Research Design in Political Science
POLS4011/POLS8058

Richard Frank
24 March 2026

WEEK 5: CASE SELECTION AND SCOPE, PART 1

PART 1: OVERVIEW

This is the first of a two-week block on case selection and scope. The previous two weeks focused on *what* we study: how concepts are defined, structured, and measured. This week and next we turn to *where* and *how many*: which cases to observe, how many you need, and what you can legitimately claim based on the cases you choose. The logic connects directly to the concept-focused weeks. Once you know what you are measuring, the next research design decision involves deciding which cases to measure it in, and whether those cases allow you to draw the inferences your research question requires.

Case selection is one of the most consequential design decisions you will make, and it is one of the most commonly misunderstood ones. Scholars often treat case selection as of mostly practical convenience. They study a case/country because they speak the language or know its history, or they use a dataset because it is available for download. There is nothing inherently wrong with either of these reasons, but they are not design justifications. A design justification for case selection explains *why* the chosen cases allow you to answer your research question validly, *what* inferential leverage those cases provide, and what the choice of cases *rules out*. This week's readings provide the conceptual tools for making and defending those choices.

The title of this two-week block, "case selection and scope," is deliberate. Case selection and scope conditions are two sides of the same coin. Selecting cases means deciding what is *inside* your study; specifying scope means being explicit about what is *outside* it and why. A well-designed study does both. An under-designed study selects cases without thinking about scope, which means the researcher cannot tell you what their findings generalise to, or more importantly, what they do *not* generalise to. This connects directly to the external validity questions that arose in Week 2 when we discussed Hyde (2007). Her Armenian natural experiment has strong internal validity, but what was its broader scope?

Next week continues the case selection and scope discussion. While this week focuses on the logic of case selection and the problem of selection bias, Week 6 discusses scope conditions and external validity more directly. The motivating question next week is basically once you have selected your cases and produced your findings, what can you claim about the world beyond those cases? The readings will nudge us further into questions of generalisability, the boundaries of causal inference, and the relationship between internal and external validity.

Plan for today

1. Overview: connecting Weeks 3 and 4 to Weeks 5 and 6
2. Readings: case selection strategies, selection bias, and applied examples
3. Group activity: diagnosing case selection in your own projects
4. Discuss your upcoming critical reviews

Key topics this week

- Case selection is a design decision, not a convenience decision.
- What are the major case selection strategies (typical, diverse, extreme, deviant, most similar, most different, and crucial cases)
- Why the cases you choose affect the answers you get. Selection bias is an inferential threat.
- The relationship between case selection and the kind of inference (descriptive, causal, or theory-building) you can make.
- Scope conditions: what your findings generalise to, and what they do not
- The connection between conceptual clarity (Weeks 3 & 4) and case identification (this week): you cannot select cases until you know what counts as an instance of your concept

The differentiated expectations continue. Honours students should be able to ***identify*** what type of case selection their project uses and ***justify*** why that choice is appropriate for their research question. MA/PhD students should be able to ***evaluate alternative*** case selection strategies, explain what each would and would not allow them to conclude, and articulate the inferential trade-offs involved in their choice.

PART 2: READINGS

Required readings

1. Seawright, Jason and John Gerring (2008), "Case Selection Techniques in Case Study Research: A Menu of Qualitative and Quantitative Options," *Political Research Quarterly* 61(2): 294–308.
2. Geddes, Barbara (1990), "How the Cases You Choose Affect the Answers You Get: Selection Bias in Comparative Politics," *Political Analysis* 2: 131–150.
3. Tannenwald, Nina (1999), "The Nuclear Taboo: The United States and the Normative Basis of Nuclear Non-Use," *International Organization* 53(3): 433–468.

1. Seawright and Gerring (2008)

This article is the main theoretical reading for this week. Seawright and Gerring (2008) provide a systematic typology of case selection techniques, each defined by its relationship to the broader population of cases and each suited to different inferential goals. The central argument is that case selection in qualitative research should not be ad hoc or convenience-driven but should follow from the research objective. Different questions demand different types of cases, and the choice of case determines what kind of inference is possible.

Seawright and Gerring (2008) identify seven case selection techniques and helpfully demonstrate how they can be implemented in practice. A *typical case* is representative of the broader population, and it sits near the centre of the distribution on relevant variables. Typical cases are useful for exploring causal mechanisms in a "normal" setting, but they are not good for identifying new variables or testing the boundaries of a theory. A *diverse case* strategy selects cases that span the range of variation on relevant dimensions, maximising the representativeness of a small-n study. This is useful when the goal is descriptive breadth rather than causal depth. An *extreme case* sits at the far end of the distribution on the outcome or key

variable of interest, which can be useful for understanding unusual outcomes but risks over-generalising from outliers.

A *deviant case* is one that defies expectations because it deviates from the pattern predicted by existing theory or from the general cross-case relationship. Deviant cases are valuable precisely because they challenge what we think we know and can generate new hypotheses or reveal omitted variables. A *most similar* design (also called Mill's method of difference) selects cases that are as alike as possible on all variables except the one of interest, which helps us isolate the effect of that variable. A *most different* design (Mill's method of agreement) selects cases that differ on as many variables as possible but shares the same outcome, which suggests that the shared factor is causally important. Finally, a *crucial case* is one where a theory is most or least likely to hold. If the theory fails in its most-likely case, it is seriously weakened. If it holds in its least-likely case, it is strongly supported. Crucial cases provide maximum inferential leverage from a single observation.

The key insight in this reading is that each case selection technique answers a different question. If you want to understand how a causal mechanism works in a representative setting, select a typical case. If you want to challenge an existing theory, select a deviant case. If you want to test a theory under the most difficult conditions, select a crucial case. Students who say "I am doing a case study of X" without specifying which type of case X represents in relation to the broader population have not really made a design choice yet. Seawright and Gerring (2008) is useful because it gives us the vocabulary and logic to make that choice clear.

This article also makes an important point about the relationship between qualitative and quantitative approaches to case selection. Several of their techniques (typical, diverse, deviant) are defined relative to a cross-case pattern, which means that even qualitative case studies benefit from quantitative information about the broader population. You do not need to run a regression before selecting a case, but you do need to know something about where your case sits relative to other cases. This connects to KKV's (1994) argument from Week 1 that the logic of inference is shared across traditions, and to Gerring and Seawright's (2022) chapter about using statistical tools to find interesting puzzles.

Reading questions

Honours students

1. Seawright and Gerring (2008) identify seven case selection techniques. Choose two that seem most relevant to your own project and explain why. What kind of inference would each allow you to make?
2. What is the difference between an extreme case and a deviant case? Why does this distinction matter for what you can conclude from your study?

MA/PhD students

1. Seawright and Gerring (2008) argue that case selection techniques should follow from the research objective. In practice, researchers often select cases for pragmatic reasons (language skills, data availability, personal connections) and then construct a post hoc justification. Is this always problematic, or can pragmatic case selection be defensible? Under what conditions?

2. The crucial case design provides maximum inferential leverage from a single observation. What assumptions does it require about the theory being tested, and when might those assumptions fail? Can you identify a published study in your field that uses a crucial case design?

2. Geddes (1990)

Geddes (1990) addresses what is probably the single most common methodological pitfall in comparative politics: selecting cases on the dependent variable. Her argument is direct and powerful. If you want to explain why some countries democratise but you only study countries that democratised, you cannot distinguish the factors that caused democratisation from factors that are present in both democratising and non-democratising countries. Put simply, the cases you choose affect the answers you get, and choosing cases poorly can produce systematically wrong answers.

The core problem is straightforward. If you select only cases where the outcome of interest occurred (revolutions, democratic transitions, civil wars, successful policy reforms), you have no variation on the outcome. Without variation, you cannot assess which factors are associated with the outcome and which are simply background conditions present in all cases regardless of outcome. Geddes (1990) illustrates this with examples from the comparative politics literature where scholars like Skocpol (1979) studying revolutions examined only cases of revolution, scholars studying economic development examined only successful developers, and so on. In each case, the research design could not distinguish genuine causes from correlates because the comparison category was missing.

Geddes (1990) argues that selection bias in case selection persists for several reasons. Positive cases are more visible and more interesting than non-cases. Revolutions are dramatic; the absence of revolution is not. Like most people, scholars are drawn to explaining things that happened, not things that did not happen. Additionally, when working with a small number of cases, researchers may not realise they have restricted their sample in a way that can bias their findings. The error is not always obvious, particularly in qualitative research where the "sample" is not drawn from an explicit population in the way a survey sample would be.

The solution Geddes (1990) proposes is conceptually simple: include cases where the outcome did not occur. If you want to explain democratisation, you need to study both countries that democratised and countries that did not (but were otherwise plausible candidates). This does not mean you need a large-n study; it means you need variation on the dependent variable. Seawright and Gerring's (2008) most similar and most different designs are also responses to this problem. They build in variation on either the explanatory variable or the outcome to avoid the trap Geddes (1990) identifies.

I think this article is useful for students at your research project stage. Many honours students, when asked "why are you studying this case?" will answer "because it is interesting" or "because it is where the thing I am interested in happened." Geddes (1990) shows why that answer is insufficient as a *design* justification, even if it is a perfectly reasonable starting point for *finding* a project (recall Gerring and Seawright's [2022] back in Week 1).

Reading questions

Honours students

3. In your own words, explain what "selection on the dependent variable" means. Can you construct a simple example from your own area of interest where this error would lead to misleading conclusions?
4. Geddes (1990) argues that including negative cases (cases where the outcome did not occur) is essential. For your own project, what would a negative case look like, and would it be feasible to include one?

MA/PhD students

3. Geddes (1990) focuses on selection bias in comparative case studies, but the logic applies more broadly. Do you think selection on the dependent variable can happen in large-n quantitative research? If so, how, and what tools exist to address it?
4. Some scholars argue that selection on the dependent variable is acceptable for certain purposes, such as theory generation or process tracing within a single case. Evaluate this claim. Under what conditions, if any, can studying only positive cases produce valid inferences?

3. Tannenwald (1999)

Tannenwald's (1999) article on the nuclear taboo serves as this week's applied example. The article addresses an important puzzle: why has the United States not used nuclear weapons since 1945, despite possessing them and facing situations where their use was arguably militarily useful? Tannenwald (1999) argues that a normative taboo against nuclear use developed over time, constraining decision-makers even when strategic calculations pointed toward use. I am asking you to read this article not because its substantive topic is directly relevant to most of your projects, but because its case selection logic is instructive.

The United States is a crucial case for Tannenwald's argument. If a nuclear taboo constrains even the most powerful nuclear state (the one with the most extensive arsenal, the greatest capacity for use, and a history of actually having used nuclear weapons) then the taboo is likely to constrain weaker nuclear powers as well. The US is a least-likely case for the taboo hypothesis in some respects (a superpower with strategic interests in nuclear use) and a most-likely case in others (a democracy with a normative culture that could generate taboo dynamics). Tannenwald (1999) is explicit that the case is chosen not because it is convenient but because it provides maximum inferential leverage.

The article also uses within-case variation over time. Tannenwald (1999) traces US decision-making across multiple wars to show how the taboo strengthened over time. This temporal variation within a single case is a form of process tracing that complements the cross-case logic of crucial case selection. It is worth noting how Tannenwald (1999) combines the crucial case logic (why the US?) with process tracing (how did the taboo evolve?) to build an argument that is stronger than either technique alone.

Tannenwald (1999) illustrates several of this week's themes. The US is a crucial case (Seawright and Gerring's [2008] typology). The article avoids selection on the dependent variable (Geddes 1990) because Tannenwald (1999) does not simply study cases where nuclear weapons were not used. She examines cases where use was actively considered but rejected,

which provides variation in the deliberation process even though the outcome (non-use) is constant. And the case selection is justified by its theoretical properties, not by convenience (Gerring and Seawright 2008).

Reading questions

Honours students

5. Why is the United States a good case for testing the nuclear taboo argument? What type of case is it in Seawright and Gerring's (2008) typology?
6. Tannenwald (1999) examines multiple crises within a single country over time. How does this within-case variation strengthen her argument compared to studying only a single decision point?

MA/PhD students

5. Evaluate Tannenwald's (1999) case selection using the frameworks from the other readings this week. Does she avoid the selection bias problems Geddes (1990) highlights? How would Seawright and Gerring (2008) classify her cases, and does that classification hold up under scrutiny?
6. Tannenwald's (1999) outcome variable (nuclear non-use) is constant across all her cases. Does this create a selection-on-the-dependent-variable problem, or does her process tracing approach resolve it? What are the limits of process tracing as a solution to selection bias?

Overall reading questions for the entire class

1. Seawright and Gerring (2008) provide a typology of case selection techniques. Geddes (1990) warns about selection on the dependent variable. How do the two arguments relate? Does every technique in Seawright and Gerring's (2008) typology avoid the problem Geddes (1990) identifies, or are some techniques more vulnerable than others?
2. Think back to Hyde (2007). She selected Armenia's 2003 election as her case because it provided quasi-random variation in monitor assignment. Using this week's readings, how would you classify Hyde's (2007) case selection strategy? What are its strengths and limitations from a case selection perspective?
3. Weeks 3 and 4 showed that how you define a concept determines what counts as an instance of it. This week shows that case selection determines which instances you study. Trace the chain from concept definition to case selection in your own project. If you changed your definition of your core concept (as discussed in Weeks 3–4), would your case set change? And if your case set changed, would your likely findings change?
4. Gerring and Seawright (2008) argue that case studies can provide causal insight if the case is well-chosen. KKV (1994) argued that all research shares one logic of inference. Geddes (1990) shows that poorly chosen cases produce biased inferences. What emerges from putting these three arguments together? Is there a general principle about when small-n research can and cannot produce valid causal inferences?

PART 3: GROUP ACTIVITY

This activity applies this week's readings to your research projects. Answer the following questions yourselves (I will bring paper instead of those darn sticky cards).

1. What is your case a case of? Define the population of cases to which your case belongs. If you cannot define the population, you may have a scope problem.
2. Using Seawright and Gerring's (2008) typology, which type of case is yours: typical, diverse, extreme, deviant, most similar, most different, or crucial? Can you justify this classification?
3. Does your project have variation on the dependent variable? If you are studying only cases where the outcome occurred, what would a relevant negative case look like, and is it feasible to include one?
4. What is your scope claim? What types of cases do you believe your findings will generalise to, and what types are outside your scope? What is the basis for that boundary?
5. If someone applied Geddes's (1990) critique to your project, what would the charge be, and how would you respond?
6. How does your case selection connect to your conceptual definitions from Weeks 3 and 4? If you changed your definition of a core concept, would different cases qualify?

Once you have written down your answers, work in groups of 2–4 students as in previous weeks. Each student should present their answers as succinctly as possible, and the other students should provide helpful feedback.

The above should take about 20 minutes. I will then go around the room like last week to ask each group to share a few interesting takeaways, challenges, or difficulties in answering the above questions. As in previous weeks, the goal is to normalise the process of refining design choices through self-analysis and peer critique. A student who discovers that their case selection is convenience-driven rather than design-justified is not in trouble. Rather they have identified something concrete to refine and improve on for their research design paper.

PART 4: CRITICAL REVIEWS

The critical review is due at the end of Week 6. You should now have chosen the peer-reviewed article you plan to review. This week's readings give you a set of tools for evaluating an article's research design. When selecting your article, consider the following:

- Does the article make an explicit case selection choice that you can evaluate using Seawright and Gerring's (2008) typology?
- Is there a potential selection bias problem that Geddes's (1990) framework would help you identify?
- Does the article specify its scope conditions, or does it implicitly claim generality without justification?
- Can you connect the article's case selection to its conceptual framework using the tools from Weeks 3 and 4?

A strong critical review does not just say an article is good or bad. It diagnoses specific design strengths and weaknesses, using the vocabulary and frameworks we have been developing. It identifies the most consequential inferential risk and explains why it matters. And it proposes a feasible redesign that would address the weakness. The tools you now have from five weeks of readings should equip you to do this. As should the first draft writing guide I am emailing you with these notes. Do feel free to give me any feedback you might have on this critical review writing guide.