BIG DATA AND PUBLIC POLICY
PPHA 3059 / SSA 4059

Course Planning Guide

Learning Goals, divided among
1. Readings
2. Recorded lectures
3. Class meeting

Goals of Lectures:
1. Clarify difficult points from the readings
2. Draw connections across the various readings
3. Fill in gaps/context of what we didn't read

<u>Week 1: Introduction</u>
In the introduction to the course, we raise big, overarching questions about the role of evidence in policy making, including how evidence stacks up against other forms of policy making knowledge, what types of evidence are used for addressing different kinds of policy problems, and how the emergence of big data and new kinds of analytic methods might be changing how we use evidence in policy making. We explore these topics in greater detail in the later weeks of the course.

Readings:
- Kingdon, John W. 2003. Agendas, Alternatives, And Public Policies. Chapters 4 & 6.
- Kitchin, Rob. 2014. The Data Revolution: Big Data, Open Data, Data Infrastructures & Their Consequences. Chapter 7.
- Mayer-Schoenberger, Viktor, and Kenneth Neil Cukier. 2013. "The Rise of Big Data." Foreign Affairs.

<u>Learning Goals for this Week</u>
1. What does "evidence driven policy" mean in a big data world?
2. What are the alternatives to evidence-driven policy?
3. How does evidence fit into the overall policy-making process? What are the barriers to evidence-based policy?
4. Why big data policy now?

<u>How We Cover Each Learning Goal</u>
1. What does "evidence driven policy" mean in a big data world?
    a. Reading: Kingdon
    b. Reading: Kitchin
    c. Lecture: We give some high-level thoughts on this, e.g.,there's this new interest in big data to inform policy, and that's what this course is going to focus on

    d. Small-group activity: students provide definitions, and we discuss in class
2. What are the alternatives to evidence-driven policy?
    a. Lecture: History of alternative approaches to policy making, which have been significantly displaced by evidence driven policy
    b. Lecture: History of evidence driven policy
3. How does evidence fit into the overall policy-making process?
    a. Reading: Kingdon
    b. Lecture: Unpacking of Kingdon specifically on whether policymakers really care about evidence, and what they also care about
4. Why big data policy now?
    a. Reading: Mayer-Schoenberger & Cukier
    b. Reading: Kitchin
    c. Lecture: This is the end of the history of evidence-based policymaking

Agenda for Lecture

1. Bear in mind that as you're developing your expertise in data crunching, the answers you get from your analysis will not determine how policy gets made. There's a bunch of other things involved in the policymaking process, evidence is only one of them
2. It wasn't always this way, that we thought we should make policy based on quantitative evidence -- a little history
    a. Some examples of what were thought to be good ways to make policy:
        i. Professional policymakers or politicians
            1. Expertise
                a. substantive knowledge related to a policy area (You can be an expert without having experience)
                b. Policy knowledge about a substantive area
            2. Feasibility
                a. Implementation capacity
                b. Money
                c. Politics -- how to get folks to agree on something
            3. Ideology (worldview)
            4. Politics -- party politics or conservative/liberal?
        ii. Other kinds of policy evidence (other than quantitative, experimental)
            1. Policy implementation research
            2. Ethnographic research on conditions we want to affect with policy
                a. Tally's Corner, Making Ends Meet, etc.
3. Tell the history of evidence-based policy, and the story of the Harris School: takeaway is the development of the ideas behind evidence-based policymaking (Angrist & Pischke)
    a. Progressive Era and the Bureaus of Municipal Research
        i. Effort to displace machine politics and replace with data-driven administration (Nicole)
    b. Economics and the Harris School (Chris)

        i. 1986 paper by Bob Lalonde shows quasi-experimental analysis can't recover correct experimental results
        ii. Economics boundary-setting: this is a policy-relevant thing for economics, but it's not economics, it's "public policy" or "program evaluation"
        iii. Credibility revolution paper: program evaluation actually becomes the dominant paradigm for economics overall
    c. The "net impact" triumph in policymaking (Nicole, as told by Breslau)
4. How does evidence fit into the overall policy-making process?
    a. Do policymakers really care about evidence? Of any kind? (Kingdon)
        i. Kingdon: garbage-can model shows that just because you have a good policy idea doesn't mean that it actually becomes policy
5. Why big data policy now? (Kitchin, Mayer-Schoenberger & Cukier)
    a. Popularization of data science; use of big data in business (connect to prior iterations of business techniques imported to government, e.g., New Public Management)
    b. Computational power
    c. Data availability
    d. Government already primed for evidence-based policymaking
    e. Spread of big data from the corporate world into the government world
    f. Politicians coming from business background, or they think consulting firms are more credible than academics
    g. Movement of analytics from economics to computer science, stats
    h. Psychological need to address our massive uncertainties :)

Agenda for Class Meeting

1. Personal Introductions: everyone introduces themselves, including Nicole and Chris - 20 minutes
    a. Name
    b. Program and year of study
    c. Why you're interested in Big Data and Public Policy
    d. Policy interests
2. Course introduction by Nicole and Chris - 20 minutes
    a. How the course came to be
    b. Goals
    c. Reading schedule
    d. Assignments
3. Lecture on the stuff above - 30 minutes
    a. Prepared slides, high-level; riff on them in class
4. BREAK - 5 minutes
5. Small group activity -- 20 minutes

    a. Kingdon's model is about problems, policy alternatives, politics. It was written in the 1980s, prior to the "big data" revolution. How does Kingdon's framework need to be revised, if at all, for the contemporary policy environment?
    b. Does the big data revolution increase the likelihood that evidence-based policy alternatives will be selected to solve particular problems? Why or why not?
6. Report out - 30 minutes
    a. Begin with one group reporting, then ask for additions, corrections, different ideas
7. What's on tap for next week?

<u>Week 2: Developing Ethical Competencies for Data Science</u>

In the ethics introduction, we lay a foundation for thinking about ethical problems in general, and explore ideas about ethics specifically related to the use of big data in policymaking. We believe that while regulation may be important in determining how big data might be used for policymaking and in general, that the technology is always going to develop more rapidly than regulation. As such, students must be encouraged to develop their own ethical compass to guide them as they engage in evidence-based policymaking.

Readings

- Cathy O'Neil. 2016. Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy. Introduction.
- Salganik, *Bit by Bit,* Chapter 6.
    - https://www.bitbybitbook.com/en/1st-ed/ethics/principles/
- Hasselbalch, Gry. 2019. "Making Sense of Data Ethics. The Powers behind the Data Ethics Debate in European Policymaking"
- Brayne, Sarah. 2017. "Big Data Surveillance: The Case of Policing." *American Sociological Review* 82: 977-1008.
- Refresher: *Political Economy for Public Policy,* Chapter 1.
- Recommended: Kahn, Falaah Arif & Julia Stoyanovich. "Mirror, Mirror (Data Responsibly, #1)." Comic book in pdf.
- Recommended: Kitchin ch 10

<u>Learning Goals for Week 2: Foreshadow at a high level some of the issues that come up in later weeks</u>

1. Understand that ethics is not a single answer but rather a series of decisions about goods and bads, and a coherent position from which to reason through them.
2. Understand that ethics as applied to policymaking has some existing commitments: equity, effectiveness, efficiency, others?
    a. The aims of public policy: what is the justification for government intervention, for the state, why do people voluntarily give power to the state to coerce them?
        i. Tragedy of the commons, free riding, etc.
    b. We give up some of our autonomy to the state so it can solve some of these problems and improve social welfare

    c. What is the idea of social welfare provided by the state?
        i. Rawls
        ii. Utilitarianism: net gain
        iii. Maybe consequentialism?
    d. Ethan: No final answer, but you have to pick one!
        i. The Perils of Quantification (article): often, in formulating policy we do some kind of cost-benefit analysis; we provide the gain to the winners, but we seldom follow up with compensating the losers (which is what makes it ethical!)
3. Understand the ethics of research, which is an important underpinning of policymaking.
4. Understand that the discussion and pursuit of ethics in research has the additional function of working to stabilize research findings in order to act and make policy.
5. Something about surveillance and control -- an ethics related to big data in particular?

<u>How We Cover Each Learning Goal</u>

1. Understand that ethics is not a single answer but rather a series of decisions about goods and bads, and a coherent position from which to reason through them.
    a. Readings: Salganik
    b. Lecture
    c. Class
2. Understand that ethics as applied to policymaking has some existing commitments: equity, effectiveness, efficiency, others?
    a. Lecture
3. Understand the ethics of research, which is an important underpinning of policymaking.
    a. Reading: Salganik
4. Understand that the discussion and pursuit of ethics in research has the additional function of working to stabilize research findings in order to act and make policy.
5. Something about surveillance and control -- an ethics related to big data in particular?

<u>Questions</u>

1. Different questions

Agenda for Lecture

1. High-level concepts about the ethics of datafication in society that apply across a number of areas.
    a. Duty-based (deontological): “rights” comes from here; morality is based on real truths (“deon” in Greek means “duty”)
    b. Consequentialist: morality is based on actual impacts of our actions
        i. Utilitarianism: big in CS world? Doesn’t look at individual, looks at collective
    c. Virtue: maybe not useful for us

    d. Oxford book: Our attempts to do things for the good of others has been messed up, so now we focus on a human rights framework; focus on individual rights; how do we see if algorithms really benefit individuals (ecological fallacy type thing)
2. Identify specific ethics concepts that relate to policy making (connect to Week 1)
3. Understand the idea that when we make decisions about how to do policymaking (including the use of data analysis) we are making claims about what is important to us as a society.

Agenda for Class Meeting

1. Check-in/community builder: 10 minutes
    a. If we were having a class potluck, what would you bring?
2. 1-minute summary of small group discussion topics:
    a. Teacher value-added (O'Neil)
    b. Alert-based policing practices (Brayne -- see table on p. 986)
    c. Child welfare risk scores for investigating families (lecture)
    d. NYC fire inspector allocation (pp. 35-36 of Mayer-Schoenberger & Cukier)
3. Small group discussion part I -- 20 minutes
    a. How would we articulate "the good" and "the right" for this topic?
    b. Should we be doing this analysis?
    c. Can we do this analysis right?
4. Large group discussion -- 30 minutes
5. Small group discussion part II -- 20 minutes
    a. Based on a-c, would you do something differently, relative to how it was done in our example, and if so what?
    b. What does this imply about the practice of data science in this policy context, or more generally? Is there a way to build structures that support ethical data science practice in this area?
6. Large group discussion -- 20 minutes
7. Closer -- 10 minutes

8. Points we want to make sure come across in large group discussion
    a. Develop a list of principles for "the good" and "the right"
        i. Are any of these principles new or unique to data science/big data?
            1. Privacy?
            2. Transparency?
                a. Understanding a system means you can game it
            3. Accuracy?
            4. Accessibility?
            5. Fairness?
            6. Feasibility?
    b. Disparate impact on different groups of doing this analysis?
    c. Is there a way to build structures for ensuring ethical practice of data science in public policy?

        i. Hasselbalch

<u>Weeks 3-4: Datafication and Data Production</u>

The section on datafication directs students' attention to the many ways in which new forms of data are produced and become available for analysis. The course interrogates the claim made by some big data evangelists that present-day volumes and velocities of data allow a truthful, frictionless representation of the world. We consider this claim, as well as its critiques, encouraging students to be critical consumers of datasets. The process of data production is at bottom a human and social one, and thus shot through with many social constructions and biases. Users of data must be aware of this fundamental reality, as well as be equipped to consider how to address it prior to analysis.

Readings:

- Kitchin, Rob. 2014. The Data Revolution: Big Data, Open Data, Data Infrastructures & Their Consequences (Chapters 1-5, 8)
- Rieder, Gernot and Judith Simon. 2016. "Datatrust: Or, the Political Quest for Numerical Evidence and the Epistemologies of Big Data." Big Data and Society.
- Adcock & Collier on measurement validity

<u>Learning Goals for Week 3</u>

1. What are different sources and types of data?
2. What are general features of data production that are true of all datasets?

<u>How We Cover Each Learning Goal</u>

1. What are different sources and types of data? [use Kitchin's typology]
    a. Reading: Kitchin, chapters 1-4
    b. Lecture:
2. What are general features of data production that are true of all datasets?
    a. Reading: Kitchin, Chapters 1-4
    b. Lecture: Different theoretical perspectives for thinking about data; human intermediation in data production

Agenda for Week 3 Lecture

1. What are different sources and types of data (typology)? [use Kitchin's]
    a. Official statistics, Administrative microdata, Etc.
    b. Numeric, Text, Image, Etc.
2. What are general features of data production that are true of all datasets?
    a. Different theoretical frameworks for thinking about data
    b. Data is always intermediated by humans
        i. Data are never "raw"
        ii. Data attempts to represent something about the world, but how well is it doing?
    c. Is all data useful?

Student Presentation Questions -- Relate to the readings and lecture material, but focus on a particular dataset you are familiar with or using in your work; student answers only questions 1-3; questions 4-5 will be presented during class in small group activity

1. How does your data fit into the typology?
2. Why was the data produced? What was their original purpose?
3. What choices had to be made in that process that affect what data is made available?
4. If about social processes, what do you know about how those social processes unfold, which might affect the data itself?
5. Identify aspects of data production process that are opaque

Agenda for Week 3 Class Meeting

1. Check-in/community-builder - 10 minutes
2. Nicole and Chris each take an example of a dataset and walk through the five questions - 30 minutes
3. Small group activity - 30 minutes
    a. Three groups: one for each of the datasets discussed in student presentation
    b. Expand on student presentation to answer questions 4-5 for each dataset: what more can you learn about the data production process for the student's dataset?
4. Report out to large group - 15 minutes (5 minutes per group)
5. Open discussion - 30 minutes
    a. Are there aspects of the typology that should be added?
    b. Are there other aspects of the data production process that we should be interrogating?
6. Closing chat statement: what is a question you still want to have answered?

<u>Learning Goals for Week 4</u>

1. How do we evaluate the "quality" of different kinds of data?
2. What are the legal/ethical implications of data production and use?

<u>How We Cover Each Learning Goal</u>

1. How do we evaluate the "quality" of different kinds of data? [Probably are general papers on this -- **CB to look for-- Maybe Adcock and Collier on Measurement Validity?**]
    a. Reading: Kitchin
    b. Lecture: Aspects of data quality
2. What are the legal/ethical implications of data production and use?
    a. Reading: Kitchin, Chapters 5 & 8
    b. Reading: Rieder & Simon
    c. Lecture: data as surveillance and control; classification as control

Student Presentation Questions -- Relate to the readings and lecture material, but focus on a particular dataset you are familiar with or using in your work; student answers only questions 1-3; questions 4-6 will be presented during class in small group activity

1. Evaluate your data according to the main dimensions of data quality
2. What features of the world is the data meant to represent?
3. Is your data representative? Of what population?
4. Is this the population you're interested in?
5. If the data are not representative, how might this cause problems for the analysis that might be done with it? [maybe this is for analysis/interpretation?]
6. Are there any legal/ethical issues that would affect how you are able to use the data? How you should use the data?

Agenda for Week 4 Lecture

1. How do we evaluate the "quality" of different kinds of data? [Probably are general papers on this -- **CB to look for**]
    a. Maybe different for different types of data
    b. Representativeness: to what population? Are population statistics known?
    c. Coverage (representative of a sub-population?)
    d. Measurement error: absolute; relative to construct
    e. Consistency (across time and location)
2. What are the legal/ethical implications of data production and use?
    a. Data integration
    b. Data surveillance
    c. Data ownership

Agenda for Week 4 Class Meeting

1. Check-in/community-builder - 10 minutes
2. Chris' reaction to Nicole's lecture -- 5 minutes
3. Any responses from students -- 10 minutes
4. Introduce spreadsheet for small group activity; walk through how to fill it out using Ridge's presentation topic as an example -- 20 minutes
5. Small group activity - 30 minutes
    a. One person talks, another writes, the rest listen and help with notetaking
6. Report out and discussion - 30 minutes
7. Open discussion: Reflection on legal and ethical issues; some examples -- 20 minutes
    a. Tax data is required as the source for property assessments -- nothing you can on that
    b. COVID data -- you can't compel people to take tests, at least not in the U.S.; you can in other countries, and there are benefits and costs to that
    c. Chicago gang database? -- Who knows why you're in it, and you can't get out of it, but it's being used for policing/legal actions
8. Remarks on next week's readings and our pivot -- Chris
9. Closing: put in the chat: what's a question you still have?

Everyone think about the topic you've chosen for your paper. Go around the group and:

1. Think of a dataset you know is used, or you could imagine being used, to study your topic

2. Name one issue of representativeness in this dataset that would affect an analysis using these data **OR** Name one issue of measurement validity in this dataset that would affect an analysis using these data
3. Brainstorm with the group: is there any practical way to reduce/eliminate the source of bias you identified?
4. Given whatever realistic improvements you can make, is this data still useful/informative for studying the topic? Compared to what?

More on measurement validity -- what we've done is very high level
More on representativeness -- clear from statistical perspective; is there a different perspective we want to consider?
Is this the population you care about? How would you know if the data are representative of the population you care about?

1. What is the population for which you want to make inferences?
2. How does the sample of data you have relate to that population?
    a. This can be a hard question to answer! It might even be impossible to answer.
    b. Identify which features relate to this question, and on which you have information in your dataset
    c. Can you make the sample you have look like the population? At least on the observables, you might be able to. But it's the unobservables that are the problem.

Weeks 5-6: Epistemologies

The section on epistemologies of quantification focuses on the forms of knowledge generated within the two analytical paradigms that dominate contemporary policy analysis: causal inference/program evaluation, which is focused on estimates of causal effects, and machine learning/data science, which is focused on prediction and classification. We ask what can be learned with each approach, what sorts of policy questions or problems can be addressed within each, and also emphasize the limitations of the sort of knowledge that can be attained.

Readings:

- Breiman, Leo. 2001. "Statistical Modeling: The Two Cultures (with Comments and a Rejoinder by the Author)." Statistical Science [explicating the algorithmic/prediction view and its desirability]
- Hofman, Jake M., Amit Sharma, and Duncan J. Watts. 2017. "Prediction and Explanation in Social Systems." Science [comparing predictive and explanatory models and arguing that they should work together. Short article]
- Holland, Paul W. 1986. "Statistics and Causal Inference." Journal of the American Statistical Association [classic statement of potential outcomes model, links to philosophy and statistics, different notions in medicine and social science]

- Meyer, Bruce D. "Natural and Quasi Experiments in Economics." 1995. Journal of Business Economics and Statistics [discussing major research designs and what they estimate]
- Deaton, Angus. 2020. "Randomization in the Tropics Revisited: A Theme and Eleven Variations." National Bureau of Economic Research. [non technical critique of RCTs in terms of knowledge production and policy implications]
- Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan, and Ziad Obermeyer. 2015. "Prediction Policy Problems." American Economic Review [definition of prediction policy problrms, distinguished from causal, argument for importance. Short article]
- Titiunik, Rocío. 2015. "Can Big Data Solve the Fundamental Problem of Causal Inference?" PS: Political Science & Politics 48 (01): 75-79. [critique of large n and large p data from causal inference perspective]
- Shmueli, Galit. 2010. "To Explain or to Predict?" Statistical Science 25 (3): 289-310. [comparison of explanatory and predictive modeling goals, procedures, and argument for clarity on what approach is appropriate for a given problem]

<u>Week 5 Learning Goals</u>

1. Define epistemology and understand its use in data analysis and policy making.
2. Understand the program evaluation/causal inference paradigm: what is its epistemology, what is the relevant family of statistical methods, what are its advantages and disadvantages
3. Understand the predictive/data science paradigm: what is its epistemology, what is the relevant family of statistical methods, what are its advantages and disadvantages

<u>How We Cover Each Week 5 Learning Goal</u>

1. Define epistemology and understand its use in data analysis and policy making.
    a. Lecture
    b. Readings:
2. Understand the program evaluation/causal inference paradigm: what is its epistemology, what is the relevant family of statistical methods, what are its advantages and disadvantages
    a. Lecture
    b. Readings: Holland, Meyer,
3. Understand the predictive/data science paradigm: what is its epistemology, what is the relevant family of statistical methods, what are its advantages and disadvantages
    a. Lecture
    b. Readings: Breiman, MAYBE NEED ANOTHER READING ON BASIC METHODS--CB TO FIND
    c.

Agenda for Week 5 Lecture

1. Chris to give a version of existing lecture of causal vs. predictive

Week 5 Student Presentation Questions -- Relate to the readings and lecture material. Focus on an existing analysis (by someone else or yourself).

1. What statistical approach/technique was used to analyze the data?
2. What was the analytical goal: prediction, causal inference, something else?
    a. Are the authors clear about which one of these they are doing?
3. What are the findings?
4. What is the policy question being answered?
    a. Is there a theory behind the way the question is posed or a hypothesis behind the answer?
    b. What policy decision hinges on the answer?
    c. Is there a way that combining causal and predictive approaches to the question offers a better policy decision?
5. Are the data and techniques appropriate for answering the question? Explain
    a. Is there another approach that could have been used?

Week 5 Class Agenda

1. Check-in/community builder -- 10 minutes
2. Nicole reaction to lecture -- 5 minutes
3. Time for questions about readings and lecture -- 20 minutes
4. Paired discussion exercise -- 20-30 minutes
    a. Start by articulating a policy question related to your paper topic.
    b. Reason through whether this is a causal or predictive question.
    c. As you reason through this, you may find that you need to change the way you stated the question: that's ok!
    d. By the end of this exercise, you should have one well-stated policy question each, and be able to say if it's causal or predictive.
    e. If you have extra time, start this cycle again with a new question on your same topic.
5. Report out: everyone states (d) -- 30 minutes
    a. If we don't get to you, feel free to come to office hours
6. Closing chat statement

<u>Week 6 Learning Goals</u>

1. Understand what kind of policy problems can be productively analyzed via each approach.
2. Understand how to frame a question to make it more suitable for either approach.
3. Knowing when a problem is not suitable for either.

<u>How We Cover Each Week 6 Learning Goal</u>

1. Understand what kind of policy problems can be productively analyzed via each approach.
    a. Lecture

    b. Readings: Hofman et al., Deaton, Kleinberg et al., Titiunik, Shmueli
2. Understand how to frame a question to make it more suitable for either approach.
    a. Lecture
3. Knowing when a problem is not suitable for either
    a. Lecture
    b. Readings: Hofman et al., Deaton, Kleinberg et al., Titiunik, Shmueli

Week 6 Student Presentation Questions -- Relate to the readings and lecture material. Start with a policy problem, e.g., "I care about education policy, and I want to know how to reduce high school dropout."

1. State your policy problem as a researchable question.
2. Is it a prediction policy problem or a causal policy problem, or neither? (If neither, go back and try again!)
3. Is there a way that the other approach can add value in answering your question? Explain why or why not.

Agenda for Week 6 Lecture

1. Chris' distillation of key points for Learning Goal 1 from Week 6 readings
2. Nicole's points on how to frame questions about policy problems for either causal or predictive approaches (Learning Goal 2)

Agenda for Week 6 Class

1. Check-in/community builder -- 10 minutes
2. Questions from readings or lecture?
3. Quick review of 9th grade on track thing
4. Whatever your question was from last week (either causal or predictive), find a related question from the other perspective
    a. Look at Figure 2 from Shmueli and identify the places where these two versions of your question require the most different approach
    b. You might want to make a table: rows are Shmueli elements, columns are causal & predictive
5. Large group discussion: go through as many of the student examples as we can
6. Closing chat statement -- Thinking about the readings for today, what is an unanswered question you have?

<u>Week 7: External Validity: Understanding Implications of Evidence for Policy</u>

The final section of the course focuses on interpretation. Even when the internal validity of a study is solid, policy implications do not always follow straightforwardly. This week we focus on issues of extrapolating from a study to a set of policy implications, focusing on issues of external validity. External validity has received much attention in the literature on experiments, but we also consider analogous issues in predictive analysis.

Readings:

- Czibor et al, The Dozen Things Experimental Economists Should Do (More of)
- Dehija, Experimental and Non-Experimental Methods in Development Economics: A Porous Dialectic
- Lucas, Theory-Testing, Generalization, and the Problem of External Validity
- Muller, Causal Interaction and External Validity: Obstacles to the Policy Relevance of Randomized Evaluations
- Pritchett, Context Matters for Size: Why External Validity Claims and Development Practice do not Mix

Questions:

1. Can you give an account of the substantive mechanisms that generate the findings? Do you need to in order to address the policy question?
2. To whom do the findings apply? To whom do they not apply?
3. What are the policy implications of the findings? What policy action is supported by them, if any?
    a. If predictive analysis, what is the scope for extrapolation?
    b. If causal, what is the scope of external validity?
4. What is the relevant notion of external validity in this case and what are the major threats to external validity?

<u>Week 7 Learning Goals</u>

1. Understand the notion of external validity as applied to causal analyses (RCTs and observational methods)
2. Understand the notion of external validity as applied to predictive analyses
3. Understand different threats to external validity in causal and predictive analysis: data production, measurement validity, the average effect, mechanisms
4. Understand how external validity is related to heterogeneous treatment effects
5. Understand how external validity is related to theory
6. Understand how to evaluate external validity of a study (causal or predictive)
7. Connect this unit's ideas to the broader course question of how analysis is used in the policymaking process.

<u>How We Cover Each Week 7 Learning Goal</u>

1. Understand the notion of external validity as applied to causal analyses (RCTs and observational methods)
    a. Readings
2. Understand the notion of external validity as applied to predictive analyses
    a. Lecture
3. Understand different threats to external validity in causal and predictive analysis: data production, measurement validity, the average effect, mechanisms
    a. Readings, exp Czibor, Muller, Pritchett
4. Understand how external validity is related to heterogeneous treatment effects

    a. Readings, esp Muller and Pritchett
    b. Lecture
5. Understand how external validity is related to theory
    a. Readings, esp. Lucas, Dehija
6. Understand how to evaluate external validity of a study (causal or predictive)
    a. Lecture
7. Connect this unit's ideas to the broader course question of how analysis is used in the policymaking process.
    a. Muller reading
    b. Lecture

Agenda for Week 7 Lecture

1. Granting that you have internal validity, what do you need to consider when planning policy action?
2. Discuss each of the threats to external validity in causal and predictive analysis
    a. Take lessons on these issues from causal literature and see how they apply to predictive analysis
3. Tell students to look for examples of these issues in the papers assigned for the week

Agenda for Week 7 Class Meeting

1. Check in/community building
2. Example of good external validity: vaccines
3. Introduce a new causal example and a new predictive example where external validity threats can be seen
    a. Causal: Moving to Opportunity -- Nicole
        i. Describe what was done and what was found
        ii. Questions of external validity to discuss
            1. What do those findings generalize to? To all low-income people? To ?? Why or why not?
            2. What policy decision would be supported by this evidence?
    b. Predictive: Gentrification -- Chris
        i. Describe what was done and what was found
        ii. Questions of external validity to discuss
            1. What do those findings generalize to? Chicago today? Another city today? Why or why not?
            2. What policy decision would be supported by this evidence?
4. Discuss in small groups: what are the threats to external validity?
    a. 2 groups do causal example
    b. 2 groups do predictive example
5. Report out
    a. Causal example
    b. Predictive example

6. Large group discussion: What did you learn from this discussion that applies to your own paper?
    a. If you don't see how this applies to your own paper topic, speak now! Or come see us in office hours.
7. Closing chat statement

<u>Week 8: Fairness</u>
Most discussion of internal and external validity focus on average

<u>Week 8 Readings</u>

- Kleinberg, Mullainathan, Raghavan. "Inherent Trade-Offs in the Fair Determination of Risk Scores"

- Mitchell, Porash, Barocas. "Prediction-Based Decisions and Fairness: A Catalogue of Choices, Assumptions, and Definitions"

- Paulus and Kent. "Predictably unequal: understanding and addressing concerns that algorithmic clinical prediction may increase health disparities"

- Kravitz, Duan, and Braslow. "Evidence-based medicine, heterogeneity of treatment effects, and the trouble with averages"

- Optional for students who find the more technical readings challenging: Barocas, Solon, and Andrew Selbst. 2016. "Big Data's Disparate Impact." California Law Review. Read part I and skim the rest. (One of the few non-technical summaries of algorithmic fairness and discrimination)

[Still needed: keep looking at articles on how heterogeneous treatment effects relate to bias and discrimination. Kravitz et al is a placeholder. Something on NCLB could still be good/more interesting.]

[<u>Reading Notes:</u> Aziz's Duke Law Journal article is very useful but maybe too law-y? Good stuff for lecture, if nothing else.]

Questions for student presentations

1. What are the substantive concerns about bias or discrimination in your application?
2. How do these concerns map onto the technical definitions given in the readings?
3. What is the right definition/standard? Is there more than one?
4. Is the standard satisfied in this case; i.e., is there bias or discrimination in the analysis? How do you know?
5. What are the tradeoffs involved in satisfying this standard?

<u>Week 8 Learning Goals: Bias, Fairness, and Legality</u>

1. Understand various technical concepts of fairness and discrmination in predictive analysis and how they relate to ethical and legal concepts
2. Understand the tensions and tradeoffs between these various concepts
3. Understand how concerns about fairness and discrimination map onto causal analysis
4. Connect this unit's ideas to the broader course question of how analysis is used in the policymaking process.

<u>How We Cover Each Week 8 Learning Goal</u>

1. Understand various technical concepts of fairness and discrmination in predictive analysis and how they relate to ethical and legal concepts
   Lecture
   Readings: Kleinberg, Mitchell, Paulus, Barocas

2. Understand the tensions and tradeoffs between these various concepts
   Lecture
   Readings: Kleinberg, Mitchell, Paulus, Barocas

3. Understand how concerns about fairness and discrimination map onto causal analysis
   Lecture
   Readings: covered briefly in Mitchell (section 5), Kravitz; maybe need one more reading

4. Connect this unit's ideas to the broader course question of how analysis is used in the policymaking process.
   Lecture

Agenda for Week 8 Lecture

This week's readings are very technical. The key for the lecture is to (a) give a non-technical summary/substantive interpretation of the key concepts; and (b) give a synthesis and framework for understanding how it all relates. May be a longer than usual lecture.

Agenda for Week 8 Class

1. Nicole: reaction to lecture
2. Chris: Summarize three buckets of concepts and allow students to ask questions
   a. Confusion matrix
   b. Non-algorithmic conceptions of fairness
   c. "Causal pathways" -- these are conceptual, not actual estimation of causal effects
   d. Heterogeneous treatment effects -- as related specifically to estimation of causal effects
      i. Be sure to specify that these are not usually discussed in parallel with algorithmic fairness/bias, but it's important to realize that RCT/causal evidence has this type of problem as well

3. Walk through one application -- maybe from healthcare, say controlling for race in prediction (use controversial among us health paper on race-correction in clinical treatment)
    a. Look at student presentations: walk through the four buckets?
4. Paired exercise on students' own topics: consider the three/four buckets of fairness concepts; how do any or all of them relate to your topic?
    a. Everyone should end with one insight or question to discuss with the larger group
5. Large group discussion regarding insights and/or questions from paired discussions
6. Foreshadowing next week
    a. No pre-recorded lecture -- expect to stay the whole 3 hours of class
    b. ProPublica article generated a lot of inquiry on the use of ML for policy applications
    c. This example draws together the whole course: data production, ethics, epistemology, fairness/bias -- think about it all as you read
    d. Order:
        i. Angwin et al.
        ii. Feller et al.
        iii. Then any order for remaining

<u>Week 9: Bail Decisions and Recidivism Prediction</u>

<u>Readings</u>
- Julia Angwin, Jeff Larson, and Surya Mattu and Lauren Kirchner. 2016. "Machine Bias." Text/html. ProPublica. May 23, 2016
- Rudin, Cynthia, Caroline Wang, and Beau Coker. 2020. "The Age of Secrecy and Unfairness in Recidivism Prediction."
- Kleinberg et al, "Human Decisions and Machine Predictions." 2018. QJE.
- Alexandra Chouldechova, The accuracy, fairness, and limits of predicting recidivism, 2018. Science Advances
- Feller at al, A computer program used for bail and sentencing decisions was labeled biased against blacks. It's actually not that clear. Washington Post. 10/17/2016. (non-technical summary of some of these issues)
- Jackson et al, "Bail reform analysis by Cook County chief judge based on flawed data, undercounts new murder charges." 2/13/2020. Chicago Tribune.

Questions for Student Presentations: NA

Learning Goals
1. Understand basics of COMPAS model and critiques of it
2. Understand procedures used in absence of machine learning, and critiques of them
3. Make connections between main topics of the course and this example: ethics, datafication, epistemology, algorithmic fairness

How we cover each learning goal

1. Understand basics of COMPAS model and critiques of it
    a. Readings, all except Jackson
2. Understand procedures used in absence of machine learning, and critiques of them
    a. Readings, incl. Jackson
3. Make connections between main topics of the course and this example: ethics, datafication, epistemology, algorithmic fairness, role of evidence in policymaking
    a. Readings and lecture

Agenda for lecture
Provide context of ProPublica article and the various responses it provoked; explain some of the more technical issues in the Chouldechova and Kleinberg articles; link to Cook County and alternatives to ML via Jackson article; give high level connections with big course topics per learning goal #3

Agenda for class
Allow students to ask any clarifying/technical questions
Discussion question: Should jurisdictions continue to use (or start using) COMPAS? If not, what should they do instead? (Should this be done in small group break out, or whole class? I think whole class)
Open discussion for end-of-quarter reflections

END Winter 2021 Quarter

Week X (for a 10-week quarter): Communicating Research Findings
A week on communicating results to policymakers/administrators/the public in a responsible way

Learning Goals
1. Understand the work that is done to stabilize scientific findings.
2. Understand how to responsibly communicate findings and policy implications from causal and predictive analysis.
3. Understand the role that uncertainty plays in policy recommendations and implementation.
4. Understand the importance of communicating the analytic tradeoffs you made in your analysis.
5. Identify how uncertainty affects the analyst’s role in policymaking: does communicating uncertainty and tradeoffs limit the effect you might have on policymaking? Is that something you should care about?

Readings A: on Communicating Uncertainty

- Manski, Charles F. 2019. "Communicating Uncertainty in Policy Analysis." Proceedings of the National Academy of Sciences
- Van der Bles, Anne Marthe et. al. "Communicating Uncertainty about Facts, Numbers and Science." Royal Society Open Science
- Thompson, Erica L., and Leonard A. Smith. 2019. "Escape from Model-Land." Economics.
- Rudin, Cynthia, and Joanna Radin. 2019. "Why Are We Using Black Box Models in AI When We Don't Need To? A Lesson From An Explainable AI Competition." Harvard Data Science Review

Leftover reading ideas not used

• Barocas, Solon, and Andrew Selbst. 2016. "Big Data's Disparate Impact." California Law Review (Argues that algorithmic bias is difficult to address not only via technical means, but also, in the area of employment discrimination, via legal or political means. Not sure the paper is appropriate for our purposes given how much of it focuses on the legal and political issues. But maybe the technical discussion (Part I) by itself is useful?) READ AT LEAST PART I

• Obermeyer, Ziad, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. "Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations." Science (Shows that health care costs are not an effective proxy for health across racial groups, because some racial groups (here, blacks) have less access to health care and thus lower health care spending; thus, lower health care spending does not equate to better health for blacks. Example of measurement validity, leading to racial bias)

• Vyas, Darshali A., Leo G. Eisenstein, and David S. Jones. 2020. "Hidden in Plain Sight - Reconsidering the Use of Race Correction in Clinical Algorithms." New England Journal of Medicine (Argues that race is not a reliable proxy for genetic difference, and thus algorithms that use race as a predictor for individual clinical decisions perpetuate the notion that race does reliably proxy genetic difference. Question: Does this argument fail to appropriately distinguish prediction and causality?) [Consider moving to prediction vs causation week]

•