidence on whether a certain step in the putative causal path conformed to expectations is simply unobtainable. Also, theories frequently do not make specific predictions on all of the steps in a causal process, particularly for complex phenomena. When data is unavailable or theories are indeterminate, process verification can reach only provisional conclusions.

Another potential problem for process tracing is that there may be more than one hypothesized causal mechanism consistent with any given set of process tracing evidence. When this problem of indeterminacy arises, there is no absolute standard for excluding alternative hypotheses that may be spurious. (Njolstad, 1990; Achen and Snidal 1989:156-157)) There are only empirical, methodological, aesthetic, analytical, and sociologicalcriteria for selecting hypotheses, testing them, strengthening or infirming them on the basis of the evidence, and modifying them (Bennett and George, 1997e). There are no absolute means of proving or disproving theories or establishing that changes in these theories are progressive rather than regressive. Still, even if it is not possible to exclude all but one hypothesis as a potential explanation for a particular case, it may possible to exclude at least some hypotheses and draw inferences that are useful for theory-building and/or policymaking.

Process Tracing as a Corrective for the Limits of Mill's Methods

Process tracing also addresses an additional misconception about case study methods. Several kinds of case study research designs bear a resemblance to John Stuart Mill's "method of agreement" or his "method of difference." However, this superficial similarity has caused many commentators wrongly to attribute to case studies all of the well-known limits that Mill and others have identified regarding these methods. In fact process tracing can compensate for many of the limits of Mill's methods, even though it cannot eliminate them entirely.

In general, this is because case comparisons, congruence testing, and process tracing are all set up by prior theories, and their results should be weighted or discounted by our existing level of confidence in these theories. This, together with methodological standards for progressive theorizing, such as Lakatos's insistence that theories must explain not only existing anomalies but "new facts," provides safeguards against the potential pitfalls of Mill's methods (Lakatos, 1970; Bennett and George, 1997e).

The key limitation of Mill's methods, which Mill himself identified, is that they cannot work well in the presence of equifinality. Put another way, Mill's methods can work well at identifying underlying causal relations only under three very strong conditions. First, the causal relations being investigated must be deterministic regularities involving conditions that are either necessary or sufficient for a specified outcome. Second, all causally-relevant variables must be identified prior to the analysis. Third, there must be available for study cases that represent the full range of all logically and socially possible causal paths. (Little, 1996; Lieberson 1994, George and McKeown, 1985)

Clearly, these strong assumptions seldom hold true. However, typological theorizing and case study methods do not require such stringent conditions. Let us consider each condition in turn. First, typological theories address the problem of equifinality directly, acknowledging and even taking advantage of the fact that there may be different causal paths to similar outcomes. The inductive development of typological theory attempts to map out the different causal paths, while the deductive development of typological theory attempts to provide theoretical reasons why particular conjunctions of variables lead to particular outcomes. Case study methods do not require causal relations of necessity and sufficiency, although case study methods -- like all methods -- offer stronger inferences on the existence of such relations than on that of equifinality or probabilistic causality. (Dion, 1997)

In addition, as long as all relevant variables are included in a typology, that typology inherently reflects interactions effects, even when those effects are not fully identified or understood by the researcher. Some critics of case study methods have suggested otherwise, arguing that these methods cannot incorporate interactions effects. (Lieberson, 1992:109-113) In fact, the logic of case study methods and the notions of causality associated with them have made case study researchers very attentive to interactions effects. (Ragin, 1987). If there are no measurement errors and there are deterministic or very high probability processes involved, admittedly two big assumptions, then two typologically similar cases, or cases with highly similar values on their independent variables, will have the same outcome, even if the interactions among the variables are caused that outcome are not fully understood or specified. Thus, we can have accurate predictions without accurate explanations, or the problem of spuriousness. For some research objectives, such as policy-makers' use of typologies, this may be acceptable, while for others, such as explanation by reference to causal mechanisms, it is not. Typological theorizing, as opposed to the simple use of typologies, pushes theorists to try to anticipate and explain interactions effects, although there is no guarantee that they will do so adequately. Process tracing and cross-case comparisons, though still fallible, may help identify which interactions are causal and which are spurious.

The second limitation of Mill's methods, the problem of left-out variables, is a potential threat to all methods of causal inference. Some critics have argued that the omission of causal variables is more likely to result in spurious inferences in case study methods than in research using other methods. (Lieberson, 1992:113). The problem with this critique is that it conflates Mill's methods with case study methods, and it does not acknowledge that process-tracing can test whether seemingly causal variables are spurious and to uncover supposedly unrelated variables that may in fact be causal. Moreover, the likelihood that relevant variables will be left out is lower for case studies than for statistical methods. Case study methods allow for the inductive identification of variables as well as their deductive specification, and they do not face a narrowly-defined degrees of freedom problem on how many independent variables to include. Indeed, one of the most visible and important many independent variables to include. Indeed, one of the most visible and important contributions of case study methods has been to identify causal variables that have been left out or insufficiently examined by studies relying on purely deductive theories or correlational methods. This is evident in the litaratures noted above on deterrence, where case studies have added variables on psychological dynamics and domestic politics, and alliance burden-sharing, where case studies have added domestic political variables and given greater weight to states' security dependence on an alliance leader.

The third limitation on Mill's methods, the requirement of having available for study cases representing all logically and socially possible causal paths, is a more binding constraint on case study methods. Causal inferences are indeed stronger when extant cases cover more of the typological space. Even so, having all possible types of cases available for study, while desirable, is not necessary. Not all research designs or research objectives require a fully inhabited typological space. Single cases, if they are most likely, least likely, or especially crucial cases, can be quite revealing about the strength of a theory. Comparisons of a few cases, if they are most similar or least similar, can also be revealing. Some cases provide more information than others on the theoretical issues of interest to a particular researcher. Moreover, for some research objectives, there may be cases for study representing most or even all of the possible types. The extant cases may also provide diverse causal paths even if the cases for any one causal path are not numerous enough for statistical methods.

Perhaps the most important difference between Mill's methods and case study methods in all three of the areas discussed immediately above is that case study methods can use within-case analyses, particularly process tracing, to ameliorate the limits of Mill's methods. Process tracing can identify different causal paths to an outcome, point out variables that otherwise might be left out, check for spuriousness, and allow causal inferences on the basis of a few cases or even a single case. These potential contributions of process tracing make case studies worthwhile even when sufficient cases exist for the use of statistical methods. Sophisticated critiques of case study methods acknowledge the value of process tracing. For example, Daniel Little, while more pessimistic than we are on the possibilities for typological theorizing, notes that such theorizing can be strengthened by the use of empirically supported social theories to establish hypothesized causal linkages that can then be tested through process tracing. As noted above, and as Little also points out, Theda Skocpol's work on social revolutions, in addition to traditional comparative analysis based on Mill's methods, uses established social theories in this manner. (Little, 1995:54)

Similarly, process-tracing provides an alternative way of making a causal inference when it is not possible to do so through the method of controlled comparison alone. An ideal case comparions based on Mill's method of difference requires identification of two cases that are similar in all but one independent variable and that differ in the outcome. When this requirement can be met, the comparison is "controlled" and provides the functional equivalent of an experiment. Variation in but one variable permits the investigator to employ experimental logic in making a causal inference regarding the impact variance in that variable has on the outcome (dependent variable). However, when the requirement needed for a perfectly-controlled comparison is not met, as is often true, process tracing can help. Process tracing can test whether each of the potentially causal variables that differ between two closely but imperfectly matched cases can or cannot be ruled out as causal. If all but one of the independent variables that differ between two cases can be ruled out as explanations for these cases' differing outcomes, and process tracing cannot rule out the last variable on which the cases differ as a cause of their differing outcomes, the case comparison provides a stronger basis for inference than either of the cases alone (for a good example, see Ray, 1995: 158-200).

Inductive and Theory-Testing Uses of Process Tracing

There remains the practical question of which theories to draw upon and which hypothesized causal paths to trace. This issue bridges the epistemological and methodological dimensions of process tracing. Depending on the research objectives --theory testing, theory development, plausibility probes, and so on -- this may be an exclusively theory-driven process or a partly or fully inductive one. For testing theories through process tracing, or process-testing, there are three sources of causal theories for potential testing. First, there may already exist within the relevant scholarly community well-developed alternative theories on the phenomenon at hand, which not only allows process verification but raises the possibility that it can assess their validity, either as alternative theories, where only one or a few is likely to survive the test, or as different causal paths to similar outcomes (ie, equifinality). This requires that these theories provide, explicitly or implicitly, well-defined causal mechanisms and predicted causal processes for the case at hand, or, more often, that the researcher uses the theory to specify the predicted process, hopefully in its entirety but at least in important respects, through which it should operate if it is to explain the case. Hubert Blalock offers some useful advice in this respect, suggesting, "at the risk of being accused of professional heresy," that when a theory is too vague to permit a definite relationship, the researcher should "forget what the theorist intended -- even though [s]he be a very renowned scholar" and insert the linkages that the researcher believes best fit the theory. (Blalock, p. 29) The researcher must also specify or at least consider the possibility that more than path, or more than one value at a certain point in a path, may be consistent with the theory, and whether these paths might also be consistent with alternative theories.

A second potential source of theories for process verification is the implicit or explicit theories that substantive experts have proposed to explain the particular case at hand or similar cases, including, for example, the writings of regional specialists, historians, ethnographers, and issue-area or functional experts. These must be translated into the terms of existing or new theories if they are to be accorded the same a priori status of these theories. If they do not relate to existing theories, they can still be rendered in terms of generalizable theoretical variables and treated as hypotheses which can be tested through a "plausibility probe." Theories codified in this manner cannot attain the same legitimacy as established theories until they have similarly attained further empirical confirmation and broader acceptance within the scholarly community.

A third source of theories for process verification, namely, the explicit or implicit causal explanations of participants in the social process under study, can also be restated in the terms of existing theories or generalizable variables. The difference here is that the researcher has to be careful to take into account of the information processing and motivational biases of the participants who are the sources of these theories, including their potential instrumental purposes in proposing particular explanations (an injunction that many no doubt find useful also in treating the theories proposed by scholars and experts). There has been considerable debate in sociology over whether case study researchers should attempt to recreate actors' experiences of the world, or impose their own theoretical framework on these actors.(On this point, see George and McKeown, p. 35; update with more recent arguments on this). This depends greatly, of course, on the research objective more recent arguments on this). This depends greatly, of course, on the research objective and the outcomes, behaviors, or factors to be explained. "Understandings" or ideas themselves may be the dependent variable, or they may be an independent variable. Even when this is the case, however, researchers will usually find it useful to render at least some variables in theoretical terms -- if the purpose is causal explanation in any of its broad senses, it is hard to imagine any explanation that leads from understandings to understandings without some underlying or intervening theoretical process.

With theories from all these sources, the researcher may have seemingly too many theories to choose from rather than too few. The researcher can limit the list of alternative theories to those that address those dimensions of the dependent variable that are of interest. In contrast to research using statistical methods, case study research should in general err on the side of including too many theories and variables rather than too few. In statistical research, the emphasis is usually on keeping as few independent variables as are absolutely necessary in a model in order to maintain high degrees of freedom. In the tradeoffs among left-out-variable bias, measurement or specification error from collapsing variables into a single variable, and insufficient degrees of freedom, statistical researchers may sometimes have to live with the first two of these problems to avoid the third. In case study research, however, as argued above, each theory brings its own set of process tracing predictions. In practice, there will be limits on whether data is accessible to test all of these process predictions, but the costs and benefits of including an additional hypothesis and its variables are different for case study methods than statistical ones. On the cost side, there is extra effort in collecting more process tracing data, and there may be gaps in such data for any given theory. The benefits of including additional theories and variables, however, are greater ability to deal with the problem of equifinality, reduced chances of spurious findings, and greater awareness of interactions effects. Whether an additional theory brings with it an abundance or a dearth of accessible process evidence is a practical matter related to the particular case, not part of an inexorable statistical law. Even theories for which no process tracing evidence is accessible must be considered in the analysis of a case if there are reasons to believe that they are relevant. The researcher must report on the existence of such untestable theories and qualify their conclusions accordingly regarding the theories that are tested, making these conclusions somewhat provisional as long as all alternative explanations could be tested.

A useful technique in charting out the expected process predictions of a particular theory is to construct a flow chart or diagram of its hypothesized underlying causal mechanisms, and drawing the causal web all the way from factors to effects. In addition to the usual boxes or letters indicating variables and arrows indicating hypothesized causal processes, causal diagrams should indicate the expected magnitude as well as the sign of an expected effect, and the nature of the expected relationship (linear, non-linear, (non)additive). (Blalock, Theory Construction, p. 29) Diagrams can also represent interactions effects, delayed linkages and possible pre-emptive/expectational behaviors, and suspected equifinality. [insert and explain examples from Blalock 35-47, others; add in equifinality etc).

With appropriately stated theories and their respective process predictions in hand, and keeping in mind the extent to which these theories have or have not survived previous empirical tests and won wide acceptance, the researcher can then engage in process verification. Lawrence Mohr has given a useful account of this method, following Michael Scriven's "modus operandi" method and his metaphor of a detective:

> when X causes Y it may operate so as to leave a 'signature,' or traces of itself

> that are diagnostic. In other words, one can tell when it was X that caused Y, because certain other things that happened and are observed unequivocally point to X. At the same time, one knows the signature of other possible causes of Y and one may observe that those traces did *not* occur. By using this technique, one can make a strong inference that X either did or did not cause Y in a certain case. For the present purpose, moreover, one notes in passing the affinity of this approach for the study of a single case. The kind of example of the modus operandi approach that is frequently given reminds one of the work of a detective or a diagnostician. (Mohr, 1985, 82-83, citing Scriven, 1976)

Several qualifications are worth emphasis. As Mohr notes, a process may leave and observable "signature," but it also may not, or the evidence may be hard to attain or inconclusive. Moreover, proving the negative, and demonstrating that a process did not occur, can be notoriously difficult. Both detectives and researchers face these difficulties, but a third problem, that theories are not always sufficiently specified to allow one to "know" the causal processes that they would predict, is more pervasive for researchers studying social phenomena than for detectives studying physical evidence.

As the "detective" metaphor suggests, when well-specified theories are available, process-testing can proceed either forward, from potential causes to effects, or backward, from effects to their possible causes, or both. Process verification should also ordinarily involve attempts to test and eliminate several alternative causal aprocesses. Thus, the detective pursues both "suspects," usually several, and "clues," constructing possible chronologies and causal paths both backward from the crime scene and forward from the last known whereabouts of the suspects. With theories, as with suspects, the evidence might not be sufficient to eliminate all but one. When theories are complementary, more than one may be consistent with the process-tracing evidence, and several may have added to the observed effect or even over-determined it. When theories make competing process predictions, the process-tracing evidence may be incomplete in ways that do not permit firm conclusions on which fits better. As the detective's colleague the District Attorney would remind us, a potential causal path cannot explain a case if it does not establish an uninterrupted causal path from the alleged cause to the observed outcome. The inaccessibility of evidence at one point in this path does not disprove the cause, but it does make it harder to eliminate competing theories beyond a reasonable doubt.

It is best if the researcher specifies theoretically well-defined alternative causal paths in advance for the purpose of process verification. This can guard against the common cognitive bias to see patterns where none exist (George and McKeown, pp. 37-38). However, it may not always be possible to specify hypothesized causal processes in advance of performing research on a case. If scholars, experts and participants have not proposed any useable theories, or if such theories have already been shown to be inapplicable to the case in question, the investigator may engage in process induction, setting out with the inductive purpose of finding one or more potential causal paths which can then be rendered as more general hypotheses for testing against other cases. This approach is also suited to the study of "deviant" cases, or cases that have outcomes that are not predicted or explained adequately by existing theories. In contrast to process verification, process induction proceeds mostly backward from effects to possible causes, though it could also involve forward tracing from a long list of potential causes that have not yet been formalized as theories or widely tested in other cases. Research on newly-discovered infectious diseases or new outbreaks of such diseases, for example, sometimes proceeds along this path.

proceeds along this path.

A second role for process induction concerns what we term "factor-centric" case studies. Most of the case study literature is "outcome-centric" in that it focuses on explaining variance in outcomes in terms of the the causal variables and contingent conditions that account for this variance. It is also possible that researchers or policy-makers may be interested in assessing the causal powers of a particular factor, particularly if it is a variable that policy-makers can manipulate, and exploring the contingent conditions under which the same values on this independent variable lead to different outcomes. For example, a central banker may pose the question: of all the half-dozen or so instances in my country in the last ten years (stipulations that control for many variables) where the central bank has instituted a sudden 0.5% increase in the interest rates, what have been the effects on the stock market at the next opening, and how did the economy do (along a specified range of measures of interest-rates sensitive activities like housing starts) in the next one, three, and six months? What contingent variables account for any variance in outcomes at these time points? To address such questions, a researcher can examine each of the cases in question using inductive process tracing. This would not be an entirely inductive process, as it would begin with theoretical guesses on which variables may be irelevant and worthy of process-tracing (such as the level of real interest rates prior to the increase, the real interest rates of other countries with strong economic ties, the prior level of unemployment, and so on). But process tracing would be largely inductive, looking to the interpretations of market and industry analysts and economists at the time, examining unusual conjunctions between process variables, and so on. The researcher could repeat this approach with bigger and/or smaller changes in the interest rate, exploring how sensitive outcomes were to the size of the rate change. The researcher might thus begin to establish a typology of interest rate increases, with categories defined by such variables as the magnitude of the rate change, surprise versus anticipated increases, increases that catch up to or rise above market rates, and so on. The goal is a better understanding of the factor in question, of the contingent variables that modify its effects, and of the past situations that may be most analogous (typologically similar) to the current policy context.

It is also possible that the investigator will not set out with inductive purposes in mind, particularly in outcome-centric case studies, but that in the course of testing the processes suggested by existing theories, they may incidentally uncover an unanticipated causal path that fits the process tracing evidence. This may occur even if one or more existing theories also appear to fit this evidence, but it is especially likely if none of the prior theories appears to fit this evidence well and the investigator, puzzled by this anomaly, searches more intensively for unexpected processes. It is useful for the purposes of theory development that researchers retain the option of following heretefore unexpected potential causal paths that are suggested by the evidence in the case.

The use of process induction to build testable hypotheses need not degenerate into an atheoretical and idiographic enterprise. When a researcher uncovers a potential causal path for which there is no pre-existing theory, there are several possible approaches for converting this a-theoretical finding into an analytical result couched in termos of theoretical variables. If deductive logic or inductive inquiry into other cases suggests a generalizable theory within which this causal factor fits, the theory can be specified, operationalized, and tested in other cases as a plausibility probe.

Another possibility is that a researcher will uncover a potential causal path that is not consistent with any of the alternative explanations specified a priori, but upon further reflection, this path appears to be an exemplar of an existing theory that the researcher had not previously thought to be relevant, or perhaps had not known about at the time that she or he specified the alternative hypotheses. In such cases, this theory assumes a status almost the same as if it had been specified a priori for testing, that is, it should be given weight proportional to its success in previous empirical tests, its general acceptance by scholars, and its ability to contribute to progressive research programs.

If a researcher cannot specify a suitable theory from an inductively-derived causal path, or find an existing theory of which the path is an example, the putative path can merely be reported as a finding consistent with the evidence but ungeneralizable until the researcher or another scholar can relate it to a useful theory or find other cases in which the path applies.

Process tracing is sometimes misunderstood to be applicable only to decisonmaking processes and other venues that involve preferences, expectations, intentions, motivations, beliefs, or learning at the individual and organizational levels.

In fact, process tracing is applicable to any hypothesized causal process, including not only readily observable constructs and processes at the macro level in the social sciences, particularly in sociology and economics. In economics, for example, theoretical relationships among variables like stock values, inflation expectations, and interest rates can involve long and complicated causal paths that can be tested empirically. Process tracing can also apply to causal mechanisms in the physical sciences, such as biology, microbiology, and epidemiology. It is true, however, that process tracing offers particular advantages relative to other methods in the study of intentional behaviors below the level of the state. This is true in part because process tracing of one or a few cases, as opposed to large N statistical analyses, allows sufficient time for the very detailed empirical inquiry that is necessary for studying these behaviors by obtaining documents, interviewing subjects, performing content analysis on documents and statements, and for establishing precise sequences of who knows and does what when.

Process tracing also lends itself to the study of the intentional behavior of individuals and organizations because this often involves the use of qualitative variables that are difficult though not necessarily impossible to quantify in a fruitful way. Even for qualitative methods like case studies, the difficulties of collecting and interpreting data on cognitions and intentions are well known, and they are indeed daunting. These difficulties are no doubt one reason that some researchers prefer to treat cognitive variables as a "black box," or dealing with them by making certain restrictive assumptions about individual behavior on the basis of game theories, rational choice models, or other simplifying approaches. In addition, many notions of causality assume that causes must precede effects, but this temporal distinction can begin to break down when expectations can create self-fulfilling or self-denying prophecies, whether those expectations are modelled in accordance with rational actor or psychological theories. Some have even questioned whether it is possible to construct causal models that involve variables as difficult to define and measure as human intentions, but this mistakes methodological difficulties for ontological ones. As one observer has pointed out, it is inconsistent to argue that intentions and understandings do not play a role in causation, because the very act of argument expresses a belief that arguments can change ideas and that ideas matter. (Sayer, p. 111) Moreover, because the decisionmaking at any level beyond the individual is a social process, it necessarily leaves behind at least some kinds of evidence -- documents, participant recollections, public communications -- even though this evidence may be far from complete or unbiased.

(George and McK., p. 37)

Because they are suited to studying micro level intentional behaviors, process tracing methods can be particularly useful in advancing the current debate between proponents and critics of rational choice models. Some critics argue that the proponents of rational choice theories have modified them in regressive ways through post-hoc theorizing, the proliferation of theoretical entities, and the selective use of evidence. On the other hand, proponents of such theories argue that such measures "might actually be instrumental or even necessary to the development of what Lakatos calls a progressive scientific research program." What both schools agree on is that the efficacy of rational choice theories must be judged in part on an empirical basis that emphasizes testing these theories against the decisionmaking processes that are actually observed. In other words, proponents of simplifying assumptions about individual behavior have to show that the benefits of these assumptions, in terms of simplicity, predictive power, and so on, must outweigh the costs of the concomitant loss of empirical richness. As two proponents of rational choice theories have stated, "because society is composed of human beings, social science explanations have to be compatible with psychological processes. This means both that it is physically possible for people to act as the social explanation requires, and to hold or form the relevant beliefs and desires . . . the question of what form such attribution should take is an empirical one." Process tracing is one useful method for examining such empirical issues, but not the only one. Statistical methods and even simulations can be used to study process tracing and evaluate deductive models.

Finally, it is worth noting that there is great interest in causal mechanisms, both as sources causal inference and windows of opportunity for policy interventions, not only among social scientists, but among those in the physical sciences as well. Epidemiologists and microbiologists, for example, are highly interested in tracing the causal mechanisms behind the transmission and debilitating effects of cancer, AIDS and other illnesses. The shift from interest in correlations to causal mechanisms, has been ennabled in part by the advent of technological means for tracing causal mechanisms on the microbial level, and is evident in the changing focus of epidemiological research (cite examples).

Conclusions

Process tracing on causal mechanisms is no panacea. It can require enormous amounts of information, and it is weakened when data is not accessible on key steps in a hypothesized process. In a prticular case, limited data or uderspecified theories, or both, may make it impossible to eliminate plausible alternative processes that fit the available evidence equally well. The validation of causal paths through process tracing, like the use of case study or statistical methods to establish covariation, must also address the demanding standards of internal and external validity (Cook and Campbell). Process tracing is no guarantee that a study can establish internal validity, or that it will uncover only relationships that are truly causal. Both false positives, or processes that appear to fit the evidence even though they are not causal in the case at hand, and false negatives, processes that are causal but do not appear to be so, are still possible through measurement error or under-specified or mis-specified theories. External validity, or the ability to generalize results to other cases, also remains a difficult standard. The findings of single case studies can only be contingent generalizations that apply to typologically similar cases, but even then, cases that appear to be typologically similar may differ in an as-yet unspecified causal then, cases that appear to be typologically similar may differ in an as-yet unspecified causal variable that leads to different outcomes. Disagreements over measurement of qualititative variables can also limit the cumulation of case study results, just as disagreements over how to define and quantify variables can limit the cumulation of statistical findings.

Despite these limits and problems, process tracing is useful method for generating and analyzing data on causal mechanisms. It can be used in studies of a single case and those involving many cases whose processes can be traced individually. It can greatly reduce the risks of the many potential inferential errors that can arise from the use of Mill's methods of comparison, from congruence testing, and from other methods that rely on studying covariation. Process tracing is particularly useful at addressing the problem of equifinality by documenting alternative causal paths to the same outcomes, and alternative outcomes for the same causal factor. In this way, it can contribute directly to the development of differentiated typological theories. Finally and most generally, process tracing is the only observational means of moving beyond covariation alone as a source of causal inference. Whether it is pursued through case studies, correlations, experiments, or quasi-experiments, it is an invaluable method that should be included in every researcher's repertoire.

***:** This paper was presented at the MacArthur Foundation Workshop on Case Study Methods, Belfer Center for Science and International Affairs (BCSIA), Harvard University, October 17-19, 1997. Much of this paper is drawn from three separate papers by Alexander George and Andrew Bennett: "Developing and Using Typological Theories in Case Study Research," Research Design Tasks in Case Study Methods," and "Lost in the Translatioin: Big (N) Misinterpretations of Case Study Research." These papers are a preliminary discussion of material that will appear in a book, *Case Studies and Theory Development,* by Andrew Bennett and Alexander George. This book is forthcoming with the BCSIA studies in International Security, published with M.I.T. Press, in 1998. These papers and additional information are available at: http://www.georgetown.edu/bennett.

Process tracing is a method used to explain change and causation, and to evaluate and develop theories in the social sciences. It has been used in psychology, political science, or usability studies. In process tracing studies, multiple data points are collected in comparison to simple input-output methods, where only one measurement per task is available.