

Using Big Data To Solve Economic and Social Problems

Professor Raj Chetty
Head Section Leader Rebecca Toseland


Photo Credit: Florida Atlantic University



Using Social Norms to Reduce Electricity Consumption: OPower Scale-Up

- OPower has mailed Home Energy Reports to millions of electricity customers in markets across the U.S.
- Mailings were randomly assigned, permitting experimental estimates of causal effects
- Allcott and Rogers (2014) evaluate impacts of these mailings by comparing electricity usage of treated and control households

Treatment Effects of Monthly and Quarterly OPower Home Energy Reports on Electricity Consumption




First home energy report received

Differential Effects of Social Norms Across Markets

- OPower gradually expanded its presence to additional markets over time
- Were impacts similar to those achieved in initial markets?
- Allcott (2015) estimates impact of home energy reports for 111 markets separately, using data on nearly 9 million customers
- In each market, construct an estimate of causal effect of home energy report by comparing treated and control households

Effects of OPower Home Energy Reports Across Utilities, by Program Start Date



Site Selection Bias in Experimental Estimates

- Allcott (2015) finds that effects of OPower treatment are larger for “early adopter” utilities
 - Likely because utilities with more environmentally-oriented customers signed up for the program in its early stages
- Illustrates a broader methodological lesson: experimental/quasi-experimental estimates in social science are not necessarily stable across settings
- Important to continue to conduct studies across areas and time periods to understand policy impacts, rather than assuming that effects will generalize

Summary: Policies to Mitigate Climate Change

- Climate change and pollution have large social costs
 - But costs are not infinite, partly because human beings can adapt
 - Economic approach of quantifying environmental costs of policies and comparing them to their economic benefits is therefore very useful
- Environmental damage can be mitigated through policies such as corrective taxes and changing social norms
 - New data are helping us to measure what policies are necessary to achieve optimal balance of social costs and economic benefits

Summary: Policies to Mitigate Climate Change

- Key problem going forward: addressing externalities in implementing the necessary policies at a national and global scale
 - Unlike other problems we have discussed in this course, cannot be tackled at individual or local level
 - Pollution emitted in California affects atmosphere in the rest of the U.S. and the world
 - Policy challenge: need to develop global coalitions to tackle environmental problems

Outline of Remaining Lectures

- Primary methodological theme of remaining lectures: applications of machine learning to prediction
- Today: discrimination and criminal justice

Next Tuesday: Guest Lecture by Sendhil Mullainathan



Professor of Economics, Harvard University
Co-Author of *Scarcity* and MacArthur “Genius”

Next Thursday: Guest Lecture by Matthew Gentzkow



Professor of Economics, Stanford University


Recipient of the John Bates Clark Medal

Outline of Remaining Lectures

- Primary methodological theme of remaining lectures: applications of machine learning to prediction
- Today: discrimination and criminal justice
- Next Tuesday: guest lecture by Sendhil Mullainathan on using big data for prediction
- Next Thursday: lecture by Matt Gentzkow on political polarization
- Tuesday June 6: Concluding lecture and review of methods by Rebecca

Discrimination and Criminal Justice

Black Men's Median Earnings vs. White Men's Median Earnings Among Full-Time Workers



Source: Current Population Survey

White Women's Median Earnings vs. White Men's Median Earnings Among Full-Time Workers



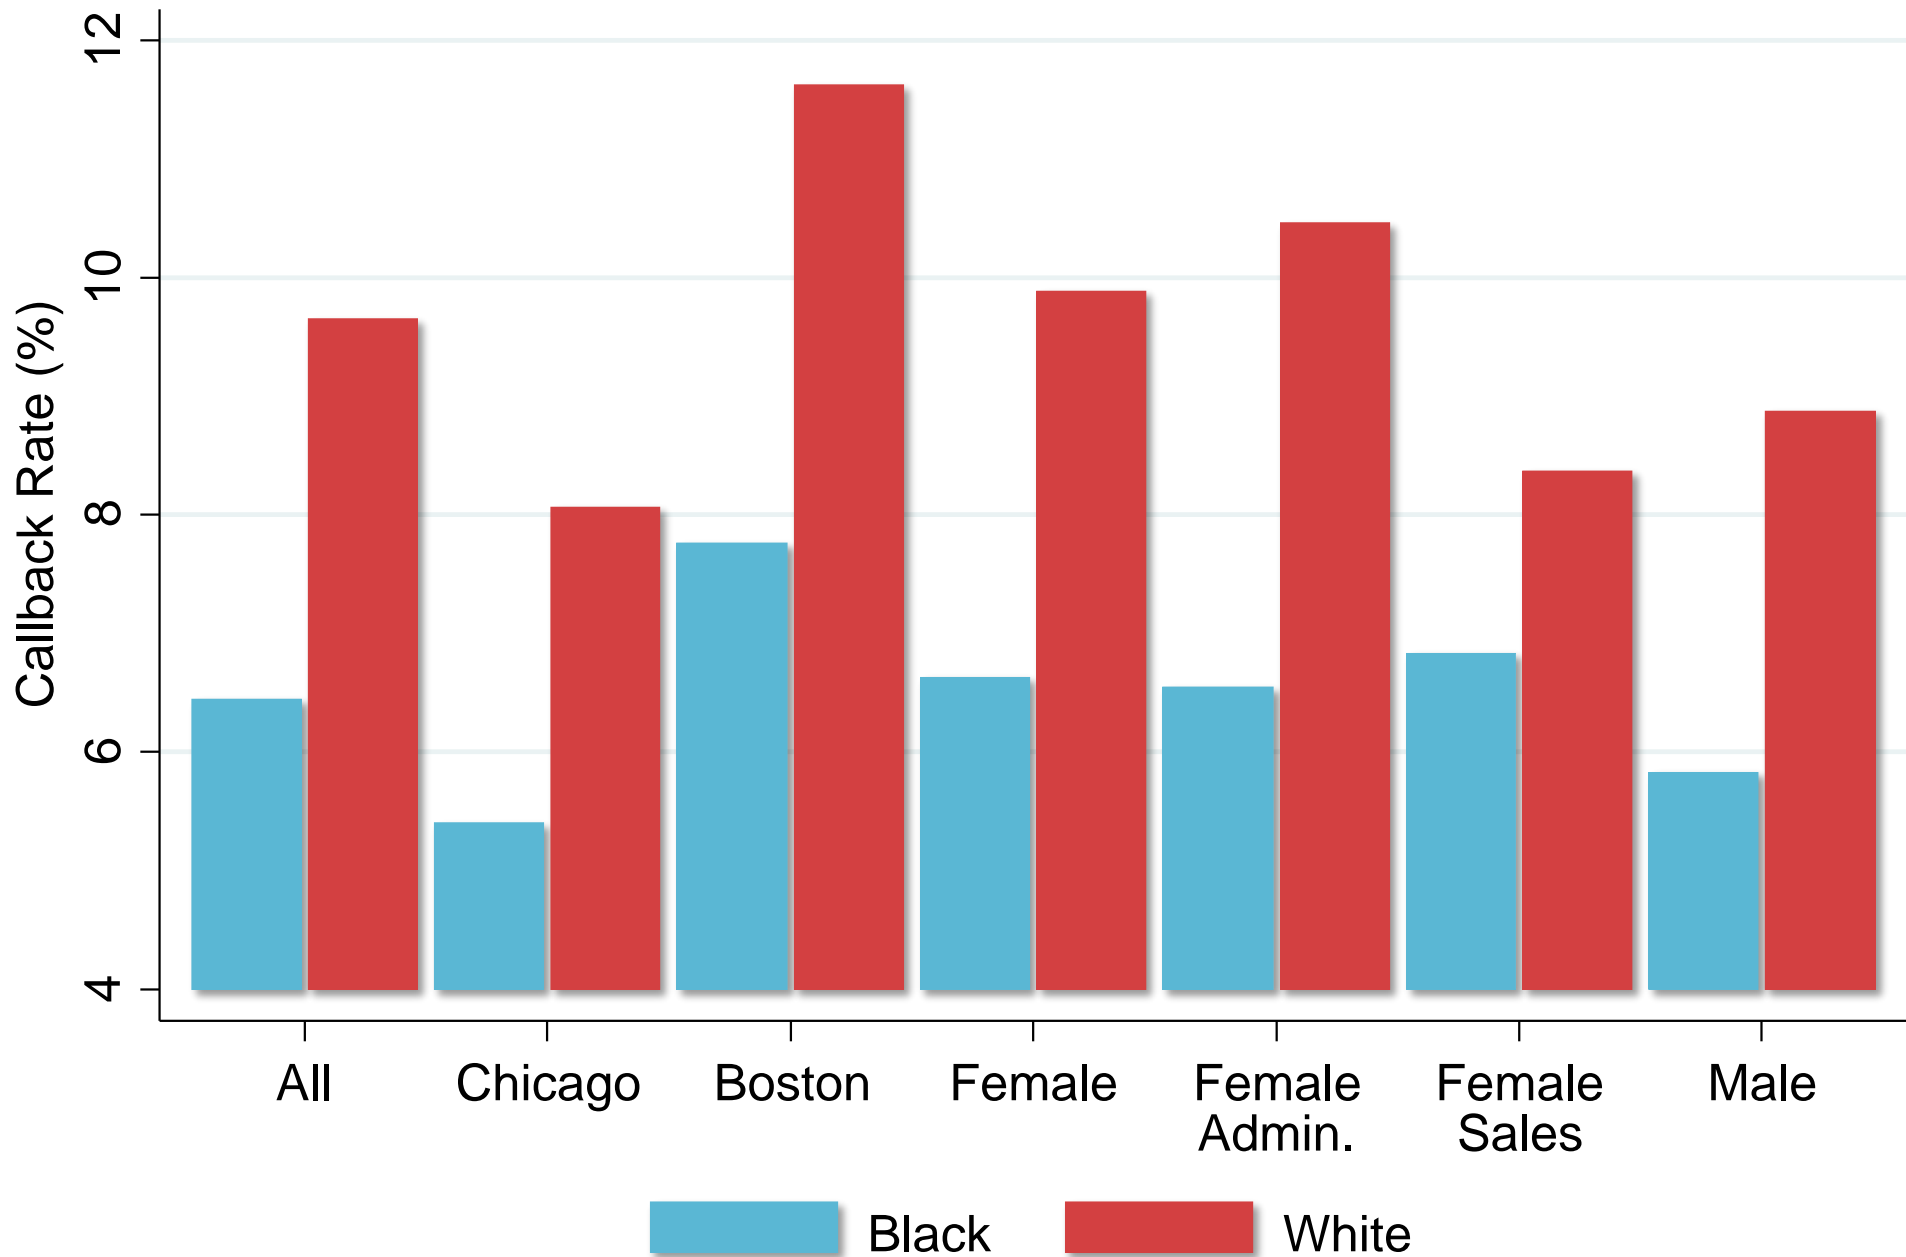
Discrimination: Motivation

- Large differences in earnings and other outcomes by race and gender in America
- Partly driven by differences in economic opportunities, schools, etc. as we have discussed previously
- But could also be partly driven by discrimination: differential treatment by employers and other actors based on race or gender
- A series of recent studies have used experimental and quasi-experimental methods to test for discrimination

Racial Discrimination in Hiring

- Bertrand and Mullainathan (2004) send out fictitious resumes with identical credentials in response to real job ads
- Vary name of applicant to be “white-sounding” (e.g., Emily Walsh or Greg Backer) vs. “black-sounding” (e.g., Lakisha Washington or Jamal Jones)
- Send 5,000 resumes in response to 1,300 help-wanted ads in newspapers in Boston and Chicago
- Analyze call-back rates in response to application

Job Callback Rates by Race for Resumes with Otherwise Identical Credentials




Source: Bertrand and Mullainathan (AER 2004)

Racial Discrimination Among Airbnb Hosts

- Edelman, Luca, and Sversky (2017) use an analogous approach to study discrimination among Airbnb hosts
- Send out fictitious requests to book listings in response to real postings
 - Set up 20 independent accounts and sent 6400 messages to hosts from these accounts in 2015
 - Randomly vary name of applicant across accounts to be white vs. black sounding (and do not include profile pictures)
- Note: study was done independently using experimental accounts, *not* in collaboration with Airbnb

AirBnB Host Response Rates by Race for Individuals with Otherwise Identical Profiles



Source: Edelman, Luca, and Svirsky (AEJ 2017)

Racial Discrimination Among Airbnb Hosts

- What happens to listings where offer was not accepted?
 - Roughly 40% of these listings remained vacant on proposed date
 - Rejecting black guests → loss of \$65-\$100 of revenue on average
- Publication of Edelman, Luca, and Sversky (2017) paper pressured Airbnb to acknowledge that discrimination was prevalent on the site
- Airbnb took steps to address the problem: more “Instant” bookings, reducing prominence of guest photos, increased support for discrimination complaints
- But stopped short of making names and photos invisible

Statistical vs. Taste-Based Discrimination

- Two different theories of discrimination that could in principle be consistent with these patterns:
 1. Taste-based discrimination: preference for employees/clients of a given race for reasons unrelated to their performance
 2. Statistical discrimination: use of race as a signal of other aspects of background that may matter for performance (e.g., language, inferences about effort, or skill)

- Both forces harmful for qualified minorities, but distinguishing between them is important to understand the solution
 - If discrimination is statistical, then providing more accurate information would help; if it is taste-based, need to change people's preferences/attitudes

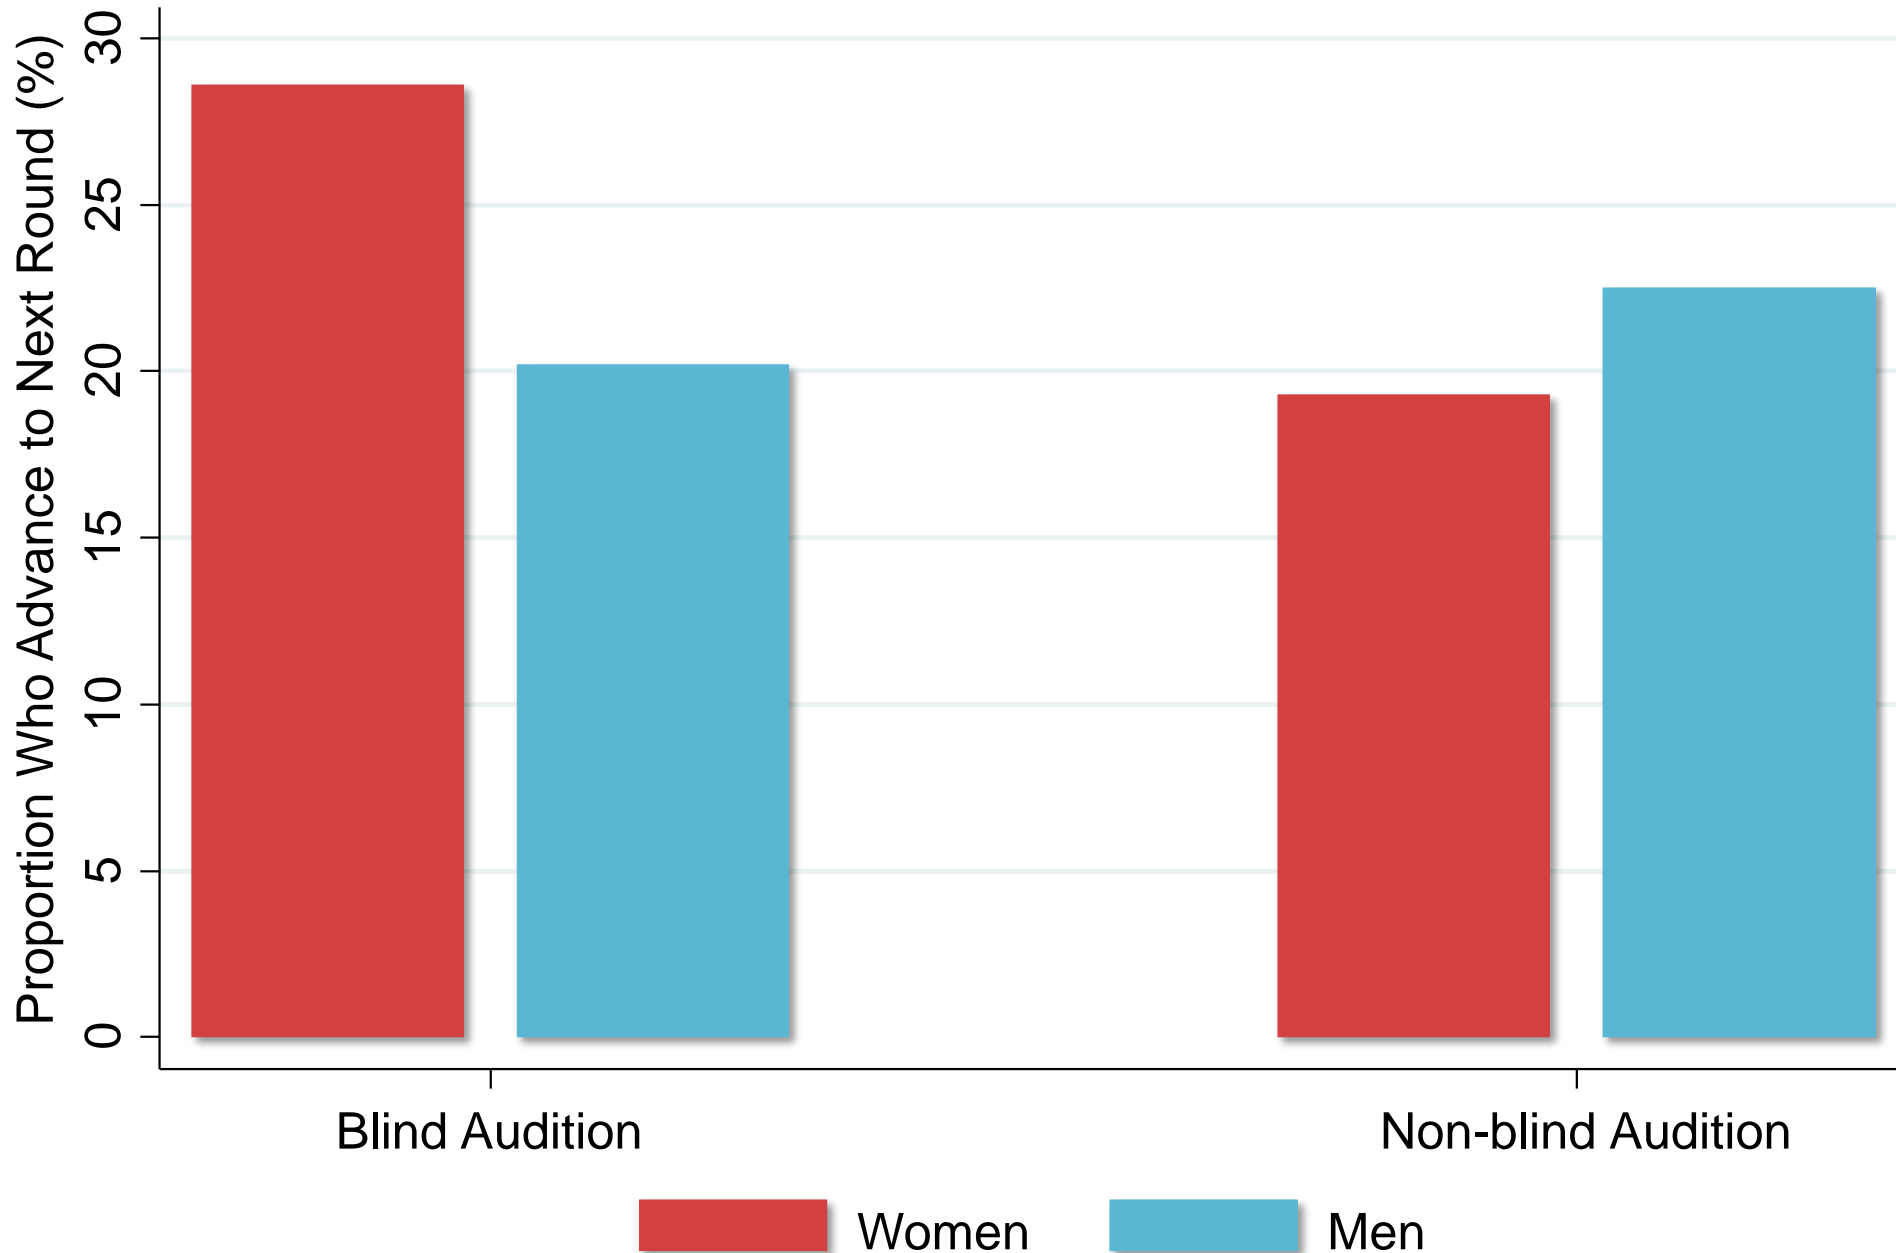
Statistical vs. Taste-Based Discrimination

- One way to test between these two explanations is to use measures of performance/output
- Do those who get screened out due to discrimination actually have lower performance?
- Difficult to study directly in most settings because we don't observe outcomes of those who don't get the job or the Airbnb booking

Sex-Based Discrimination in Orchestra Auditions

- One setting where people can see productivity: auditions for orchestras
- Goldin and Rouse (2000) compare outcomes of auditions that were “blind” (behind a curtain) vs. non-blind
 - Data from eight major symphony orchestras from 1950-1995
 - 14,000 auditions, some of which were blind and some of which were not, depending upon orchestra policy

Average Success Rates at Orchestra Auditions by Sex and Type of Audition



Sex-Based Discrimination in Orchestra Auditions

- Goldin and Rouse results suggest that discrimination is partly taste-based (or based on biased beliefs about performance)
- If all one cares about is quality of music performance, then orchestras are sacrificing performance quality by giving men a preference


Biases due to Decision Fatigue

- Discrimination/suboptimal choice is not driven entirely by deep-rooted beliefs
- Decisions also vary greatly based on factors unrelated to substantive features of the issue at hand
- Danziger et al. (2011) demonstrate this by analyzing data on judges' decisions to grant prisoners parole

Studying Decision Fatigue in Parole Decisions

- Data: 1,100 judicial rulings on parole for prisoners in Israel over 10 months
- Judges review about 20 cases on average each day in succession
 - Ordering of cases depends upon when attorney shows up and is essentially random
- Judges can decide to grant parole or reject (delay to a future hearing, maintaining status quo)
- Key institutional feature: two breaks during the day for meals
 - 10 am for late-morning snack (40 mins)
 - 1 pm for lunch (1 hour)

Proportion of Rulings in Favor of Prisoner by Time of Interview During the Day



Can Machines Help Humans Overcome Biases?

- Most recent literature asks whether big data and machine learning can help people make decisions that are less biased
- Idea: develop algorithms to predict outcomes of interest and use these to guide or replace human decisions
- Analogous to the way in which Netflix presents movie recommendations based on history of movies viewed, demographics, etc.
- These algorithms are not necessarily subject to human biases, but do they outperform humans' decisions overall or have other biases/shortcomings?

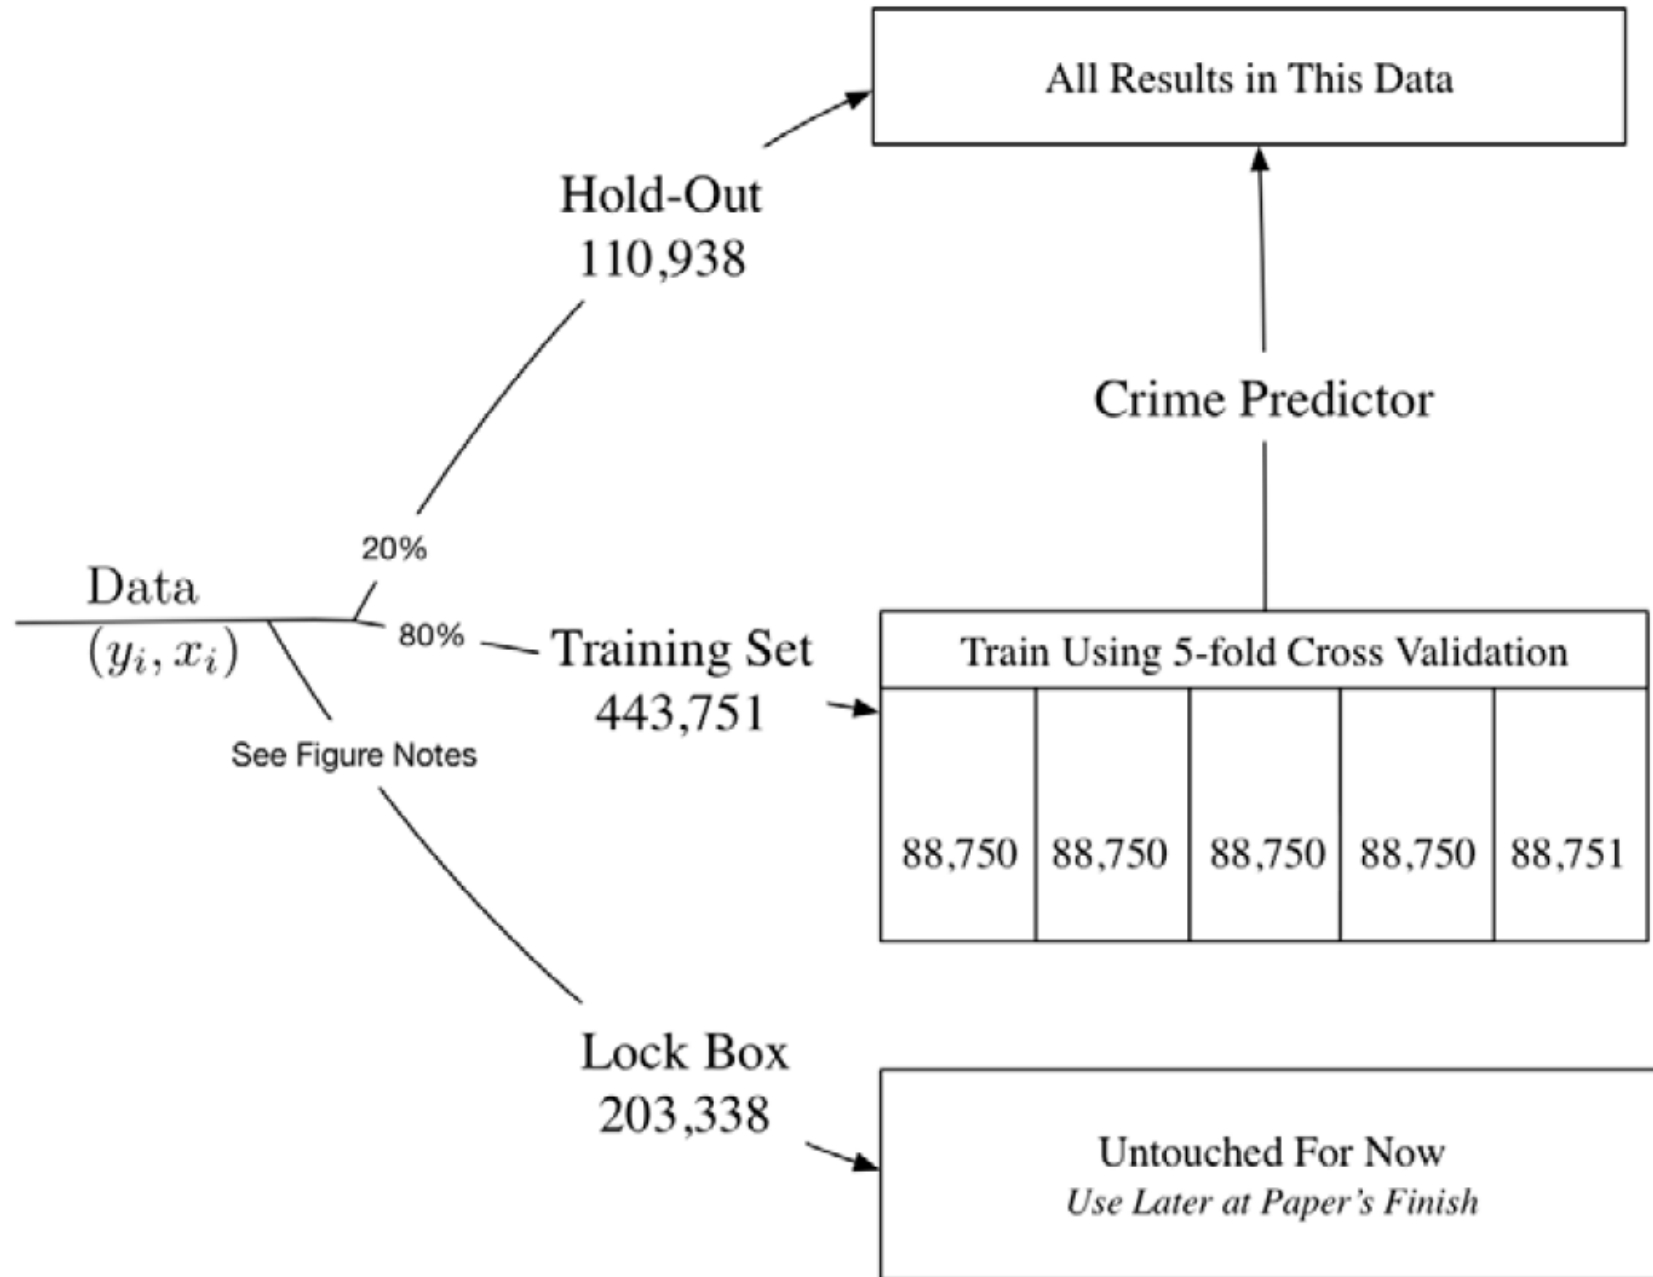
Decisions to Jail vs. Release Defendants

- Kleinberg et al. (2017) compare the accuracy of human decisions and machine predictions in the criminal justice system
- Every year, ~10 million people are arrested in the U.S.
- After arrest, judges decide whether to hold defendants in jail or let them go
- By law, decision is made with the objective of minimizing risk of flight (failure to appear at trial)
- Kleinberg et al. compare machine learning predictions and judges' actual decisions in terms of performance in achieving this objective

Decisions to Jail vs. Release Defendants

- Data: 750,000 individuals arrested in New York City between 2008-2013
 - Same data on prior history that is available to judge (rap sheet, current offense, etc.)
 - Data on subsequent crimes to develop and evaluate performance of algorithm
- Define “crime” as failing to show up at trial; objective is to jail those with highest risk of committing this crime
 - Other definitions of crime (e.g., repeat offenses) yield similar results
- First divide data into three separate samples

Data Used for Empirical Analysis



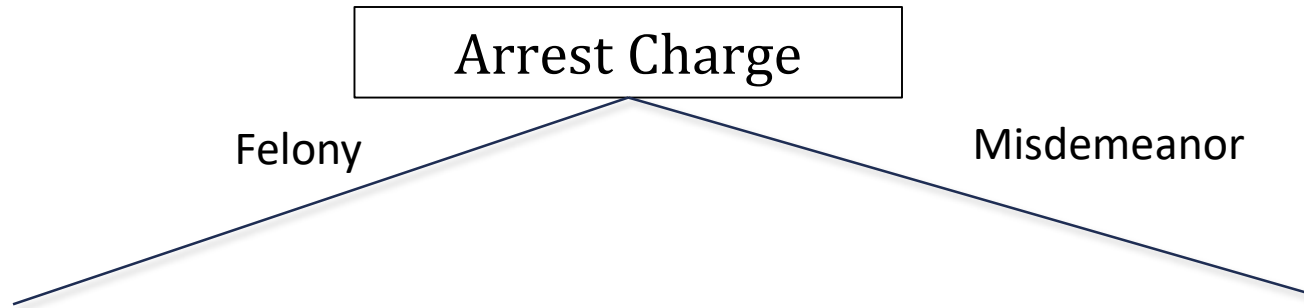
Methodology: Machine Learning Using Decision Trees

- Predict probability of committing a crime using a machine learning method called decision trees
- Main statistical challenge: need to avoid overfitting the data with large number of potential predictors
 - Can get very good in-sample fit but have poor performance out-of-sample (analogous to issues with Google flu trend)
 - Solve this problem using cross-validation, using separate samples for estimation and evaluation of predictions

Methodology: Machine Learning Using Decision Trees

- Three steps to develop predictions using decision trees
 1. Split the data based on the variable that is most predictive of differences in crime rates

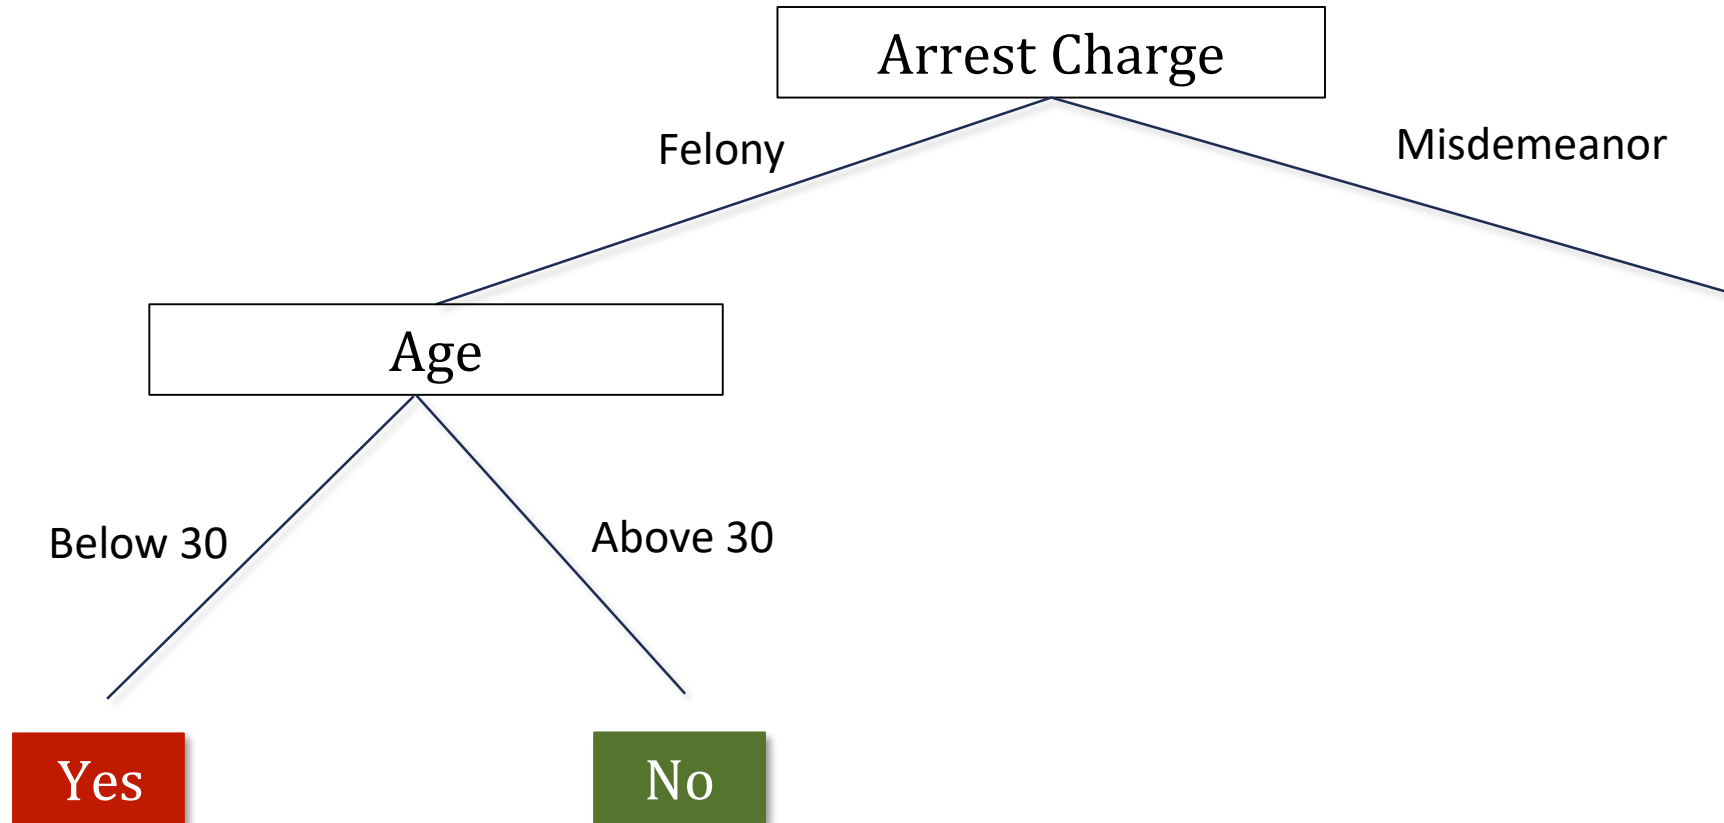
Hypothetical Decision Tree for Decision to Jail Defendant



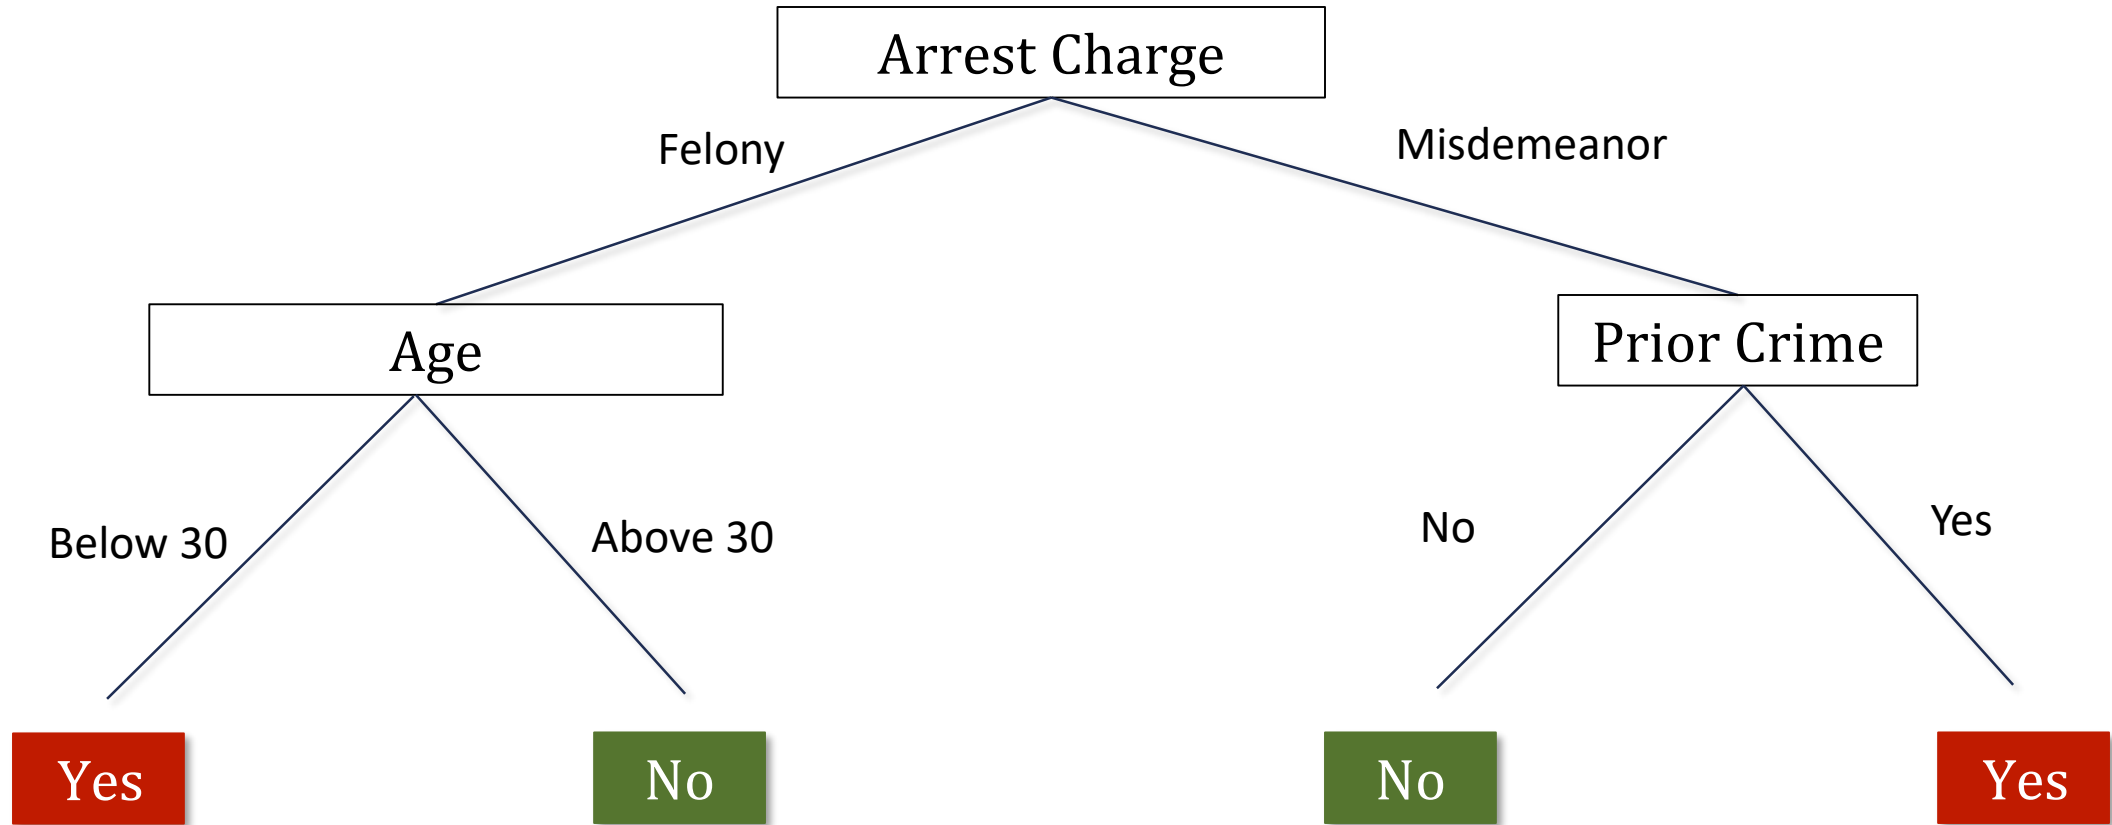
Methodology: Machine Learning Using Decision Trees

- Three steps to develop predictions using decision trees
 1. Split the data based on the variable that is most predictive of differences in crime rates
 2. Grow the tree up to a given number of nodes N

Hypothetical Decision Tree for Decision to Jail Defendant




Hypothetical Decision Tree for Decision to Jail Defendant



Methodology: Machine Learning Using Decision Trees

- Three steps to develop predictions using decision trees
 1. Split the data based on the variable that is most predictive of differences in crime rates
 2. Grow the tree up to a given number of nodes N
 3. Use separate validation sample to evaluate accuracy of predictions based on a tree of size N
- Repeat steps 1-3 varying N and choose tree-size N that minimizes average prediction errors


Prediction Errors vs. Size of Decision Tree




Comparing Machine Predictions to Human Predictions

- Applying this method yields predictions of crime rates for each defendant
- Machine-based decision rule: jail the defendants who have the highest predicted risk
- How does this machine-based rule compare to what judges actually do in terms of crime rates it produces?
- Answer is not obvious: judges can see things that are not in the case file, such as defendant's demeanor in courtroom

Judges' Release Decisions vs. Machine Predictions and Crime Risk

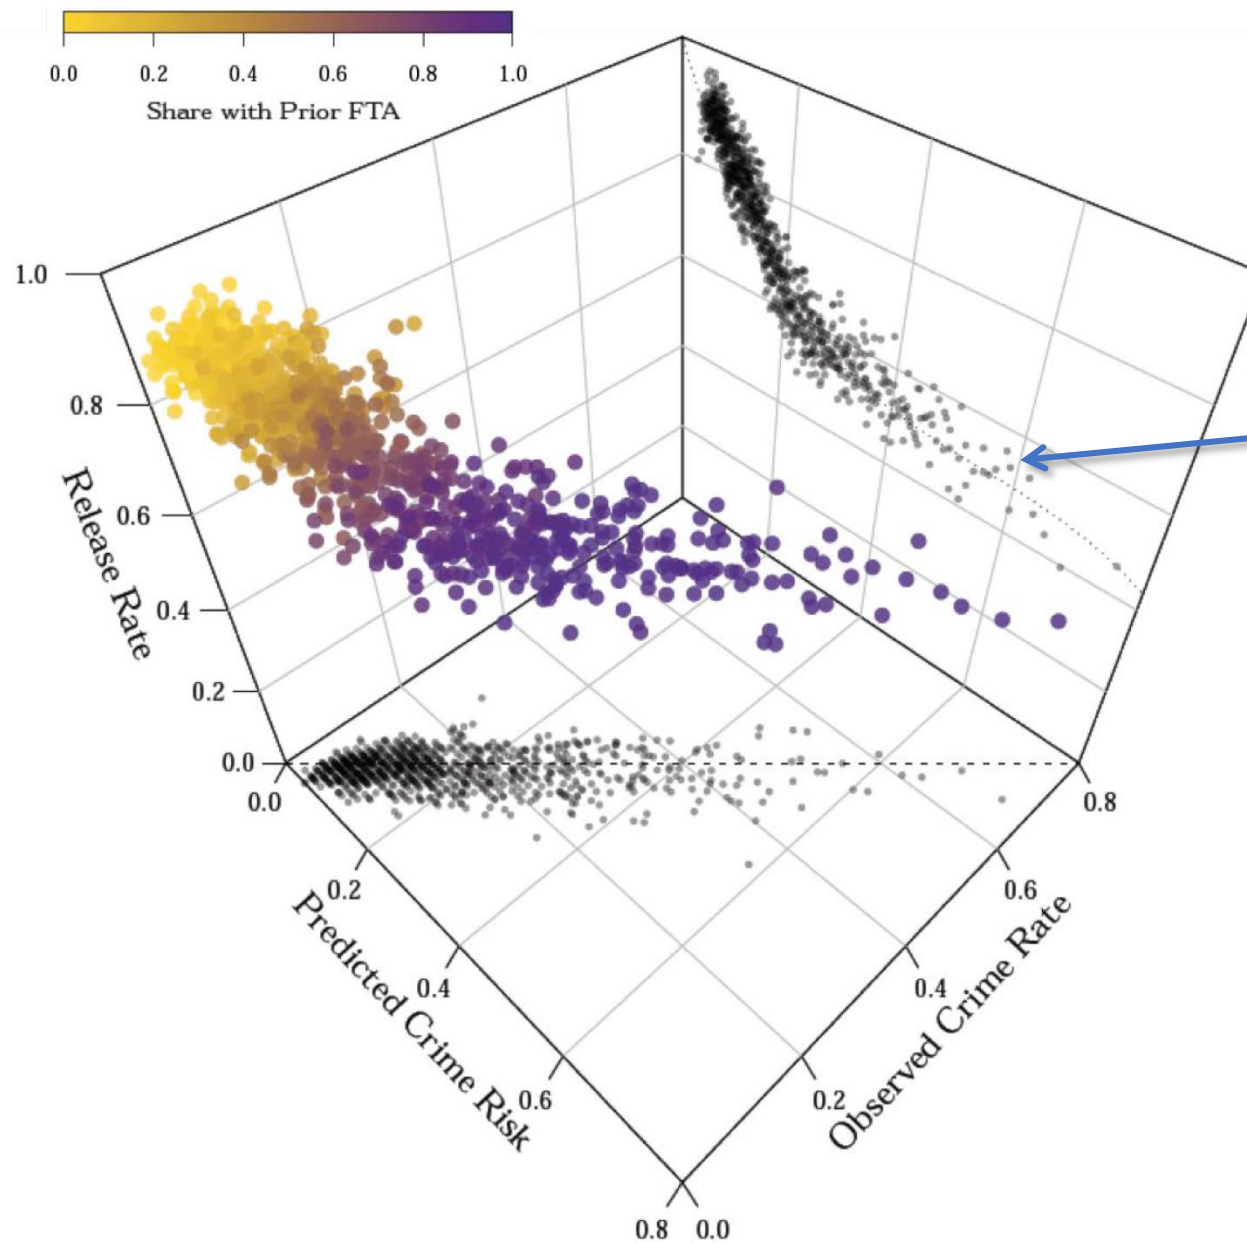


Judges' Release Decisions vs. Machine Predictions and Crime Risk



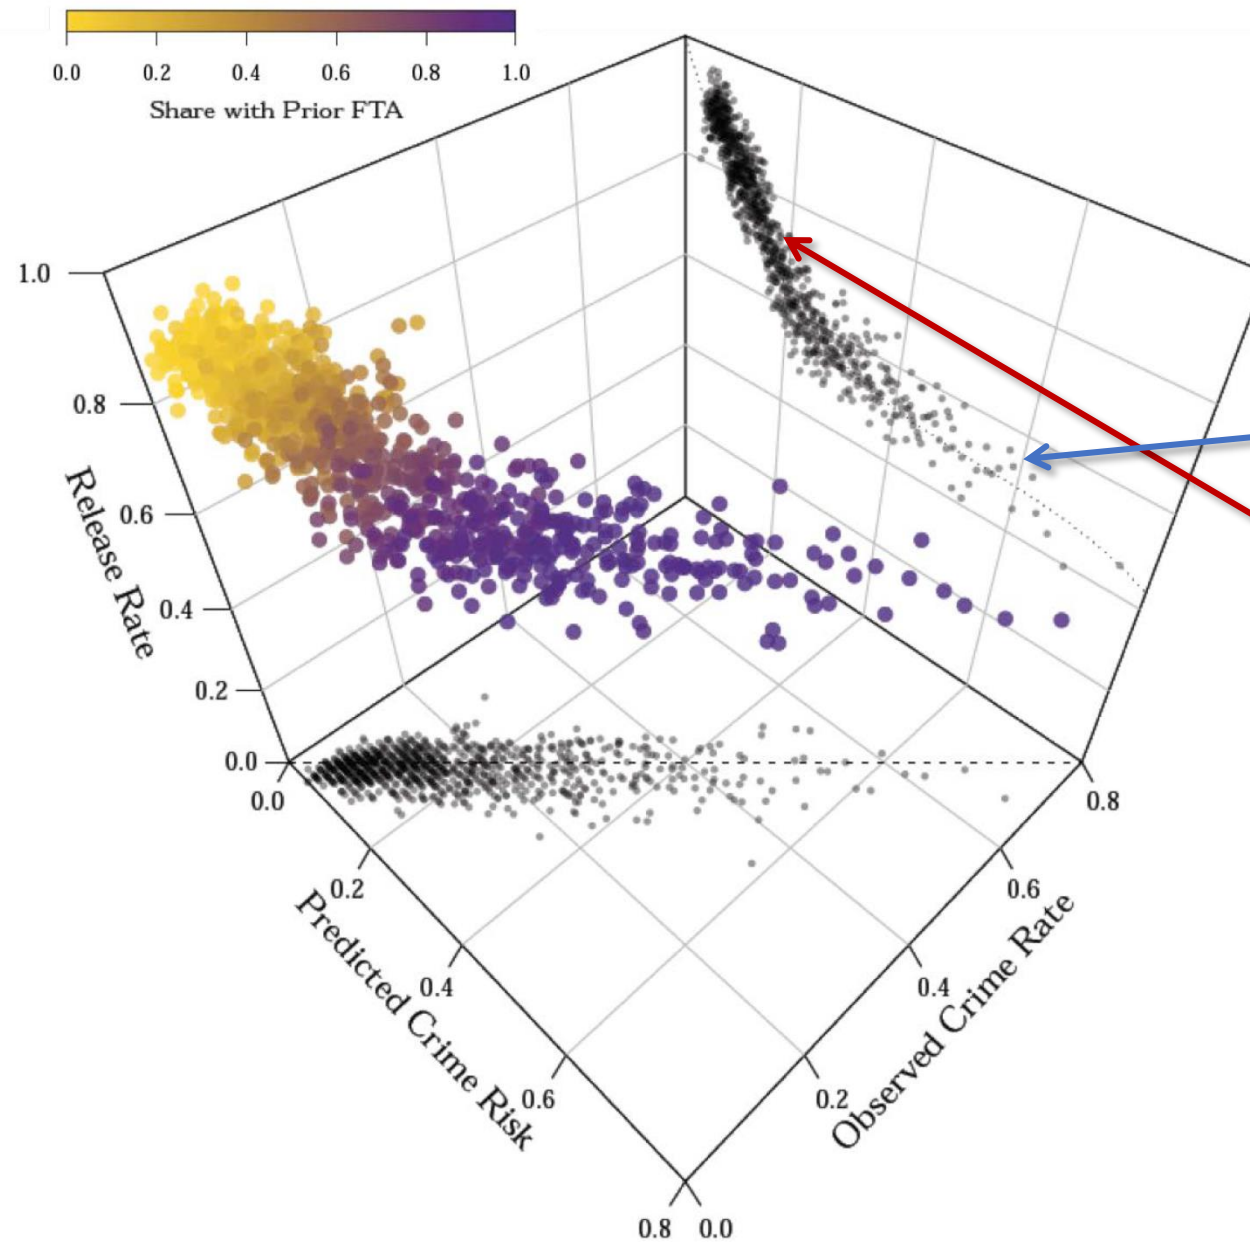
Predicted crime risk is highly correlated with observed crime rates (for defendants who are released)

Judges' Release Decisions vs. Machine Predictions and Crime Risk



Judges release 50% of defendants whose predicted crime risk exceeds 60%

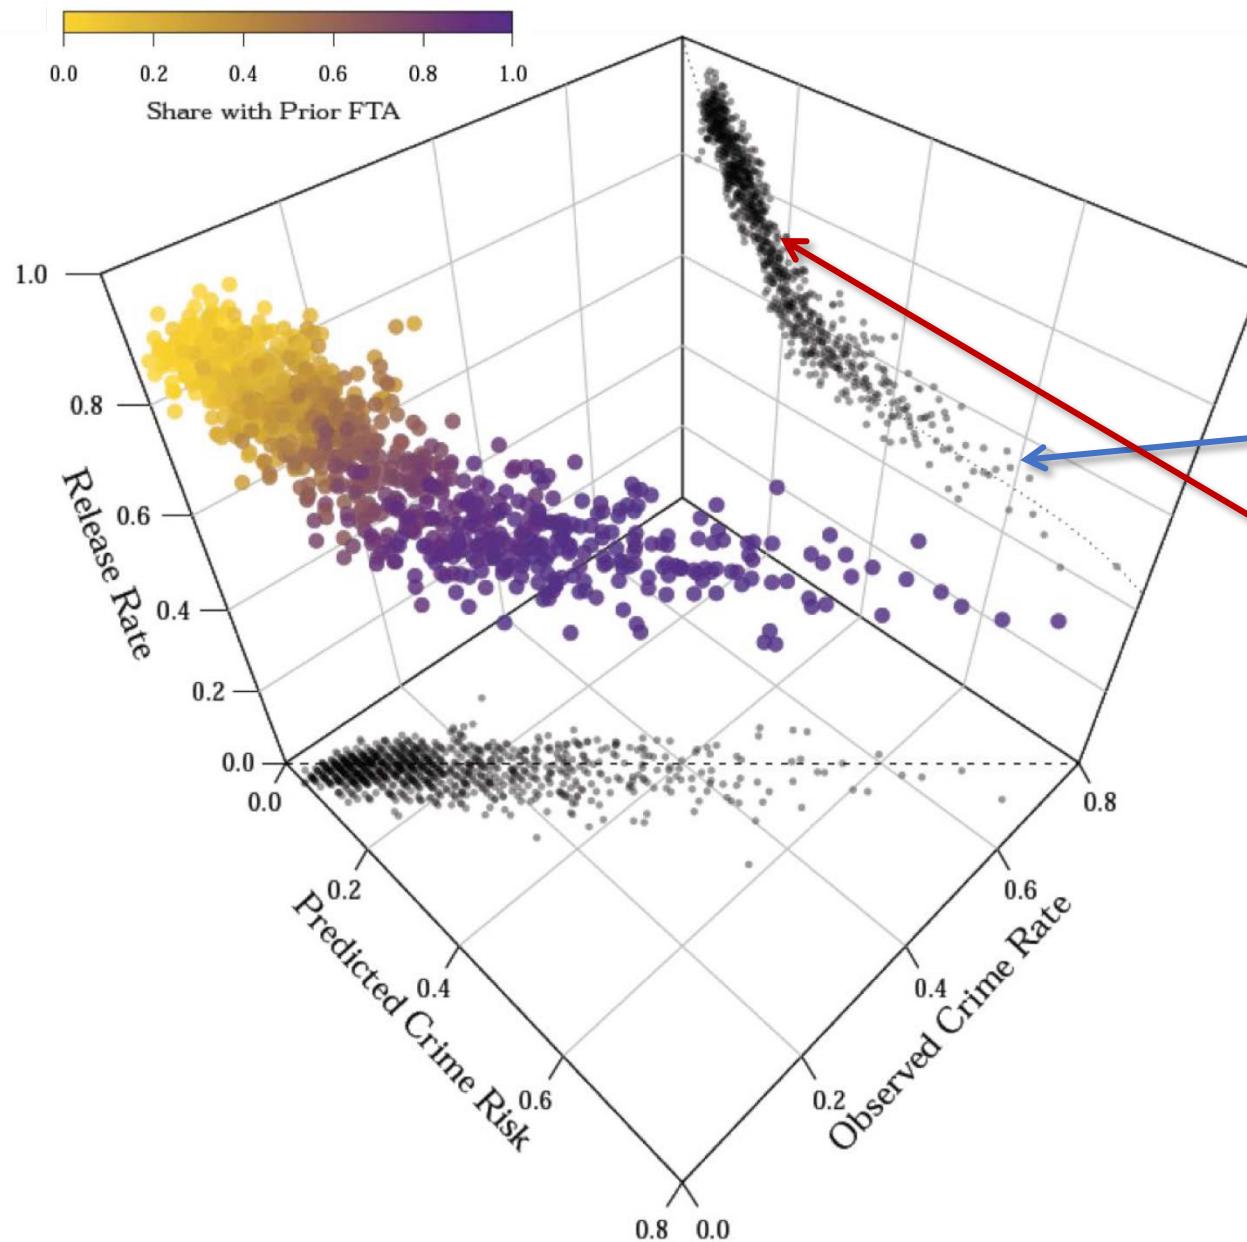
Judges' Release Decisions vs. Machine Predictions and Crime Risk



Judges release 50% of defendants whose predicted crime risk exceeds 60%

Yet they jail 30% of defendants with crime risks of only 20%

Judges' Release Decisions vs. Machine Predictions and Crime Risk



Judges release 50% of defendants whose predicted crime risk exceeds 60%

Yet they jail 30% of defendants with crime risks of only 20%

Swapping low-risk and high-risk defendants would keep jail population fixed while lowering crime rates


Comparing Machine Predictions to Human Predictions

- How large are the gains from machine prediction?
 - Crime could be reduced by 25% with no change in jailing rates
 - Or jail populations could be reduced by 42% with no change in crime rates
- Why? One explanation: Judges may be affected by cues in the courtroom (e.g., defendant's demeanor) that do not predict crime rates
- Whatever the reason, gains from machine-based “big data” predictions are substantial in this application

Predictive Policing

- Another active area of research and application of big data in criminology: predictive policing
 - Predict (and prevent) crime before it happens
- Two approaches: spatial and individual
 - Spatial methods rely on clustering of criminal activity by area and time

Times Between Burglary Events Separated by 0.1 Miles or Less



Predictive Policing

- Another active area of research and application of big data in criminology: predictive policing
 - Predict (and prevent) crime before it happens
- Two approaches: spatial and individual
 - Spatial methods rely on clustering of criminal activity by area and time
 - Individual methods rely on individual characteristics, social networks, or data on behaviors (“profiling”)

Ethical Debate Regarding Predictive Analytics

- Use of big data for predictive analytics raises serious ethical concerns, particularly in the context of criminal justice
- Tension between two views:
 - Should a person be treated differently simply because they share attributes with others who have higher risks of crime?
 - Should police/judges/decision makers discard information that could help make society fairer and potentially more just than it is now on average?
- Further research on costs and benefits of predictive analytics and study of ethical issues required before widespread application

Using Big Data to Solve Economic and Social Problems

- What we've learned in teaching this class
 - What we see in the data are reflected in your stories
 - Teaching the tools of social science through applications is valuable
 - Couldn't have done it without Rebecca, Christos, Evan, Jason, Juan, and Will
- What do we want you to take from this course?
 - Methods to analyze real-world problems
 - The value of taking a data-based, scientific approach to policy questions
 - An interest in focusing on the social good

An Opportunity and a Challenge

Area	Chance of Rising from Bottom to Top Fifth
Dubuque, IA	17.9%
San Jose, CA	12.9%
Washington DC	10.5%
Chicago, IL	6.5%
Memphis, TN	2.6%

