

Process-tracing the Effects of Ideas

Alan M. Jacobs
Department of Political Science
University of British Columbia
C425-1866 Main Mall
Vancouver, B.C. V6T 1Z1
Canada
jacobs@politics.ubc.ca

Forthcoming in Andrew Bennett and Jeffrey T. Checkel, (eds.), *Process Tracing in the Social Sciences: From Metaphor to Analytic Tool*, Cambridge University Press.

This chapter examines the uses of process tracing for empirically testing ideational explanations and theories of political decision making.¹ Ideational mechanisms have characteristics that make them especially difficult to study, as compared to materially driven causal processes. Ideas are unusually difficult to measure and are often highly correlated with other plausible causes of political outcomes. Moreover, key mechanisms of ideational influence operate within a “black box” of unobservability from the perspective of the historical researcher. These challenges of ideational analysis motivate this chapter’s arguments in two respects. On one level, the chapter seeks to demonstrate that process-tracing represents an especially powerful empirical approach for distinguishing between ideational and material effects. At the same time, the chapter reckons with the considerable challenges that the study of ideational causation presents, even for careful process tracing.

The chapter offers ideational analysts a set of process-tracing strategies as well as guidance in identifying the conditions under which each strategy can be fruitfully applied. Broadly, the paper emphasizes three hallmarks of effective tracing of ideational processes. The first of these is expansive empirical scope. It is tempting for analysts testing ideational explanations to zero in on key moments of political decision, on the handful of elite actors who were “at the table,” and on the reasons that they provided for their choices. However, for reasons outlined below, a narrow focus on critical choice points will rarely be sufficient for distinguishing ideational from alternative explanations. To detect ideational effects, our analytic field of view must be expansive in terms of both

¹ The author thanks Justin Shoemaker for invaluable research assistance and the volume’s editors and participants at the Georgetown Authors’ Workshop for helpful comments.

temporal range and *level of analysis*. A well-specified theory of ideas will imply predictions not just about individual elites' statements and behavior at key moments of choice, but also about continuity and change, sequences of events, flows of information, and movements of actors across institutional settings over time.

Second, in outlining, illustrating, and assessing a set of empirical strategies, the paper emphasizes the importance of careful and explicit reasoning about *the processes that generated the data* under analysis. As in all inferential endeavors, analysts seeking to trace ideational processes must relentlessly confront their interpretations of the data with plausible alternatives. In ideational analysis, this means paying especially close attention to the ways in which the institutional and political contexts of choice generate *strategic incentives*. These incentives include pressures for actors to speak, behave, or keep records in ways that occlude, rather than reveal, the considerations motivating their decisions.

Finally, the chapter underlines the role of theory-specification in process tracing. Tightly specified theories with detailed mechanisms can substantially enhance the discriminating power of process-tracing by generating relatively sharp and unique empirical predictions. In the realm of ideational analysis, analysts can often fruitfully draw more detailed causal logics from psychological theories of how individuals process information and form beliefs. At the same time, the chapter points to the risk that an *overly* narrow specification of mechanisms may lead analysts to miss ideational processes that are in fact present.

As the editors indicate in their introductory chapter, process tracing is a versatile analytic approach that can be put to different kinds of knowledge-generating purposes. The analysis below is primarily focused on the *deductive* testing of claims about

ideational effects, rather than the inductive generation of hypotheses. The tools assessed here, however, may be equally applied to the testing of *general* theories as to the testing of explanations of *specific* cases. The causal processes of concern here, moreover, operate at multiple levels of analysis. Viewed narrowly, the effect of ideas on decision-making may play out on a very “micro” scale, at the level of individual-level cognition and short-run governmental processes. Yet, as I have foreshadowed, the chapter will argue that substantial empirical leverage can be gained from a more macroscopic approach: from the analysis of patterns of behavior and interaction among individuals and across organizations over extended stretches of time.

The remainder of this chapter proceeds in four sections. The first substantive section lays conceptual foundations by defining an ideational theory and distinguishing it from alternative logics of explanation. The second section then outlines three acute empirical challenges that afflict the testing of ideational claims. Next, taking into account these challenges, the third section outlines, illustrates, and assesses several types of process-tracing tests of ideational influence. These tests involve a variety of forms of data and logics of inference, including the analysis of communication; the examination of within-unit covariation (both over time and cross-sectionally); the tracing of paths of ideational diffusion; and analysis of the substantive content of decision outcomes. The chapter closes with reflections on the core analytical investments scholars must make if they are to effectively trace ideational causation in politics.

Defining an ideational theory

As the volume's editors point out in their introductory chapter, good process tracing involves, first, casting a wide net for plausible alternative accounts and, second, being as empirically "tough" on one's primary explanation as on the alternatives. Testing a theory against its competitors, however, first requires a clear *conceptual* distinction between alternative causal logics. In this section, I offer a definition of an ideational causal theory and logically distinguish ideational theories from non-ideational alternatives.

I conceptualize an ideational theory (or explanation of an outcome) as *a causal theory (or explanation) in which the content of a cognitive structure influences actors' responses to a choice situation and in which that cognitive structure is not wholly endogenous to objective, material features of the choice situation being explained.*²

The first part of this definition is straightforward: an ideational theory posits a causal effect of the content of actors' cognitions on their choices. These cognitions may include normative commitments, causal or descriptive beliefs about the world, or mental models or analogies from which actors draw specific beliefs or policy prescriptions.

It is the second part of the definition, however, that distinguishes an ideational theory from most alternative lines of explanation. It is a common feature of most theories of political choice that actors' choices flow causally from their cognitions. In standard game-theoretic accounts, for instance, actors' choices of strategy result from (among other things) their beliefs and their preferences. Nearly all theories of choice could, in this trivial sense, be considered "ideational."

² Herein, I refer interchangeably to ideational theories and explanations; the arguments I make are intended to apply to both.

How, then, can we conceptually distinguish ideational theories from non-ideational alternatives? In this chapter, I refer to non-ideational explanations of choice, broadly, as *materialist* explanations. We can conceptualize one key difference between ideational and materialist explanations by thinking about how each accounts for *variation* in actors' choices. In a materialist logic of explanation, variation in choices is caused by variation in *the objective, material parameters of actors' choice situations*. Material causes may include differences across cases in the relative material payoffs of the alternatives, arising from variation across those cases in the causal relations linking options to material outcomes. Material causes may also include differences in the menu of feasible alternatives (or strategies), arising from differing material capabilities or differing institutional or technical constraints. Rationalist institutional theories; theories grounded in class-based, sectoral, or geographic economic interests; and neo-realist theories of strategic interaction in international relations are among the more common forms of materialist explanation in political science.

In an ideational theory, by contrast, variation in choices across cases is explained by reference to variation in the *content of actors' cognitions*. This may include variation in the relative value that actors place on different material outcomes (i.e., goals or normative commitments); differences in actors' mental maps of the causal relations linking alternatives to outcomes (i.e., causal beliefs); or differences in actors' descriptive beliefs about the state of the world. A requisite feature of an ideational account, moreover, is that this variation in cognitions must not be purely a function of material conditions. The ideas in question, that is, must have a source *exogenous* to material

features of the present choice situation.³ Such prior causes may include exposure to ideas held by other actors through policy networks or processes of political socialization. Alternatively, actors' beliefs may arise from the lessons they draw from a disproportionately formative historical experience. Whatever the idea's prior cause, however, a claim of ideational causation necessarily implies that decision makers' beliefs or goals are *not fully determined by the material parameters of the choice being explained*.

Thus, an account in which actors in different cases hold different causal beliefs because the *true* causal relations objectively differ across those cases would *not* be an ideational explanation: the ultimate cause here would be the material conditions of choice. On the other hand, an account in which actors operating in environments governed by similar true causal relations act on different *beliefs* about those causal relations – beliefs which were shaped by something other than the objective causal relations themselves – would be an ideational explanation. As should be clear, ideational accounts are fully compatible with an instrumentalist logic of choice in which actors select the goal-maximizing option given their causal beliefs. The key distinguishing feature of an ideational theory is that those goals and beliefs can *vary* independently of objective material conditions, generating differing decisions.

This extended definition now allows us to delineate the empirical task of testing an ideational theory. In particular, the definition implies three elements that must be operationalized in order to establish ideational causation. Any test of an ideational

³ One may be able to trace the origins of many ideas to some set of material conditions: e.g., the past economic or sociological circumstances of their original formulation and dissemination. The key requirement here is that the ideas cannot be endogenous to material features of the choice situation *which is presently being explained*.

explanation must seek evidence that: (1.) decision makers possessed particular cognitions (a measure of the independent variable), (2.) that those cognitions shaped their choices (evidence of a mechanism of influence), and (3.) that those cognitions were not simply reducible to material features of the circumstances of choice (evidence of exogeneity of the independent variable).

The challenges of testing ideational theories

Attempts to adduce evidence of these three elements – to empirically distinguish ideational from material influences – confront a distinct set of challenges. I identify here three hurdles to ideational analysis, which roughly parallel the three evidentiary tasks identified above: the unusual difficulty of observing the independent variable; the difficulty of observing key mechanisms of influence; and a frequently close alignment between actors' ideational commitments and their material incentives.

First, the independent variable in an ideational theory – the ideas to which political decision-makers subscribe – is particularly difficult to observe. Error in the measurement of ideas can arise from the fact that the most readily interpretable manifestation of actors' cognitive commitments – their own verbal expressions of their ideas – is often a systematically biased indicator. As the volume's editors point out in Chapter 1, evidence that is provided by political actors themselves is subject to bias whenever those actors have incentives to conceal their true motives. Politics generates strong pressures for actors to employ verbal communication to strategically misrepresent the reasoning underlying their choices (Goldstein 1993; Shepsle 1985). In particular, officeholders or interest-group leaders, seeking to broaden support coalitions and advance

their careers, have strong incentives to occlude many of the material and self-interested motives that might underlie their policy positions. They likewise have incentives to exaggerate the importance of “good policy” motives and broad social benefits. They will in turn select “good policy” justifications that conform to widely embraced normative frameworks and causal models connecting chosen policies to valued goals.

The result will often be systematic measurement error and a tilt of the inferential scales in favor of ideational explanations – in particular, those centered around “pro-social” or widely accepted ideas – and against material explanations based on a logic of decision-maker self-interest.⁴ Importantly, this problem is not limited to utterances made *at the time* of decision: in recounting decisions in later memoirs and interviews, actors may face similar incentives to forge reputations for disinterested, civic-minded leadership.

Second, even where ideas can be well measured, analysts will face difficulty in assembling evidence of the mechanisms through which those ideas influence choices. Consider the mechanisms through which other commonly studied independent variables – such as institutions or the organization of interests – shape political outcomes. Many of these mechanisms operate at the level of *social interaction*. Institutional models of policymaking – such as theories of veto points or veto players – posit efforts by opponents of policy change to exercise influence at points of institutional opportunity, and efforts by proponents to bargain their way to winning coalitions across institutional venues. While some of this activity may be (strategically) hidden from view, much of it

⁴ By the same logic, it may also generate bias against explanations centered around “anti-social” ideas (e.g., racist ideas).

will be at least in principle observable by virtue of the fact that it involves communication and behavioral interaction *among* individuals and organizations.

Far more of the causal action in an ideational theory, by contrast, is *intrapersonal*, taking place *inside the minds* of individual decision-makers, as their pre-existing conceptual frameworks lead them to prioritize particular goals, attend to particular pieces of information, or employ particular causal logics. The challenge here is one of connecting independent variable to outcome: even if the analyst can establish that actors *hold* certain beliefs or goals, the intrapersonal nature of much of the causal process makes it more difficult to establish that actors *applied* those ideas to the choice being explained.

Finally, ideational analysis will often confront a challenge of multicollinearity. Competitive theory-testing is much easier when the analyst can observe suspected alternative causes varying independently of one another across cases. In politics, however, actors' ideas and their material circumstances are not independently "assigned." In fact, common patterns of political interaction will often select for ideas that push actors' choices in the same direction as their material incentives. One important selection process derives from the logic of delegation. Many influential actors in politics – from elected officials to agency directors to interest-group leaders – owe their positions of authority to an act of delegation by one or more principals (e.g., voters or legislators). These agents often face, on the one hand, strong material incentives to make choices that satisfy their principals (e.g., the threat of electoral punishment). Yet, whenever principals have a choice among agents, they will seek to reduce the risk of "agency loss" by *selecting* agents who share their goals (Bendor et al. 2001). Wherever an effective agent-selection mechanism is operating, the result will tend to be a high correlation between the

principal's demands and the ideational worldview of the agent. The result is a causal confound: the agent's material incentives to satisfy the principal will tend to dictate similar choices to those implied by the agent's own ideas. So, for instance, members of the U.S. Congress who take conservative stances on social issues are more likely than those taking liberal stances to (a.) sincerely hold conservative social attitudes and (b.) come from districts in which a large share of the voting public holds conservative social attitudes. While this may be good news for democratic representation, it is bad news for causal inference: if the former fact supports an ideational explanation of roll-call voting patterns, the latter will suggest an equally plausible office-seeking motive. In sum, in many political contexts, processes of agent-selection will deprive analysts of independent variation in ideational and material causal variables, making it harder to sort out potential causal confounds. In addition, a high correlation between actors' ideas and their material circumstances makes it harder for the analyst to establish that the former are exogenous to the latter.

I have argued that testing an ideational theory requires looking for evidence that decision makers' choices were influenced by the content of their cognitions and that those cognitions are not reducible to material parameters of the choice situation. I have now contended that cognitive content is difficult to observe without bias; that mechanisms of individual-level cognitive influence are unusually elusive; and that cognitions and the material conditions of choice will often be highly correlated. How, in light of these challenges, should the testing of ideational theories proceed?

Strategies of process tracing ideational effects

In the remainder of this paper, I describe and assess a set of process-tracing strategies for discriminating between ideational explanations of political choice and plausible materialist alternatives. In some ways, process tracing methods are ideally suited to addressing the challenges of studying ideational causation. For instance, the detailed, context-sensitive analysis of cases allows scholars to closely examine the strategic incentives generated by particular choice situations and to exploit variation at multiple levels of analysis and over time. At the same time, the nature of ideational causation creates unique challenges for process-tracing. The difficulty of detecting the operation of individual-level cognitive mechanisms is particularly problematic for an analytic approach that is so dependent on mechanism-related evidence. In crafting research designs based around process tracing, we must therefore think carefully about the ways in which ideational mechanisms might leave behind observable clues at *higher* levels of aggregation: in interpersonal interactions and communication, in organizational dynamics, and in the substance of the outcomes chosen.

In this section, I consider a set of empirical tests centered on the core elements of the definition of ideational causation introduced earlier in the chapter. Each empirical test contributes to one or more of the evidentiary tasks deriving from that definition:

1. **Measuring the independent variable:** identifying decision makers' sincere ideational commitments,
2. **Establishing the exogeneity of independent variable:** identifying an ideational source external to the choice situation being explained, and

3. **Finding evidence of a causal mechanism:** establishing that the relevant ideas were applied to the choice being explained.

In addition, certain tests discussed below complement the first three tasks by:

4. **Reducing multicollinearity:** identifying and exploiting independent variation in possible material and ideational causes.

In discussing each test below, I do four things. First, I elaborate the logic of inference underlying each test, specifying the observable implication (of an ideational theory) that it examines. Second, I identify the probative value of each test. The tests contribute differentially to the four evidentiary tasks identified above. Moreover, they vary in the degree to which they refer to unique evidence for an ideational theory (that is, in their *sufficiency*) and in the degree to which they test for a certain prediction of that theory (that is, in their *necessity*) (see Bennett and Checkel, this volume, and Van Evera 1997). I thus characterize each test according to its degree of necessity and sufficiency: the degree to which a test's failure impugns an ideational theory and to which its passage adds to the theory's credibility. Third, the probative value of each test depends on certain assumptions about the processes generating the data. For each strategy, I therefore outline key conditions that determine the strength or validity of the test. Fourth, I provide illustrations of each strategy drawn from prominent studies of the role of ideas in politics.

To structure the exposition, I group the empirical tests roughly according to the kinds of data on which they draw. In particular, I consider tests that draw on:

- the analysis of communication,
- the examination of within-unit covariation *over time*,
- the examination of within-case covariation *cross-sectionally*,

- patterns of ideational diffusion, and
- the substance of decision outputs.

Table 1 summarizes the tests, the evidentiary tasks to which they contribute, and the assumptions on which they hinge.

Throughout, the discussion emphasizes key themes foreshadowed in the chapter's introduction: the advantages of expanding the scope of inquiry both temporally and across levels of analysis; the importance of careful reasoning about processes of data-generation, including actors' strategic incentives; and the benefits of theoretical specificity.

Analyzing (mostly private) communication

The most legible manifestation of an idea will sometimes be its verbal expression. Often, the tracing of ideational causal processes relies heavily on an analysis of the things that decision-makers say and write. Indeed, among the most intuitive observable implications of most ideational theories of influence is the expectation that *we should observe communication, during the process of decision making, that is congruent with the idea*. Under favorable conditions, testing for this implication can serve two evidentiary purposes: it can provide a measure of the independent variable – revealing what ideas actors hold – and provide evidence of the operation of an ideational mechanism, suggesting that actors *applied* a particular set of values, beliefs, analogies, etc. to the decision in question.

For reasons outlined above, verbal communication by strategic political actors can be misleading. As this volume's editors point out in Chapter 1, the analyst must interpret

actors' statements with careful attention to the motives or incentives that the speaker may have had to say particular things. Among the determinants of those incentives is the *context* in which utterances are made. I unpack here the implications of one specific element of context that Bennett and Checkel discuss: the speaker's audience. More particularly, I explore here the implications of *privacy*: whether statements are made in to a small circle of fellow elites or to the general public.

Analysts of ideational effects often privilege statements delivered in more private settings – e.g., discussions within cabinet or correspondence between officials – over public statements. There is good reason to make this distinction. In more public settings, political elites will, in general, have stronger incentives to justify pre-determined decisions in socially acceptable terms. In private settings, on the other hand, decision-makers can let down their guard. Especially where actors with similar goals are deliberating together, it is more likely that they will understand themselves to be engaged in the collective pursuit of optimal (from their shared perspective) choices. In such a setting, actors are more likely to candidly reveal their goals, their causal beliefs, and their lines of reasoning in order to maximize the effectiveness of deliberation. Where an assumption of “collective deliberation” is justified, privately communicated statements can be a rich source of data on actors' cognitive commitments and their sources.⁵

One of the most striking uses of private communication to test an ideational argument appears in Yuen Foong Khong's (1992) study of U.S. decision-making during the Vietnam War. The ideas posited as influential in Khong's study are *analogies*

⁵ Public statements may also be revealing for some evidentiary purposes: for instance, where the analyst is interested in the kinds of policy justifications that public audiences find legitimate.

between past historical events – particularly, the appeasement of Hitler at Munich and the Korean War – and current choice situations. In testing his analogical theory, Khong relies heavily on quotations from correspondence, meeting minutes, and other primary documentation of closed-door deliberations over Vietnam among top U.S. officials. These communications reveal actors repeatedly reasoning about the risks and potential benefits of military options in Vietnam by reference to events in Europe in the 1930s and the Korean peninsula in the 1950s. Khong shows actors engaging in this process of selective historical inference repeatedly, across numerous contexts, and often in great detail.

In some cases, records of private deliberations may also be revealing for their silences. The analysis of reasoning in which actors do *not* engage plays an important role in my own study of governments' long-term choices in the field of pension policy (Jacobs 2011). The study seeks explain the choices that governments have made between two alternative methods of financing public retirement schemes: pay-as-you-go, or PAYGO, financing (the collection of enough tax revenue each year to match annual spending) and pre-funding (the accumulation of a fund to meet long-run pension commitments). Among the propositions tested is the claim that policymakers' choices were influenced by the “mental model” that they employed to conceptualize pension arrangements: in particular, by whether they understood a state retirement program as (a.) a form of insurance, analogous to private insurance or (b.) a social mechanism for the redistribution of resources. While the insurance model was expected to tilt actors' preferences toward pre-funding, a redistributive understanding was expected to yield preferences for PAYGO financing. Further, these ideational effects were theorized to arise through an *attentional*

mechanism: a given mental model was expected to direct actors' attention disproportionately toward those particular lines of reasoning logically implied by the model, and away from logics extrinsic to it.

The case of the design of the world's first public pension scheme, in Germany in 1889, yields especially clear verbal evidence of this effect (84-90). On the one hand, archival records show actors in closed-door settings drawing repeatedly on an understanding of public pensions as a form of "insurance" and articulating actuarial lines of reasoning that flow from this private-sector analogy. Equally revealing, however, is the *absence* of any record that officials considered key lines of reasoning that were inconsistent with the model. For instance, in their tight focus on the actuarial logic of commercial insurance, Bismarckian officials never spoke about the *political* consequences of fund-accumulation: in particular, the possibility that a pension fund accumulated in state coffers might be misused or diverted by future governments. This silence is particularly revealing – as evidence of biased information-processing – by comparison to two further observations. First, actors in other cases analyzed – where the redistributive model was dominant – referred frequently to the political considerations that German officials ignored. Second, the political risks to fund-accumulation appear to have been objectively present in the German case: within 30 years of the program's enactment, its fund had been wiped out by political misappropriation.

What is the probative value, for an ideational theory, of a test for private communicative evidence? How necessary is the discovery of such evidence to the survival of the theory? And how sufficient is such evidence for concluding that ideas had

an effect on the outcome? The answers to these questions depend on the assumptions that we can plausibly make about the process generating the data in a given case.⁶

We would seem to be on most solid ground in characterizing communicative evidence as *necessary* for the survival of an ideational explanation: it would seem hard to credit such an explanation if we had looked hard and failed to find significant verbal references to the ideational constructs hypothesized to have been influential. The wrinkle, as this volume's editors point out in Chapter 1, is that an absence of evidence cannot always be interpreted as evidence of absence. For many political and policy decisions, a sufficiently complete and reliable set of records of actors' closed-door deliberations may not exist or be available to the researcher, especially where actors were intent on keeping their discussions secret. Moreover, some widely held beliefs may never be voiced by actors during deliberations precisely because they are understood to be common knowledge.

Following the Bayesian intuition that the editors outline in this volume's introduction, the degree to which communicative evidence can serve as a "hoop test" – high in necessity – depends on the likelihood that we *would* have found verbal evidence of a set of ideas if actors had in fact held and applied those ideas to the decision. When assessing an absence of evidence, we must ask several questions about the data-generating process, including: Do we have evidence of deliberations in the *venues* within which actors would have been likely to apply and give voice to the idea in question? How *complete* is the available record of the deliberations in those venues? Would actors have had an *incentive* to voice the idea during deliberations if they subscribed to it? In my

⁶ The succeeding discussion draws upon Bennett and Checkel (this volume, 20-21), George and Bennett (2005), and Trachtenberg (2006).

study of German pensions, the absence of evidence of certain lines of reasoning is made more compelling because the data are drawn from (a.) relatively comprehensive transcripts, (b.) across several deliberative venues, (c.) containing participants who, if they *had* thought of the unmentioned considerations, would have had clear incentives to draw on them because the arguments would have bolstered the case for their desired outcome.

What about the sufficiency of the test? When is verbal evidence sufficient to establish actors' ideational commitments or that actors applied those ideas in reasoning about the choice? One threat to the sufficiency of communicative evidence is the fact that actors' statements in internal deliberations may – despite their privateness – be affected by strategic dynamics. Even in closed-door settings, political elites may frame arguments for the purpose of coalition-building, rather than open-minded deliberation, selecting lines of reasoning to maximize the persuasive effect on fellow decision-makers. Moreover, available records of deliberations may have been created or released strategically by participants in the decision-making process; records revealing less pro-social material motives may tend to be suppressed. As George and Bennett (2005) emphasize, assessing the probative value of archival evidence thus requires knowledge of the broader context within which deliberations unfolded: the role of a given discussion and deliberative venue within the larger decision-making process; the incentives and pressures that actors faced; and the procedures by which records were kept, stored, and declassified in the political context under analysis. The sufficiency of verbal evidence will be higher to the extent that we can, through empirical and logical argumentation, rule out strategic motives among both speakers and record-keepers.

The examples above also suggest that we can increase the sufficiency of the test – that is, the uniqueness of the empirical predictions – by increasing the specificity of the theory itself (see also Bennett and Checkel, this volume, 23). Effective causal-theory-testing via process-tracing always depends on a clear specification of the causal logic or mechanisms underlying a causal effect (Collier et al. 2004; Hall 2003; George and Bennett 2005). And the payoffs to relatively high theoretical specificity are apparent in both Khong’s and my own analyses of communicative evidence. Both studies set out to test ideational claims grounded in relatively detailed cognitive mechanisms, drawn from psychological models of mental representation and information-processing. These theories do not posit simply that a given set of ideas will influence decisions: they also supply a more specific set of predictions about the ways in which ideas should shape the *processes* through which actors arrive at those decisions, yielding a substantially harder test of ideational claims.

Drawing on schema theory, Khong, for instance, predicts not just that actors will make use of analogies but that they will ignore or discount information inconsistent with the analogy and interpret ambiguous information in ways that support the analogy. In my study of German pension politics, the theory yields the “risky” prediction that actors on both sides of an issue will display the same allocation of attention across considerations: thus, even opponents of a policy option should fail to attend to some considerations (those outside the dominant schema) that would speak strongly against the option. Such observations would be hard to reconcile with a strategic account of deliberation. By generating predictions that are less likely to be observed under alternative theories, a better-specified theory increases the sufficiency of supporting evidence.

At the same time, the analyst should weigh an important risk of crisp specification of mechanisms: while rendering ideational accounts more falsifiable, positing a particular cognitive mechanism of causation raises the probability of *falsely rejecting* an ideational explanation. In my own study, it was possible that ideas influenced German policymakers' choices through a cognitive mechanism *other* than the attentional mechanism that I theorized (say, by shaping actors' underlying goals). Deductive process tracing based on my tightly specified attentional theory would then have led me to understate the importance of ideas in shaping the outcome. How should the analyst manage this tradeoff between Type I and Type II errors? One way to guard against the danger of false negatives by theorizing multiple cognitive mechanisms, though this tactic will reduce the sufficiency of the tests. A strong familiarity with the relevant findings in cognitive and social psychology can also help rule out the least-plausible mechanisms. Moreover, the analyst should consider leavening deduction with induction. As Bennett and Checkel (this volume, 38-39) explain, a key advantage of process tracing is that in-depth engagement with cases provides opportunities for uncovering evidence of causes and mechanisms that had not been previously theorized. Thus, the researcher might begin with one tightly specified ideational mechanism; if no evidence for that mechanism is found, inductively search for clues of other ideational processes; and if another ideational logic is suggested, derive empirical predictions from that new logic and collect additional evidence to test them.

Examining covariation over time

For reasons outlined above, material pressures and actors' ideational commitments will often be systematically correlated. However, analysts can enhance their prospects of finding independent variation in suspected causes by studying decision making over time. Suspected causes that push in the same direction at the level of a case may diverge (a.) over stretches of time extending *beyond* the case or (b.) across temporal stages *within* the case. The analyst can exploit such independent variation to test for the distinct over-time correlations predicted by alternative theories. Temporally structured evidence can, further, permit inferences about both the exogeneity and the sincerity of actors' apparent ideational commitments. I discuss here two types of tests drawing on over-time covariational evidence: one grounded in the analysis of ideational stability and change across decision-making episodes in a single unit; another based on the inspection of sequences within a single case of decision-making.

Covariation over time: Analyzing ideational stability and change

Observation of the behavior of key decision-makers over substantial stretches of time can help distinguish ideational from material causes by uncovering independent variation in these two sets of factors. One strategy of longitudinal analysis exploits the fact that cognitive commitments are typically slow to change and that beliefs are robust to new information (see, e.g., Nickerson 1998). By analyzing decision making over an extended time horizon, the analyst can test the following observable implication of many ideational theories: *that, because cognitive constructs are relatively resistant to change, we should see evidence of relative stability over time in both actors' ideas and in the choices that are hypothesized to result from them, even as material conditions change.*

In effect, this test multiplies the number of cases available for analysis within a single unit (e.g., a country) by taking in a stretch of time covering a series of decisions. This will often mean extending the temporal scope of analysis prior to or beyond the decision(s) initially of central interest to the investigator (see also Bennett and Checkel, this volume, 33-35). The analyst then applies a longitudinal form of Mill's (1868) Method of Agreement to rule out alternative causes. If actors' statements and choices remain consistent with a hypothesized ideational commitment at multiple points in time, even as material pressures shift, then those material factors become less plausible as an explanation of actors' decisions. Furthermore, the case for both the exogeneity and the sincerity of actors' stated ideational commitments is considerably strengthened if those apparent commitments do not change with material conditions. If suspected "ideas" shift with the material winds, they are more likely to be endogenous or insincere post hoc justifications of choices that are actually driven by those material forces.

Students of ideational effects have frequently engaged in long-term longitudinal analysis to exploit this logic. Judith Goldstein (1993), for instance, in her landmark study of U.S. trade policy, examined decision-making over the course of more than a century. This timeframe included two decades-long periods during which a single idea – protectionism in one period, free-trade liberalism in the other – was dominant. In examining decision-making across several episodes in each period, Goldstein demonstrates that commitments to protectionism and free trade, respectively, were little moved by changes in economic conditions to which, under a materialist explanation, they ought to have been highly sensitive. During the postwar era of liberal dominance, for instance, Congress and the President continued to reduce tariff barriers even as the

country's trade position dramatically worsened and well-organized interests lobbied hard for protectionism (167-69).

Sheri Berman (1998), in her comparative study of social democratic parties, similarly leverages a longitudinal design to examine the presence and effects of specific ideas. She demonstrates the cognitive grip of Marxist doctrine on German social democrats by outlining party leaders' rigid adherence to it over time; most strikingly, during the Weimar period the party refused to broaden its appeal beyond the working class or embrace Keynesian responses to unemployment despite strong electoral incentives and problem pressures to do so.

To summarize, the longitudinal, within-unit Method of Agreement can lend support to an ideational theory to the extent that expressed ideas and observed choices remain constant as possible material incentives vary. Such an observed pattern lends support to the claims that (1.) actors truly hold the beliefs that they profess, (2.) those beliefs are not merely a function of (changing) material circumstances, and (3.) the material factors that vary are not the explanation. A variant of this logic is to examine whether actors' positions are consistent with their expressed ideas across *issue* areas. Observing an actor who supports a social welfare program from which she happens to benefit, but defends that program with egalitarian arguments, one can ask whether the actor also supports a redistributive program from which she derives no pecuniary or electoral advantage.

The *necessity* of a longitudinal Method-of-Agreement test is greater the less change there is in the matrix of material payoffs in the period under analysis: if actors' verbal reasoning or choices shift frequently with relatively modest change in material

conditions, then the ideational theory is seriously impugned. The *sufficiency* of the test is greater the more change there is in material payoffs: the more material pressures change while actors' choices and statements remain the same, the more decisively those incentives are ruled out. Analyzing longer stretches of time can thus tend to increase sufficiency.

Even so, evidence of consistency over long periods of time will not always imply support for an ideational theory. For some ideational theories – particularly, those that allow for learning – evidence of *change* in ideas and outcomes under particular circumstances can provide crucial support. For instance, when existing ideas and their policy implications *fail* in spectacular fashion, a theory of learning might expect actors motivated by “good policy” reasoning to reconsider prior understandings and adjust course. Here again, a clear specification of an ideational theory's mechanisms becomes important. At a cognitive level, are actors understood to engage in such strongly confirmatory reasoning that we should expect consistency over time *regardless* of the outcome? Or are there conditions under which actors are expected to attend to discrepant information and revise their ideas – i.e., learn? Are there sociological processes through which old ideas and their adherents get replaced by new?

Often, ideational theories do not explicitly answer these questions, but some do. Berman and Hall (1993) usefully adopt relatively clear – and differing – positions on the conditions for ideational change and, as a result, look for distinctive kinds of longitudinal evidence. Berman emphasizes the biasing effect of ideas on how actors process information, arguing that “ideas play a crucial role in structuring actors' views of the world by providing a filter or channel through which information about the external

environment must pass” (30). Given this model of cognitive self-reinforcement, Berman seeks evidence of over-time ideational and policy *rigidity*, even in the face of failure and seemingly clear objective indications that other options might be preferable.

Hall, in contrast, sees prior ideas as constraining only up to a certain point. Actors will tend to draw by default on existing paradigms, even in the face of considerable policy failure. But when failures sufficiently accumulate – and if they are inexplicable in the terms of the old paradigm – then social learning may occur.⁷ In support of this argument, Hall demonstrates, on the one hand, rather remarkable consistency in British policymakers’ adherence to Keynesian principles and prescriptions – despite their ineffectiveness – through the stagflation of the 1970s. At the same time, he shows that Keynesian doctrine lost credibility and was replaced following persistent failures that were incomprehensible from the standpoint of Keynesian theory. In Hall’s argument, that is, it is precisely because British policy *does* change in response to a strong form of objective feedback that the case for a particular kind of ideational influence receives support.⁸

As these examples illustrate, the longitudinal pattern for which analysts should go looking depends strongly on their theoretical priors about the conditions under which ideas change. It is worth noting that not all claims about ideational change and persistence are equally falsifiable. A prediction of strict rigidity is relatively easy to test for: any evidence of significant ideational change undermines the theory. The predictions of a learning mechanism are much harder to specify and operationalize (see Levy 1994).

⁷ The process of learning that Hall documents appears to be more sociological than cognitive, driven as much by shifts in the locus of authority as by individual-level information-processing.

⁸ For a related argument, see Culpepper (2008)

If learning can occur in the wake of dramatic failure, what counts as “dramatic”? If repeated failure is necessary, how much repetition is required? When exactly does an unexpected failure become an anomaly that forces ideational revision? Moreover, different theories might make different predictions about *which* actors will be most likely to change their minds: for instance, those with a material stake in the policy outcome, or those most directly exposed to information about the failure? Without well-crystallized theoretical accounts of the mechanisms through which learning operates, empirical tests based on a logic of learning can only be relatively weak “straws in the wind” (Van Evera 1997).

Finally, important considerations flow from the reliance of this test on the inspection of covariation between independent and dependent variables. Much of the recent literature on qualitative methods has drawn a sharp contrast between the logics of causal inference underlying process tracing, on the one hand, and correlational analysis (whether small-n or large-n), on the other hand. This contrast, for instance, underlies Brady, Collier, and Seawright’s (2004) distinction between a correlational “dataset observation” (DSO) and a “causal process observation” (CPO). As scholars have pointed out, many canonical methodological principles (most prominently expressed, in King et al. 1994) are drawn from a logic of covariation and apply differently or not at all to the analysis of CPOs.

In practice, small-n case study research partakes of both logics, blending causal-process analysis with correlational analysis. As in the test described here, case analysts often unpack cases into multiple sub-cases (temporally or cross-sectionally) and analyze the correlation of suspected causes and outcomes across those sub-cases. And whenever

they are drawing leverage from the inspection of covariation – whatever the level of analysis – the standard assumptions required for drawing unbiased causal inferences from correlations must be defensible. Crucially, familiar concerns about omitted variables and endogeneity apply in full force to these “within-case” covariational strategies.

The case-study researcher will be in an especially strong position to rule out endogeneity: indeed, a number of the tests explored here are specifically aimed at establishing the exogeneity of ideas. But scholars employing tests based on covariation – including other covariation-based tests described below – must think especially hard about the threat of omitted confounding factors. In particular, they must ask: are there other material conditions that remained constant alongside ideas (or that covaried with ideas) that might also have influenced the outcome? If there are, then the analyst will need to employ additional tests (which may themselves draw on CPOs) to rule out those variables’ confounding influence on the outcome.

Covariation over time: Examining the sequence of decision-making

As just discussed, process tracing over time may mean examining covariation across decisions within a given unit. Yet the analyst can also leverage useful variation across the sequence of steps *within* a single decision-making process. Sequential analysis can take advantage of the fact that different actors and different venues are likely to play an important role at different stages in processes of policymaking or institutional design. Sequential analysis begins by examining a decision-making trajectory to determine a stage in the process, *S*, at which a plausible alternative was removed from the menu of viable options. The analyst can then inspect most closely the motives of actors at and

prior to S , relative to the motives of actors involved after that watershed moment had passed. This test relies on the following empirical prediction: if an option was removed from the menu of active alternatives for ideational (or material) reasons at stage S , then *we should be able to observe actors who plausibly held that idea (or who had that material interest) centrally engaged in the policymaking process at or before S* . This test contributes to causal inference by generating independent variation – over time within a decision-making episode – in material and ideational factors that are correlated at the level of the episode taken as a whole.

In my analysis of pension policymaking (Jacobs 2011), I seek to distinguish between electoral and ideational motives in governments' choices between PAYGO financing and pre-funding. In general, PAYGO financing tended to be the more appealing option in electoral terms because it imposed the lowest costs on constituents and delivered the largest pensions in the near term. At the same time, prominent ideas about the political economy in some of the cases analyzed also favored PAYGO financing, particularly the notion that elected governments cannot credibly commit themselves to saving large reserves for future use. Cases in which pro-PAYGO ideas were dominant are thus especially difficult to decipher because material pressures (electoral incentives) and ideas push in the same direction.

The study's analysis of British pension politics illustrates how sequential evidence can help pry apart correlated potential causes. The outcome to be explained in this case was British ministers' decision to place their pension system on a PAYGO basis in 1925 (104-107). As secondary histories and archival records make clear, Conservative ministers in Britain initially designed a scheme with full pre-funding. This blueprint was

then sent to an influential interdepartmental committee of civil servants for vetting and, according to an internal report, was rejected by this committee on the grounds that elected officials could not be trusted to resist short-run political pressures to spend the fund – a view with a long pedigree within Whitehall. After this stage, there is no evidence in the historical or archival record of pre-funding having been considered further by elected or unelected officeholders. These temporally ordered data are revealing on two points: (a.) that those actors with the *strongest* electoral motivations (ministers) placed the less electorally appealing option on the agenda and (b.) that that option no longer appeared on the menu after those actors with the *weakest* electoral motivations (career bureaucrats) – and a strong set of cognitive commitments running counter to the plan – had rejected it. In short, the observed sequence is far less consistent with an electoral than with an ideational explanation.

Tightly assembled sequential evidence can prove quite decisive against either ideational explanations or rival hypotheses by helping to eliminate, as potential causes, the beliefs or motives of downstream actors (whether ideationally or materially generated). At the same time, temporal orderings must be interpreted with caution. If political actors are even moderately strategic, they will frequently take positions and make choices *in anticipation* of other actors' reactions. Perhaps British civil servants simply discarded an option that they knew their political masters would, if presented with it, later reject. Or perhaps ministers sent the plan to committee precisely in the hope that senior bureaucrats would kill it.⁹ In social causation, temporally prior observations of political behavior can be endogenous to subsequent (expected) outcomes. Sequential

⁹ What makes both possibilities unlikely in the present example is the prior step in the sequence: the initial design and proposal of the idea by ministers themselves.

analysis should thus be informed by evidence or reasoning about the incentives that actors involved early in the process might have had to pander to the preferences of those who would arrive on the scene later.

Examining within-case cross-sectional covariation

We have considered the use of over-time within-unit variation to cut against multicollinearity of ideational and material forces. A similar logic also applies to the disaggregation of cases *cross-sectionally* – across subunits within a case. Some ideational theories, for instance, may usefully imply predictions about the positions that *individual actors* should be observed to take on the issues up for decision.¹⁰

The logic of inference here closely follows the familiar logic of analyzing cross-case variation, but at a lower level of aggregation. Actors within a case (individuals or organizations) will display varying degrees of exposure to experiences, information, or argumentation that might shape their beliefs, goals, or conceptual toolkits. They will also vary in their material stakes in the choice. This information will be analytically useful whenever those two patterns diverge: when the cross-actor distribution of ideational exposure is only weakly correlated with the distribution of material stakes. If the relevant ideational and material influences and actor positions can be well measured, the resulting test approaches “double decisiveness.” That is, it would seriously impugn *either* an ideational explanation or the materialist alternatives if well-measured variation in actors’ stances on the issue did not correspond to variation in their exposure to ideational influences or to their material stakes in the issue, respectively.

¹⁰ Related strategies could involve unpacking a country-level case into subnational units or institutional settings across which suspected causal conditions and actor positions vary.

Andrew Bennett's (1999) study of Soviet military interventionism in the 1970s and 1980s makes substantial use of this method. Bennett seeks to explain why the Soviet Union (and, later, Russia) chose to intervene in some times and places but not others. His prime theory yields an ideational explanation in which Soviet and Russian leaders' beliefs about the effectiveness of military intervention derive from personal experiences: personal involvement in a successful intervention is expected to reinforce actors' beliefs in the efficacy of the use of force while involvement in a failed intervention is theorized to lead to learning and belief change.

At the level of the cases as a whole, the outcomes to be explained appear overdetermined: we observe the presence of both material and ideational forces that could explain Soviet and Russian policy choices. Several of Bennett's most decisive tests thus leverage variation *within* the state in actors' exposure to formative experiences from which lessons could be drawn. This differential exposure derives, for instance, from generational differences and variation in whether or not actors were directly involved in the Soviet Union's disastrous intervention in Afghanistan. Crucially, moreover, Bennett attends closely to points of divergence between the distribution of material stakes and the distribution of learning opportunities. So, for instance, while the military as an organization had a material stake in an expansion of its turf and resources, not all members of the organization had personal experience of the Afghan war; veterans of that failed intervention were, in turn, among the fiercest uniformed opponents of the war in Chechnya. This within-case correlational pattern helps to carry the ideational explanation through a critical hoop while simultaneously casting significant doubt on an important materialist alternative.

The probative value of this type of test depends on a number of conditions. First, the certainty of the prediction, for either an ideational or materialist theory, depends on how well actors' sincere positions on the issue can be measured. And, like measurement of ideas themselves, measurement of actor positions needs to take into account potential strategic dynamics: the possibility that actors may have had strategic reasons to take public positions that differ from their sincere preferences.

Second, theoretical clarity is once again crucial. In particular, the method relies on a clear specification of what kind of "exposure" is causally important: what kind of stimulus ought to generate or transmit a given set of ideational commitments? Bennett draws heavily on cognitive psychology to identify the ways in which particular kinds of experience and information ought to translate into actor beliefs. Only with this specification in hand can he determine which actors have been "exposed" and which have not. On the other hand, as discussed above, the choice of degree of specification presents a dilemma. The more precise the specification, the more closely this test approaches a smoking gun for an ideational explanation. Yet unwarranted precision also risks setting up an overly restrictive hoop test and a false negative finding. (And the same issues apply to postulating the sources of actors' material interests.)

Third, because it draws on the inspection of covariation, this test is vulnerable to familiar threats to correlation-based inference. To avoid omitted variable bias, for instance, the analyst must be careful to account for all plausible influences on actors' positions that are also correlated with their ideas.

Tracing ideational diffusion

In their account of process-tracing in this volume's introduction, the editors evoke the image of a radioactive medical dye: the dye's observed path through a patient's body can help discriminate among alternative biological processes that are themselves unobservable. Likewise, while individual-level cognitive structures may be difficult for the political analyst to observe, discriminating evidence often lies in the observable pathways along which ideas travel through a political system. I turn now to three tests for ideational influence that center on paths of ideational diffusion. I discuss tests for (1.) the origins of ideas, (2.) the transmission of ideas across actors, and (3.) the movement of ideational "carriers" across institutional settings.

Identifying ideational origins

Establishing that ideas mattered in a decision making process requires establishing that they are *exogenous* to the material circumstances of choice. If an ideational framework is indeed exogenous, then the following prediction should usually hold: *there should be evidence of a source for the idea that is both external and antecedent to the decision being explained.* Where the exogeneity assumption is valid, such evidence will usually be easy to find: typically, proponents of new issue understandings or ideological frameworks *want* to transmit them in order to influence the course of social events – and are thus likely to make and disseminate statements of their views. This strategy, in most cases, is thus a hoop test: without a demonstration of prior intellectual ancestry, the case for ideational influence should usually be considered weak.

Such demonstrations are, unsurprisingly, quite common in ideational accounts. Berman exhaustively documents how the Swedish Social Democrats' programmatic

beliefs emerged from the thinking of early party leader Hjalmar Branting (38-65) while those of the German SPD emerged from the thinking and argumentation of theoreticians Friedrich Engels and Karl Kautsky (66-95). Goldstein traces the free-trade ideas that dominated the postwar era back to work being done in economics departments at U.S. universities decades earlier (88-91). And Hall can readily establish that monetarist ideas had become well established within the U.S. economics profession and been, subsequently, taken up by British right-of-center think tanks and journalists prior to the policy shifts that he seeks to explain.

Demonstrations of antecedent origins do not, by themselves, establish exogeneity. One reason is that actors within the decision-making episode being explained could have “cherry-picked” – from among the pre-existing ideas available in their environment – those that were most compatible with their material interests. The ideas employed during the decision making process would, in such a situation, be endogenous “hooks” for policies chosen on other grounds. Moreover, not just any intellectual antecedent will satisfy the hoop test. The source must have been sufficiently prominent and credible to have influenced the intellectual environment in which the case is situated.

But should we always consider the search for an ideational antecedent to be a hoop test? What if the causally important idea is the “brainchild” of the episode’s key decision maker, who never had occasion to express this belief prior to the choice being explained? In such a situation, there might be no observable intellectual antecedent, even if an ideational explanation is right. A crucial implication is that not all ideational claims are equally amenable to empirical analysis. The idiosyncratic beliefs of lone individuals

will usually be harder to study, and claims about them harder to falsify, than arguments about the influence of socially *shared* cognitions with identifiable origins.

Tracing paths of ideational transmission

A prior source for an idea is itself insufficient to sustain an ideational account: the analyst should also be able to demonstrate that the idea was *available* to decision-makers prior to the decision being explained. In this subsection and the next, I suggest two types of evidence that may, independently, help satisfy this hoop test of ideational influence. First, the analyst could identify *a pathway – an organizational structure or a social interaction – through which information or argumentation was likely to have been transmitted to authoritative actors.*

Alastair Iain Johnston (1996), in his case study of Chinese security policy, examines an ideational explanation of China's apparent shift toward a more constructive engagement in arms-control. One form of evidence for which Johnston looks is indications that Chinese officials were exposed to new, more dovish security ideas through transnational communities of experts. He uncovers evidence of several pathways of dissemination, finding that considerable numbers of Chinese officials spent time at Western security institutes and took part in bilateral meetings and training programs with U.S. organizations committed to arms control – much of this, prior to the policy shift being explained (43-46). These data help keep an ideational explanation in contention. (For a reason that I outline below, however, Johnston ultimately concludes that the transmitted ideas were not responsible for the policy change observed.)

Erik Bleich (2003), in his study of race politics in Britain and France, similarly provides evidence of transnational contacts as a pathway of ideational dissemination. He shows that an influential group of British Labour Party politicians were exposed to new understandings of racism – as a problem of access and discrimination, best handled through civil penalties and administrative procedures – both through visits to North America and through the study of U.S. and Canadian models of race relations (53-56).

Identifying mobile “carriers”

As just discussed, one way to establish ideational availability is to find evidence that actors in positions of institutional authority came into contact with the relevant ideas. Alternatively, the availability can be established with evidence of the *movement of individuals* – individuals reliably known to hold a given set of ideas – into decision-making institutions. In this second version of the test, *changes in outcomes should follow the entry of identifiable “carriers” of the relevant ideas into key loci of political authority* (beyond examples discussed below, see Checkel 1997).

The institutional analysis of mobile carriers is central to Margaret Weir’s (1989) explanation of the differing fates of Keynesian policy prescriptions in the United States and Britain. Weir begins with the observation that Keynesian policies were introduced earlier in the United States but proved less enduring than in Britain. She accounts for this temporal pattern, in large part, by reference to differing patterns of recruitment and distributions of power in the two political systems (for a parallel argument in another context, see Risse-Kappen 1994). Staffed by a large number of political appointees, the U.S. bureaucracy is a relatively porous environment characterized by rapid turnover in

personnel and without a single, centralized locus of policymaking authority. With high turnover across a fragmented system, disciples of Keynesian thought gained relatively quick entry to federal economic councils and agencies in the 1930s. The dispersion of authority, however, limited their ability to enact the type of coordinated policy responses that Keynesian theory prescribed. Moreover, serving at the pleasure of the president, Keynesian advisors never achieved a stable and secure foothold within government. The result is the quick adoption of, but unsteady commitment to, countercyclical macroeconomic management.

Weir also documents, by contrast, the far more regimented environment of the UK Treasury: not only was the department dominated by career bureaucrats (making turnover slow), but recruitment procedures and lines of authority severely limited the entry or influence of carriers of new ideas. The Treasury's virtual monopoly of economic policymaking authority within the state further restricted access opportunities for ideational upstarts. It took the national emergency of World War II to pry the system open: Treasury authority was temporarily diluted, and Keynesian economists (including Keynes himself) were brought into government to help manage the wartime economy. Following the war, the same organizational rigidities and concentration of authority that postponed the Keynesians' entry then secured their position within the state, leaving them ensconced in career positions at the Treasury. Keynesian principles came to dominate British fiscal and economic policymaking for the next 30 years (on a similar point, see Blyth 2002).

Analyses based on personnel movements across institutions hinge on a few important assumptions. First, we must be able to reliably identify the carriers' ideational

commitments. Indeed, what makes a carrier analytically “useful” is that her cognitive commitments are more readily knowable than those of other actors involved in decision-making, especially elected officials. Carriers’ belief systems can often be inferred by reference to their sociological context – such as their embeddedness within professional networks or the site of their training – or from past verbal communication. In this respect, the most “useful” carriers will have a prior track record of activity *outside* of politics – i.e., in an intellectual or professional setting in which the incentives for strategic misrepresentation of beliefs are limited. Second, for their ideas to have explanatory power, the carriers must not only take up residence within major loci of authority; they must have sufficient influence *within* a venue for their ideas to shape its outputs.

Finally, the analyst must dispense with an alternative explanation: that the carriers were *selected* by a set of political principals in order to provide intellectual cover for an option that was appealing those principals for reasons of material interest. Where experts are hand-picked for political convenience, these carriers – and their ideas – are epiphenomenal. One response to this quandary is to employ the carriers as an explanation of *longer-term* rather than immediate choices: even if politicians chose carriers strategically, those carriers may exert long term influence if they remain in place – like entrenched Whitehall bureaucrats – long after their political masters have departed the scene.

Where key assumptions can be met, analyses of ideational availability – based on either transmission paths or mobile carriers – can aid causal inference in a few ways. First, they can contribute to an unbiased measure of decision makers’ ideational commitments by helping to establish that actors had access to the relevant ideas or (in the

case of mobile carriers) providing evidence of their views under reduced strategic pressures. Second, in doing so, these tests can lend support to claims of ideational exogeneity. Third, as Weir's study demonstrates, the analysis of ideational movement can exploit distinctive temporal variation in ideational availability: it can demonstrate that a given idea was "on the scene" when congruent policy change occurred, the absence of policy change prior to the idea's arrival, and the fragility of policy change after mobile carriers' exit. Such over-time patterns can significantly undermine the sufficiency of non-ideational alternatives, suggesting that the availability of the relevant idea was *necessary* for the outcome to occur. Of course, since this logic draws on patterns of covariation between potential causes and outcomes, the standard cautions about correlational inference (discussed above) apply here.

Unpacking the substance of decision outputs

In large-n analyses, scholars are usually forced to code decision outputs relatively crudely – along a single dimension or using a very small number of categories. Small-n analysis, in contrast, affords the opportunity to attend much more closely to qualitative features of actors' decisions, and such scrutiny can sometimes produce evidence with substantial potential to discriminate among possible motives. The analyst can usefully ask the following question of a policy or institutional choice: is this precisely the way actors would have constructed the policy or institution if they had been motivated by a given normative commitment or causal belief? A detailed examination of the "fit" between the outcome and alternative lines of reasoning can contribute to a demonstration of the

mechanisms at work: in the best case, it can help discriminate among the possible considerations or motives that actors might have applied when making the decision.

Inspection of decision outputs can test an observable implication of the following form: *if a choice was driven by policymakers' commitment to Goal X, then the output should be take a form that, given the state of knowledge at the time, was likely to effectively promote Goal X.* A casual example will help illustrate: President George W. Bush's tax cuts of 2001 were partly sold to the public as much-needed stimulus for a slowing economy. A detailed inspection of the package's provisions would cast doubt, however, on an explanation based on a commitment to quickly boosting the economy. For instance, a large majority of the revenue cuts were both substantially delayed in time and targeted to those (the wealthy) least likely to spend the additional disposable income (Hacker and Pierson 2005). On the other hand, these policy details are highly congruent with an electoral logic of redistribution toward wealthy supporters.¹¹

In his study of Chinese security policies, Johnston undertakes a systematic unpacking of the dependent variable and test for ideational fit. The starting point for Johnston's analysis is an apparent shift toward greater cooperativeness in Chinese arms-control policy: the question is whether this shift is generated by a strategic interest in improving China's image or by a new set of more internationalist ideas about the sources of global security. By closely examining the specific international agreements to which China has been willing to accede, Johnston is able to derive considerable discriminatory leverage. In particular, he finds that Chinese leaders have largely cooperated with international arms-control efforts when those efforts would exact a low *cost* to China's

¹¹ They might also be congruent with a goal of long-run growth combined with a supply-side belief about the beneficial effects of tax cuts on the rich.

military capabilities but walked away from efforts that would impose substantial, binding constraints (49-57).

A test for the “fit” between outcome and motive may serve as a hoop test not only for an ideational explanation but also for materialist alternatives. One can ask, that is, is the outcome precisely what an actor with a hypothesized political or economic motive would have chosen? Were the material benefits of the policy, for instance, surgically directed toward those constituencies most critical to the government’s reelection prospects or more broadly diffused? Moreover, as the Johnston example illustrates, *costly* features of the output are often the most illuminating. The fact that a choice imposes costs on decision makers or their constituents in order to more effectively advance Goal X is an especially informative signal about the importance of Goal X to decision makers.

The application of this test also confronts an important complication: because political decisions are usually collective choices, they often involve compromise among actors with divergent beliefs or goals. Deviations from an ideational (or material) logic may, therefore, reflect not the absence of that logic’s operation but the comingling of that logic with other motivations. This complication is not intractable, however; indeed, it can be turned into a testable hypothesis. By closely examining the decision-making process alongside the details of the outcome, the analyst should be able to determine how well any departures from the prescriptive logic of an idea held by one set of actors “fit” the demands of other actors with veto power or strong bargaining leverage.

IV. CONCLUSION

The process-tracing strategies explored here require, on the whole, three types of analytical investment. The first is an investment in *breadth of empirical scope*. In measuring politicians' and policymakers' ideational commitments, analysts might begin by examining actors' statements at or just prior to the critical moment of choice. But an ideational theory's observable implications can be readily multiplied, and their uniqueness enhanced, by expanding the inquiry both temporally and across levels of analysis. Establishing the exogeneity of actors' ideas almost always requires expanding the historical scope of inquiry to periods prior to the choice being explained. By examining extended stretches of time, analysts can also make discriminating observations about the degree of stability of, or the timing of change in, actors' statements and issue positions, relative to change in the material context of choice. Likewise, by shifting the focus from the individual level toward larger patterns of social interaction, scholars can track the movement of ideas and their adherents across organizational settings and institutions. Substantial leverage can also be gained by disaggregating episodes to inspect within-case correlations across both participants and sequenced steps in the decision-making process.

At the same time, this paper has suggested that none of this is straightforward: each of these strategies can only be credibly employed when key assumptions can be made plausible. To put the point another way, the sufficiency of these empirical tests – for substantiating an ideational account – depends on the analyst's ability to rule out alternative interpretations of the evidence. Hence the second analytical commitment required of good process tracing of ideational effects: close attention to the *assumptions* required for drawing valid inferences from evidence. In part, this means close attention to

the incentives generated by the institutional, organizational, and societal context. Analyzing these incentives means deploying case-specific knowledge of formal and informal institutional structures, patterns of political competition, economic and social conditions, and details of the substantive issue at hand. For those tests that rely on assessing within-case covariation, the analyst must also attend carefully to the assumptions necessary for drawing causal inferences from correlations, including accounting for potential confounds.

Third, I have emphasized the value of *richly theorized mechanisms* for effective process-tracing. Ideational mechanisms can be fruitfully specified both in terms of individual-level cognitive processes and sociological processes through which ideational frameworks are disseminated, embedded within organizations, and replaced over time. Theoretical refinement can go a long way toward rendering ideational theories more falsifiable. Yet, as I have cautioned, tight specification of mechanisms also has risks: a deductive search for evidence of narrowly theorized mechanisms may render ideational effects more elusive, generating more false negatives. Thus, where more than one mechanism is plausible for a given effect, the analyst will often want to test multiple ideational logics, trading a measure of falsifiability for an increased chance of picking up ideational effects.

Though it has not been a focus of this essay, a final word of caution is in order about the selection of cases. As Bennett and Checkel point out in Chapter 1, any inferences drawn from process tracing must take into account whether the case examined is most or least likely for the theories being tested. Whether it makes sense to choose cases that are most or least likely loci of ideational influence depends on the analyst's

goals. Most likely cases will be especially informative for the inductive *building* of ideational theories, as they are good places to observe ideational mechanisms unfold (see, e.g., Parsons 2002). Inductive theory-generation may be aided by the selection of cases across which outcomes appear to align poorly with material conditions or in which actors faced high uncertainty about their material interests – and, thus, where there was greater room for the operation of ideational processes (see, e.g., Berman 1998; Blyth 2002). However, if the analyst seeks to test a *general* theory of political decision making, a selection procedure that selects “most likely” cases from a domain of decisions will, on average, lead us to overstate the overall influence of ideas in that domain. A more balanced test would involve a selection rule uncorrelated with the likelihood of ideas mattering – for instance, choosing for wide variation in the outcomes to be explained. Alternatively, if the analyst wants to subject an ideational theory to an especially *hard* test, then she should seek out “least likely” cases, such as those in which actors’ material stakes were high and pushed strongly in favor of the observed decision (George and Bennett 2005). Evidence that actors’ decisions in such cases were nonetheless shaped by their particular cognitive commitments would offer especially strong support to a general claim of ideational influence.

Bibliography

- Bendor, J., A. Glazer, and T. Hammond. 2001. "Theories of Delegation." *Annual Review of Political Science* 4 (235-269).
- Bennett, Andrew. 1999. *Condemned to repetition? : the rise, fall, and reprise of Soviet-Russian military interventionism, 1973-1996*. Cambridge, Mass.: MIT Press.
- Berman, Sheri. 1998. *The social democratic moment: ideas and politics in the making of interwar Europe*. Cambridge, Mass.: Harvard University Press.
- Bleich, Erik. 2003. *Race politics in Britain and France : ideas and policymaking since the 1960's*. Cambridge, UK ; New York: Cambridge University Press.
- Blyth, Mark. 2002. *Great Transformations: Economic Ideas and Institutional Change in the Twentieth Century*. New York: Cambridge University Press.
- Checkel, Jeffrey T. 1997. *Ideas and international political change : Soviet/Russian behavior and the end of the Cold War*. New Haven, Conn.: Yale University Press.
- Collier, David, Henry E. Brady, and Jason Seawright. 2004. "Toward an Alternative View of Methodology: Sources of Leverage in Causal Inference." In *Rethinking Social Inquiry: Diverse Tools, Shared Standards*, ed. H. E. Brady and D. Collier. Lanham, MD: Rowman and Littlefield.
- Culpepper, Pepper D. 2008. "The politics of common knowledge: ideas and institutional change in wage bargaining." *International Organization* 62 (1):1-33.
- George, Alexander L., and Andrew Bennett. 2005. *Case Studies and Theory Development in the Social Sciences*. Cambridge, MA: MIT Press.
- Goldstein, Judith. 1993. *Ideas, interests, and American trade policy*. Ithaca: Cornell University Press.
- Hacker, Jacob S., and Paul Pierson. 2005. *Off Center: The Republican Revolution and the Erosion of American Democracy*. New Haven: Yale University Press.
- Hall, Peter A. 1993. "Policy Paradigms, Social Learning, and the State." *Comparative Politics* 25 (3):275-96.
- . 2003. "Aligning Ontology and Methodology in Comparative Politics." In *Comparative Historical Analysis in the Social Sciences*, ed. J. Mahoney and D. Rueschemeyer. New York: Cambridge University Press.
- Jacobs, Alan M. 2011. *Governing for the Long Term: Democracy and the Politics of Investment*. New York: Cambridge University Press.

- Johnston, Alastair Iain. 1996. "Learning Versus Adaptation: Explaining Change in Chinese Arms Control Policy in the 1980s and 1990s." *The China Journal* (35):27-61.
- Khong, Yuen Foong. 1992. *Analogies at War: Korea, Munich, Dien Bien Phu, and the Vietnam Decisions of 1965*. Princeton, N.J.: Princeton University Press.
- King, Gary, Robert O. Keohane, and Sidney Verba. 1994. *Designing Social Inquiry: Scientific Inference in Qualitative Research*. Princeton: Princeton University Press.
- Levy, Jack S. 1994. "Learning and foreign policy: sweeping a conceptual minefield." *International Organization* 48 (2):279-312.
- Mill, John Stuart. 1868. *A System of Logic*. London: Longmans.
- Nickerson, Raymond S. 1998. "Confirmation bias: a ubiquitous phenomenon in many guises." *Review of General Psychology* 2 (2):175-220.
- Parsons, Craig. 2002. "Showing ideas as causes: the origins of the European Union." *International Organization* 56 (1):47-84.
- Risse-Kappen, Thomas. 1994. "Ideas do not Float Freely: Transnational Coalitions, Domestic Structures, and the End of the Cold War." *International Organization* 48 (2):185-214.
- Shepsle, Kenneth A. 1985. "Comment." In *Regulatory Policy and the Social Sciences*, ed. R. Noll. Berkeley: University of California Press.
- Trachtenberg, Marc. 2006. *The Craft of International History*. Princeton: Princeton University Press.
- Van Evera, Stephen. 1997. *Guide to methods for students of political science*. Ithaca: Cornell University Press.
- Weir, Margaret. 1989. "Ideas and Politics: The Acceptance of Keynesianism in Britain and the United States." In *The Political Power of Economic Ideas: Keynesianism Across Nations*, ed. P. A. Hall. Princeton: Princeton University Press.

TABLE 1. STRATEGIES OF PROCESS TRACING IDEATIONAL EFFECTS

Empirical test	Evidentiary task to which it contributes	Assumptions or limitations
Analyzing (mostly private) communication	<ul style="list-style-type: none"> - Measurement of independent variable, by observing statements under reduced strategic pressure - Establishing causal mechanism: application of ideas to decision 	<ul style="list-style-type: none"> - Requires relatively complete deliberative record - Must take into account internal (e.g., intra-governmental) strategic motives for persuasion - More decisive when specific psychological mechanisms theorized and evidence of those specific mechanisms sought
<p>Examining covariation over time</p> <p><i>Analyzing ideational stability and change</i></p> <p><i>Examining within-case sequences</i></p> <p>Examining within-case cross-sectional covariation</p>	<ul style="list-style-type: none"> - Reducing multicollinearity: materialist factors vary over time while ideational measures and outcomes remain constant; or ideas and outcomes change as predicted by learning-based theory - Establishing exogeneity of ideas - Measurement of independent variable: stable beliefs more likely to be sincere than rapidly changing ones - Reducing multicollinearity: temporally separates out potential influences - Reducing multicollinearity: tests for within-case covariation between actors' issue positions and their exposure to ideas and material incentives 	<ul style="list-style-type: none"> - Requires establishing wide variation in material incentives; greater material change → greater sufficiency - Requires theoretical specificity about conditions under which ideas should persist or change - Learning-based theory will have less certain and unique predictions than theory of ideational persistence - Must rule out strong role of strategic anticipation of preferences of actors involved later in process - Must rule out correlational confounds (omitted variables, etc.)

Tracing ideational diffusion		
<i>Identifying ideational origins</i>	<ul style="list-style-type: none"> - Establishing exogeneity of ideas - Measurement of independent variable: helps validate communication as an unbiased measure of sincere ideas 	<ul style="list-style-type: none"> - A “hoop test” for ideational theory, but low in uniqueness - Must rule out strategic “cherry-picking” of ideas - Test more applicable to shared cognitions than individuals’ idiosyncratic beliefs
<i>Tracing paths of ideational transmission</i>	<ul style="list-style-type: none"> - Measurement of independent variable - Establishing exogeneity of ideas - Reducing multicollinearity: variation in transmission over time tests for causal necessity of ideas 	<ul style="list-style-type: none"> - Generally, a “hoop test”: ideas must be available to actors hypothesized to have used them - May also be satisfied with evidence of “mobile carriers”
<i>Identifying mobile “carriers”</i>	<ul style="list-style-type: none"> - Measurement of independent variable - Establishing exogeneity of ideas - Reducing multicollinearity: can test for covariation between policy and mobile carriers entry/exit 	<ul style="list-style-type: none"> - Carriers’ ideational commitments must be readily identifiable (e.g., based on professional affiliation) - Carriers must move into key loci of authority - Must rule out strategic <i>selection</i> of carriers by political principals at time of outcome to be explained
Unpacking the substance of decision outputs	<ul style="list-style-type: none"> - Establishing causal mechanism: costs and benefits of decision outputs, and their visibility, helps discriminate between strategic and ideational motives 	<ul style="list-style-type: none"> - Must take into account collective nature of decision making and potentially diverse motives