Paper by Numbers

James Fenske
V4.0. October 2023[1]

Abstract [150 words]. State your question and why it is important. State your main result. State what it means, and what mechanisms explain it. Describe how you arrive at it (e.g. main data source and identification strategy). State your contribution to the literature (i.e. why this is new and important). State any bonus or additional results above the main result. Yes, 150 words is not many.

For a job market paper, 150 words may be too short. Many job market papers have long abstracts. Unless you are submitting to a journal that demands a 100 word limit, there is no need to make your abstract that short.

I. Introduction

1. **The Hook**
- *Part 1:* State clearly what the question is that you seek to answer. Sometimes I do this in the first sentence of the paper.
- *Part 2:* The space for a paper. This is – I think – the hook. There is no one way to state that a paper is needed, but here are a few standard hooks:
    - Establish equipoise: i.e. show that this question is unanswered, because theory is ambiguous, because empirical results are conflicting, or both. *Equipoise* is an overly fancy word you learn if you take ethics training from e.g. the National Institutes of Health. Since all medical trials can harm participants, it is unethical to have a trial if we already know the answer to the research question.
    - Point to the existence of a puzzle that needs to be explained and that your paper will explain.
    - Note that a policy that resembles your X variable is being debated by policymakers, or in the news, but that we lack credible estimates of its benefits and costs. This can be tricky given how long the publication process is, so focus on a more general policy debate you're sure will still be active in 5 years.
    - Argue that there is a conventional view that policy X works/variable X has a known effect in context A, but there is substantial reason to believe it may be very different in context B (e.g. land titling has been shown to improve investment, but is this true where the financial system is underdeveloped and so land can't be used as collateral?). It is even better if this difference in context helps us test between models or understand mechanisms.

I am updating this guide for an audience principally of MSc and PhD students who are often unsure how to write up their dissertations or thesis chapters. The gap this guide fills is that it is a numbered, paragraph-by-paragraph guide to what needs to be in your paper. This template is not the only way to write a paper. But if you've never written a paper before, this template might be easier than figuring out a structure on your own.

- o Note that, while the outcome you study is very important, little is known about how to affect it. Avoid using this since it is a lazy crutch of last resort. Show your work – if you just assert that "there is little on X", I can almost surely prove you wrong with an ISI Web of Science search. Or with work outside economics. If you're going to say "little is known" then say why little is known – e.g. lack of data, lack of natural experiments. Then you have to say that your paper solves the problem that kept us with little knowledge. Indeed, this reason little is known is likely to be a better hook than simply claiming that little is known.
- *Part 3:* Importance
  - o Make it clear why your topic is important, and why anyone would want to read it/publish it/give it a good grade.
  - o Explain why knowing the answer is important both in the real world and from the perspective of the editor of a general-interest journal.
  - o You need an answer to the question "Why should I care?" where "I" is an editor, referee, or grader.
  - o If your paper is part of an already very full literature (e.g. microcredit), you need to establish very early that this is not just one more paper on the pile. So: here is good. What's the novel twist? The novel insight?
  - o If your paper is not obviously "economics," this is one of two chances you have to not lose the reader. Frame your question as an economic question of interest to economists.
- These three parts can come in any order, but they should all be in the first paragraph.
- Parts 2 and 3 are related, but they are not the same. Just because a question is unanswered and challenging to answer doesn't mean it is worth answering.

2. The contribution
- Give a broad overview of what you do, framed in terms of how it fills the gap established in the first paragraph.
- That is: complement the hook. If the hook was "knowing the fiscal multiplier is really important but experimental evidence is non-existent," then the contribution may be "and so we solve that problem using a randomized controlled trial in Kenya."
- Another example: If the hook was "land titling may not encourage investment in the absence of a developed financial system," the hook may be "this paper evaluates the effect of land titling in a context where the financial system is poorly developed."

Given that the Economic Journal has, in the past, said that their ideal paper length is 8,500 words, a 67 paragraph "checklist" for an empirical paper should allow any student to use this as a template and adjust it to a desired word count. Each paragraph should be about 100-150 words. Paragraphs more than 200 words long should always be split apart. Paragraphs less than 50 words should be rethought. 8,500 words is a rough guide. Papers in journals are getting longer. I read one recently, forthcoming in the JPE, that was 16,000 words before the bibliography. Some "short" papers are also now quite long – AER: Insights wants 7,000 words and many RESTAT "short" papers are 8,500 words. So: you will need to cut or add depending on your audience.

This is a template for a reduced-form empirical paper. That is, I expect your paper to estimate the effect of X on Y using some sort of natural experiment (e.g. a policy change), probably using one of: differences in differences, an event study, instrumental variables, synthetic control, or a regression discontinuity. A theory-heavy paper may have a different structure, e.g. intro, model, data, empirics. A structural paper, macro paper, or trade paper may have a structure like: intro, data, preliminary graphical evidence and stylized facts, model, estimation, counterfactual simulations. I don't write papers like those so my advice on that type of paper would be even worse than my advice here.

- State why it was challenging to make this contribution. Audiences won't appreciate that you have done a paper-sized amount of work if you don't tell them.
- An "In this paper, I …" statement. Sometimes this is in the first sentence of the second paragraph.
- Acknowledgement: This hook-contribution structure is from Dan Bernhardt's "Miscellaneous Thoughts on How to Write Papers." That set of slides has several concrete examples I won't repeat here.

3. Methods: What you do
- Make the data source clear.
- Make the identification strategy clear. What is the source of exogenous variation in X?
- This paragraph differs from the previous paragraph in its concreteness. For example, if the hook is "knowing the fiscal multiplier is really important but experimental evidence is non-existent," and the contribution is "and so we solve that problem using a randomized controlled trial in Kenya," this paragraph would tell me that you give $1000 to 10,000 households in 600 villages, collected data on their consumption expenditures after 6 months and 18 months, etc …
- In a science paper, this would be the "materials and methods" section of the extended abstract.

4. Methods: Why they are correct, or at least adequate
- Make it clear why this is the "right" analysis to run. e.g. what endogeneity problem does it solve?
- Why does your context offer exceptional internal validity (why is the answer from this context right)?
- Why does your context offer exceptional external validity (why is the answer from this context useful elsewhere)?
- If there is an alternative approach, why is yours better? Why don't you use the alternative approach?

5. Results
- What are your main findings? What is the main "punchline" of the paper?
- How big is this result? Don't just tell me whether it's significant or not. Is it big or small?
- What does this magnitude mean from a policy perspective? In comparison to other estimates in the literature?
- If you have written a thorough paper with lots of results this outline of results may take a few paragraphs. But don't lose the reader. You can't have ten "main" results.
- Results that come from a model, such as policy counterfactuals, belong later, in the mechanisms paragraphs.
- This paragraph is also a nice place to describe a picture in the introduction that illustrates your results. It may be the only picture anyone remembers from the paper.

And, as usual: Nothing I have to say will be more important than John Cochrane's writing tips! If you want more general advice on how to write, there are plenty of other, much better, sources available. Chaubey's "Little Book of Research Writing." David Levine's "Cheap Advice." Thomson's "Advice for the Young Economist." Thomson's "The Young Person's Guide to Writing Economic Theory." Bellemare's "Doing Economics." Weisbach's "The Economist's Craft." Williams' "Style: Toward Clarity and Grace." Everything on Masa Kudamatsu's Tips4Economists website. You should also read King's "On Writing" and Snyder's "Save the Cat!" Writing a logline is a lot like writing an abstract. Creating a beat sheet is a lot like making an outline for a paper.

For a nice example, see Figure 2 in Grosfeld, I., Sakalli, S. O., & Zhuravskaya, E. (2020). Middleman minorities and ethnic violence: anti-Jewish pogroms in the Russian empire. *The Review of Economic Studies*, *87*(1), 289-342.

6. Mechanisms: Story
- Your "story" can't wait until later in the paper, so tell it here.
- What theory/model/channels account for your result? Provide a brief summary of your model or conceptual framework.

7. Mechanisms: Evidence
- What is the empirical evidence you have to support the story from the previous paragraph?
- What alternative theories or explanations do you rule out and how?

8. Robustness
- I try to anticipate the biggest objections that will come to the minds of referees and dismiss them in this section. That is, I also preview my most important robustness checks.
- If there are standard checks for your method (e.g. McCrary density tests with an RDD) they go here.

I.I. Contribution
i.e. "Related Literature" with a more purposeful sub-section title.

- The idea of this section is to convey the following: "The existing literature has found A, but leaves B unanswered. Luckily, my paper answers B."
- I call this section "contribution" rather than "related literature" precisely to keep this purpose in mind. Even if this naming choice adds the confusion of two separate "contributions" in the introduction.
- I make this a sub-section of the introduction, rather than its own section, solely because that is the current fashion in the economics journals that count for tenure and promotion in my department. You need to follow the conventions of your field. By following them, you show that you know what they are.
- If your paper is not obviously "economics," this is the second of two chances you have to not lose the reader. Show your paper contributes to a literature published in mainstream economics journals, even if your paper is interdisciplinary.

X. Optional paragraph: Policy Implications
- Some policy-driven papers will outline the policy implications of the paper before describing the contribution to the academic literature.
- Note: this cannot be weak speculation about policy, or it will be grounds for a referee to reject the paper.

9. What is known in literature 1

- Cite literatures that are well published, recently published, and full of potentially sympathetic referees. You want a referee who finds your methods familiar and convincing. Editors choose referees in part by looking at your bibliography. They decide if your paper is important based on what it contributes.

- Be strategic. This is the section where you tell your editor who you would like to have as a referee. Cite heavily from referees/literatures where your approach is standard, not where it will be considered strange or not credible. For example: if you're using a revealed preference approach to test a theory in IO without making parametric assumptions, it may be better to cite papers in consumer theory that take the same nonparametric approach than to heavily cite an IO literature built on the structural estimation of parametric models.
- If you are writing on topic A and submitting to journal B, you should really cite some work on topic A published in journal B.
- Don't write a set of one sentence summaries. Group papers e.g. by finding, by outcome, by method, by setting… e.g. "While there is considerable evidence that microcredit increases entrepreneurship (Smith, 1999; Dasgupta, 2001), there is little evidence that it increases wealth in the long run (Stevens, 2002; Xu, 2003)."
- Don't compile your references manually. Use BibTeX. If you do compile them manually, make sure everything cited in the text is cited in the bibliography, and vice versa.
- Check every working paper in your bibliography. It may now be forthcoming or published. Or it may be in a newer working paper series. You don't want the editor and referees to think you've fallen behind the literature.

10. What you add to literature 1

- While you don't have to break "what is known" and "what I add" into two paragraphs, this division is intended to make it clear that you can't just summarize existing work. You only review the literature for the purpose of explaining why your work fills a gap in it.
- One of my pet peeves in seminars is people who put up a "related literature" slide that is just a list of existing papers. You have to explain why what you do is new and a substantial advancement of our knowledge.
- My formula for this: Identify 2-3 seminal papers in the literature and 10-15 recent well-published or famous unpublished papers. For each, write down in one sentence a) the main lesson, b) what is left unanswered that you answer. Each paragraph in this section can list all the a's, then all the b's. Then collapse sentences together. It is bad to have the structure "Smith found A. Jones found B. Adeboye found Z." It is better to have "The literature emphasizes A (Smith), B (Jones) and C (Adeboye)."
- Your contribution should not be "I use a new tool" or "I look at a different aspect of the question." It should be how your results make us rethink what we know. What new economic knowledge do you add that is of importance beyond your study context?

11. What is known in literature 2
- See above.

12. What you add to literature 2
- See above.

13. What is known in literature 3
- You can get away with two, rather than 3 literatures, so long as you discuss them in substantial depth. The "contribution" section as a whole should be at least 600 words.

14. What you add to literature 3
- See above.

15. The structure of the paper
- "In section 2, I explain why labour market institutions should be linked to growth. In section 3, I outline my sources of data…"
- Some people don't like these sections, but I find them useful.

X. An alternative introduction structure from Josh Angrist has 6 paragraphs:
- Research question and motivation
- How you answer the research question
- Results
- Mechanisms, heterogeneity, and policy implications
- Literature
- …. but that only ads up to 5. You may need to split 1. Hence 6. So in this view of the world, 15 is too many. But I think introductions are getting longer, and the "contribution" subsection replaces the "literature review" you might see in other fields.
- … also: pay attention to parallels. You shouldn't have two research questions and three main results.

X. More generally:
- The introduction should focus only on what is important.
- Why should others care?
- Why do you do what you do?
- The focus should be on the big picture, not on detail.
- Enumerate the most important results, not all of them.

II. Context and Identification Strategy

X. On putting the identification strategy before the data:
- This is not how most people structure their papers. I have three reasons for jumping straight from the introduction to the identification strategy in many of my papers.
  - The principal contribution of my paper is usually an empirical result. Sometimes I can only narrow things down to a couple of mechanisms, or the model is really simplistic, so I start with the strong bits.
  - I want to get to my main contribution as fast as possible.
  - This structure frames how the reader reads what follows. Knowing what the data will be used for, they'll hold it to the standard of "adequate for what I will do with it" not "the data are perfect." So: you shut down irrelevant objections by doing this. In seminars, this ordering similarly keeps the discussion on topic and useful to you.
- Other people like to put this as its own section, after their data description. Others merge it with the results section. I find that audiences are more willing to stay awake through the data section if you tell them up front what you're actually going to do with the data.
- Essentials that the reader needs to know to even get through the intro can go in the intro. But the intro shouldn't be a background section.

16. Overview of the context and identification strategy section

- Outline the main points of the rest of the section

II.I. Context

17. Context: The setting
- If you are going to have a "background" or "context" section, I think it should be here. Why? If the context is needed to understand the natural experiment you exploit, it should come just before it. If this context is not needed to understand your empirical approach, it isn't needed at all. Don't include background for the sake of background.
- Here's where you can introduce your setting. But make sure to keep it focused on issues of interest. If you're working on e.g. sex-biased educational investments in India, don't start off by telling me that India is a large country in Asia. Begin with the relevant context e.g. gender differences in wellbeing in India.
- Calling this section "institutional background" might help you focus it even more narrowly on only what is relevant and nothing more.

18. Why this is the right setting
- Stress why this is the right setting for this paper – why does it offer exceptional internal and external validity?
  - Since every context is different, external validity may be impossible to establish. But you can point out similarities between your context and several other environments.
- Stress that the relationship between X and Y is particularly salient or important in this context.
- Stress why this setting not only gives you an unusually good ability to evaluate whether X affects Y, but also why it lets you examine the mechanism behind this relationship better than in other settings.

19. Context: Y in your setting
- Tell us about the nature of your outcome variable in your setting, and what determines it.
- For example, if I had a paper on "does maternal education reduce child mortality in Nigeria? Evidence from Universal Primary Education (UPE)," I might describe levels of child mortality in Nigeria, what the main causes are of child death, and which populations (e.g. rural, minorities) are most at risk.
- This may draw heavily on literature not from economics, e.g. government publications, articles in other fields, books. Indeed, showing a thorough reading of this material shows you are serious and that you know your context well.
- Determinants of Y other than your main X variable that are introduced here should be included as controls or mechanisms (or made irrelevant by your identification strategy) later on in the paper. Allude here to how these variables affect your empirical approach.

20. Context: X in your setting
- Do for X what you just did for Y.
- For example, if I had a paper on "does maternal education reduce child mortality in Nigeria? Evidence from UPE," I might describe levels of education in Nigeria, particularly among women, what the main sources are of education (government schools? missions?), what the nature is of Nigerian education (is it in English? what is

in the curriculum that might affect child mortality?) and which populations (e.g. rural, minorities) have more education than others.

- Again, this will motivate some of the controls and mechanisms that are discussed later in the paper. Allude here to how these variables affect your empirical approach.

21. Context: M in your setting
- Do for your mechanisms what you just did for Y. And if the mechanisms are measurable variables, this is easier.
- For example, if I had a paper on "does maternal education reduce child mortality in Nigeria? Evidence from UPE," I might describe the channels through which education can improve child survival. Do more educated women have more bargaining power? Do more educated mothers know more about vaccines and hygiene? Do they have more money to invest in child health? Do they migrate to environments where there are more services?
- Base this on the mechanisms you actually test later in the paper – there should be a tight connection between the two sections, and if you discuss mechanisms in the same order it will be even clearer. Allude here to how these variables affect your empirical approach.

22. Context: Why the natural experiment you exploit exists
- A reduced-form empirical paper is the application of an identification strategy to a known dataset. If this does not describe your paper, rethink your paper.
- Tell me about the natural experiment you exploit – when did it happen, why, and what form did it take?
- For example, if I had a paper on "does maternal education reduce child mortality in Nigeria? Evidence from UPE," I might describe how and why UPE was enacted, making clear its timing and spatial coverage.
- Because this section will stress the determinants of your natural experiment, think about whether any of these are confounders and need to be controlled for. If the answer is yes, tell the reader here how you will control for them.

23. Context: Why natural experiment you exploit is an ideal experiment
- What is the natural experiment you exploit? Is it the rollout of a new policy? Is it an exogenous event like a commodity price shock or drought? Is it an assignment rule given by a certain law? Is it a border discontinuity? That is, map the policy into one of the workhorse tools discussed by e.g. Cunningham's "Causal Inference: The Mixtape."
- Explain what the natural experiment is that you exploit and why it is plausibly exogenous.
- For example, if I had a paper on "does maternal education reduce child mortality in Nigeria? Evidence from UPE," I might focus on how its timing is essential for defining "post" and how its differential effect on some states but not others is essential for defining "treatment." This description will make it clear even before the main estimating equation appears that the correct approach will be a difference in differences or event study.

X. *A general note on context*
- This section does not exist to just give background. For your paper to "tell us something" rather than being a rote presentation of statistical results, you need to have a "story" that runs through it. The story you introduced no later than the "mechanisms"

paragraph in the introduction. The elements needed to understand that story must all be introduced in the background section. And if there is something in the background section that isn't used later on to tell your story, remove it. This is a principle of storytelling called "Chekhov's Gun" (Wikipedia: "a dramatic principle that states that every element in a story must be necessary, and irrelevant elements should be removed").

- At each step, tell the reader how the background shapes your empirical approach. The background section tells me what I need to know to understand and believe your methods and results. Not for its own sake.

II.II. Identification Strategy

24. The main estimating equation
- Say what regressions or other tests you will perform with the data. Actually write out the equation, even if it is a simple OLS regression.
- What controls will you add in your baseline, and why?
- Explain how you will interpret the results (e.g. "a negative coefficient on beta implies X, a positive coefficient implies Y").
- How will you cluster your standard errors?
- Make it clear what endogeneity problem you solve. Why are you not just running OLS?
- Justify what you're doing: why does this give us a causal effect, and how do you control for confounding factors?
- Explain every single piece of notation in the equation. What is each subscript? What is each fixed effect?
- Explain why every single thing is in the equation – what it is and what it does. What, for example, does each fixed effect control for?
- … this may take several paragraphs, not just one.

25. The sources of identifying variation
- Step away from the math and tell the reader in words what comparison is being used to identify your effect of interest. For an RDD, for example, you'll be comparing two units just on either side of the cutoff. For an event study, you'll be comparing how much the gap between treated and untreated units changes after treatment begins.
- For example: "We estimate the impact Nigeria's UPE policy, which expanded the availability of primary schools after 1976 in certain regions. A significant aspect of this policy is that it was targeted towards regions outside the country's Western region. Thus, the policy is well suited to a difference in difference analysis. Our identification strategy exploits variation within states and within birth cohorts, and exploiting the within-state variation in exposure driven by the differential expansion of primary schools across states."
- Make sure to give the intuition for what you're doing. Don't just state that you e.g. use a Bartik instrument without telling us how it works and where variation now comes from.
- … this may take several paragraphs, not just one.

26. Threats to identification and how you will handle them
- State what the identifying assumption is – what does the error term have to be orthogonal to?

- What sorts of standard omitted variables problems are no longer an issue because of your identification strategy?
- What sorts of omitted variables could still be a problem? Why are these sorts of problems unlikely?
- Preview the main exercises that you will conduct in the "robustness" subsection of the results section.

27. Outline of any additional estimating equations
- If you have any other major empirical exercises you conduct, outline them here.
- Depending on how clearly these sections read, you may show equations that you will estimate in the robustness or mechanisms section here, or you may leave them until then.

III. Data
    28. Overview of the data section
- Outline the main points of the rest of the section
- I find that it confuses readers less if I add sub-subsection headings for the outcomes, treatment of interest, and controls.
- Make sure to only discuss characteristics of the data that readers need to know to understand the results. Other detail can go in the data appendix.
- Were the data especially hard to get? E.g. were they digitized from a 17^{th} century collection of handwritten bills of exchange? Primary surveys done in a Ugandan refugee camp? Emphasize the novelty and difficulty from the start.

    29. The shape of your dataset
- You are likely merging together many datasets from many sources, collected at different levels. Start by telling the reader your ultimate goal, e.g. a panel of firm-by-year observations in which you have profits, the percentage of the directors who are Parsi, etc… If you tell me how you use dividends from quarterly financial statements, I may mistakenly think the regressions are at the quarter-by-firm level. If you tell me you have individual lists of directors, I may mistakenly think the regressions are at the director level…
- It is possible to combine this with the previous paragraph.

    30. Data on the dependent/left hand side (LHS) variable: Collection

- What data source do you use for your dependent variable?
- If these come through a third party (e.g. the US census, via IPUMS) make that clear.
- If these come from another paper, provide a citation. Indeed, almost all datasets have a citation involved. Don't just say "we use ACLED data" without citing Raleigh et al.
    - You may even want to add hyperlinks for data sources that are on the web.
- At what level are the data originally available (e.g. individual)?
- How are these data collected? What does this data collection process tell us about the reliability or representativeness of this measure? E.g. "Data were collected by the OECD as part of its blah blah blah project. These cover 12 countries over the years 1980 to 2013 and include manufacturing firms with at least 50 employees… Because these data are collected by the labour regulator, participation is mandatory and so firms cannot select out of the sample if they are above the employment threshold… "

- What are the possible errors in data collection (e.g. topcoding of incomes, missing wages for nonrespondents) and how do you account for them?

31. Data on the LHS variable: Operationalization
- Define the actual variable you use, preferably with a formula. How is the dependent variable defined, precisely? E.g. "The variable "$\ln(wage\_\{it\})$" is the natural logarithm of the mean hourly wage for full-time workers in industry i in year t, in pounds…"
- At what level do you use the data (e.g. municipality by year)? This is important if the original resolution of the data is not the same as the level of observation in your regressions. If you do need to aggregate or merge the data, explain how you do so.
- If there are any decisions that need to be made (e.g. defining "full time" as above 35 hours), justify them. And tell the reader now that you will use alternatives for robustness.
- Why this dependent is variable good to use given the broader theory of change you have in mind (e.g. in a regression of wages on education, why are wages a good measure of income, rather than e.g. annual earnings or total hours worked?).
- How does this variable compare to others used in recent well-published papers?
- Are there any important transformations you need to make e.g. deflating by CPI, removing seasonality, taking the inverse hyperbolic sine, or winsorizing? Why do you make them? Again, tell the reader now that you will show robustness to alternatives.
- Err on the side of giving a formula rather than only describing the variable in words. If a variable is a ratio, give me the formula – ratio of what to what? You can't assume your writing is clear on its own. Specialists may know that a "rising sex ratio" means that the population is becoming more male, but non-specialists won't. "A rising ratio of boys to girls" is much clearer.
- What patterns e.g. over time and space in this variable exist that lead us to anticipate your main results?
- What alternatives do you use for robustness?
- … as you can tell, this may take a few paragraphs, not just one.

32. Data on the principal right hand side (RHS) variable: Collection
33. Data on the principal right hand side (RHS) variable: Operationalization
- Provide details like those for the LHS variable.

34. Data on the key mechanisms
- Provide details like those for the LHS variable.
- Also explain why you use these variables. Typically they will be variables used for heterogeneity analyses, other outcomes, or mediators.
- Like with the LHS and RHS variables, you should make the differences between "here is how the data are collected" and "here are the variables I define" clear, but you may not have a simple two paragraph order, especially as these data may come from multiple sources.
- To prevent this section from dragging on forever, I sometimes discuss these variables in the "mechanisms" section. But I warn the reader about it here, and I do have to have that discussion in the mechanisms section.
- Check this section against what you have written about the mechanisms in the introduction. You need to write enough in the introduction for this section to make sense, since you won't go into depth on mechanisms until later in the paper.

- But: the "mechanisms" in your paper is a story, not a list of variables. Don't lose sight of the story.

35. Data on other important variables
- Provide details like those for the LHS variable. This paragraph is probably about your controls.
- Also explain why you use these variables. Typically they will be controls for confounders. That is: for each control variable, give a) a definition, b) a source, and c) a reason for including it. You can alternatively explain why you use these variables when you set out your identification strategy.
- To prevent this section from dragging on forever, I sometimes discuss any controls not in the main result table later on as they are introduced in the "robustness" section. But I warn the reader about it here, and I must have that discussion in the robustness section.

36. Any other data issues
- Do you need to explain any other issues to the reader? For example, if you are linking parents to children across census waves, do you have to explain the algorithm? If you are making districts consistent over time, do you have to explain how? If you are merging GIS data to survey data, do you need to explain how? How do you take lots of different things measured at lots of different levels and turn them into a coherent row-by-column dataset you can run regressions with?

37. Summary statistics
- Put a table of summary statistics here. Format all your tables in the same format used by the journal you are targeting.
    - Also: use the journal you are targeting as a guide for what to include, e.g. what to report in the summary statistics, whether to report standard errors or p values, whether to use asterisks, how many decimal places to report.
    - Write numbers like an American. Commas are for thousands not decimals. Commas separate thousands, not lakhs and crores.
- This table is useful a) to spot possible issues with the data (do any variables have implausible means or ranges?), b) to help readers understand magnitudes in later tables, and c) to remind us of the context (e.g. are average farms in this sample really big? Really small?).

38. Foreshadowing your findings
- This may not be the final paragraph of this section, it may come earlier, as appropriate.
- Find a way in this section to make us anticipate your main result and mechanisms. For example, show maps of your X and Y variables side by side, and note the similarities. Or split your summary statistics by treatment and control before and after treatment, commenting on the similarity of the baseline means and the differences that emerge after treatment.
- This will make your results more credible – if they're visible to the naked eye, we become less worried that they are an artefact of some fancy econometric wizardry.

IV. Results

39. Overview of the results section
- Outline the main points of the rest of the section

40. Outline of main results table
- What results do you find? Are they statistically significant?
- If they are not significant, are they precise nulls? If the confidence intervals are wide, you may need to redesign your study so that it is either conclusively zero or conclusively non-zero, rather than under-powered and inconclusive.
- Are they what you expected? If not, go back and re-write your introduction and mechanisms section. Your results are useless if they cannot be understood. You can't write a paper that finds effects but is puzzled by them; you need a story you can explain.

41. Magnitudes of results
- Are your results big?
- How do they compare to similar effects in the literature?
- How do they compare to other policies that affect your outcome of interest that other researchers have estimated? Are they policy-actionable? Would the policy be cost effective compared with alternatives? If the cost-effectiveness calculation takes more than a paragraph, move it later in the paper and make it a sub-section. That approach is common with randomized controlled trials.
- Does the treatment of interest explain a lot of the variance in the outcome?
- What if the policy you study had never been enacted… would the world look very different?
- How big a policy intervention would be needed to solve the "problem" measured by your Y variable?
- Avoid the trap of looking like you care only about significance and not magnitude. Magnitude matters.

42. Importance of the results
- Here is another place to remind your reader of the story you are telling. How do these results fit?
- What do these results imply for theory or policy?

43. Outline of any additional results
- This is also a place to put results that directly support the main results (so, not quite a mechanism, not quite a robustness check). For example, if you are showing that there is an effect of some X on firm productivity, you can report alternative results with e.g. different measures of TFP or output per worker or revenue productivity here.
- But: too many results will distract a reader. Most papers have one table of main results. If your main results don't fit into a single table, reconsider whether these additional results should instead go in the robustness section, mechanisms section, or in another paper.

X. A general note:

- This sub-section should not only tell me what your results are, but how they support the overall "story" of the paper.

IV.I. Robustness

X. I often make robustness its own section, after the mechanisms, depending on the paper.

44. Principal robustness checks 1

- Key point: what are the main threats to identification, and how do you address these? Some people actually call this section "threats to identification."
- Outline a major threat to identification.
- Outline what empirical exercise you perform to show that it is not a concern
- For many methods, these come from a standard checklist that you can figure out by looking at other papers. For example, if you're using an RDD, you're going to want to show McCrary density tests, placebos with predetermined variables on the LHS, and robustness to changes in bandwidth.
- You find out what threats to validity there are and what robustness checks to run by a) discussing your paper with multiple people and b) reading a large number of papers with similar methods. For ideas, look at recently published papers in top journals that use approaches similar to yours. Both those that look at similar questions and those that look at completely different questions but use similar methods. Make a list of their robustness checks, and figure out which ones you can apply to your data.
- Some classic examples:
  - Do your results survive adding additional controls? Do they seem to hold for all the data, or only for certain industries, or countries, or years?
  - Are your results driven by outliers? Are they robust to discarding certain sub-samples? To alternative measurements of the variables of interest? To alternative means of clustering?
  - What about unobserved variables, or other interpretations of your findings? Can you dismiss these?
  - If your story truly explains the result, what similar treatments shouldn't matter (placebos) and what other samples should be untreatable (falsification exercises)?
  - Can you find extremely narrow comparisons in which your results should still hold (e.g. two children of the same mother, two counties on opposite sides of a state border…). Note: this would use a different source of variation than your main result.

45. Principal robustness checks 2

- You want to have at least three major empirical exercises or three paragraphs worth of evidence on robustness, otherwise this section will appear thin, and the reader will think you aren't very careful, or thorough, or that you don't read very much.

46. Principal robustness checks 3

- You want to have at least three major empirical exercises or three paragraphs worth of evidence on robustness, otherwise this section will appear thin, and the reader will think you aren't very careful, or thorough, or that you don't read very much.

47. Additional robustness checks

- After a point, however, the audience will think the section is dragging on. Construct as long a "to do" list as possible, but put everything that isn't so high priority that it comes in the first three paragraphs into the appendix. Have one paragraph in the text summarizing this appendix material.

- How do you decide what is "high priority"? Generally, any particular method has the standard, mandatory checks – e.g. McCrary density tests with a regression discontinuity. And there will be context-specific issues that are key in your study, and that come up every time you talk about the paper. Lead with those.

X. A general note:

- While some parts of this sub-section cannot help but read like rote statistical exercises, it becomes easiest to read if you frame it around ruling out alternatives to your story. Motivate each exercise. Why might the failure of each robustness check invalidate your story?

V. Mechanisms
48. Overview of mechanisms section
- Outline the main points of the rest of the section
- Briefly re-capitulate the "story" you tell, and how this whole section supports it.

V.I. Model/Conceptual Framework
49. Related theories
- What theories of Y (or of the effects of X) exist? How do they relate to your conceptual framework?
- This paragraph serves as motivation: How does your result teach us about "economics", i.e. broader explanations of human behaviour proposed by economists?
- It also serves as signalling: you demonstrate you've done your due diligence and thoroughly read the economics relevant to your topic at hand: you're a deep thinker and not a lazy student so you know where your theory fits in the literature.
- It also serves as acknowledgement of the literature: no theory or conceptual framework emerges out of the ether, so you explain how past contributions have led you to this.

50. Conceptual framework or model: Environment
- Many authors prefer to make the conceptual framework or model a whole section, rather than part of the mechanisms section. And they prefer to have it come before the empirical results. You should decide this based on who your audience is. Do referees in this field believe that an empirical paper principally has value based on which theories it allows us to support or reject?
- When I was a graduate student, every paper needed to have a model for the sake of having a model. Partly, a model forces you to think formally about why X may lead to Y. In terms of signalling for a thesis or paper, it also shows a) you're not mathematically incompetent, b) you're a "real economist", whatever that is, and c) you learned something in all your modules. Not all papers have models, though, so I generally see three approaches:
  - *Model:* This usually won't accomplish anything a conceptual framework can't, but audiences still want to see one. Obviously, for a more theory-heavy paper or a paper framed as a test of a theory, this should have its own section that precedes the data. It will serve the benefits of making your argument explicit and clarifying what you mean. It can also make counterintuitive results clearer and help you check if your intuition is in fact logically consistent.
  - *Conceptual Framework:* Just because you don't write a model in math doesn't mean your story doesn't have to be internally consistent. So, drawing on the

literature, can you tell a story that explains your results? Putting this in one place will force you to be logically consistent. Nunn and Wantchekon's AER paper on mistrust has a good example of a "conceptual framework" section.

- *Possible explanations for the main result, discussed in sequence:* If there are 2 or 3 plausible explanations that your data can't distinguish, but all of which are interesting, you can discuss them in turn, stating how well supported each is given the limited evidence you do have. This is often the weakest way to write this section, because it doesn't show that you've thought about anything deductively or formally, and several of these explanations may be inconsistent with each other. If you can't test between them, it is seen as a weakness of your paper (you can't "explain" the result convincingly).

- I have set this subsection up as "the model in five paragraphs". But this structure is artificial. Your model or conceptual framework will take as long as it takes. But it needs to be logically ordered, hence this guide.
- The conceptual framework also needs to be sufficiently well thought out, even if it is verbal, that it could fill five paragraphs.
- This first paragraph is one of "setup." Here, you explain what the constraints are that people face. For example: what do the utility functions, production functions and budget constraints look like?

51. Conceptual framework or model: Incentives
- This paragraph is another "setup" paragraph. Now that we know the environment, put it all together to formulate the agent's incentives. For example, you might write the Bellman equation here.

52. Conceptual framework or model: Behaviour
- This paragraph is half way to the results. How do agents behave? For example, what are the first order conditions, and what trade-offs do they highlight?

53. Conceptual framework or model: results
- Now you state the results. What, for example, are equilibrium wages and prices?
- If you are actually doing this as a model in math, then this will be the part where you list your propositions. Any proposition in math must also be stated in words. Indeed, any math that contains no economic insights can be put in the appendix.

54. Implications of the framework or model for data
- Any story you tell will have more implications in data than what you showed as your main result. For example: if you claim that India's midday meals program increases educational attainment because better-fed kids can pay more attention in class, then you should be able to see test scores and learning outcomes increase as well.

X. A note on the conceptual framework
- A major problem I see with the "mechanisms" section in many student papers is that it is a long list of additional results that have no clear purpose. The purpose of the mechanisms section is to tell me your story and to show me you have evidence in favour of that story and against alternative stories. You can't do that if you don't start with a model or conceptual framework.
- So, here's one brute force way to write the conceptual framework.
  - Make a numbered list of every result you're thinking of putting in this section.

- Start two paragraphs. The first paragraph begins with the sentence "The reason X affects Y in our data is …". The second paragraph begins with "There are several alternative explanations of the link between X and Y. The first is…"
- Go back to the first paragraph. Finish that first sentence. For example, "The reason maternal education affects child survival in Nigeria is that it increases women's ability to bargain for a higher share of household resources to be spent on their children."
- Then follow it up with a sentence that starts with "As evidence, we show that…" For example, "As evidence, we show that children in households affected by UPE receive more vaccinations and miss fewer meals; further, their mothers are more likely to report that they have decision-making power on several dimensions of household spending."
- For additional results in the numbered list that support your story, add two additional sentences. The first sentence tells me how the result furthers your story. The second tells me the result. For example: "This bargaining power effect should be greatest for women whose bargaining power would have been low without UPE. Consistent with this, we show our effects are greatest in ethnic groups with more patriarchal gender norms."
- Now finish the second sentence from the second paragraph. For example: "The first is that education raises women's awareness of available medical treatments."
- Then follow it up with the results from the numbered list that rule out that story. For example: "We show, however, that women in our sample who were affected by UPE show no greater knowledge of recommended vaccines."
- Go through every numbered result. If you can't tell me how it furthers your story or rules out an alternative, it does not belong in the mechanisms section. Make sure that the relevance of every result you present to your overall story is contained in your conceptual framework.
- Now, go back and fill in the gaps in the logic. In the example above: why, exactly, would education increase bargaining power? Do these women have higher status? Are they better at negotiating? Does the "why" have any testable implications you can document in data?
- If you fill in enough gaps in the story and order the discussion logically, you will come close to the five paragraph approach above.
- How do you know in which order to introduce your numbered results? For those that support your story, you can put proximate causes first (e.g. education increases child survival because children of educated mothers get more vaccines), deep causes second (e.g. these children get more vaccines because their mothers have more bargaining power), and supporting outcomes third (e.g. the effect is biggest among more patriarchal groups).

V.II. Mechanisms: Evidence

55. Evidence on mechanisms 1
- The four most common ways that I see papers provide evidence on mechanisms are:
  - Heterogeneity:
    - Who responds most to treatment? Showing in the data who responds most helps tell us why you find what you find. This exercise will be useless if it is not used to support any particular story, so keep your story in mind while doing this.

- For example, if you regress investment on land rights and claim that "restrictions on land rights affect investment incentives because they make it hard to use land as collateral," then the link between land rights and investment should be stronger where there are banks. Without banks, nobody can use land as collateral, so the correlation should be nil.
  o Additional outcomes:
    - Put other variables on the LHS. Sometimes these will come from the same dataset and will be in the same sample, sometimes they won't, depending on data availability.
    - For example, if you regress investment on land rights and claim that "restrictions on land rights affect investment incentives because they make it hard to use land as collateral," then you should find that households with more land rights have more loans from banks.
  o Mediation analysis:
    - If your mechanisms can be measured in the same dataset and sample as your main result, you can feed them through a "mediation analysis" (e.g. see the documentation for medeff or ivmediate in Stata). Does controlling for them "explain" the main result?
    - For example, if you regress investment on land rights and claim that "restrictions on land rights affect investment incentives because they make it hard to use land as collateral," then you should find controlling for loans from banks should make the coefficient on land rights smaller.
  o Additional empirical exercises:
    - Sometimes there are other empirical results that don't fit neatly into the above but support your story. For example, one paper (an older version of Iyugin et al.) claims that potatoes reduced conflict in medieval Europe. They claim that one mechanism is that potatoes were more drought resistant and so reduced the sensitivity of conflicts to droughts. They basically run a regression of conflicts on drought, potatoes, and potatoes X drought, showing the interaction is significant and reduces the "main effect" of drought.
  o A warning
    - Nothing in this section is for its own sake. Don't split the sample by gender if we don't learn anything about mechanisms due to the split. For every piece of empirical evidence you present here, make it clear how it supports your "story." This section does not exist for you to just do more stuff.

56. Evidence on mechanisms 2
- You want to have at least three major empirical exercises or three paragraphs worth of evidence on mechanisms, otherwise this section will appear thin, and the reader will think you are making up a story or explanation without evidence.

57. Evidence on mechanisms 3
- You want to have at least three major empirical exercises or three paragraphs worth of evidence on mechanisms, otherwise this section will appear thin, and the reader will think you are making up a story or explanation without evidence.

58. Additional evidence on mechanisms

- After a point, however, the audience will think the section is dragging on. Construct as long a "to do" list as possible, but put everything that isn't so high priority that it comes in the first three paragraphs into the appendix. Have one paragraph in the text summarizing this appendix material.

V.III. Alternative Mechanisms

59. Alternative mechanism 1
- There are alternative explanations of your result. Show how you rule them out in data.

60. Alternative mechanism 2
- You want to have at least three major empirical exercises or three paragraphs worth of evidence on mechanisms that you rule out, otherwise this section will appear thin, and the reader will think you haven't thought very carefully about alternatives or have been lazy in ruling them out.

61. Alternative mechanism 3
- You want to have at least three major empirical exercises or three paragraphs worth of evidence on mechanisms that you rule out, otherwise this section will appear thin, and the reader will think you haven't thought very carefully about alternatives or have been lazy in ruling them out.

62. Additional alternative mechanisms
- After a point, however, the audience will think the section is dragging on. Construct as long a "to do" list as possible, but put everything that isn't so high priority that it comes in the first three paragraphs into the appendix. Have one paragraph in the text summarizing this appendix material.

A general note: Again, this section becomes easiest to read if you frame it around ruling out alternatives to your story.

Other things that can go in the mechanisms section:

- *Qualitative evidence:* This will be particularly relevant for you if you're doing economic history as a field. Plan on also visiting archives and obtaining information on people affected by the event you study. You can't incorporate this information into your main results, but it will help readers believe your story if detailed descriptions of some cases sound like your story and are consistent with the main results.
- *Case Studies:* If you are doing something on a large cross-section of heterogeneous places, it can be hard to think of a "one size fits all" story. If the observations that have the most statistical influence all have some common features according to the secondary literature, you can make sure your story is consistent with these features, and with the experiences of these cases. For example, if you were studying the effects of drought on civil war, you could find three conflicts in your data that were preceded by droughts and then describe what the secondary literature says about them.

VI. Conclusion
63. Summary of what you did
- Briefly reiterate your setting and identification strategy.

64. Summary of what you found

- Briefly reiterate your main results, what they are robust to, the mechanisms that explain them, and the empirical evidence for these mechanisms.
- Re-capitulate your story here.

65. Why what you found is important for the literature

- Why did no one write this paper before? An answer to this question helps explain why what you do is novel and important and needed right now.
- This paragraph is partly a summary of what you say in the "contribution" section of the introduction, but you can be a bit "bigger picture" here. For example, we have learned more about the general theoretical/conceptual literature that you have discussed at the start of the "mechanisms" section.
- Do your results require us to revise our existing views? Do they support one set of theories, but not another?

66. Why what you found is important for policy

- If your results have no policy implications, then policymakers won't fund this sort of research justify your existence. What is the real world importance of what we've learned, and how should we act on it?
- But, be careful. Don't just go out and endorse a policy. There's always costs and benefits not included in your paper. And: you're not a partisan hack. I hope.
- Instead, note that your results imply that XYZ is a likely outcome of policy ABC.
- If you're not evaluating a policy, note instead that you've identified a relationship (or lack of one) between X and Y. Note there are policies that propose either changing X, or changing Y. Note what your results imply about these policies.
- That said: this discussion should not be weak speculation or it can be grounds for a referee to reject. These policy implications will require rigorous thought.

67. What you couldn't do, but later papers should

- What would you have liked to do ideally, but that the data did not permit?
- This isn't "things I could have done better," but "things it is impossible to do better given current data availability." You don't want to give the grader or referee ammunition for marking you down or rejecting your paper. Instead, you want to a) show that you are not over-claiming and b) set up the rest of the paragraph.
- Suggest that later work can overcome some of the limitations above.
- If you have an ongoing agenda, this is a good chance to advertise your works in progress as contributions that will soon appear that help overcome these limitations.
  - Downside: referees or an editor may ask you to add it to the current paper.
- What are the next steps for future research? Try not to end on a cheesy set of words like "for future research" or "in the future."
- That said: this discussion should not be weak speculation or it can be grounds for a referee to reject. These recommendations for future research will require rigorous thought.

A typical running order of tables:
- Table 1: Summary statistics
- Table 2: Main Result
  - A nice graph of the main results is very important. This is often dictated by the method (e.g. the event study with a DID, a discontinuity plot with an RDD). But

a) with these methods and b) with other methods, if your main result is clearly visible in a simple depiction of data, in as close to a "raw" form as possible, you should put the figure in the introduction.

- Table 3: Most important robustness checks (the rest go in the appendix)
- Tables 4-6: Mechanisms and Other Cool Results
- There's no good excuse for having more than 9 tables. I think 5-7 is the norm. Appendices can have infinite length. Save the tables and figures in the paper for what's important.

On your tables and figures: you should make sure that they are self-explanatory. Things like units of measurement should be there. Many readers will skip the text and just look at the tables and figures. Look at how detailed the table and figure notes are in recent papers in top journals. Yours should be equally detailed.

Help! My paper is too short!

If your paper is less than 8,500 words, it will stand out for being too short. So: here's a bunch of questions that you should answer in your paper. Skim these. Pick any you haven't answered and put the answer in your paper. These answers ought to get you over 8,500 words. And if your paper is already 8,500 words, you should still make sure you've answered as many of these questions as you can.

Questions for the introduction
1. What is the question you ask?
2. Why is that question important to answer?
3. Why hasn't the question been answered before?
4. What do you do?
5. Why is your approach the right one?
6. What do you find?
7. Are your results big?
8. Are your results robust?
9. What explains your results?
10. Why is this economics?

Questions on the literature
11. What is the closest study to yours and how is your study a contribution over and above what it has done?
12. What results are already known in this literature?
13. What gaps remain in this literature?
14. What difficulties encountered by others explain why there are gaps in this literature?
15. What gaps in the literature do you fill?
16. Why did no one do what you are doing before? That is: how did you overcome the difficulties others faced?
17. How does your work build on papers by the people who are most likely to be your referees?
18. How does this paper build on work that has recently appeared in your target journal?
19. How does this paper build on work by members of the editorial board of your target journal?
20. Thinking beyond your narrow subfield, what do we learn about "economics" from this study?
21. How does this study connect to the rest of your coherent research agenda?
22. What do other fields have to say about your question of interest? Are you contributing anything to these literatures?

Questions on the context

23. What was the motivation for your treatment of interest (e.g. why did the government pass the law you are using in your diff in diff? why was the eligibility cut-off that you are using in your RDD defined at that level?)?
24. What are the main drivers of the outcome variable in your study (e.g. "In India, access to clean water is much poorer in rural areas, especially those without paved roads, and among the lower castes.")?

25. What is the institutional framework that governs your outcome variable? (e.g. "In India, improved water pumps are largely purchased by households using their own funds, though 23% of Gram Panchayats have organized communal pumps in the village square.")
26. The same two questions as above, but this time for your right hand side variable.
27. What are the important differences between your context and the US (where most of your referees are) that need to be explained for referees to understand the paper?
28. How does the context motivate your choice of confounders?
29. How does the context motivate your thinking about mechanisms and how you measure them?
30. What confounders or mechanisms might be relevant in other contexts but do not apply in your context? Why?
31. Why should we expect results obtained in your context to be generalizable to other contexts?

Questions on the identification strategy

32. Why is an identification strategy needed? What would the biases be from vanilla OLS, or from something simpler?
33. How does your identification strategy approximate an experiment?
34. What is your estimating equation?
35. What does every term in this estimating equation mean?
36. What forms of omitted variables bias or other threats to identification does your identifications strategy address? e.g. "Classroom by year fixed effects control for any unobservable factors that are constant across children within a given year, such as teacher quality."
37. What are your identification assumptions? What are they, intuitively speaking?
38. What would a violation of the identifying assumptions look like? How likely is it? Can you rule it out?
39. Why do you include the controls that you do?
40. How do you cluster your standard errors?
41. Why do you use the timing that you do? Should the RHS variable be lagged one period? Should the LHS and RHS variables be first-differenced? Should contemporaneous shocks matter, or a sequence of them?
42. Who are the "compliers" in your natural experiment and how do they differ from the general population? How does this affect its external validity? How does it affect interpretation?
43. Are there any alternative ways to approach this same question? Will you use them for robustness? Is your approach better?
44. What other studies are closest to yours in terms of identification strategy? Do you make any innovations or improvements over the standard in the literature?

Questions on the data

45. How were the data on the outcome variable collected, either by you or by the data provider? If someone else collects them… what are they collected for?
46. Why should we believe these data are accurately reported?
47. How are these data available? From an archive? From IPUMS? From the FAO website?
48. How are your key variables defined? Why? Is this the way other studies measure the same concepts? Include equations.

49. Who else has used these data? Why aren't they more widely used?
50. What work went into making these data suitable for a regression?
51. Why is this the best data source to use for this question? How do these data compare to other relevant datasets?
52. How representative is this sample of the general population?
53. What are the possible errors in these data and how do you deal with these potential problems?
54. Do you need to focus on any particular subsample? If so, what are the benefits and costs of doing this?
55. What alternative data sources exist? Will you use them for robustness? If not, is there some advantage of your data?
56. Why do you assume the functional form that you do (linear, log, etc…)?

Questions on the results

57. If I were to take a single number away from this paper as the headline result, what would it be?
58. How do your magnitudes compare to those in other studies?
59. How "big" are your results?
60. How "precise" are your results? What are the biggest and smallest effects you can rule out?
61. Can you restate the magnitude of your results in another way to make the intuition or thought experiment behind them clearer?
62. Are your results the same across major subsamples of the data?
63. How important is X? That is, what would the counterfactual work look like if treatment had never happened?
64. Why is this result not obvious? Why did we need a study to tell us this?

Questions on robustness

65. How sensitive are your results to controls?
66. Are your results driven by a single subgroup of the data (a set of countries, a set of districts, a single gender, a single time period, a set of outliers)?
67. How sensitive are your results to functional form assumptions?
68. How sensitive are your results to the data source or variable definition?
69. How sensitive are your results to alternative measures of your standard errors (e.g. randomization inference, alternative clustering, spatial dependence…)?
70. Have you run the robustness checks that are standard to your identification strategy (e.g. for an RDD the McCrary density test, sensitivity to alternative bandwidths, putting controls on the LHS; for IV, the first stage F stat, a placebo exercise, an overidentification test of you have more instruments than endogenous variables).
71. Could there be errors in your variables definitions or in the data itself that systematically bias you towards finding your result (e.g. RHS measurement error if you find a null effect; if X increases reporting rates, could greater reporting make it look like X increases Y when it really decreases Y?).

Questions on the mechanisms

72. Which theories do your results support and which do they not support? Base your answer on implications of theoretical models.

73. If you could write a model that predicted your main result as an outcome, what would the assumptions be and what would the logic be that gets us from these assumptions to the result?
74. What is the secondary evidence in favour of your interpretation of the results? This will come from the literatures in other disciplines. This might be qualitative if others who study the phenomenon you study are e.g. historians, anthropologists, political scientists... If, instead, you're on the border of medicine, this evidence may be quantitative.
75. What is the primary qualitative evidence in favour of your interpretation of the results? This can come from many sources: newspapers, participant interviews, internal email records of the firms you're studying, transcripts of speeches in parliament, letters by the colonial governor to the district officer... This is how other disciplines do research and you will look more credible by using what they consider evidence rather than reporting a bunch of correlations and interpreting them as if you are an armchair theorist.
76. Why should we believe your story? What hard evidence do you have for this interpretation rather than any other?
77. How robust is this evidence? Draw on questions from the previous section.
78. What are the other plausible interpretations of your results, and how can we rule them out?
79. Why are your results context-specific? If you argue, for example, that clans mitigated risk in historical China and so spurred population growth, why don't we see clans that do the same thing in Europe?

Questions for the conclusion

80. What is the main takeaway from your study?
81. How do we need to revise our current understanding based on your results?
82. What are the limitations of your approach?
83. What unanswered questions do you leave for future work?
84. What are the implications of your results for policy?
85. How will later researchers be able to use your methods, results, data, etc...?

A Cookie Cutter Approach to Applied Microeconomics

This section is intended to provide an overview of the "must do" exercises for three workhorse methods of applied microeconomics: instrumental variables, regression discontinuity, and differences in differences.

Instrumental Variables

Intuition
- Use an instrumental variable "Z" to isolate exogenous variation in your treatment of interest, "X." Z must predict X, but must be uncorrelated with the unobserved determinants of your outcome, "Y."

Overview
- Undergraduate level: Wooldridge, "Introductory Econometrics: A Modern Approach." Chapter on Instrumental Variables Estimation and Two Stage Least Squares.
- Postgraduate level: Hansen, "Econometrics." Chapter on Instrumental Variables.
- Recent guide to best practice: Cunningham, "Causal Inference: The Mixtape," chapter on Instrumental Variables."
- Stata guide: Baum, C. F., Schaffer, M. E., & Stillman, S. (2007). Enhanced routines for instrumental variables/generalized method of moments estimation and testing. The Stata Journal, 7(4), 465-506.

Three Recent Papers
- Heldring, L. (2021). The origins of violence in Rwanda. *The Review of Economic Studies*, *88*(2), 730-763.
- Alsan, M. (2015). The effect of the tsetse fly on African development. American Economic Review, 105(1), 382-410.
- Heblich, S., Trew, A., & Zylberberg, Y. (2021). East-Side Story: Historical Pollution and Persistent Neighborhood Sorting. Journal of Political Economy, 129(5), 1508-1552.

Mandatory Exercises
- Discuss separately a) why the instrument is as good as random, conditional on controls, and b) why it is unlikely that there are alternative channels, other than X, through which the instrument should matter.
- Compare the OLS and IV coefficients, justifying any differences in terms of measurement error, omitted variables bias, or differences between the LATE and ATE, while ruling out violations of the exclusion restriction.
- Show the first stage and reduced form (e.g. Heldring (2021, Table 4)).
- Show the appropriate first stage F statistic, usually Kleibergen Papp, which must be above 10 (e.g. Heldring (2021, Table 4)).
- Show a "placebo": there is no correlation between Z and Y where Z cannot affect X (e.g. Alsan (2015, Table 5)).

Helpful Exercises
- Show the instrument is uncorrelated with observables (e.g. Heldring (2021, Table 3)).
- Conley's "plausibly exogenous" test: show small violations of the exclusion restriction cannot eliminate the main result (*plausexog* in Stata).

- Report confidence intervals robust to weak instruments (*weakiv* in Stata).
- Test the exclusion restriction directly, following D'Haultfœuille et al. (2021) (*testex* in Stata).
- If you are using IV in the context of another method with weak compliance (e.g. in an RDD, not everyone on the other side of the cutoff is treated; in a diff in diff, not everyone in the treatment group receives treatment in the post period), then subject your reduced form specification to all the robustness checks that are standard for that other method.

Regression Discontinuity

Intuition
- Use the sharp change in X induced by a running variable crossing a cutoff in order to identify the effect of X on Y.

Overview
- Undergraduate level: Cunningham, "Causal Inference: The Mixtape." Chapter on Regression Discontinuity.
- Postgraduate level: Hansen, "Econometrics." Chapter on Regression Discontinuity.
- Recent guide to best practice: Cattaneo et al., "A Practical Introduction to Regression Discontinuity Designs: Foundations."
- Stata guide: Calonico, S., Cattaneo, M. D., Farrell, M. H., & Titiunik, R. (2017). rdrobust: Software for regression-discontinuity designs. *The Stata Journal*, *17*(2), 372-404.

Three Recent Papers
- Montero, E. (2022). Cooperative property rights and development: Evidence from land reform in El Salvador. *Journal of Political Economy*, *130*(1), 48-93.
- Lowes, S., & Montero, E. (2021). Concessions, violence, and indirect rule: Evidence from the Congo free state. *The Quarterly Journal of Economics*, *136*(4), 2047-2091.
- Dell, Melissa, and Benjamin Olken. The Development Effects of the Extractive Colonial Economy: The Dutch Cultivation System in Java. Review of Economic Studies 87, no. 1 (2020): 164-203.

Mandatory Exercises
- Show several figures depicting the RDD (*rdplot* in Stata, e.g. Lowes and Montero (2021, Figure 4)).
- Use a local linear polynomial and the optimal, MSE-minimizing bandwidth (*rdbwselect* in Stata) as the baseline; show robustness to alternative polynomials and bandwidths (e.g. Lowes and Montero (2021, Table 2)).
- Do McCrary density tests (*rddensity* in Stata) that show there is no bunching around the cutoff.
- Report balance tests that show that predetermined observables do not change at the cutoff (e.g. Lowes and Montero (2021, Table 2)).

Helpful Exercises
- If compliance is imperfect at the cutoff, estimate a fuzzy regression discontinuity with the dummy for being on the "treatment" side of the cutoff as the excluded instrument (e.g. Montero Table 4).
- Check: In a binned scatterplot with no confidence intervals and no visual indication of the cutoff, is the discontinuity obvious?
- Use bounding exercises to show that selective sorting around the cutoff would need to be severe to explain the result (e.g. Lowes and Montero (2021, Appendix)).
- Use alternative cutoffs as placebos (e.g. Lowes and Montero (2021, Appendix)).
- If the RDD is spatial, Conley standard errors (*acreg* in Stata).
- Discuss how "local" the estimate is. If the cutoff were at a different level, would the treatment effect be the same? For example, if a birthweight of 3,000 grams rather than 1,500 grams were used to assign newborns to intensive care, the effect on later

outcomes might be smaller. This involves understanding why the cutoff was set as it was.

- Show robustness to changing the kernel (triangular, uniform).
- Do a "donut hole" dropping observations very close to the cutoff.

Differences in Differences

Intuition
- Compare the change over time in treated units to the change over time in untreated units to assess the effect of treatment.

Overview
- Undergraduate level: Wooldridge, "Introductory Econometrics: A Modern Approach." Chapters on "Simple" and "Advanced" Panel Data Methods.
- Postgraduate level: Hansen, "Econometrics." Chapters on Panel Data and Difference in Differences.
- Recent guide to best practice: Kahn-Lang, A., & Lang, K. (2020). The promise and pitfalls of differences-in-differences: Reflections on 16 and pregnant and other applications. *Journal of Business & Economic Statistics*, *38*(3), 613-620.
- And another: Miller, D. L. (2023). An Introductory Guide to Event Study Models. *Journal of Economic Perspectives*, *37*(2), 203-230.
- And another: "Two-Way Fixed Effects and Differences-in-Differences with Heterogeneous Treatment Effects: A Survey" by Clément de Chaisemartin and Xavier D'Haultfoeuille (2022): not yet published, but submitted to The Econometrics Journal.
- Stata guide: One example is Villa, J. M. (2016). diff: Simplifying the estimation of difference-in-differences treatment effects. *The Stata Journal*, *16*(1), 52-71, but you may be better off just looking at the help file for *reghdfe*.
- Stata guide for event studies: Clarke, D., & Tapia-Schythe, K. (2021). Implementing the panel event study. *The Stata Journal*, *21*(4), 853-884.

Three Recent Papers
- Anderson, D. M., Brown, R., Charles, K. K., & Rees, D. I. (2020). Occupational licensing and maternal health: Evidence from early midwifery laws. *Journal of Political Economy*, *128*(11), 4337-4383.
- Giorcelli, M., & Moser, P. (2020). Copyrights and creativity: Evidence from Italian opera in the Napoleonic age. *Journal of Political Economy*, *128*(11), 4163-4210.
- Bau, N. (2021). Can policy change culture? Government pension plans and traditional kinship practices. *American Economic Review*, *111*(6), 1880-1917.

Mandatory Exercises
- With more than 2 pre-treatment periods: Do an event study plot, interacting "treatment" with the time periods, with the final pre-treatment period as the omitted category. Insignificant coefficients before treatment show that trends were parallel before, and so support the assumption that they would have remained parallel were it not for treatment (e.g. Bau, 2021, Figure 3). It should be clear from the figure when treatment begins, even without a vertical bar at zero.
- If not a randomized controlled trial: include fixed effects for the cross-sectional units at which treatment varies, and for time periods (e.g. individual and year fixed effects).
- In a randomized controlled trial with both pre-treatment and post-treatment data, simply controlling for "Treatment" and "Post" may be sufficient due to the randomization. The full set of dummies can increase precision by reducing the residual variance, or lower it by reducing the degrees of freedom.
- Cluster by the cross-sectional units at which treatment varies (e.g. states in Anderson et al, 2020, Table 2).

- o Though, note that Abadie et al. (2023, QJE) suggests this may be too conservative and proposes a CCV estimator instead
- Show that future lags of treatment do not predict changes in the outcome (e.g. Anderson et al, 2020, Table 2).

Helpful Exercises
- Drop non-switchers (e.g. Anderson et al, 2020, Table 3).
- Use tighter-than-necessary fixed effects, e.g. state-by-year rather than year in an application where you have a panel of counties and X varies across counties within a state in a given year.
- "Falsification": Show that outcomes that should not respond to treatment do not (e.g. Anderson et al, 2020, Table 4).
- Balance: Show treatment and control are similar before treatment (e.g. Giorcelli and Moser, 2020, Table 1).
- Add linear trends for cross-sectional units (e.g. Giorcelli and Moser, 2020, Table 3).
- With a staggered difference in differences (treatment starts or ends at different times for different units), there is now a recent literature arguing that a simple event study approach conflates several different treatment effects. de Chaisemartin and D'Haultfœuille 2020; Goodban-Bacon 2019; Abraham and Sun 2020; Callaway and Sant'Anna 2020 are the key references. Some of these have Stata packages (*did_multiplegt*, *eventstudyinteract, eventstudyweights*) or R packages (*did*).
- Employ synthetic control (*synth* in Stata).
- Use differences in differences with matching (*diff* in Stata).
- Use "doubly robust" difference in differences (*DRDID* in R).