

Moreover, process predictions are often unique--i.e., no other known theories predict the same patterns--hence process tracing often offers strong tests of a theory. For example, predictions about the motives that actors should reveal in private discussion are often quite singular, hence confirmation of these predictions strongly corroborates the theory.²⁹ If it supplies abundant and reliable data that bears on unique process predictions of this sort, a single case can provide a very strong test of a theory. As noted above, the investigator will still be unsure what antecedent conditions the theory may require to operate, and discovering these conditions remains an important task. They can be found only by exploring other cases. However, the validity of the theory, and its ability to explain at least one case, are strongly confirmed.

Process tracing is a close cousin of congruence procedure type 1. A process trace analysis starts by dividing the theory's explanation into many small explanatory hypotheses. It then explores the congruence between predicted and observed values on the independent and dependent variables of each of these explanatory hypotheses. Thus process tracing could be thought of as a series of sequential type 1 congruence procedures. One key difference lies in the point of comparison for assessing congruence with predictions: in a process trace observations are most often compared to previous within-case values on the study variable of the moment, while type 1 congruence procedures compare values observed in the case to values in typical external cases (as noted above). Another difference lies in the elements of theory that each method best illuminates.³⁰ Congruence procedures (both types) are often better at testing a theory's prime hypothesis, i.e., its hypothesized relationship between the IV and DV. Process tracing usually sheds relatively more light on the more fine-grained hypotheses that comprise a theory's explanation. In so doing, however, process tracing also assesses the prime hypothesis, since a proof of the elements of the theory adds up to a proof of the prime hypothesis, and disproof of elements can cast doubt on the prime hypothesis.

²⁹In other words, process tracing often provides "smoking gun" tests--see this memo, below, section VII.

³⁰Of course there is no need to choose between these two procedures, since both can be done at the same time on the same case.

How strong are the theory-tests that case studies pose? Scientists tested Einstein's general theory of relativity with what was, in essence, a single real-time congruence procedure type 1 "case study": the observation of the May 29, 1919 solar eclipse. The passage of this single test brought the theory wide acceptance (although it was re-tested many times) because the tested prediction was unique--there was no plausible competing explanation for the result that the theory predicted and astronomers observed, hence the passed test was very strong.³¹ Any case study that tests equally unique predictions can offer equally decisive results. Social science case studies will seldom be decisive, but this problem stems from the messy nature of social science data and the complexity of social phenomena, not the inherent weakness of the case method.

III. CREATING THEORIES WITH CASE STUDIES

Case studies can serve five (5) main purposes: testing theories, creating theories, identifying antecedent conditions, testing the importance of these antecedent conditions, and explaining cases of intrinsic importance.³² The previous section discussed theory testing. This section covers theory-making,³³ and the last three functions are covered in the following three sections.

To infer new theories from cases we start by searching cases for associations between phenomena, and for testimony by people who directly experienced the case (e.g., a participant in the case) on their motives and beliefs about the case. These associations and participant accounts offer clues on cause and effect. Then we ask: "of what more general phenomena are these specific causes and effects examples?" Once candidate causes and effects are framed in general terms the investigator has theories that can be tested against other cases and applied to other cases.

³¹A synopsis of these events is Albert Einstein, *Relativity: The Special and the General Theory: A Popular Exposition*, trans. Robert W. Lawson (New York: Crown Publishers, 1961), pp. 123-132. Einstein's theory predicted that gravity would deflect the path of light toward a gravity source by a specific amount; hence it predicted that stars lying behind the sun would be next to it during a solar eclipse, and predicted where they would appear to be. No other theory made this prediction. Hence the May 29, 1919 eclipse, which generally confirmed Einstein's prediction, won the theory broad acceptance.

Scientists conducted this case study in real time, studying the eclipse as it occurred, but they could just as well have studied a past eclipse if one were photographed showing the positions of nearby stars during the eclipse.

Another decisive real-time case study from physics is found in observation of the return of Halley's comet in 1758-59, which powerfully supported Isaac Newton's theory of gravity. In 1705 astronomer Edmond Halley applied Newton's theory to predict the comet's return in 1758-59; it appeared almost exactly on schedule. Donald K. Yeomans, *Halley's Comet: The Chronological History of Observation, Science, Myth, and Folklore* (New York: John Wiley, 1991), pp. 118-119. Opponents of Newton's theory "were silenced" when this and another test, the 1737 measurement of the flattening of the earth, "confirmed Newton's theory beyond a shadow of a doubt." J. Lévy, "The Solar System," in René Taton, *The Beginnings of Modern Science: From 1450 to 1800* (New York: Basic Books, 1964), pp. 437-455 at 438.

³²These purposes overlap and several (e.g., explaining cases, creating theories and testing theories; identifying antecedent conditions) often can and should be pursued simultaneously.

For another typology of case-study formats see Lijphart, "Comparative Politics and the Comparative Method," which distinguishes six types of case study (p. 691): (1) atheoretical; (2) interpretive; (3) hypothesis-generating; (4) theory testing; (5) theory infirming; (6) deviant. Five of Lijphart's categories overlap with three mine: his study types #1, #2, and #3 are types of theory-testing study, his type #2 is a case-explaining study, and his types #3 and #6 are types of theory infirming study. However, Lijphart's first study-type is descriptive history, a study-type I do not consider here. On the other hand he omits antecedent-condition-identifying and condition-testing studies, which I consider below.

³³Theory-making case studies are also called "heuristic," "hypothesis-generating," and "exploratory" case studies. Eckstein, "Case Study and Theory," pp. 104-108; Lijphart, "Comparative Politics and the Comparative Method," pp. 1, 3-4.

Four basic methods are used to infer theories from case studies: controlled comparison, congruence procedure, process tracing, and Delphi method. Controlled comparison uses across-cases comparative observations to infer theories. Congruence procedure and process tracing deduce theories from within-case observations. The Delphi method consults the views of case participants.

- A. **Controlled comparison.** The investigator infers hypotheses from contrasts or similarities in aspects of several cases. This is done by adapting Mill's method of difference and method of agreement to the task of theory-making. In the method of difference the investigator explores several cases with similar characteristics and different values on the study variable (i.e., the variable whose causes or effects we seek to discover), looking for other differences between the cases. These other cross-case differences are nominated as possible causes of the study variable (if we seek to discover its causes) or possible effects (if we seek its effects). Similar cases are picked to reduce the number of candidate causes or effects that emerge: the more similar the cases, the fewer the candidate-causes. This makes the real cause easier to spot.³⁴ Likewise, in the method of agreement the analyst explores cases with different characteristics and similar values on the study variable, looking for other similarities between the cases. These similarities are nominated as candidate causes or effects of the variable.

The method of difference is preferred when the characteristics of available cases are homogeneous (i.e., most things about most cases are quite similar). The method of agreement is preferred when the characteristics of available cases are heterogeneous (most things about most cases are different.)

- B. **Congruence procedures.** The investigator explores a case looking for within-case correlation between the study variable (usually the dependent variable) and other phenomena. These phenomena are nominated as possible independent variables in new hypotheses. Three specific formats are used.

One method lies in "examining the outliers," i.e., exploring cases that are poorly explained by known causes, on the assumption that unknown causes explain their outcomes. Specifically, we look for cases where a phenomenon is present but its known causes are absent. Still-undiscovered causes must explain the phenomenon, and exploring the case may reveal them.³⁵

Second, the investigator selects cases with extreme high or low values on the dependent variable, and explores them looking for phenomena associated with the DV. When the DV is present in abundance its cause should also be present in unusual abundance and so should stand out against the case background. When the DV is absent its cause should also be prominent by its absence.

³⁴For examples of method-of-difference theory making, see Memo 1, "Hypotheses, Laws and Theories: A User's Guide," in this working paper, section IV, note 24.

³⁵Studies of outlier cases are also known as "deviant" case studies. Lijphart, "Comparative Politics and the Comparative Method," pp. 692-693. The logic of studying outliers follows the logic of John Stuart Mill's "method of residues"; on this method see Mill, *A System of Logic*, pp. 397-398.

Third, the investigator selects cases with extreme within-case variance on the dependent variable, and explores them looking for phenomena that co-vary with the DV. If values on the DV vary sharply its cause should also be prominent by its absence, standing out against the more static case background.³⁶

- C. **Process-tracing.** The investigator traces backwards the causal process from the case outcome was produced, at each stage inferring from the case what caused the cause. If this backward process-trace succeeds it leads back to the cause.
- D. **Delphi method.** The views of case participants or others who experienced the case are mined for hypotheses. Those who experience a case often have important unrecorded data that is lost to later investigators. Their memories and judgments can be the basis for making hypotheses that could not be discovered by direct observation alone.³⁷

IV. INFERRING ANTECEDENT CONDITIONS FROM CASE STUDIES

As noted above, a weakness of single case studies lies in their concealment of 'antecedent conditions,' i.e., background conditions required for theories to be true that magnify their action. However, these background conditions can be uncovered by examining selected new cases.

Four methods of inferring antecedent conditions are most useful. (These methods include all the four methods of inferring theories, outlined above.)

- A. **Controlled comparison.** The investigator infers antecedent conditions by looking for contrasts or similarities in the characteristics of several cases by adapting the method of difference to the task.³⁸ Specifically, we select and examine cases that resemble previously-studied cases except in their value on the independent variable. For example, if previously examined cases had high values on the IV and the DV, we would now examine cases with high values on the IV and low values on the DV, and a close resemblance to the previously examined cases in all other regards. Thus, if the hypothesis that "elite belief that conquest is necessary for war" has been tested using Europe 1900-1914 as a case (high values on the IV and DV), we would next look for other cases when conquest was believed necessary for war resulted, and which otherwise resembled pre-1914 Europe. If we

³⁶This discussion of congruence procedure assumes we are looking for the causes of a study variable. If we are looking for antecedent conditions, the procedure is different. First, we select cases with extreme high or low values on the independent variable, and explore them looking for phenomena associated with the IV. When the IV is present in abundance its effect should be present in unusual abundance and should stand out against the case background. When the IV is absent its effect should be prominent by its absence.

Second, we select cases with extreme variance on the independent variable, and explore them looking for phenomena that co-vary with the IV. Values on phenomena caused by the IV should vary sharply, standing out against the more static case background.

³⁷For an example of Delphi-method theory making see Memo 1, "Hypotheses, Laws and Theories: A User's Guide," in this working paper, section IV, note 26.

³⁸The method of agreement is too weak to bother with.

high-IV-low-DV cases, this suggests that the conditions required for the theory's operation are abundant, and the theory has broad applicability (or "external validity.") If we find such cases, we inspect them looking for points of difference with previously examined cases. Important antecedent conditions will appear as these points of cross-case difference.

- B. **Congruence procedures.** The investigator measures the gap between the predicted and observed values on the dependent variable and then looks for within-case correlation between the size of the gap and values on other phenomena. Phenomena that correlate with the gap (i.e., that are scarce when the value on the DV is lower than the IV value warrants, and are abundant when the DV value is higher than the IV value warrants) are nominated as possible antecedent conditions. Two format are used.

One method lies in examining outliers--i.e., cases where the theory's posited cause is present but its predicted effect is absent. If we assume the theory has validity, this pattern indicates that an important antecedent condition is also notably absent. The missing antecedent condition can be identified among conditions that are often present but are absent in the outlier.

Second, we explore cases with large within-case variance in the value on the dependent variable and constant high values on the independent variable. This pattern suggests that an important antecedent condition varies within the case. It should announce itself as a factor that co-varies with the dependent variable.

- C. **Process-tracing.** The investigator traces backwards the causal process by which the case outcome was produced, at each stage asking if we can infer from the context what antecedent conditions the process requires.
- D. **Delphi method.** The views of case participants or others who experienced the case are mined for possible antecedent conditions. They may have observed in person telltale dynamics that cannot be observed from a distance.

V. TESTING ANTECEDENT CONDITIONS WITH CASE STUDIES

Antecedent conditions, like hypotheses, should be tested before they are lent credibility. Like hypotheses they can be tested three ways: controlled comparison, congruence procedure, and process tracing.

- A. **Controlled Comparison.** The investigator explores paired observations in two or more cases, asking if values on the pairs are congruent or incongruent with the premise that the antecedent condition contributes to causing the dependent variable. For example, if values on the condition variable are higher in case 'A' than case 'B', values on the dependent variable should also be higher, relative to values on the independent variable, in case 'A' than 'B'. If possible, cases are selected according to an adaption of Mill's criteria for the method of difference: cases should have similar general characteristics, similar values on the IV, and different outcomes. If the CV has an impact its values should co-vary with values on the DV.

- B. **Congruence procedure.** Two congruence procedures are most useful for antecedent conditions. First, we study cases with extreme (high or low) values on the condition variable (CV) and a value greater than zero on the independent variable. A very high value on the CV should multiply the effects of the IV on the DV, moving their values above predicted ranges (with "predicted" meaning the value predicted by the test theory in light of the value on the IV in the case). A very low value on the CV should diminish the IV's impact on the DV, lowering their values below predicted ranges. Second, we study cases with large within-case variance on the value of the CV and little or no within-case variance on the IV. If the CV is important the DV's value should co-vary with the IV's value.
- C. **Process tracing.** The investigator explores the chain of events or the causal making process by which initial case conditions are translated into case outcomes. Antecedent conditions will leave footprints in this process: antecedent conditions refer to their importance and events will occur in a sequence that follows the appearance and disappearance.

VI. EXPLAINING CASES⁴⁰

As noted in Memo 1,⁴¹ explanations for specific cases are assessed by answering the following questions:

(1) Does the explanation exemplify a valid general theory (i.e., a covering law)? A specific explanation must exemplify a valid covering law. An explanation that does not exemplify a covering law is a false general theory falls.

(2) Is the covering law's causal phenomenon present in the case? The causal phenomena that the specific explanation posits must be present in the case. If not explained, the case falls. (Even if 'A' is a confirmed cause of 'B', it cannot explain instances of 'B' when 'A' is absent.)

(3) Are the covering law's antecedent conditions met in the case? Theories that do not explain the outcomes of cases that omit their necessary antecedent conditions.

(4) Are the covering law's intervening phenomena observed in the case? The causal phenomena that link the covering law's posited cause and effect should be evident in the case and appear in the proper order.

The logic of case-explaining parallels that of the pathologist doing an autopsy or a crime-solving detective. Specific explanations of the death (or crime) are evaluated by asking if they rest on a valid covering law, if the conditions for that covering law's operation--cause and required antecedent conditions--are observed in the case at hand, and if the causal phenomena that signal its inner workings are also observed. A case-explaining in itself does not test theories, although the evidence collected could be used to check their validity.

³⁹These two methods observe the methods of congruence procedure type 1 (comparison to typical values) and type 2 (multiple within-case comparisons, respectively). See section II B, this memo, above.

⁴⁰Case-explaining studies are also called "explanatory," "interpretive," and "disciplined-configurative" case studies. See *Case Study Research*, p. 5; Lijphart, "Comparative Politics and the Comparative Method," p. 692; Eckstein, "Case Study and Theory," pp. 99-104.

⁴¹Memo 1, "Hypotheses, Laws and Theories: A User's Guide," in this working paper, section VII.

Political scientists seldom do case-explaining case studies, partly because they define the task of case-explaining as the domain of historians. However, historians often explain cases in a softer way than political scientists would. Their explanations are left vague and the predictions they infer from these explanations are left unspecified, hence the meaning of their evidence is often ambiguous. General theories that underlie their explanations are deeply buried. As a result their explanations are hard to interpret and evaluate. Hence cases often remain poorly explained unless political scientists do the job.

VII. STRONG/WEAK TESTS; PREDICTIONS AND TESTS

Strong tests are better than weak tests, and the results of strong tests carry more weight than the results of weak tests.

A strong test is one whose outcome is unlikely to result from any factors except the operation or failure of the theory. Strong tests evaluate predictions that are *certain* and *unique*. A *certain* prediction is an unequivocal forecast. The more certain the prediction, the stronger the test. A *unique* prediction is a forecast not made by other known theories. The more unique the prediction, the stronger the test.⁴²

Four types of tests can be distinguished:

Hoop tests. Predictions of high certitude and no uniqueness provide decisive negative tests: a flunked test kills a theory or explanation, but a passed test gives it little support. E.g. "Was the accused in the state on the day of the murder?" If not, she's innocent, but showing that she was in town does not prove her guilty. To remain viable the theory must jump through the hoop this test presents, but passage of the test still leaves the theory in limbo.

Smoking gun tests. Predictions of high uniqueness and no certitude provide decisive positive tests: passage strongly corroborates the explanation, but a flunk infirms it very little. For example, a smoking gun seen in a suspect's hand moments after a shooting is quite conclusive proof of guilt, but suspects not seen with a smoking gun are not proven innocent. An explanation passing a "smoking gun" test of this sort is strongly corroborated, but little doubt is cast on explanations that fail the test.

Doubly-decisive tests. Predictions of high uniqueness and high certitude provide tests that are decisive both ways: passage strongly corroborates an explanation, a flunk kills it. If a bank security camera records the faces of bank robbers, its film is decisive both ways—it proves suspects guilty or innocent. Such a test combines a "hoop test" and "smoking gun" test in a single study. Such tests convey the most information (one test settles the matter) but are rare.

Straw-in-the-wind tests. Most predictions have low uniqueness and low certitude, and hence provide tests that are indecisive both ways: passed and flunked tests both provide straws in the wind that weigh in the total balance of evidence but are themselves indecisive. Thus many explanations for historical events make probabilistic predictions

⁴²On strong vs. weak tests and unique vs. certain predictions see Memo 1, "Hypotheses, Laws and Theories: A User's Guide," in this working paper, section VI 9.

("if Hitler ordered the holocaust, we should probably find some written records"⁴³), whose failure may simply reflect the downside probabilities. We learn nothing by testing such straw-in-the-wind predictions, but such tests are never done on themselves.⁴⁴ Unfortunately, this describes the predictions we usually work with.

This discussion highlights the need to assess the uniqueness and certitude of predictions when interpreting case study evidence. All evidence is not equal because the predictions they test are not equally unique or certain. Hence authors should consider the uniqueness and certitude of their predictions.

Interpretive disputes often arise from disputes over what outcomes theories predict and over the certitude of these predictions. Does Realism make predictions that are contradicted by the end of the Cold War? Some scholars say yes, others say no. Disputes can be narrowed by framing theories clearly to begin with, since vague and rhetorical statements leave more room for divergent predictions.

Interpretive disputes also arise from quarrels over the uniqueness of predictions. Do other theories or explanations predict the same result? If so, the passing of a test is less impressive. The Fischer school argues that the December 8, 1912 German war cabinet reveals a plot among the German elite to instigate a major war. Some critics argue that the Kaiser's mercurial personality explains his bellicose talk at that meeting—he came off steam by saying things he did not mean. In short, they point to a competing explanation for events that some Fischerites claimed was a "smoking gun" for their elite conspiracy of the war. The question then rides on the plausibility of this competing explanation.

This underlines the need to explain and justify the predictions that case studies make. Debates over predictions often divide analysts. Hence authors should surfacing these debates. They can be surfaced by explicitly framing and explaining predictions.

Can the investigator test a theory with the same case from which it was inferred? In practice is criticized on grounds that such tests lack integrity. The criticism rests on the evidence for blind testing. The assumption is that data not used to infer a theory is more known to an investigator than used data, hence the investigator using unused data is tempted to sample the data selectively.

However, prohibiting the re-use of theory-inspiring cases for theory testing is not possible in practice. (How would this rule be policed?) Moreover, important evidence can be lost by excluding theory-inspiring cases from the testing database, since such cases inspire theories are often among the best testing laboratories. Hence we need other ways to guard against test-fudging. Infusing social science professions with high standards of honesty is the best solution.⁴⁵

⁴³In fact there is no written record of an order from Hitler mandating the holocaust, yet historians agree that Hitler gave it. A discussion is Sebastian Haffner, *The Meaning of Hitler*, trans. Ewald Osers (Cambridge: Harvard University Press, 1979), pp. 133, 138-143.

⁴⁴These last four terms—"hoop test," "smoking gun test," "doubly-decisive test" and "straw-in-the-wind test"—are used to fill gaps in the lexicon.

⁴⁵For more on these arguments see Memo 1, "Hypotheses, Laws and Theories: A User's Guide," in this working paper, section VIII 3.

VIII. INTERPRETING CONTRADICTIONARY RESULTS

What should investigators do when tests produce contrary results--i.e., theories pass some tests and flunk others? Answer: investigate further. Five procedures are appropriate:

- A. Infer and test additional predictions, with a special eye toward finding "hoop" and "smoking gun" tests. Such additional tests may resolve the confusion.
- B. Double-check the accuracy of data used for past tests. Some may be wrong. If so, an unambiguous result may emerge from double-checking: all tests may now be passed or flunked.
- C. Reconsider the predictions you inferred from the theory. Were they fair? Sometimes false flunks (or false passes) are reported because false predictions are tested.
- D. Replicate your tests using new cases. Replication may produce more consistent results.
- E. Repair the theory in ways that enable it to pass flunked tests, e.g., by limiting the scope of its claims or removing flunked explanatory hypotheses. This can salvage a damaged theory (although the salvaged product is now a different theory.)

IX. CASE SELECTION CRITERIA

Practitioners of case studies have not produced a comprehensive catalogue of possible case study research designs,⁴⁶ or a comprehensive list of case selection methods. Accordingly, I have made my own list of useful case selection criteria.⁴⁷ My list (of nine criteria) does not exhaust the logical possibilities, but includes all possibilities that seem strong to me. Specifically, I recommend cases be selected for the following attributes: (1) data richness; (2) extreme values on the independent variable, dependent variable, or condition variable; (3) large within-case variance in values on the IV, DV, or CV; (4) divergence of predictions made of the case by competing theories; (5) case resemblance to current policy problems; (6) ease of controlled comparison between cases (mainly using Mill's method of difference); (7) outlier character; (8) intrinsic importance; (9) appropriateness for replication of previous tests.

This list reflects two general criteria for case selection:

>> Investigators should select cases that best serve the purpose of their inquiry. As noted above, five purposes for case studies can be distinguished: testing theories, creating theories, identifying possible antecedent conditions that theories require to operate, testing the importance of these antecedent conditions, and explaining cases of intrinsic importance. The selection criteria that is most appropriate differs from purpose to purpose, hence investigators should be clear in their purpose before they select cases.⁴⁸ Some of the

⁴⁶Noting this failure is Yin, *Case Study Research*, p. 18.

⁴⁷These criteria evolved from discussions with Jack Snyder, Steve Walt, Tom Christensen and Chaim Kaufmann and include their ideas.

⁴⁸This means you can not know what cases are best to select until you frame your questions. Decisions on case selection are premature before you know what you want to know.

following selection criteria are appropriate for most purposes, but some serve two purposes. Hence investigators should take care to match criteria and purpose. Table 1, below, for a summary of matches and mismatches between mission selection criteria.)

Case selection criteria should therefore differ with the stage at which the investigator stands. Investigators first seek to infer theories, then to test theories, then to test range (or "external validity") by inferring and testing antecedent conditions. Case selection vary across these tasks, and hence vary with the stage of the inquiry.

>> When testing theories investigators should select cases to maximize the range and number of tests they let the investigator perform. The best case selection is the most strong tests (i.e., tests of predictions that are certain and/or unique) with the least research effort.⁴⁹

1. **Select data-rich cases.** We learn more from case studies that let us answer more questions about the case. The more data we have the more questions we can answer. Hence more tests are possible, hence data-rich cases are preferred *ceteris paribus*.⁵⁰

Selecting cases for data-richness is especially appropriate if you plan to test theories using process-tracing, since process-tracing has large data requirements.

Data richness can take several forms. Abundant archival data may be available. Participants in the case may be alive and available for interviews. Other researchers may have studied the case for their own purposes and done the legwork for you. Etc.

2. **Select cases with extreme (high or low) values on the independent variable (IV), the dependent variable (DV), or the condition variable (CV).**⁵¹ In doubt, select cases in which the study variable (i.e., the variable whose effects are being studied) is present in unusually large quantities or unusually small quantities.⁵²
 - a. To test a theory, select cases with extreme values on the independent variable. Such cases offer strong tests because the theory's predictions about them are certain and unique (as noted above in this memo, section II B).

⁴⁹What follows in part repeats remarks made earlier (in sections II, III, IV, and V) on inferring and testing theories and antecedent conditions, since methods of case selection are an aspect of general methods of inquiry.

⁵⁰Concurring is Yin, *Case Study Research*, p. 40.

⁵¹Recommending this selection criteria is Eckstein, "Case Study and Theory," pp. 119-120.

⁵²This "extreme value" method of case selection is closely akin to method #6, controlled comparison. The difference is that with method #6 cases are selected to ease explicit cross-case comparisons, while here cases are selected to ease comparisons to normal conditions. Such comparisons are clearest if within-case values on IV and DV contrast with their normal values.

It is often argued that one should select cases that are representative or typical of the universe of cases. The "extreme value on the IV" method of case selection argues the opposite, that cases that are atypical in their endowment with the independent variable teach us the most.⁵³

Some also argue that selecting cases for extreme values on the IV sets up weak tests because passage of the test is likely: high IV values should produce some DV even if the theory operates only weakly, hence the test is easy to pass. However, this view rests on a mis-definition of "strong test." A strong test is one whose outcome is unlikely to result from any factor other than the operation or failure of the theory. Using this definition a test using a case selected for extreme value on the IV is a strong test. We should expect extreme results in such a test.⁵⁴ If they occur, these extreme results are unlikely to stem from other factors. If they do not occur, this is unlikely to stem from any cause other than the theory's failure. Hence cases with extreme IV values are laboratories for strong tests.

- b. To make a theory, select cases with extreme values on the dependent variable. If the DV is present in abundance its cause should also be present in unusual abundance, hence the cause should stand out against the background of the case more clearly. This makes it easier to spot. Likewise, if the DV is unusually scarce its causes should be made more striking by their greater absence.
- c. To infer antecedent conditions, select cases with extreme and opposite values on the IV and DV--specifically, with very high values on IVs and very low values on DVs. These are cases where the theory's posited cause is abundantly present but the predicted effect is notably absent.⁵⁵ Such a pattern indicates that an important antecedent condition is also notably absent. The missing condition can be identified among conditions that are normally present but are absent in the studied case.⁵⁶

⁵³Thus I chose the 1914 case to test offense-defense theory (which posits that war is more likely when conquest is believed easy) partly because 1914 was the heyday of the "cult of the offensive," a remarkable European elite belief that conquest was easy. Shortly before 1914 this belief reached heights never seen before or since. Because it was extreme it should have had extreme effects, if it ever has any effects. Hence these effects should have been clearly visible in 1914, standing out starkly from the European political landscape, and should have appeared in quantities unlikely to be produced by measurement error or by the action of other causes.

The predictions of offense-defense theory in 1914 therefore are both certain and unique. The predicted effects are too large to ascribe to measurement error or to other causes (hence the predictions are unique). The absence of these effects likewise could not be plausibly blamed on measurement error or the overriding effects of other phenomena, since these could not mask or override such large effects (hence the predictions are certain). Hence the test posed by the 1914 case is very strong. Offense-defense theory would be strongly confirmed by passing it, and badly damaged by flunking it. See Van Evera, *Causes of War*, vol. 1., chapter 7.

⁵⁴I assume there is no threshold effect, i.e. the impact of the IV on the DV does not peter out above a certain threshold. With threshold effects this selection criteria is less useful.

⁵⁵For example, to infer conditions required for economic depression to cause war, we should select cases where deep depressions occurred but no war resulted. To infer conditions required for literacy to cause democracy we should select highly literate societies with authoritarian regimes. Etc.

⁵⁶Under selection method IX 2a and IX 2b, above, cases are chosen to highlight the contrast between observed and normal values on the IV and DV. Here (method IX 2c), in contrast, cases are chosen to highlight the contrast between observed values on IV and DV. Sharp contrasts are sought in both cases, although the sought contrasts differ in nature.

- d. To test a candidate antecedent condition (i.e., a condition that requires to operate or that magnifies its action), select cases with extreme values on the condition variable (CV). If the CV is very abundant its effect on the DV should be strikingly strong. If the CV is scarce the IV's impact on IntVs and DV should be strikingly weak. In instances where the predicted results are pronounced and hence less likely to be explained away by measurement error or the actions of a third variable.⁵⁷

3. Select cases with large within-case variance in the value on the independent variable, dependent variable, or condition variable across time or space

- a. To test a theory, select cases with large within-case variance in the value on the independent variable. Theories make predictions about the DV given the IV, hence variance in the IV's value gives rise to variance in the DV's value. The more within-case variance in the IV's value, the more within-case variance in the DV's value we have to test. Such variance takes the form of diachronic change in the value on the IV, i.e., change over time within the period covered by the case, or diversity on the value of the IV across regions, groups, organizations, or individuals present in the case (i.e., synchronic variance across space).

Selecting cases for within-case IV-value variance is especially appropriate if you plan to use a congruence procedure for testing, since congruence procedures rely on observing within-case variance.

- b. To make a theory, select cases with large within-case variance in the value on the dependent variable. The cause of the DV should also vary within the case. This makes the cause easier to spot against the background. Candidate causes will announce themselves as case characteristics that vary with the DV's value--i.e., as factors present when the DV is high, missing when it is low.
- c. To infer a theory's antecedent conditions, select cases with large within-case variance in the value on the DV, and constant high values on the IV. If a case contains some observations where the relative value on the IV and DV match the theory's predictions (high on IV and DV), and some observations where they do not (high on IV, low on DV). Candidate antecedent conditions will announce themselves as factors that are more abundant when relative values match predictions, i.e., when DV values are higher.

⁵⁷This selection strategy needs adjustment if the impact of antecedent conditions follows sharp thresholds, e.g., a theory that requires some amount of the CV to cause the DV but further increases in the CV have no effects. Thus rain requires some amount of sunshine to cause grass to grow, but moderate sunshine is as good as constant sunshine. In such cases we should select cases for very low values on the CV, and test the prediction that the IV should lack causal impact on the DV in such cases. Selecting cases with very large values on the CV is unfruitful since very large values predict the DV as moderate values.

⁵⁸Obliquely recommending this selection criteria is Eckstein, "Case Study and Theory," pp. 119, 126.

- d. To test a candidate antecedent condition, select cases with large within-case variance on the value of the CV. If the CV is important the DV's value should be higher, relative to the IV's value, when the CV is abundant than when it is scarce.
4. **Select cases about which competing theories make opposite and unique predictions.**

This selection method is appropriate if you are more interested in testing the relative power of the two theories than testing a theory against the null hypothesis (that is, you prefer to arrange a Lakatosian "three cornered fight" over a "two-cornered fight.")⁵⁹

- a. If you are testing the relative power of two theories, choose cases about which they make opposite and unique predictions, e.g., a case with large and opposite within-case variance in the values on the two IVs (e.g., values on one IV fall over time, values on the other rise over time, etc.) The DV should co-vary with the stronger IV.
- b. If you are testing the relative power of two antecedent conditions, choose cases where the IV is present and the CVs show large and opposite within-case variance (e.g., values on one CV fall over time, values on the other rise over time). The DV should co-vary with the stronger CV.
5. **Select cases that resemble current situations of policy concern.**⁶⁰ A theory inferred from or tested in a case that resembles a second case will more often "travel" to that second case--i.e., operate in the second case as well. Hence policy prescriptions deduced from the first case can be more safely applied to the second. Scholars interested in offering policy prescriptions should therefore study cases whose characteristics parallel the characteristics of current or future policy problems.

A study of health policy in Minnesota yields more reliable prescriptions for health policy in Wisconsin than a study of health policy in Burkina Faso. Theories that operate in Burkina Faso may well require conditions absent in Wisconsin, hence prescriptions deduced from these theories will prove unsound in Wisconsin. This is less likely of theories that operate in Minnesota since Minnesota and Wisconsin are similar in many ways.

6. **Select cases that are well-matched for controlled cross-case comparisons.** Cases can be selected to allow their pairing for controlled comparison, i.e., for the method of difference (cases have similar characteristics and different values on the study variable) or the method of agreement (cases have different characteristics and similar values on the study variable). The method of difference, being the stronger of the two, is usually preferred.

Controlled-comparison criteria can be applied to select both single case and multiple cases. A single case can be selected with an eye on comparing it to existing case studies that others have already researched and written about. Specifically, if we plan to perform a method-of-difference comparison, we select a new case whose characteristics resemble those of an already-studied case but has different values on the study variable. Multiple cases can be selected with an eye on comparing them with each other (i.e., we select cases with similar characteristics and diverse study-variable values) or with already-studied cases (i.e., we select cases with characteristics similar to those of an already-studied case but with different values on the study variable).

- a. To test a theory using a controlled comparison method, select cases with similar characteristics and different values on the IV or the DV (i.e., for the method of difference). The theory passes the test if study cases that values on both IV and DV differ across the cases in parallel fashion. For example, e.g., if the IV has a higher value in case 1 than case 2, the DV should have a higher value in case 1 than case 2.

However, the method of difference is a fairly weak instrument for testing (and the method of agreement is even weaker.) Hence other selection criteria should have higher priority for theory-testers.

- b. To make a theory, select cases with similar characteristics and different values on the study variable (i.e., for method-of-difference comparison) or cases with different characteristics and similar values on the study variable (i.e., for method-of-agreement comparison).

Candidate causes or effects (i.e., potential independent or dependent variables) will announce themselves as differences in the characteristics of compared cases when the method of difference is used. They announce themselves as similarities in the characteristics of compared cases when the method of agreement is used.

The method of difference is preferred when the characteristics of compared cases are quite homogeneous (most things about most cases are similar). The method of agreement is preferred when the characteristics of compared cases are quite heterogeneous (most things about most cases are different).

- c. To infer antecedent conditions, select cases using variants of the method of difference or the method of agreement.

For the method of difference, choose cases with: (1) similar values on the IV; (2) similar case characteristics; (3) different values on the DV. Candidate antecedent conditions will announce themselves as differences in the characteristics of the compared cases.

For the method of agreement, choose cases with: (1) similar values on the IV; (2) different case characteristics; (3) similar values on the DV. Candidate antecedent conditions will announce themselves as similarities in the characteristics of the compared cases.

⁵⁹On two- and three-cornered fights see Memo 1, "Hypotheses, Laws and Theories: A User's Guide," in this working paper, section VI 8, note 36.

⁶⁰Jack Snyder recommends this criteria.

- d. When testing a candidate antecedent condition, select cases with: (1) similar values on the IV; (2) different values on the DV. The condition passes its test if values on the CV correspond with values on the DV across cases.
7. **Select outlier cases.** Cases that are poorly explained by existing theories are selected for study, on the assumption that unknown causes explain their outcomes, and that these causes might be identified by examining the case, giving rise to a new theory. To do this we select cases that lie far from the regression line that expresses the relationship between the dependent variable and its known causes (i.e., "outlier" cases.) Candidate new causes will announce themselves as unusual characteristics of these cases, especially those that are associated with the dependent variable within the case.
- a. To make a theory, select cases where the DV's known causes are scarce yet the DV is abundantly present. This suggests that unknown causes are operating in the case, and that study of the case may reveal them.
- b. To infer an antecedent condition, select cases where the DV's known causes are abundant yet the DV is scarce or absent. This suggests that unknown antecedent conditions are absent in the case, and that study of the case may identify them.
8. **Select cases of intrinsic importance.** Selecting cases of intrinsic human or historical importance (e.g., World War I, World War II, the Holocaust) is appropriate if our object is to explain particular human experiences. We select such cases with a nod to their data-richness (there is little point in studying cases where the record is too thin to answer our questions), but mainly in accord with the magnitude of their human consequences.
9. **Select for test replication.** Thorough theory testing requires repeating initial tests to corroborate their results. When doing this, cases are chosen for their appropriateness as laboratories to replicate previous tests.⁶¹ This approach considers multiple cases as multiple experiments. Test replication, not cross-case comparison, is the goal of later studies in the series.

Previous tests can be repeated exactly with new cases (e.g., a theory previously tested by congruence procedure can again be tested by a congruence procedure) or new test formats can be adopted, using the same case or a new one (e.g., a theory previously tested by congruence procedure can now be tested by process tracing, using the same case or a new one.) All of the first six test criteria listed above (criteria IX 1-IX 6, above) are used to select cases. (Selection criteria IX 7 and IX 8 are not followed because they are inappropriate for selecting cases for theory-testing.)

What selection criteria should we use when we select multiple cases? We have two prime choices: to select for controlled comparison or for replication. Since controlled comparison is a weak method of testing, selecting for replication is usually a better choice. (Note that selection criteria IX 1-IX 5, though designed for selecting single cases for studies, become appropriate for selecting multiple cases under this criteria).

⁶¹Making this point is Yin, *Case Study Research*, pp. 45-50.

Table 1, below, summarizes matches and mismatches of study missions selection criteria.

TABLE 1
NINE CASE SELECTION CRITERIA:
WHEN IS EACH APPROPRIATE?

Case selection criteria	Appropriate when testing theories?	Appropriate when inferring theories?	Appropriate when inferring antecedent conditions?	Appropriate when testing antecedent conditions?
1. Data richness	Yes	Yes	Yes	Yes
2. Extreme values on IV, DV, CV	Yes (on IV)	Yes (on DV)	Yes (high on IV, low on DV)	Yes (on CV)
3. Large within-case variance on values of IV, DV, CV	Yes (on IV)	Yes (on DV)	Yes (on DV)	Yes (on CV)
4. Competing theories make divergent predictions about the case	Yes	No	No	Yes
5. Resemblance to current policy-problem cases	Yes	Yes	Yes	Yes
6. Matched for cross-case controlled comparison (i.e., method of difference or agreement)	Seldom	Yes	Yes	Seldom
7. Outcome unexplained by other theories (i.e. an "outlier" case)	No	Yes	Yes	No
8. Intrinsic importance	No	No	No	No
9. Good case for replicating previous tests	Yes	Yes	Yes	Yes

MEMO 3:

WHAT IS A POLITICAL SCIENCE PH.D. DISSERTATION?

Ph.D. dissertations in political science¹ can perform seven principal missions. This gives rise to seven types of dissertation, one for each mission. Most dissertations perform several of these missions, and thus are hybrids, but it is still useful to consider possible ideal-type dissertations.

(1) A **theory-proposing** dissertation advances new hypotheses. A deductive argument for these hypotheses is advanced. Examples may be offered to illustrate these hypotheses and to demonstrate their plausibility, but strong empirical tests are not performed.²

(2) A **theory-testing** dissertation uses empirical evidence to evaluate existing theories. This evidence can take the form of large-n analysis, case studies, or both.³

Note: many dissertations are a blend of Type #1 and #2. They do some theory-proposing and some theory-testing.⁴ However, it is legitimate to focus exclusively on proposing theory, or on testing theory, as long as one contributes useful knowledge.

¹What follows is directed at students in the international relations subfield. I apologize to those in other subfields for my IR-centrism.

²Examples of theory-proposing works include Robert Jervis, *Perception and Misperception in International Politics* (Princeton: Princeton U. Press, 1976); Robert Jervis, "Cooperation Under the Security Dilemma," *World Politics*, Vol. 30, No. 2 (January 1978), pp. 167-214; Kenneth N. Waltz, *Theory of International Politics* (Reading, Mass.: Addison-Wesley, 1979); Geoffrey Blainey, *The Causes of War*, 3rd ed. (New York: Free Press, 1988); Thomas C. Schelling, *Arms and Influence* (New Haven: Yale, 1966); Thomas C. Schelling, *The Strategy of Conflict* (New York: Oxford University Press, 1960); Carl von Clausewitz, *On War* (Princeton: Princeton University Press, 1976); Hans J. Morgenthau, *Politics Among Nations*, 5th ed. (New York: Knopf, 1973). Note: hypotheses may be developed by deduction (e.g., Schelling), or by induction (e.g., Clausewitz), or both.

³Examples of works that focus on theory-testing include Richard K. Betts, *Nuclear Blackmail and Nuclear Blackmail* (Washington, D.C.: Brookings, 1987); Steve Chan, "Mirror, Mirror on the Wall ... Are the Freer Countries More Pacific?" *Journal of Conflict Resolution*, Vol. 28, No. 4 (December 1984), pp. 617-648; and Erich Weede, "Democracy and War Involvement," in *ibid.*, pp. 649-664.

⁴Examples of hybrid works that both propose and test theories include Barry R. Posen, *The Sources of Military Doctrine: France, Britain, and Germany Between the World Wars* (Ithaca: Cornell University Press, 1984); Stephen M. Walt, *The Origins of Alliances* (Ithaca: Cornell University Press, 1987); Jack Snyder, *Myths of Empire* (Ithaca: Cornell University Press, 1991); Jack Snyder, *Ideology of the Offensive: Military Decision Making and the Disasters of 1914* (Ithaca: Cornell University Press, 1984); and John J. Mearsheimer, *Conventional Deterrence* (Ithaca: Cornell University Press, 1983). Note: such works often begin as theory-testing projects; the authors begin by testing others' theories, and develop their own theories in midstream. This reflects the great difficulty of creating theory from a standing start. (I advise students not to try it. Instead, test someone else's theory. Creative lightning may strike you while you're at it, but if it doesn't you still can produce a good thesis.)

(3) A **literature-assessing** (or "stock-taking") dissertation summarizes and assesses existing theoretical and empirical literature on a subject. It asks if existing tests are valuable and existing tests are persuasive and complete.⁵

(4) A **policy-evaluative or policy-prescriptive** dissertation evaluates current and future public policies or policy proposals. Are the factual and theoretical premises of the proponents and opponents of proposed policies valid or invalid? Will the policy produce the results that its proponents promise?

It is often said that policy-prescriptive work is not theoretical, but the opposite is true. All policy proposals incorporate theoretical assumptions (i.e., assumptions about social causation). Hence all evaluation of public policy requires the framing and testing of theory, hence it is theoretical, hence it falls under the rubric of political science.

Policy prescriptive work can be focused several ways, for example, on evaluating a particular policy; on evaluating competing solutions to a given problem; or on the broader implications of a political or technical development (e.g., the nuclear revolution, or the collapse of the Soviet empire).

(5) A **historical explanatory** thesis uses theory (academically recognized theory, folk theory, or "common sense" deduction) to explain the causes, pattern, or consequences of historical cases. Such works often provide a good deal of description but focus on explaining what is described.⁷

⁵Examples of literature-assessing works include Kenneth N. Waltz, *Man, the State, and War* (New York: Columbia University Press, 1959); Benjamin Cohen, *The Question of Imperialism* (New York: Basic Books, 1973); Jack Levy, "Causes of War: A Review of Theories and Evidence," in Philip E. Tetlock, Jo L. Huberman, Robert Jervis, Paul A. Huth, and Charles Tilly, eds., *Behavior, Society, and Nuclear War*, 2 vols. (New York: Oxford University Press, 1987), pp. 1:209-333; Robert Gilpin with Jean M. Gilpin, *The Political Economy of International Relations* (Princeton: Princeton University Press, 1987).

⁶Examples of policy prescriptive work include: Jerome Slater, "Dominos in Central America: Will They Fall or Not? Matter?" *International Security*, Vol. 12, No. 2 (Fall 1987), pp. 105-134; Charles L. Glaser, *Analyzing Strategic Policy* (Princeton: Princeton University Press, 1990); Robert Jervis, *The Illogic of American Nuclear Strategy* (Princeton: Princeton University Press, 1984); Shai Feldman, *Israeli Nuclear Deterrence* (New York: Columbia University Press, 1991); Robert Art, "A Defensible Defense: America's Grand Strategy After the Cold War," *International Security*, Vol. 16, No. 1 (Spring 1991), pp. 5-53; Barry Posen, "Inadvertent Nuclear War? Escalation and NATO's Northern Flank," *International Security*, Vol. 7, No. 2 (Fall 1982), pp. 28-54; John J. Mearsheimer, "A Strategic Misstep: The Maritime Security Dilemma in Europe," *International Security*, Vol. 11, No. 2 (Fall 1986), pp. 3-57; Samuel Huntington, "Conventional Deterrence and Conventional Retaliation in Europe," *International Security*, Vol. 8, No. 3 (Winter 1983/84), pp. 405-427; Joshua Epstein, "Soviet Vulnerabilities and the RDF Deterrent," *International Security*, Vol. 6, No. 2 (Fall 1981), pp. 158; Albert Wohlstetter, "The Delicate Balance of Terror," *Foreign Affairs*, Vol. 37, no. 2 (January 1959), pp. 158-164; a Soviet policy is Richard Ned Lebow, "The Soviet Offensive in Europe: The Schlieffen Plan Revisited?" *International Security*, Vol. 9, No. 4 (Spring 1985), pp. 44-78.

⁷Examples of historical explanatory works include James C. Thompson, "How Could Vietnam Happen? An Alternative Perspective," in Morton H. Halperin and Arnold Kanter, eds., *Readings in American Foreign Policy: A Bureaucratic Perspective* (New York: W.W. Norton, 1973), pp. 98-110; Leslie H. Gelb with Richard K. Betts, *The Irony of Vietnam: The System and the Strategy* (New York: W.W. Norton, 1982); John Lewis Gaddis, "The Long Peace: Elements of Stability in the Postwar System," *International Security*, Vol. 10, No. 4 (Spring 1986), pp. 99-142; Arthur Schlesinger, Jr., "Origins of the Vietnam War," *Foreign Affairs*, Vol. 46, No. 4 (October 1967), pp. 22-52; Thomas J. Christensen, *Useful Adversity: Strategy, Domestic Mobilization, and Sino-American Conflict, 1947-58* (Princeton: Princeton University Press, 1988); and John J. Mearsheimer, *Liddell Hart and the Weight of History* (Ithaca: Cornell University Press, 1988); and, on temporary history, Barry R. Posen, "The Security Dilemma and Ethnic Conflict," *Survival*, Vol. 35, No. 1 (January 1993), pp. 27-47.

(6) A **historical evaluative** thesis evaluates the factual and theoretical beliefs that guided official or unofficial policy actors, and/or evaluates the consequences of the policies they pursued.⁸

Note: dissertations of Types #5 and #6 are rare and little admired in political science. This reflects a general bias in the field favoring the creation and testing of theory over the application of theory. However, this bias misguided. If theories are never applied to anything, then what are they for?

Moreover, scholarship of Types #5 and #6 lacks a friendly home in any other discipline, so such work won't be produced unless political scientists do the job. Some historians are averse to explicit explanation, instead preferring to "let the facts speak for themselves." Others will elaborate a preferred explanation, but they rarely compete contending explanations against one another, as one must to fully evaluate an explanation. Historians are also (with some exceptions) generally averse to writing evaluative history. However, without explanatory historical work history is never explained; and without evaluative historical work we learn less from the past about present and future problem-solving. Hence some field should accept these tasks. I nominate political science.

(7) A **predictive** dissertation applies theories to extrapolate the future world from current events or from posited future developments.⁹ A purely predictive dissertation is a risky project, because the future is constantly happening, raising the danger that the project may be overtaken by events. Therefore, students should generally steer clear of dissertations of this sort. However, this warning isn't iron-clad. Predictive work can be valuable, and can take dissertation form.

These seven varieties of thesis can be summarized as falling into four types: theory-proposing (#1); theory-testing (#2); theory-applying (#4, #5, and #6); and literature-assessing (#3).

Dissertations of Type #1 and #2--theory-making and theory-testing--have the most cachet in political science, but all seven types are legitimate if they are well-done. Be clear in your own mind about which type of thesis you are doing.

⁸Examples of historical evaluative works include Bruce M. Russett, *No Clear and Present Danger: A Skeptical View of the U.S. Entry Into World War II* (New York: Harper & Row, 1972); John Mueller, "Pearl Harbor: Military Inconvenience, Political Disaster," *International Security*, Vol. 16, No. 3 (Winter 1991/92), pp. 172-203; Paul M. Kennedy, "Tirpitz, England and the Second Navy Law of 1900: A Strategic Critique," *Militaergeschichtliche Mitteilungen*, Vol. 2 (1970), pp. 33-57; Paul Kennedy, *Strategy and Diplomacy, 1870-1945* (Aylesbury: Fontana, 1983), chapters 5 and 7 ("Strategic Aspects of the Anglo-German Naval Race" and "Japanese Strategic Decisions, 1939-1945"); Gerhard Ritter, *The Schlieffen Plan: Critique of a Myth*, trans. Andrew and Eva Wilson, foreword by B.H. Liddell Hart (London: Oswald Wolff, 1958; reprint ed., Westport, Conn.: Greenwood Press, 1979); Robert W. Tucker and David C. Hendrickson, *The Imperial Temptation: The New World Order and America's Purpose* (New York: Council on Foreign Relations, 1992), part 2 ("The Gulf War: An Autopsy"); Alan T. Nolan, *Lee Considered: General Robert E. Lee and Civil War History* (Chapel Hill: University of North Carolina Press, 1991); Paul W. Schroeder, *The Axis Alliance and Japanese-American Relations, 1941* (Ithaca: Cornell University Press, 1958), chapter 9 ("An Appraisal of American Policy"), pp. 200-216.

⁹Recent examples of predictive work include Robert Jervis, "The Future of World Politics: Will It Resemble the Past?" *International Security*, Vol. 16, No. 3 (Winter 1991/92), pp. 39-73; John J. Mearsheimer, "Back to the Future: Instability in Europe After the Cold War," *International Security*, Vol. 15, No. 1 (Summer 1990), pp. 5-56; Stephen Van Evera, "Primed for Peace: Europe After the Cold War," *International Security*, Vol. 15, No. 3 (Winter 1990-1991), pp. 7-57; Jack Snyder, "Averting Anarchy in the New Europe," *International Security*, Vol. 14, No. 4 (Spring 1990), pp. 5-41. The latter three pieces also offer policy prescriptions but their main thrust is predictive.

Finally, some words on **descriptive** dissertations are in order. Such dissertations describe political circumstances.¹⁰ They come in two types: **contemporary/descriptive** (current developments and conditions are described)¹¹ and **historical descriptive** (events and conditions are described.)¹²

A descriptive thesis is an eighth possible type of Ph.D. thesis. However, purely descriptive theses are poorly received by other political scientists. Political scientists want authors to explain or evaluate the events, policies or ideas they describe. Pure description must be combined with some making, testing, or application of theory. However, description must often precede explanation or evaluation, since phenomena that have not been described cannot be explained or evaluated. Hence students who devote heavy attention to description, giving rise to largely descriptive dissertations. This is fine as long as the student also does some explaining or evaluation.¹³

¹⁰Description establishes data points: explanation explains the structure of data that has already been described. Following statements illustrate the difference: "In January 1991 oil sold for \$40 per barrel on the world market" (pure description); and "in late 1990 the Persian Gulf crisis caused consumer fear that war might disrupt global oil supplies, causing panic oil buying, which pushed up oil prices from under \$20 to \$40" (description and explanation--the price is described and explained).

¹¹Examples of contemporary/descriptive work include articles from the 1980s describing the conventional military in central Europe: see John J. Mearsheimer, "Why the Soviets Can't Win Quickly in Central Europe," *International Security*, Vol. 7, No. 1 (Summer 1982), pp. 3-39; John J. Mearsheimer, "Numbers, Strategy, and the European Balance: Coping With Complexity in Threat Assessment," *International Security*, Vol. 9, No. 3 (Winter 1984-85), pp. 174-185; Barry R. Posen, "Measuring the Conventional Balance: Coping With Complexity in Threat Assessment," *International Security*, Vol. 9, No. 3 (Winter 1984-85), pp. 186-202. Note, however, that these works are not theory-barren. Both authors depend on causal hypotheses--the rule and the force-to-space ratio hypothesis--hence their description is theory-reliant.

Examples on other topics include Steve Fetter, "Ballistic Missiles and Weapons of Mass Destruction: What Is to Be Done?" *International Security*, Vol. 16, No. 1 (Summer 1991), pp. 5-42; Bruce G. Blair, "Command and Control (Washington, D.C.: Brookings, 1985); and Ashton B. Carter, "Assessing Command and Control Vulnerability," in Ashton B. Carter, John D. Steinbruner, and Charles A. Zraket, eds. *Managing Nuclear Vulnerability* (Washington, D.C.: Brookings, 1987), pp. 555-610. (These pieces are not purely descriptive; they also offer prescriptions but their main focus is descriptive.)

¹²Examples of works by political scientists that are largely historical/descriptive include Fred Kaplan, *The Armageddon* (New York: Simon & Schuster, 1983); Scott Sagan, "Nuclear Alerts and Crisis Management," *International Security*, Vol. 9, No. 4 (Spring 1985), pp. 99-139; and Richard K. Betts, *Soldiers, Statesmen, and Cold War Strategy* (Cambridge: Harvard University Press, 1977). Such works by historians include John Lewis Gaddis, *The Cold War: A Critical Appraisal of Postwar American National Security Policy* (New York: Oxford University Press, 1966); David Alan Rosenberg, "The Origins of Overkill: Nuclear Weapons and American Strategy, 1945-1960," *International Security*, Vol. 7, No. 4 (Spring 1983), pp. 3-71; Luigi Albertini, *The Origins of the War of 1914*, 3 vols., trans. Isabella M. Massey (London: Oxford University Press, 1952-57; reprint ed., Westport, Conn.: Greenwood Press, 1966); and Holger Herwig, "Clio Deceived: Patriotic Self-Censorship in Germany After the Great War," *International Security*, Vol. 12, No. 2 (Fall 1987), pp. 5-44. These works all provide some explanation, and several do some theory, but their focus is descriptive.

¹³A good dissertation of this kind is Peter J. Liberman, "Does Conquest Pay? The Exploitation of Occupied Economies" (Ph.D. dissertation, MIT, 1991); later published as *Does Conquest Pay? The Exploitation of Occupied Societies* (Princeton: Princeton University Press, 1996). Liberman devotes substantial space to explaining the phenomenon he describes because previous scholarship left the phenomenon he explains (the benefits of empire) largely undescribed. He proceeds to develop and test explanations for the patterns he describes.

MEMO 4:

HELPFUL HINTS ON WRITING A POLITICAL SCIENCE PH.D. DISSERTATION

I often make the following suggestions to graduate students who are launching Ph.D. dissertations.¹

I. TOPIC SELECTION

A good dissertation asks an important question. The answer should be relevant to real problems facing the real world.

Hans Morgenthau once lamented that social scientists often hide in "the trivial, the formal, the methodological, the purely theoretical, the remotely historical--in short, the politically irrelevant".² Such conduct is both a crime and a blunder. Being relevant is more fun, better for the world, and a good career move. Scholars who advance bold arguments win more praise than abuse if their work is sound. Research gains visibility largely by having college teachers assign it. Teachers assign work that frames debates. Hence work that boldly presents a side in an important debate or starts its own debate will be more widely assigned and thus more renowned. Its author will bask in academic fame and glory.

How can good topics be found? Starting yesterday, keep a "Books and Articles That Someone Should Write" file. When you form a mental picture of something you want to read, but a search reveals that it doesn't exist, record its hypothetical title and stash it in your "Books and Articles" file. Many of these absent articles won't be suitable projects for you, but some will. The rest are possible topics for your friends and future students. You do a major service by devising projects they can execute.

After each graduate school class, write an audit memo about the subject area of the course asking what was missing. What important questions went unasked? What answers did you expect to find in the literature that never appeared? What research projects could provide these answers?

Ph.D. qualifying exams offer another time to audit the field for fillable holes. You have surveyed the field's horizon: now write a memo on questions and answers that turned up missing in the literature and research that could provide the missing answers.

¹Also useful on this topic is Peter A. Hall, "Helpful Hints for Writing Dissertations in Comparative Politics," *PS: Political Science and Politics* (December 1990), pp. 596-598.

²Hans J. Morgenthau, "The Purpose of Political Science," in James C. Charlesworth, ed., *A Design for Political Science: Scope, Objectives, and Methods* (Philadelphia: American Academy of Political and Social Science, 1966), pp. 63-79, at 73. Morgenthau further complained of a "new scholasticism," in academe--the pursuit of "an intellectual exercise ... that tells us nothing we need to know about the real world." Scholars maintain their reputations by "engaging in activities that can have no relevance for the political problems of the day"; instead they substitute a "fanatical devotion to esoteric terminology and mathematical formulas, equations and charts, in order to elucidate or obscure the obvious." As a result, social science resembles "a deaf man answering questions which no one has asked him." -Ibid, p. 74; and Hans J. Morgenthau, *Truth and Power* (New York: Praeger, 1970), pp. 246, 261.

Thesis topics can also be found in public policy debates. First read up on debate you care about. Then identify the key disputes of fact or theory that drive opposing sides to their opposite conclusions. Then devise a research project that addresses one or more of these disputes. This search method locates research questions that are unsolved and germane to important public policy questions.³

II. ORGANIZATION

A good thesis has a thesis--a main line of argument, or a set of related arguments. If your thesis lacks a thesis, think it through again. If your thesis has too many threads, consider ways to organize your ideas more simply.

III. YOUR DISSERTATION PROSPECTUS

Your dissertation prospectus support your applications for research support. Your prospectus should be 5-10 pages long. It should frame the question(s) you address, the reasons why these questions are worth exploring, your working hypothesis (i.e., the answers you expect to find), your methods of inquiry, and the reasons you chose these methods.

You should footnote your prospectus as you would a research paper. Geographic footnotes to existing work on your topic are important.

Before sending it out, circulate your prospectus among friends and colleagues. Comments and criticisms.

For more on the dissertation prospectus, see Memo 5, "The Dissertation Prospectus," next in this working paper.

IV. YOUR INTRODUCTORY CHAPTER

The introduction and conclusion are the most-read parts of most theses and are read parts of many, hence their design merits special attention.

You should start your thesis with a summary introduction chapter. A summary introduction helps readers measure your evidence against your claims and arguments from the outset. This makes your thesis more readable.

Your summary introduction should answer six (6) questions:

1. What question or questions do you address? Spell them out clearly.

A thesis can propose theories, test theories, explain historical events, or evaluate or present policies or policy proposals. It can summarize and assess a body of work. It can describe contemporary circumstances or historical events. It can do several of the above. State clearly which of these missions your thesis fulfills.

³Charles Glaser recommends this search method for research topics.

⁴Purely historical or descriptive theses are exempt from this requirement, but their authors should still identify any themes or structure that emerge in the material they present.

Frame your questions in terms that call for specific answers. Questions that begin "how can we understand" (e.g. "how can we understand the meaning of the nuclear revolution?" or "how can we understand the process by which nationalism arises?") are so open-ended that vacuous non-responses ("we can understand the meaning of the nuclear revolution by reading Bob Jervis") would technically qualify as answers. Focused questions are better: "what are the consequences of the nuclear revolution?", or "what are the causes of nationalism?" Questions that inquire about "cause" or "consequence," or that pose specific descriptive tasks ("how numerous were Stalin's victims?") are better, since readers can more easily tell if you answered them.

2. Why do these questions arise--from what scholarly literature or real-world events? What previous literature has been written on these questions? What is the "state of the art" on the subject?

If your questions arise from an evolving scholarly literature, you should discuss that literature in the text of your introduction, and note ancillary or related literature in footnotes. Note any controversies in this literature, explain their origins and evolution, detail the arguments made by both sides, and summarize their current status. Note the factual or theoretical crux of any continuing disagreements. Note also the holes in the current literature. What questions have *not* been explored? (Hopefully yours is among them.) And, if you can, interpret the motives that sustain continuing controversies. What, if any, political or methodological motives are driving the disputants apart? Are these disputants honest scholars or paid polemicists? In short, explain what's been going on in the field you are entering.

If your questions arise from historical or contemporary events, detail these events, explain their significance, and explain why they give rise to the question or questions you address. Also mention any existing literature on the subject you address, and note holes in that literature.⁵

3. What answer or answers will you offer? *Clearly summarize your conclusions in your introduction.* Your summary should offer enough detail to let readers grasp the main elements of your argument by reading your introduction alone. It should run several pages at least.

The opposite strategy, of seducing readers by withholding conclusions until late in the document, merely tries readers' patience. Moreover, your argument is lost on the many readers who won't read past your introduction.

4. What competing explanations, arguments, interpretations, or frameworks will you reject or refute?⁶ Clearly identify the books, articles, and ideas that you demolish.

Connect your thesis to all the debates and literatures that it speaks to. If it speaks to several debates or literatures, flag this so participants in each debate will realize that your work matters to them. This helps them and also you: they will cite you and make you famous.

⁵Regarding sequence, a gracefully written first chapter can start with your questions or with their historical/factual context. It can work better to first frame the facts that stir your questions, then frame the questions these facts inspire.

⁶Graceful chapter construction may be served by addressing this question and question #2 at the same time.

5. How will you reach your answers? Say a few words about your sources. If you are doing case studies, explain how you selected your sources. If you are doing archival research, say so, and identify the archives and sources you used. If you are doing interviews, offer some remarks on your interview subjects and procedures. If you are doing a large-n statistical study, explain the origins and construction of the data you are using, and explain your method of analysis--in terms comprehensible to the layman among your readers who have forgotten their statistics. If you are using a complex method, explain its nature. If your approach is largely deductive, explain this.

If there are methods or sources that readers might expect you to use, but you did not use, you might note this and explain your decision. If you discovered that a method proved to be unavailable and lines of research that proved infeasible might have been pursued, note this. If there are important questions that you did not answer, identify these and explain why you couldn't answer them. Instead of writing your way around gaps in your knowledge, explain them honestly in your introduction. (But do your research in a way that avoids such lame excuses.)

6. What comes next? Provide a roadmap to the rest of the thesis: "Chapter 1 describes how I began my life of crime; chapter 2 details early arrests; chapter 3 details the journey to death row; chapter 4 offers general theoretical conclusions and policy implications. Something of that sort.

Subjects #1 ("What is your question?"), #2 ("Why does this question matter?"), and #3 ("What is your answer?") are the most important. Make sure you cover them.

Summary introductions of this sort reduce confusion about what you are doing and what you does not say. They also serve a diagnostic purpose for the author. The reader's summary can alert you to internal contradictions or other flaws in the argument. This helps you flag problems that need fixing.

Introductions are often best written last. It may work to later conclude your project you should keep your eye on how you will answer these questions as the project comes.

V. YOUR CONCLUDING CHAPTER

In your conclusion you may want to summarize your questions and answers. If your summary introduction was cursory. However, I recommend that you summarize your research only briefly and then explore its implications at greater length. Do any implications follow from your discoveries? What general theories does your research support, and which does it reinforce? What broader historical questions does your research raise? What further research is called for by your discoveries? This is the larger significance of your research.

IV. STUDY DESIGN AND PRESENTATION: OBSERVE CUMULATIVE KNOWLEDGE NORMS

Political science is often criticized because few questions are ever settled and the same issues are revisited over and over. Things will improve if social scientists follow norms that foster the cumulation of knowledge. So please follow these injunctions:

1. Have a research design before you start your research. This platitude is too often honored in the breach. "The main purpose of the [research] design is to help to avoid the situation in which the evidence does not address the initial research questions."⁷ Those who proceed without a research design risk being marooned on a mismatch between their questions and their evidence.

2. Frame your argument clearly. Knowledge accumulates only if readers know what authors have said.

If your thesis proposes, tests, or applies theories, the reader should be able to "arrow-diagram" these theories.⁸ If your hypotheses cannot be reduced to arrow diagrams, then your writing (and probably your thinking) are too muddy. Think your project through again. This advice applies to explicitly theoretical work and to policy-prescriptive work. All policy prescription rests on theories, and good prescriptive writing frames these theories clearly.

If your thesis is largely descriptive or historical, your main discoveries should be clearly summarized at least once in the thesis, preferably at the outset.

If your thesis tests theories or explanations, clearly frame their predictions (or "observable implications") before presenting evidence. Theories and explanations are tested by inferring predictions from the explanation, and then asking if the predictions are confirmed or disconfirmed by the evidence. You should explicate this process for your readers by clearly framing the predictions your evidence tests. (Most authors omit this step but that doesn't make it right.)

Frame all predictions that flow from your theory, including those that are falsified by the evidence or prove untestable. Failed predictions should be identified and their failure confessed. If some predictions are confirmed and some fail, say so and offer interpretation.

Thus your overall format should be (a) frame your theory/explanation; (b) infer predictions from it; (c) perform tests; (d) infer implications.

3. Be definitive. Your thesis should reflect a comprehensive survey of literature and evidence relevant to your subject. Your footnotes should provide a comprehensive bibliography to the important literature relevant to your topic. This requires that you gain mastery of all aspects of your subject.⁹

⁷Robert K. Yin, *Case Study Research: Design and Methods*, 2nd ed. (Thousand Oaks, CA: Sage, 1994), p. 20.

⁸On arrow-diagrams see "Hypotheses, Laws and Theories: A User's Guide," in this working paper, section I.

⁹Basic but still worth consulting for clues on doing research is Kate L. Turabian, *A Student's Guide to Writing College Papers*, 3rd ed. (Chicago: University of Chicago Press, 1976). Another guide, directed at hard scientists but useful for social scientists, is Robert A. Day, *How To Write and Publish a Scientific Paper*, 3rd ed. (Phoenix: Oryx Press, 1988). Also, ask reference librarians for help. They know many tricks.

4. Document all sources and statements of fact. This requires a good personal system for storing and retrieving your evidence. One of my rules of thumb: when in doubt, make photocopies. Copy anything and everything. This eases data retrieval and documentation of sources.

5. "Argue against yourself." Acknowledge counter-arguments that might be raised by skeptical readers, and briefly address them late in the text. Concede what you can to these arguments and explain why you won't concede more. This shows readers that you have given due thought to possible objections or alternate interpretations. It also forestalls baseless criticism of your work.

6. Do plausibility probes as the first phase of your research. In other words, try to find the answer before doing your study. The experimental science model proceeds from question to hypothesis to prediction to experiment to conclusion. This mechanistic process seldom works for us. Instead we go from question to hypothesis to prediction to exploration (plausibility probe) to revised hypothesis to prediction to larger data collection to conclusion. In short, we often "work backwards" from answer to problem. You must do this to narrow the range of possible answers we fully investigate. Otherwise you would waste energy doing full-dress tests of hypotheses that a cursory look at the evidence would refute.

7. Clearly identify works that your thesis revises, contradicts, or supersedes. If your thesis is theoretical or policy-prescriptive, identify by name authors whose work you would refute. If your thesis is descriptive or historical, identify exactly which previous work you are revising. This may annoy the superseded authors, but otherwise your readers will continue to quote outmoded work.

How can you sharpen your methodological skills? Re-read works you admire, taking an eye on how the authors executed their projects. Form an attitude on what is right and wrong, and note the methods and sources they used. Consider whether their methods or sources might be appropriate for your possible dissertation project.

VII. WRITING

A well-written thesis is more likely to be published, assigned, and quoted. So keep the following points in mind:

1. "That which is simple is also good." Your thesis should make a single main point or handful of related points. It should have a clear, simple structure.

Avoid cluttering your thesis with extra ornaments and gargoyles (as students of mine know because you researched something doesn't mean it belongs in the manuscript. It is painful--"I sweated hours over this!"--too bad! In the world of research, half your work is done to be thrown away or saved for a later project.

¹⁰Of course, if deeper study contradicts the results of our plausibility probe we report this. Good scholars go where the evidence takes them.

The logic of presentation varies from the logic of discovery. Your research followed the logic of discovery, but your write-up should follow the logic of presentation. This means it should move simply and clearly from your questions to your answers. It is seldom wise to present your discoveries in the same order in which you made them.

Pitch your writing at a level appropriate for college undergraduate readers. Do not write at a level that only your faculty supervisors can understand. Scholarship that isn't used in the college classroom has little impact; hence you should take pains to address the average student.

2. The following structure is often appropriate for thesis chapters:
 - a. your argument;
 - b. your supporting evidence;
 - c. counter-arguments, qualifications, and limiting conditions of your argument;
 - d. brief concluding remarks, which may include comments on the implications of your argument, or may note questions they raise.

3. Start each chapter with several paragraphs summarizing the argument presented in the chapter. You may cut these summaries from your final draft if they seem redundant with your summary introduction, but include them in your first drafts. They will help your supervisor and friends to read and comment on individual chapters. You may also want to keep these summaries in if they seem to fit. Finally, forcing yourself to summarize your argument in each chapter is a good way to make yourself confront contradictions or shortcomings in that argument.

Often these chapter summaries are best written after you write the chapter, but don't forget to add them at some point.

4. Start each paragraph with a topic sentence that distills the point of the paragraph. Later sentences should offer supporting material that explains or elaborates the point of the topic sentence. Qualifications or refutation to counter-arguments should then follow. In short, paragraphs should have the same structure as whole chapters.

A reader should be able to grasp the thrust of your thesis by reading only the first sentence of every paragraph.

5. Break chapters into sections and sub-sections. More subsections are better than fewer; they help your readers follow your argument. Label each section or subsection with a vivid section heading that communicates the meaning of the section.

6. Write short, declarative sentences. Avoid the passive voice.

For more advice on writing see William Strunk Jr. & E.B. White, *The Elements of Style*, 3rd ed. (NY: Macmillan, 1979); Joseph M. Williams, *Style: Ten Lessons in Clarity and Grace*, 3rd ed. (New York: HarperCollins, 1989); and Teresa Pelton Johnson, "Writing for *International Security*: A Contributor's Guide," *International Security*, Vol. 16, No. 2 (Fall 1991), pp. 171-180.

7. If you are doing case studies: it often works to write detailed chronological histories of the case before doing the case study. This helps you gain mastery of the case. Then rewrite the chronology as a case study.

VIII. STYLE

On all matters of style (e.g. footnote and citation format, etc.) consult Kate L. T. *A Manual for Writers of Term Papers, Theses, and Dissertations*, 5th ed. (Chicago: Chicago Press, 1987). Slavishly obey her instructions. Style mistakes make your script look unprofessional.

IX. VETTING

When you finish some thesis chapters, circulate them to several friends for comments and criticism. Don't be shy. The first law of scholarship is "two heads are better than one." Vetting will improve your work.

If your chapters are really half-baked--and early dissertation chapters usually are--terrible--do show some caution. It is probably best not to show them to complete strangers who may conclude from them that you are brain-dead, and that your respirator should be turned off. However, do show them to friends who can be trusted to know you are not brain-dead, even though the condition of your chapters suggests otherwise, and who will help you kick them into shape.

Conversely, when others ask you to vet their work you should make a serious effort. Helping others improve their written work is an important professional obligation. In carrying out this obligation, show mercy and compassion if your colleague's work shows early brain death--while also making clear that there is significant room for improvement and offering specific feasible suggestions.

Do not look solely to your professors for vetting or criticism. Your friends should play an equal, perhaps even larger, role.

Graduate students sometimes view their fellow students as competitors to be avoided at a distance and left unhelped. This is a serious error, for two reasons. First, it is unchristianlike. You should axiomatically, in your personal and professional life, aspire to be a mensch. Our profession and the wider world need more mensches: so be one. You and I both hope that you take this appeal to heart. We will be proud of you if you do. And mensches help their fellow students and colleagues. Second, aloofness from your fellow students is a career-management blunder. The history of social science literature is a record of triumphs and discoveries by scholars who formed empowering communities of mutual help and thereby outperformed their atomized colleagues. Those who act alone often sink to the bottom, while those who help one another excel and prosper. In Virginia, there is no conflict between collegial conduct and the imperatives of professional success. (On this matter study carefully Robert Axelrod's *Evolution of Cooperation*, pp. 63-66, which summarizes the keys to success in academic life.)

¹¹Robert Axelrod, *The Evolution of Cooperation* (New York: Basic Books, 1984).

X. YOUR ABSTRACT.

At an early stage, write a 1-2 page abstract that provides a clear, cogent summary of your thesis. Circulate this abstract when you circulate draft chapters, to help your readers grasp the general drift of what you are doing.

You should also include a provisional dissertation table of contents with provisional chapter titles when you circulate draft chapters. This helps your readers see the big picture.¹²

XI. DEALING WITH YOUR THESIS COMMITTEE

Your advisor owes you a thoughtful reaction to your thesis proposal, and some reaction as you produce chapters. However, this is your thesis, not your advisor's thesis. Your name goes on the cover. If you are really stuck--as you will be from time to time--ask for help, but don't expect anyone to hold your hand through the whole process. Your advisor has the right to expect you to solve most of your problems yourself, and to seek your own solutions before asking others to get involved.

Your committee members owe you one but only one careful read of your thesis chapters. Do not expect iterated readings. A loving advisor may give you more than one but don't expect it. Hence you should carefully choose when you want your committee members to read your drafts.

Edit chapter drafts before showing them to your committee. It takes your committee far longer to read very rough drafts, and they are less able to make useful comments. So neaten everything up before sending it around. (If you want early mid-course correction from your committee, ask your advisor to react to a detailed chapter outline, not a half-baked draft.)

Listen carefully to your advisor's advice. Most of this advice will probably be wise, some will be misguided. You needn't follow the errant advice, but do have reasons for spurning advice you reject.

Do not make your advisor repeat things twice. Remember that when you deal with your advisor your professionalism is on display.

XII. DEALING WITH YOUR HEAD, YOUR FAMILY & YOUR FRIENDS

Writing a dissertation is a difficult and lonely act that requires great force of will. The best way to summon up this will is to choose a topic that fires you up. This consideration argues for choosing a topic that stirs your passions over topics that fit current field fashions but excite you less.

The spouses, significant others, parents and friends of academics often fail to fully grasp the central importance and great difficulty of writing a Ph.D. dissertation. They

¹²You might also include your prospectus when you circulate chapters to readers unfamiliar with your project, so they can grasp what you originally set out to do.

grow impatient with the many months of strange behavior--the falling-down absentmindedness, the vacant stare, the seemingly-permanent hermit-like disappearance into murky library stacks, the vast cluttering of the apartment with clouds of note paper, mumbling to yourself when others can hear you, etc. You must be strong against these enemies of knowledge. Forgive them for their ignorance and abuse, but do not succumb to their entreaties to goof off on weekends, go to the beach, have a beer, act like a normal person, etc. Those who have not written dissertations can never understand how important it is to remain focused on the project. The best you can do is explain over and over, that your whole career rides on writing a decent thesis, and that a decent thesis is like climbing Mount Everest: it can be done, but only by careful preparation and intense focus on the task. If this doesn't work, take solace in the fellowship of dissertation-writing friends who are in the same fix, and hope the divorce papers drop before you get your degree.

XIII. HOW TO LEARN MORE ABOUT HOW TO WRITE A THESIS

Re-read several books you like that used approaches similar to yours and imitate their better aspects.¹³

¹³For example, I recommend that students doing historical case studies look at Barry Posen, *Sources of Military Power* (Ithaca: Cornell University Press, 1984), and Stephen Walt, *Origins of Alliances* (Ithaca: Cornell University Press, 1987). The theory chapters of these books are well done; when in doubt, imitate them when doing your own theory.

MEMO 5:**THE PH.D DISSERTATION PROPOSAL**

Your dissertation proposal explains your project to the world. You will use it to persuade funders and research institutions to support you, and to elicit comments and suggestions on your project from friends and colleagues.

It should frame the question(s) your dissertation will answer, and should explain how you propose to answer them. It also should persuade readers that your questions are important, and your plan of action is practical.

A complete proposal should answer five (5) questions:

1. What question or questions do you address?
2. Why does this question arise? (From what scholarly debates or real-world events?) Why does it matter? Say a few words about the origins and significance of your project.
3. What previous literature has been written on the question? Describe the "state of the art" on the subject.

If a substantial literature has already appeared on the subject you address, you should explain and distinguish majority and minority views (providing footnotes to relevant literature), and sketch the manner in which important relevant controversies have evolved.

Note: Questions #2 and #3 overlap, and can often be answered together in a single statement.

4. What working hypotheses will you explore? You can't be sure of your answer until you complete your research, but readers want to know what hunches you plan to investigate.
5. How will you reach your answers? Say a few words about the methodology you are choosing, why you are choosing it, and how you will implement it. If you are doing case studies, identify your cases and explain their selection. If you are analyzing large-n data bases, identify and describe them. If you are doing interviews or other field research, briefly explain how you plan to go about it. If you are doing survey research, briefly describe your survey sources. If you are doing archival research, explain which archives and sources you will use. If your approach is largely deductive, explain this. If there are methods that readers might expect you to use, but that for some reason you will not use, you might note this and briefly explain your decisions.

You should answer these questions in roughly 5-10 typed doublespaced pages. Footnote your proposal as you would a research paper. It may also be appropriate to append a preliminary bibliography of 1-2 pages listing some of the sources that you plan to consult.

To learn more about how to write a proposal, ask a friend who has written a reputedly good proposal if you can look at what they did. Pay close heed to proposals that were well received by others.

MEMO 6:**HOW TO WRITE A PAPER**

I often offer the following suggestions to undergraduates writing class papers.

I. GENERAL FORMAT

The following general format is often appropriate: "tell them what you're going to tell them; then tell them; then tell them what you told them."

II. INTRODUCTION FORMAT

Begin your paper with a short summary introduction. This summary introduction should answer up to five (5) questions:

1. What question or questions do you address?
2. Why do these questions arise? From what literature or real-world events?

Offer background that clarifies your questions and puts them in context.

3. What answer or answers do you offer? Summarize your bottom line in a few sentences.
4. How will you reach your answers? Say a few words about your sources and methods.

5. What comes next? Provide a roadmap to the rest of the paper: "Section I explains how I began my life of crime; Section II details my early arrests; Section III describes my trip to death row; Section IV offers general theoretical conclusions and policy implications." Something of that sort.

#1 ("What is your question?"), #2 ("Why does this question arise?"), and #3 ("What is your answer?") are essential: make sure you cover them. #4-#5 are optional.

Summary introductions of this sort help readers grasp your argument. They also help you diagnose problems with your paper. A summary introduction can be hard to write. A possible reason: gaps or contradictions in your arguments or evidence, which summary exposes. Solution: rethink and reorganize your paper.

III. CONCLUSION FORMAT

You should probably recapitulate your argument in your conclusion. However, a good summary introduction often makes a full summary conclusion redundant. If so, recapitulate quickly and then use your conclusion to explore the implications of your argument. What policy prescriptions follow from your analysis? What general arguments does it call into question, and which does it reinforce? What further research projects does it suggest?

IV. ARGUMENTATION

Four injunctions on argumentation should be kept in mind.

1. Use empirical evidence--facts, numbers, history--to support your argument. Purely deductive argument is sometimes appropriate, but argument backed by evidence is always more persuasive.

2. Clearly frame the general point(s) that your evidence supports. Don't ask facts to speak for themselves.

To sum points #1 and #2: state evidence to support your arguments, and state the arguments your evidence supports.

3. "Argue against yourself." After laying out your argument, acknowledge questions or objections that a skeptical reader might raise, and briefly address them. This shows readers that you were thoughtful, thorough, and paid due regard to possible objections or alternate interpretations.

Often, of course, the skeptic would have a good point, and you should grant it. Don't claim too much for your theories or evidence!

4. Use footnotes to document all sources and statements of fact. On footnote and citation format, consult and obey Kate L. Turabian, A Manual for Writers of Term Papers, Theses, and Dissertations, 5th ed. (Chicago: U. of Chicago Press, 1987), in paperback. (You should probably buy a copy.)

V. WRITING

Good writing is essential to clear thinking and effective communication. So bear the following points in mind:

1. Your paper should make a single point or a handful of related points, and should follow a simple organization. Avoid cluttering it with extra points. If you developed an argument that later became ancillary as you rethought your paper, drop the argument from the paper. This is painful ("I sweated hours on that idea!") but extraneous arguments drain power from your main argument.

2. Break your paper into sections and subsections. More sections is better than fewer. Sections help readers see the structure of your argument.

Label sections with vivid section headings that convey the main message of the section.

3. I recommend the following structure for sections/subsections:

- a. Your argument;
- b. Your supporting evidence;
- c. Counter-arguments, qualifications, and limiting conditions of your argument.

4. Start each section with several sentences summarizing the argument presented in the section. You may cut these summaries from your final draft if they seem redundant

with your summary introduction, but you should include them in your first drafts to show how they look. Writing such summaries is also a good way to force yourself to decide what you are and are not doing in each section, and to force yourself to confront contradictions or shortcomings in your argument.

Often these section summaries are best written after you write the section, but don't forget to add them at some point.

5. Start each paragraph with a topic sentence that distills the point of the paragraph. Later sentences should offer supporting material that explains or elaborates the point of the topic sentence. Qualifications or refutation to counter-arguments should then follow. In short, paragraphs should have the same structure as whole sections.

A reader should be able to grasp the thrust of your argument by reading only the first sentence of every paragraph.

6. Write short, declarative sentences. Avoid the passive voice. (Passive voice: "the kulaks were murdered"--but who did it? Active voice: "Stalin murdered the kulaks.")

7. Always write from an outline. Outlines are major aids to coherence and readability.

For more advice on writing see William Strunk Jr. and E.B. White, The Elements of Style, 3rd ed. (NY: Macmillan, 1979); Joseph M. Williams, Style: Ten Lessons in Clarity and Grace, 3rd ed. (New York: HarperCollins, 1989); and Teresa Pelton Johnson, "Writing for International Security: A Contributor's Guide," International Security, Vol. 16, No. (Fall 1991), pp. 171-180.

If you are doing a research paper, you might also consult Kate L. Turabian, A Student Guide to Writing College Papers, 3rd ed. (Chicago: University of Chicago Press, 1976) for more clues on how to do research.

VI. VETTING

Ask a friend or two to give your paper a look before you turn it in; and return the favor for them when they have a paper underway. Two heads are better than one, and giving and receiving comments are important skills.

VII. GENERAL BEAUTY TIPS

Take care to turn in a neat, clean paper. Run your spellchecker. A messy-looking paper suggests a messy mind.

VIII. HOW TO LEARN MORE ABOUT HOW TO WRITE PAPERS

Re-read articles you or others admire and imitate their better aspects.

MEMO 7: PROFESSIONAL ETHICS

Many professions have codes of ethics and teach professional ethics in their professional schools. For example, most students of law, business and medicine now take a course in professional ethics at some point.

The social sciences should also teach/discuss professional ethics, for three reasons. First, the wider world cannot easily hold us accountable for our general professional performance. No market forces compel us to deliver a useful research product. Hence we need extra measures to ensure that we do. Otherwise we risk degenerating into social parasites. (Groups that are accountable to no one seldom serve anyone well and usually turn parasitic or worse.) A shared body of professional ethics that defines our obligations to society can reduce the accountability gap by helping us to hold our own feet to the fire. Second, our students cannot easily hold us accountable. If we teach badly there is little they can do about it. Hence we must hold ourselves accountable to perform as teachers. A body of ethics that defines our teaching obligations can be a self-accountability mechanism. Third, the lack of shared norms among ourselves on proper conduct in all manner of personal and professional settings causes confusions and even tragedies that injure involved parties and consume the time of the rest of us who must act as judges. These confusions and tragedies could be reduced by reaching a wider and clearer agreement on professional norms of conduct toward each other.

Issues to cover in a course module on professional ethics might include:

1. What do social scientists owe the outside world? Relevance? Honesty? Anything? Do we have a social contract of some sort with wider society, or are we free to conduct ourselves however we please?

One view is:

- a. Social science has a duty to be relevant. We have an implicit contract with society: in exchange for tenure, summers off, etc. we agree to spend at least some energy answering society's more urgent questions. This does not require headline chasing. We can meet our obligation with policy research or with more abstract work that could have policy implications far down the road. But social science violates its contract if it drifts into complete irrelevance, as so much of it has.
- b. Social science has a duty to be unpopular if necessary. Many important social ideas have distributive effects that injure someone--often someone who is noisy or powerful. Taking their heat is part of our job. We should expect it and should try to remain undiverted by it. We are given the privilege of tenure in part so we can stand up to it, and we waste and abuse that privilege if we let fear of criticism divert our work.

2. What do social scientists owe their students?

One view is:

A teacher's mission is to produce learned men and women with the thoughts of their own. This requires curbing our urge to clone our own ideas, giving our students full latitude to form their own ideas, including ideas that clash with ours.

We all love our own ideas, but the place to frame them is in our writing. We should argue a clear point of view. In our teaching we should present both sides. Students should be asked to come up with their own. In short, teachers should bifurcate their work: arguing their ideas and teaching others how to think (hence imposing no answers) in their writing. We should hold ourselves to this standard all the time (I haven't) but we should

Teachers also owe the teaching of skills that are little fun to teach. But it's our job.

3. What do social scientists owe each other? Impartiality in reviewing? Fellowship applications, tenure review? Civility of discourse? Sharing of ideas? General menschlike conduct? Etc. Etc.

One view is:

- a. Diversity of methods and arguments should be maintained. A marketplace of ideas and methods produces the best results.

In hiring we are entitled, indeed obliged, to define the boundaries of our field and to exclude that which lies beyond it. Tolerance for heterodoxy that we hire and tenure witch doctors and alchemists. But dogmatism should be frowned upon. Scholars badly serve their students and field by reflexively trying to clone their methods. We should defend our field to the point of exterminating others.

In reviewing manuscripts we should exclude agreement with an author's argument as criteria to recommend for or against publication. Reviewers should endorse manuscripts that ably and importantly argue their case, even if the argument sets the reviewer against his own. We learn through debate. Debate requires that diverse views be published in print. Reviewers should render judgments that foster this.

Conversely, mutual backscratching in reviewing is a form of self-defense. Do not cut your personal friends and methodological or ideological enemies. Make any breaks. Corrupt back-scratching occurs in our field. We should make it right.

- b. Civility of discourse should be preserved. We should state our views clearly, but *ad hominem* attack should be kept to a minimum to

All members of the community should staunchly defend the political freedom of speech of all others. We have a duty to defend our bitterest opponent's right to be heard. Civil discourse collapses if we fail to do this (an important lesson of the 1960s).

- c. Menschen make the world go round and this should be recognized in field management decisions. Some value should be placed on menschhood in hiring and promotion.
4. Love, romance, etc.

A sound general rule is: no romance of any kind across power lines. Faculty and graduate students are adults, and human beings do mate with each other, often with others who share their intellectual interests. However, romance across power lines, i.e., between two people one of whom has power over the other, should be strictly taboo. When such romance seems a good idea to both parties, the power relationship should be clearly and permanently ended before romancing begins. The more powerful party should recuse themselves from all further decision making regarding the less powerful party. If this is for some reason administratively impossible then romance is a very poor idea, may in fact be sexual harassment in the eyes of the weaker party, is generally corrupting to the professional integrity of both, etc.