



# Using Big Data to Solve Economic and Social Problems

Professor Raj Chetty

Head Section Leader: Gregory Bruich, Ph.D.

Spring 2019



HARVARD  
UNIVERSITY



OPPORTUNITY  
INSIGHTS

# **Predictive Analytics in Criminal Justice**

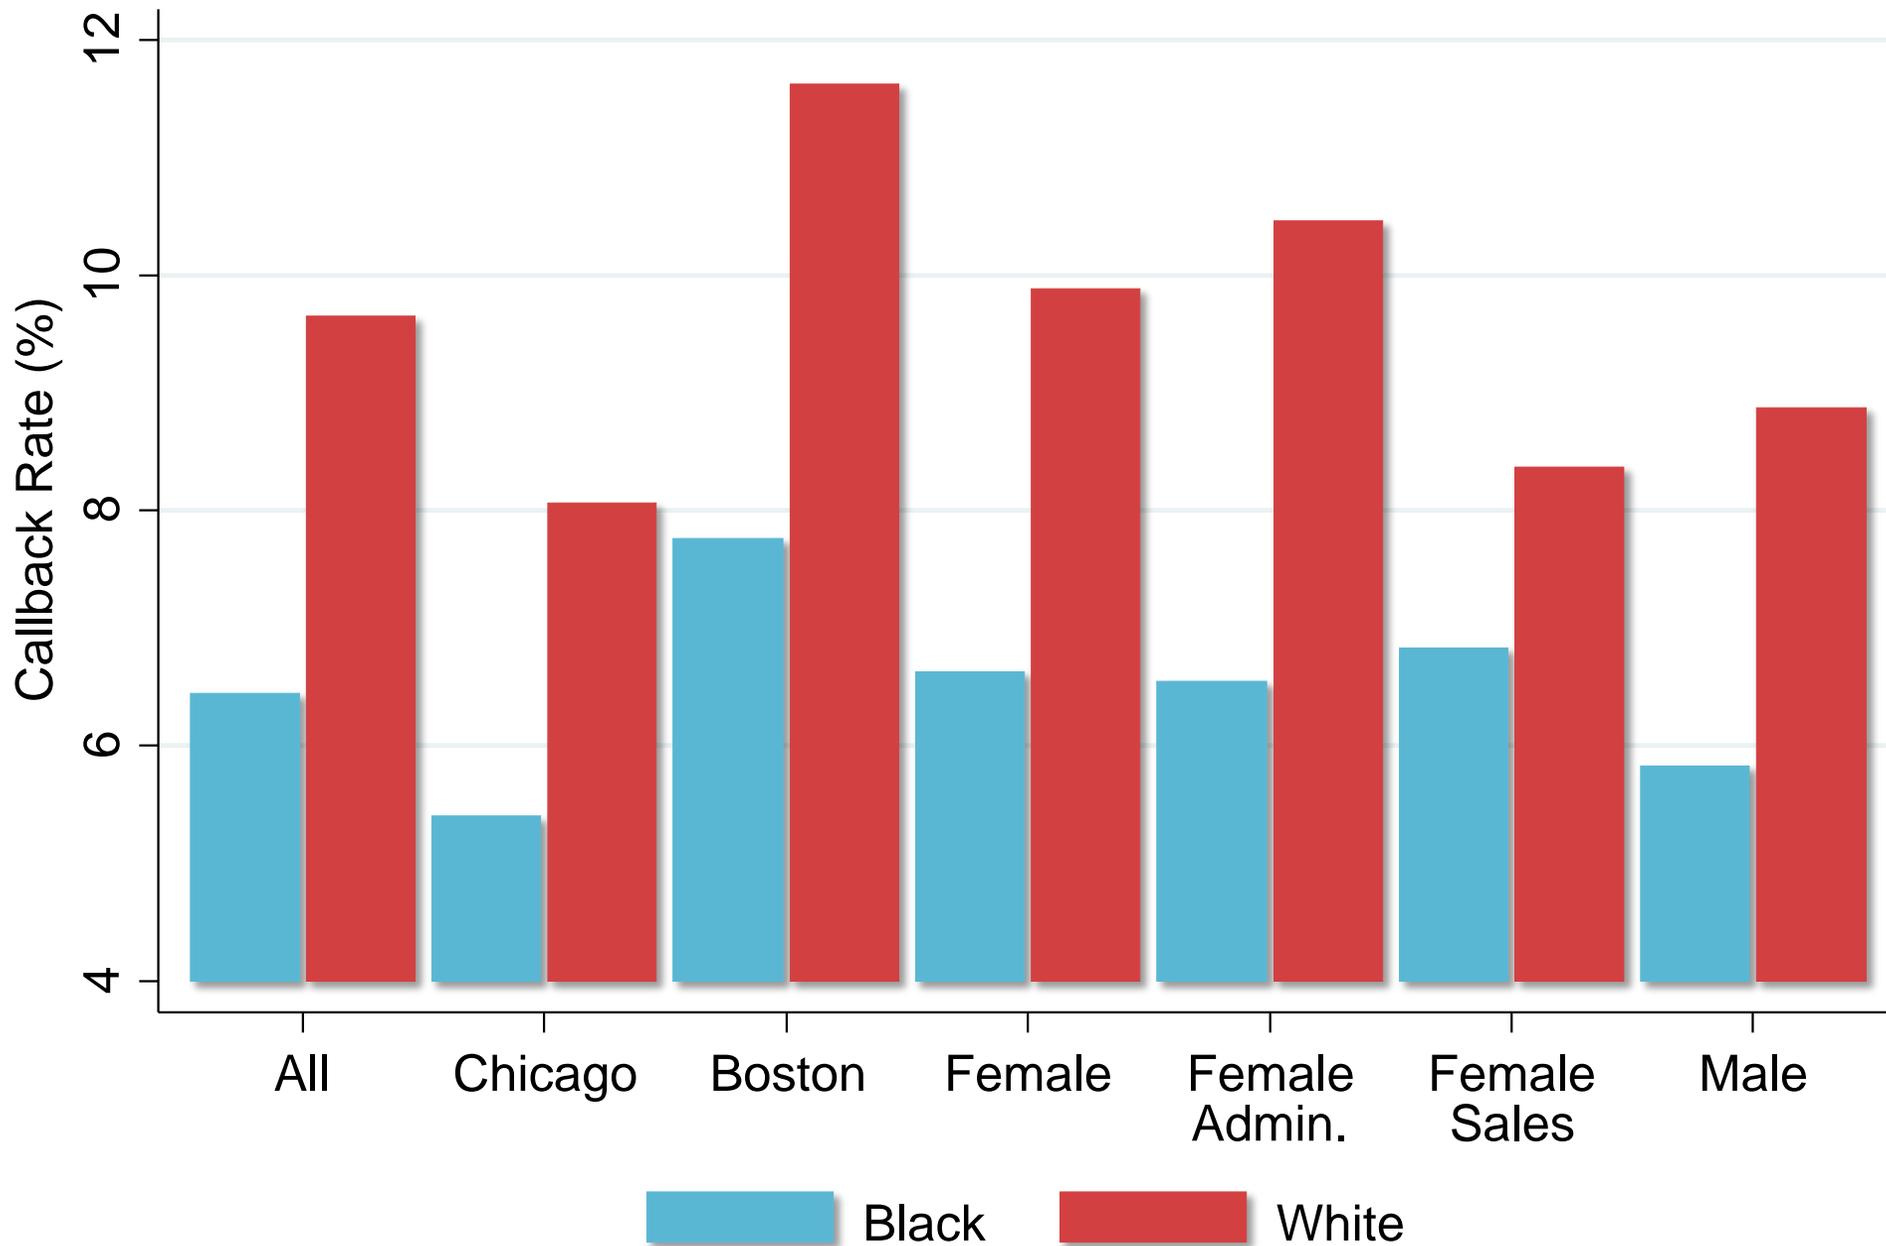
# Motivation: Biases in Human Decision Making

- Growing interest in using machine learning (predictive analytics) tools to aid decision makers, e.g. in the context of criminal justice
- Humans' decisions often exhibit substantial biases
- Begin by demonstrating this fact with a series of empirical examples

## Racial Discrimination in Hiring

- Bertrand and Mullainathan (2004) study biases in hiring by sending 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

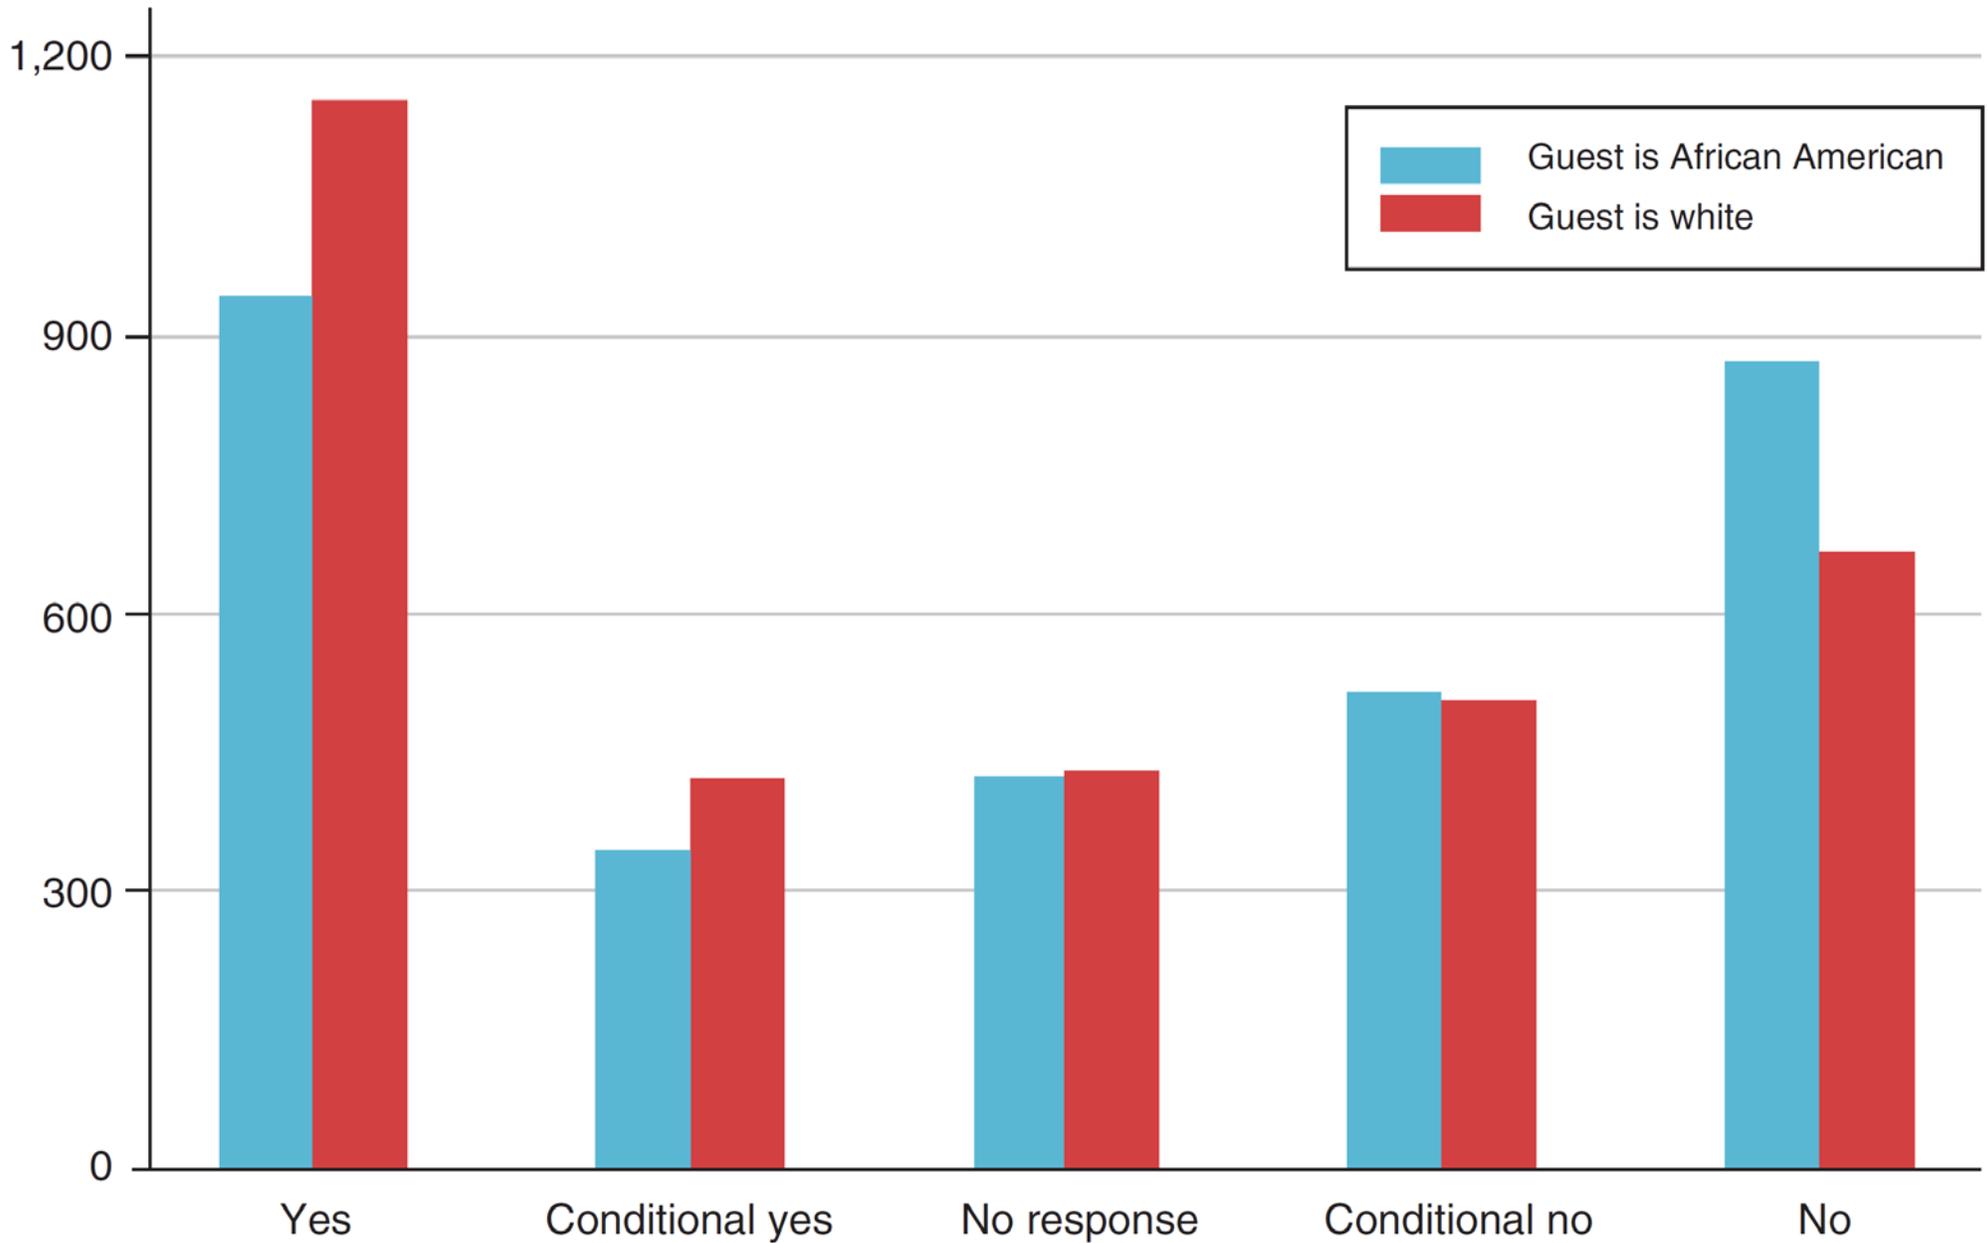


Source: Bertrand and Mullainathan (AER 2004)

# Racial Discrimination Among Airbnb Hosts

- Edelman, Luca, and Sversky (2017) use an analogous “audit study” to analyze 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)

# 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

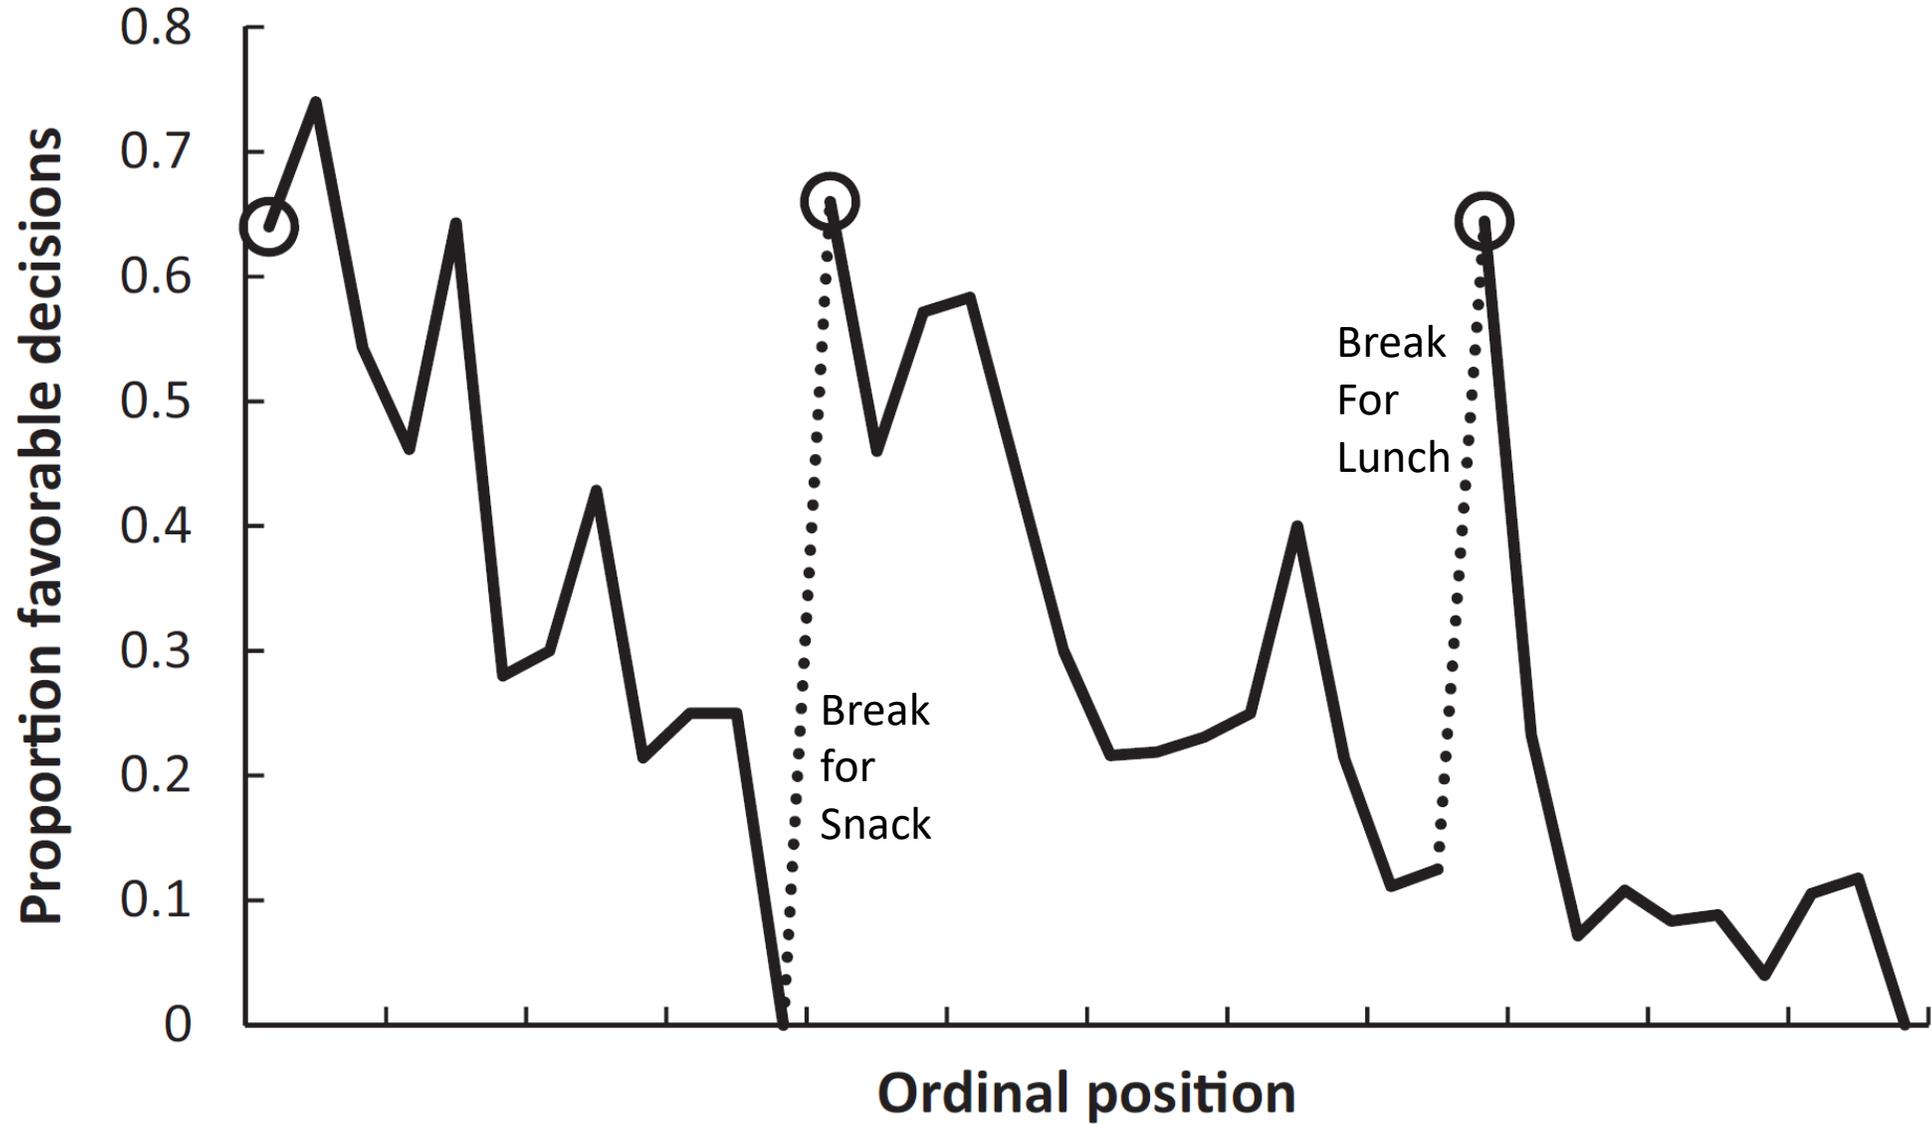
## Biases due to Decision Fatigue

- Such biases in decisions are not driven entirely by deep-rooted beliefs
- Decisions also vary greatly based on transitory 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?

- Can machine learning help us reduce such biases?
- Idea: develop algorithms to predict outcomes of interest and use these to guide or replace human decisions
- These algorithms are not necessarily subject to human biases, but do they outperform humans' decisions overall or have other biases/shortcomings?

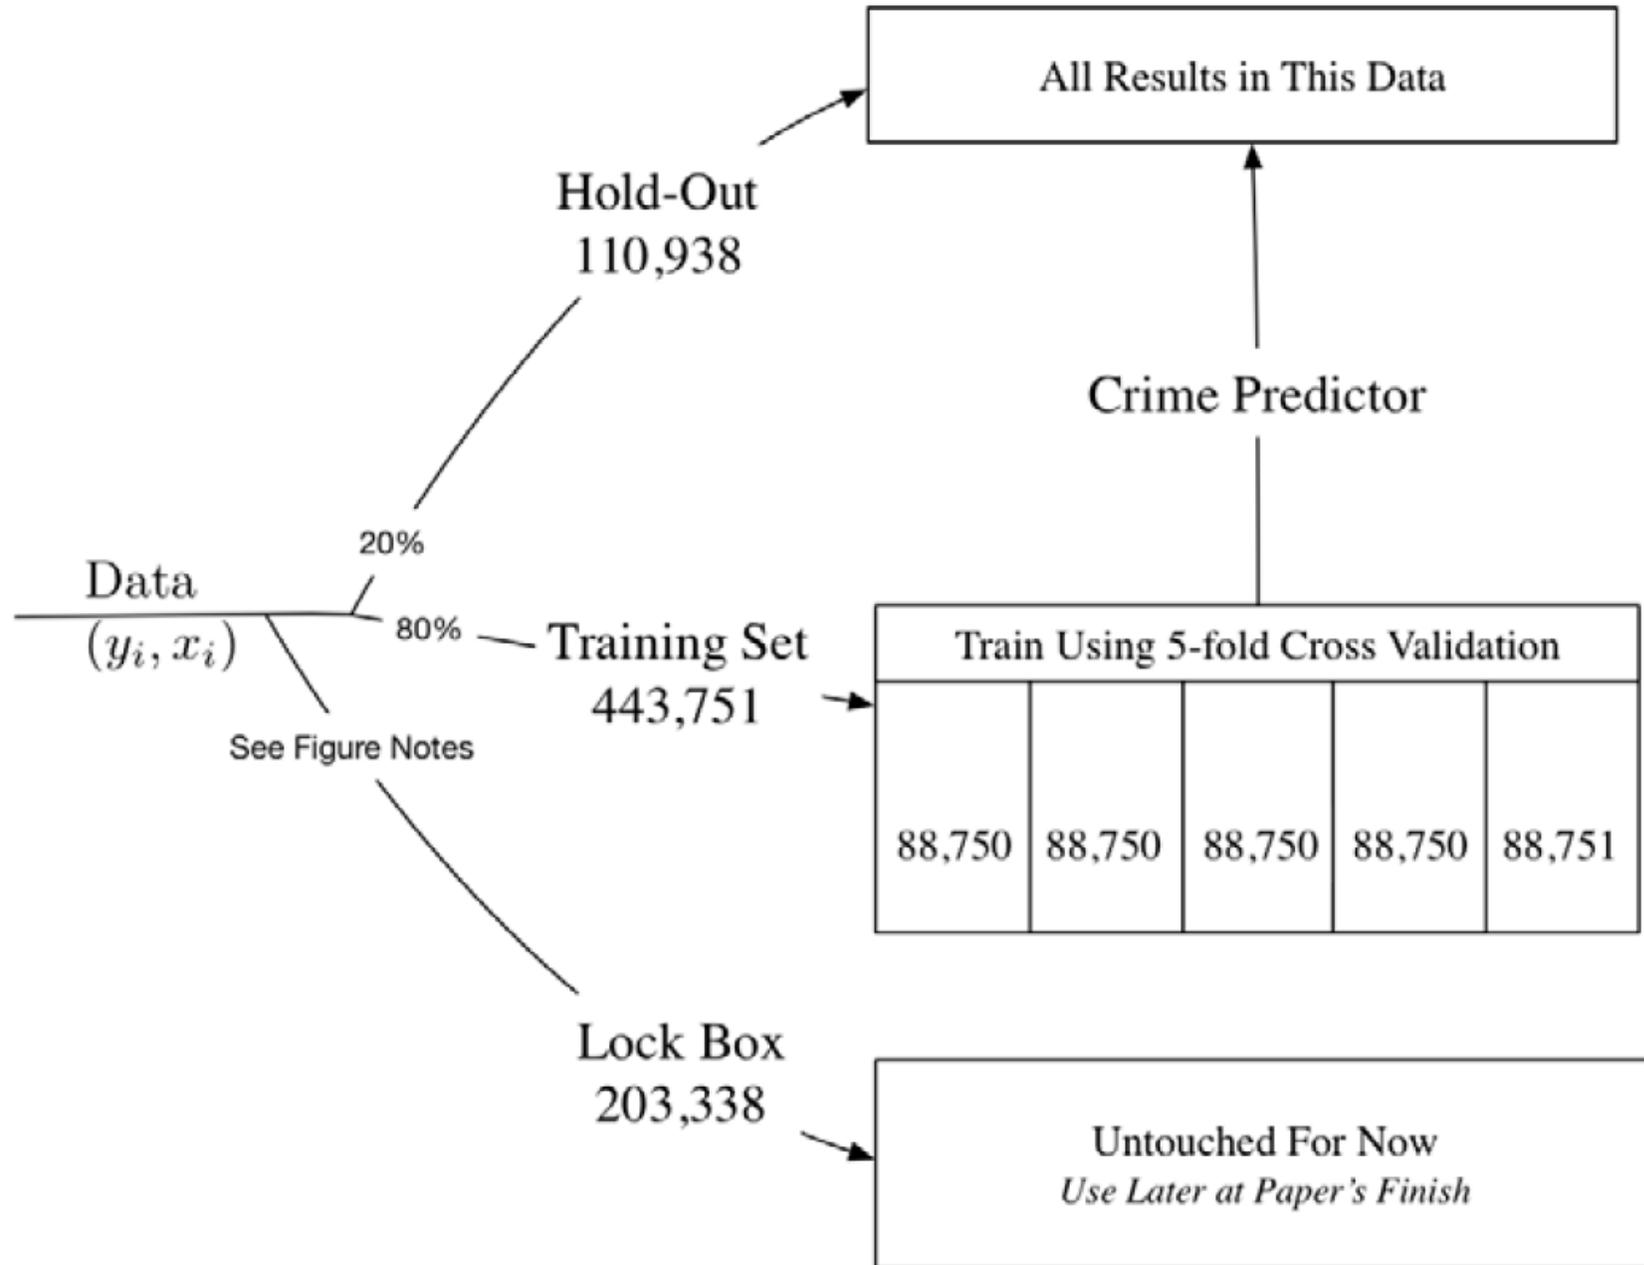
# Decisions to Jail vs. Release Defendants

- 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



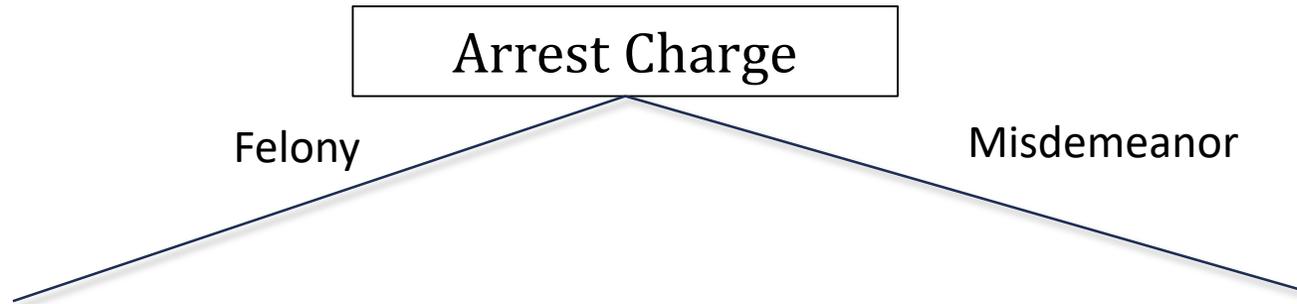
# Methodology: Machine Learning Using Decision Trees

- 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

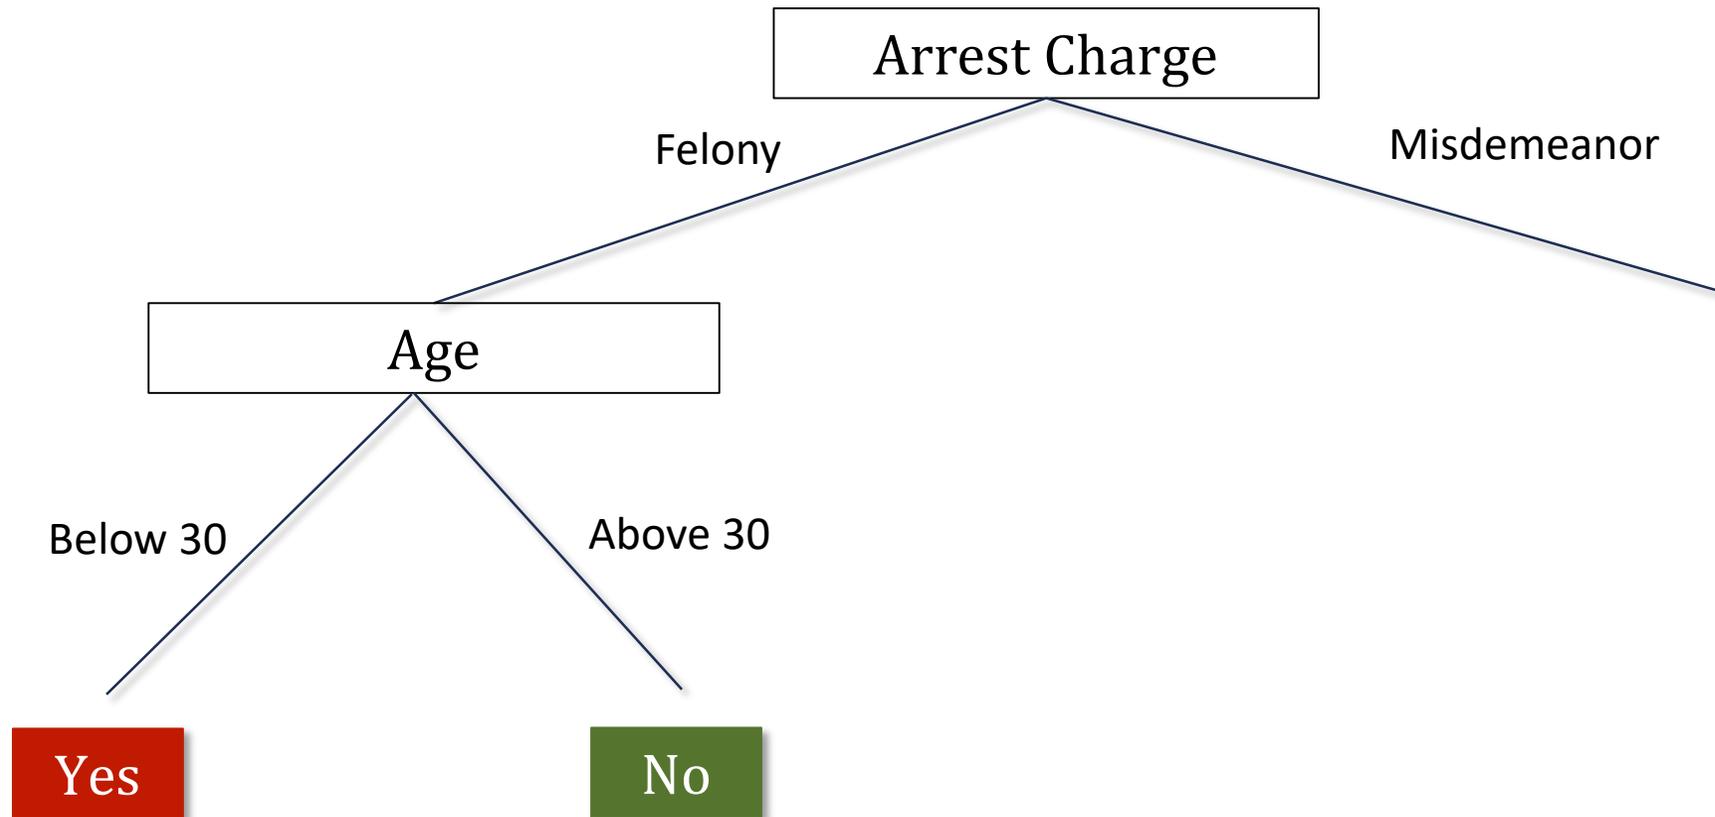
# Hypothetical Decision Tree for Decision to Jail Defendant



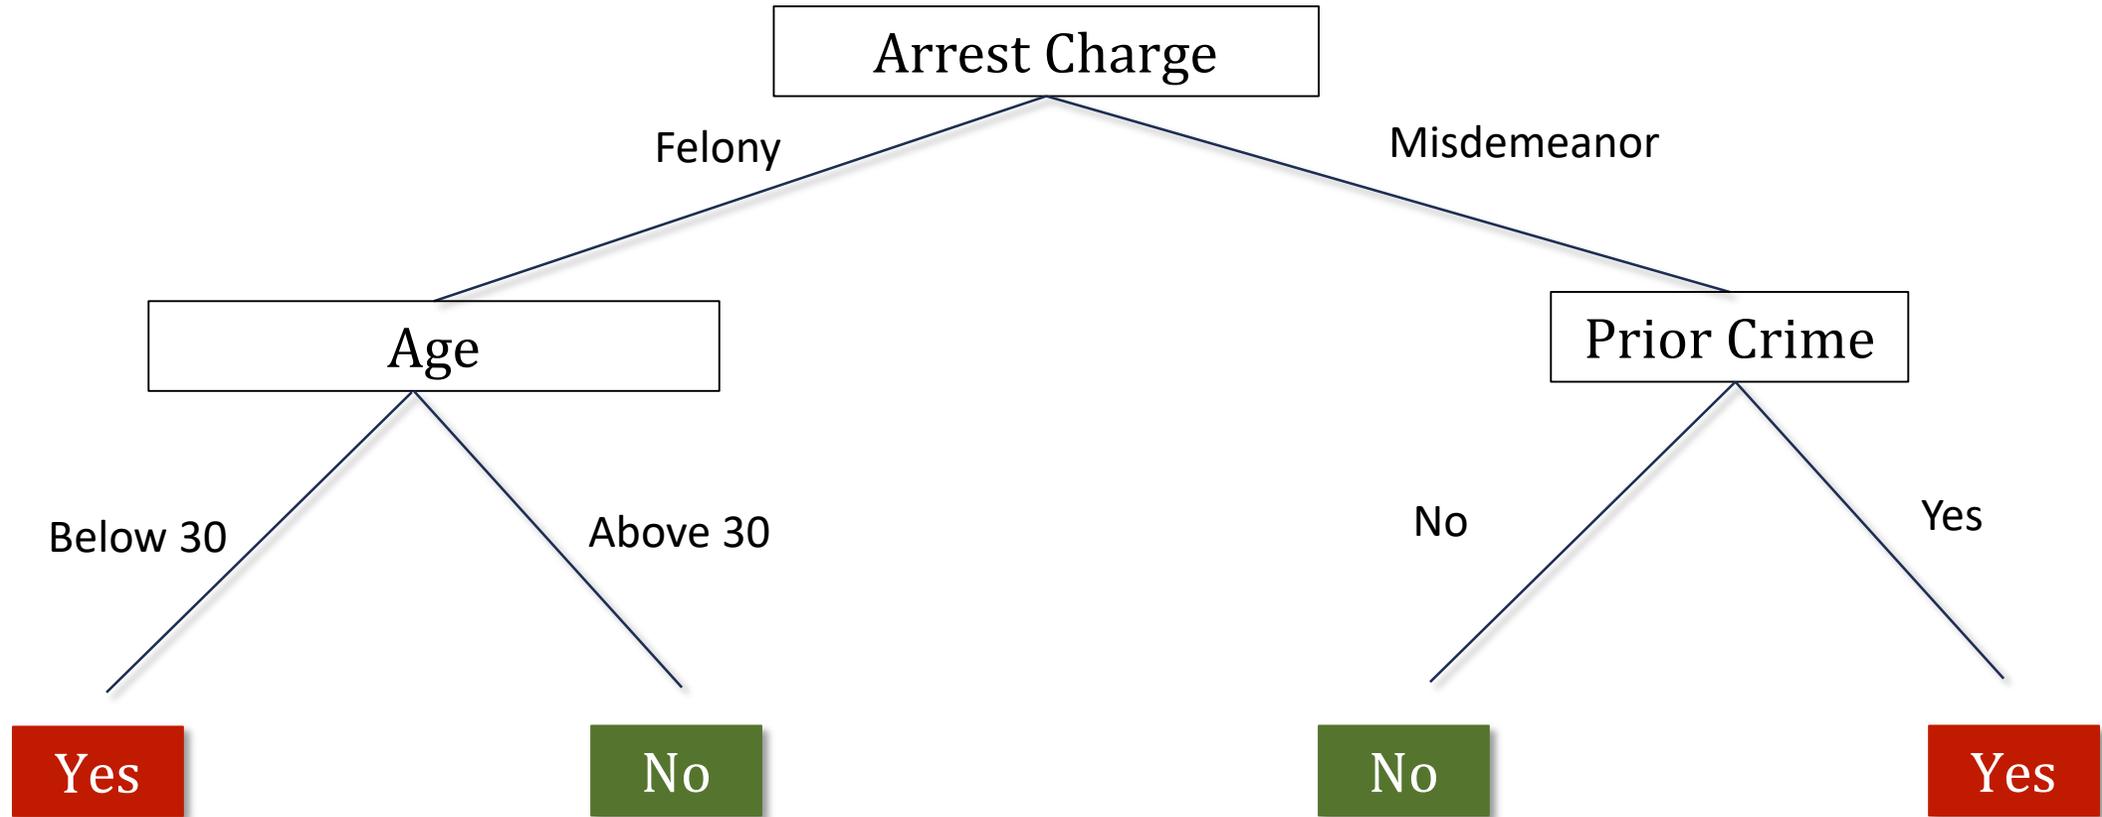
# Methodology: Machine Learning Using Decision Trees

- 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



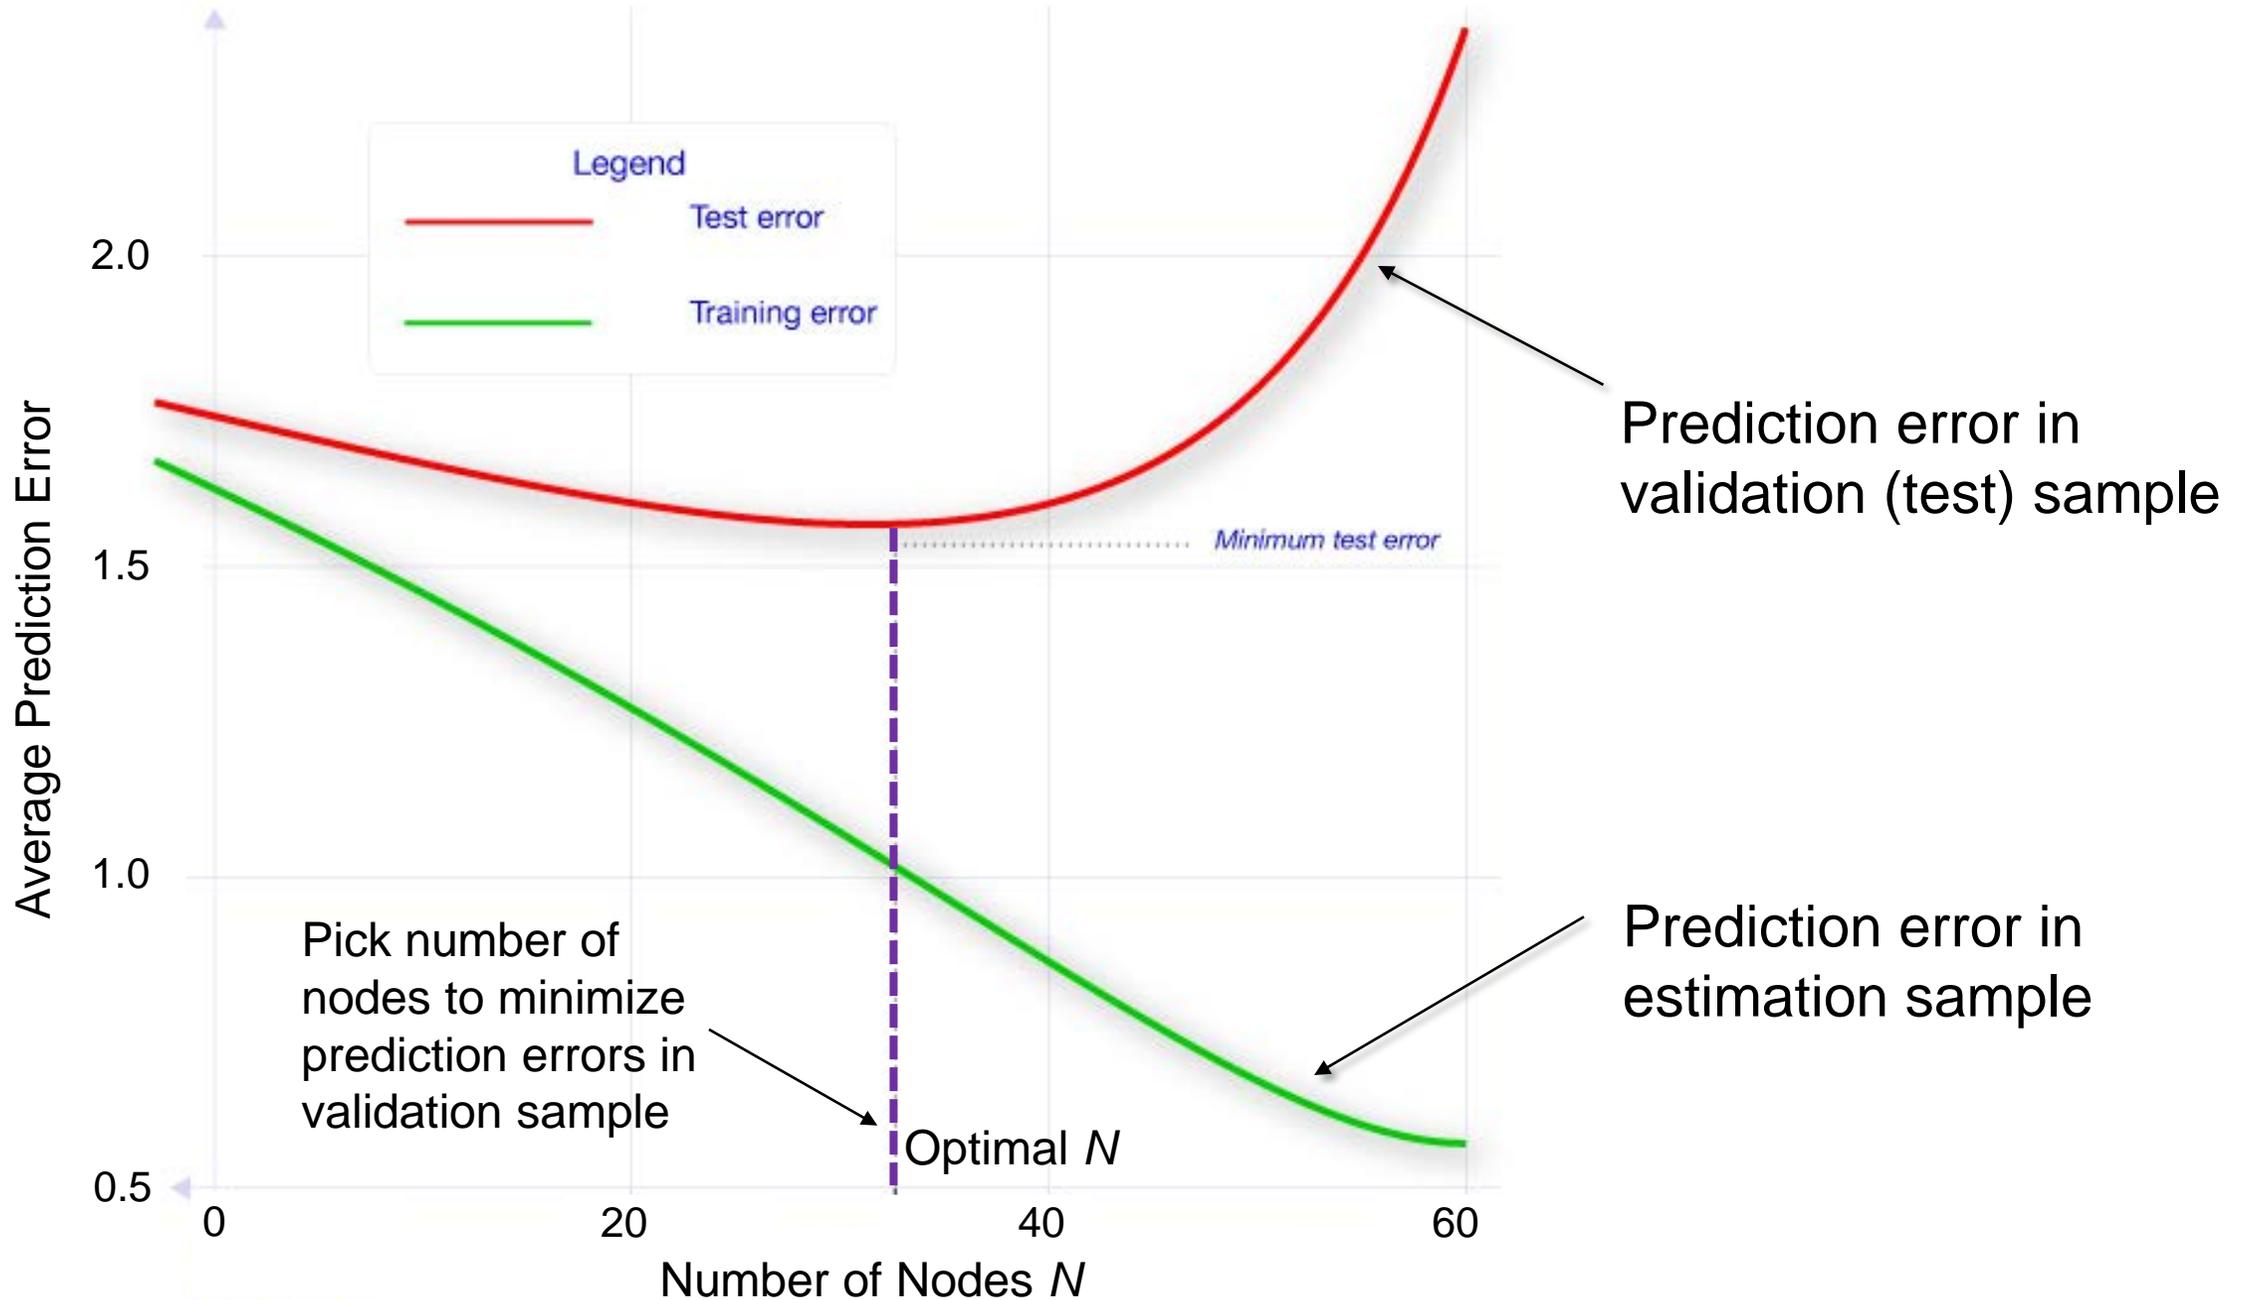
# Hypothetical Decision Tree for Decision to Jail Defendant



# 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

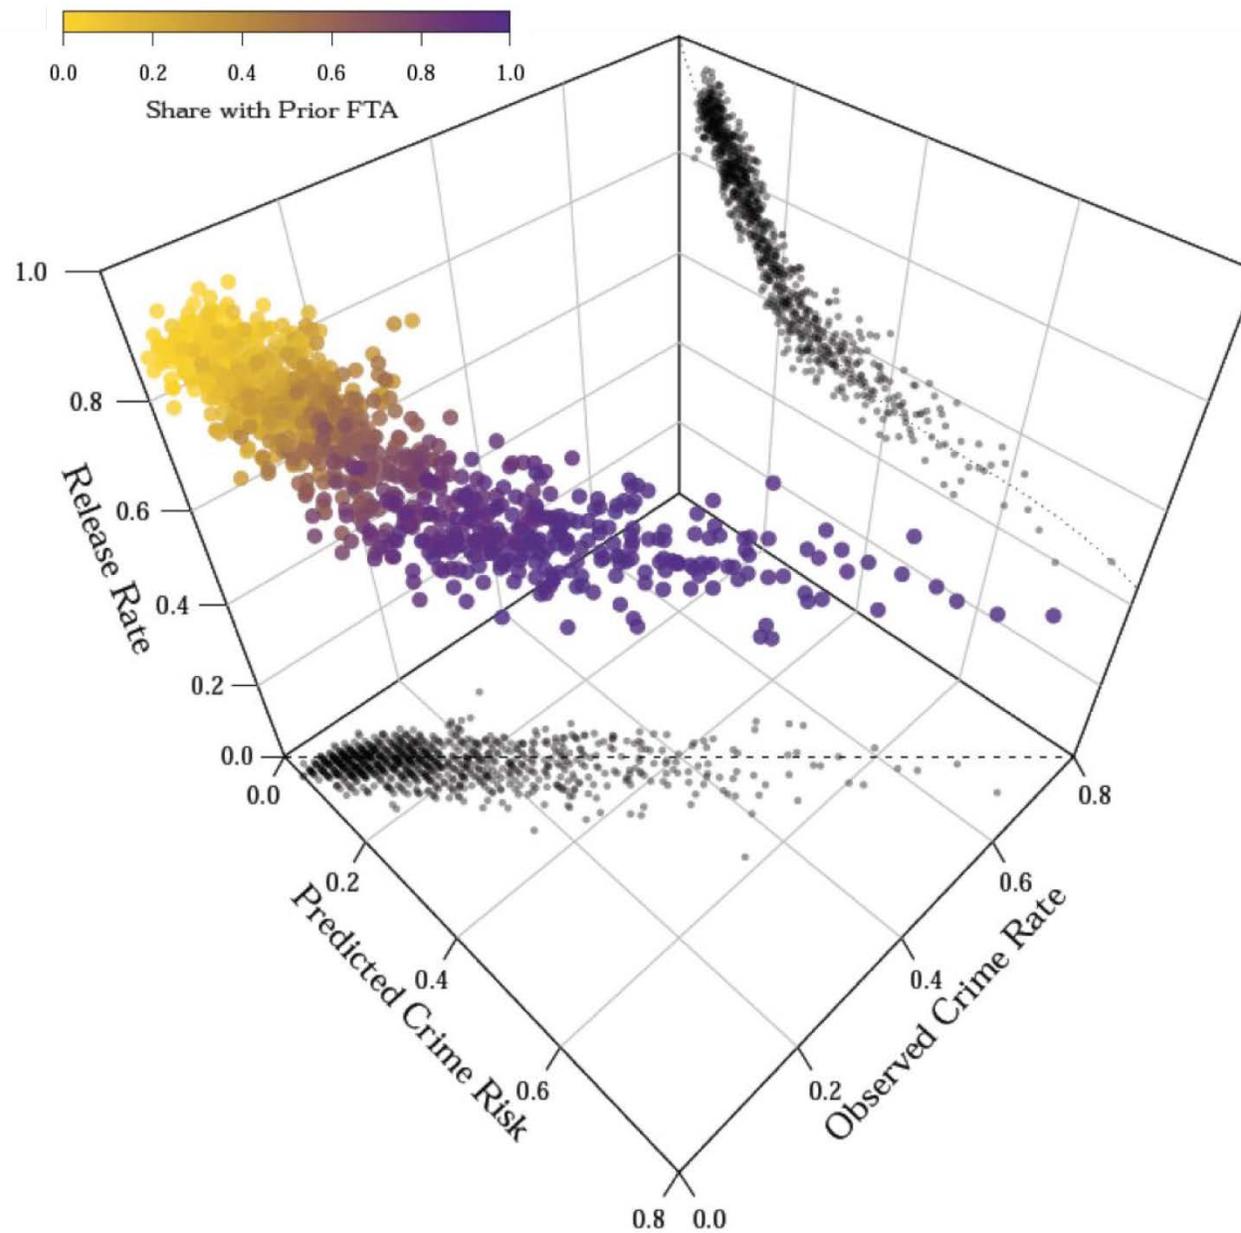
# Prediction Errors vs. Size of Decision Tree



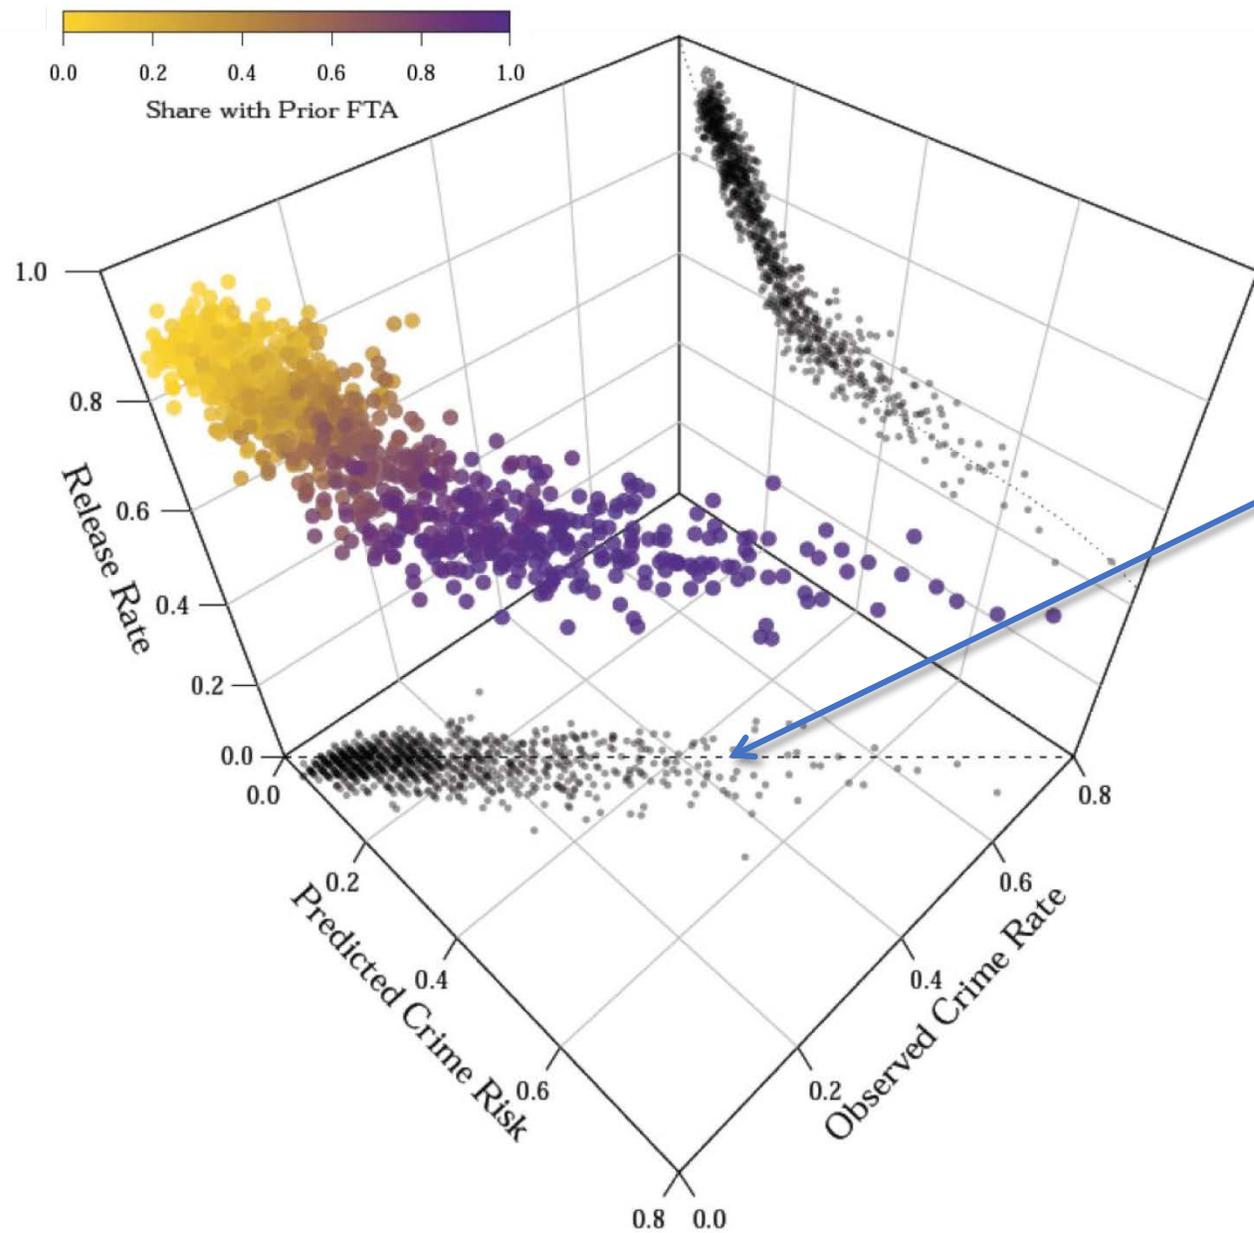
# Comparing Machine Predictions to Human Predictions

- 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

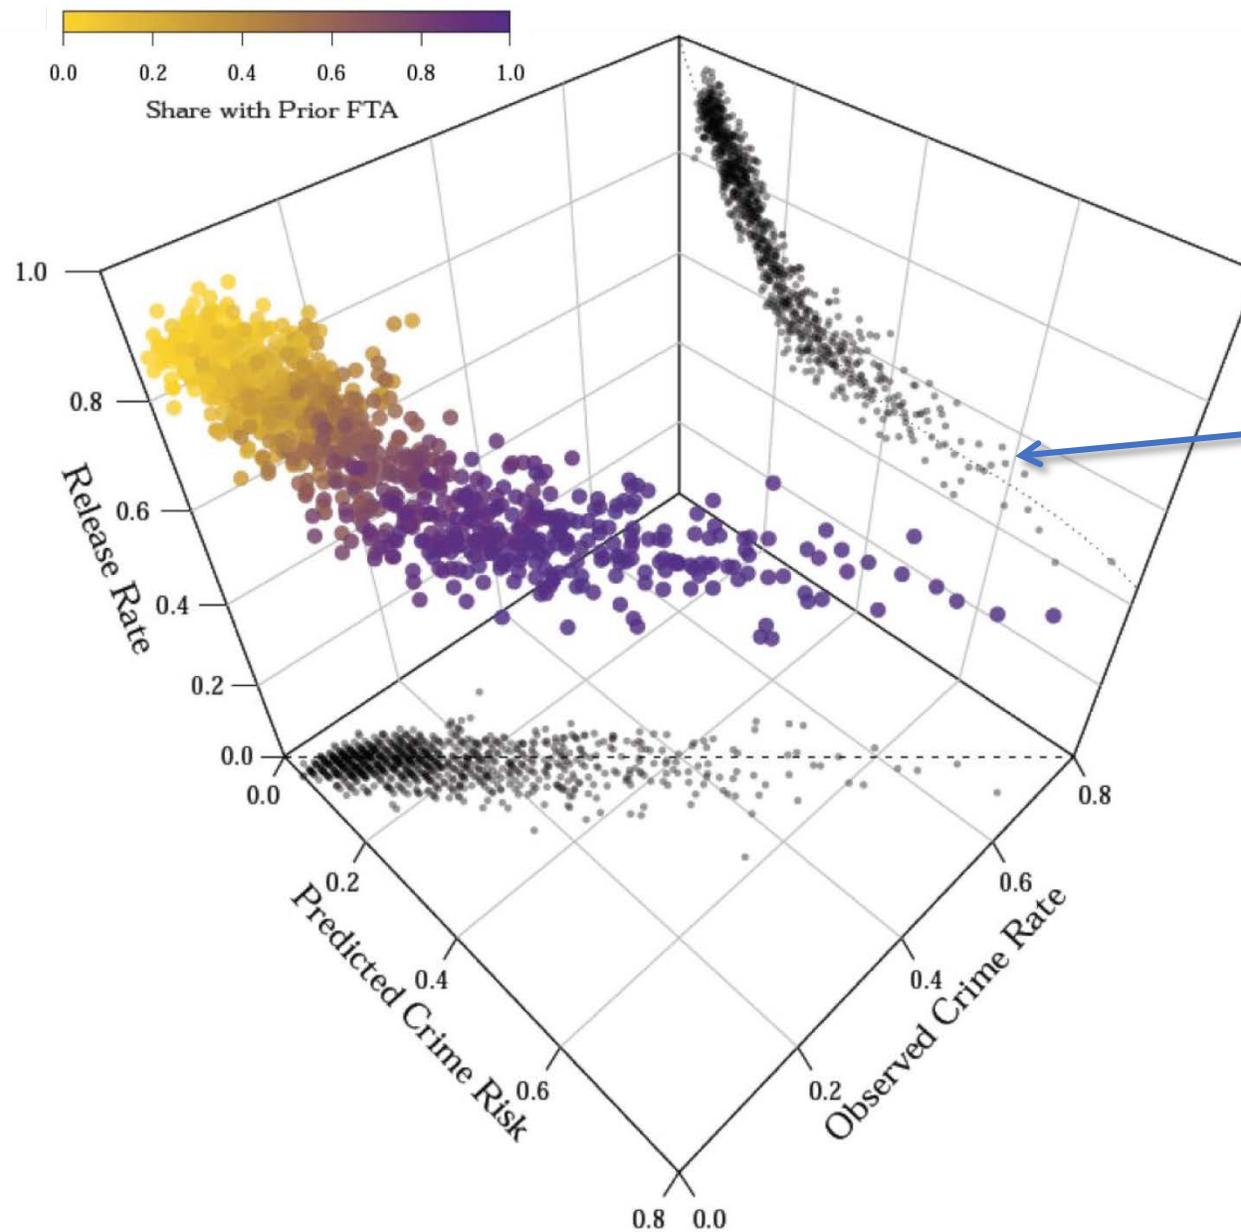


# Judges' Release Decisions vs. Machine Predictions and Crime Risk



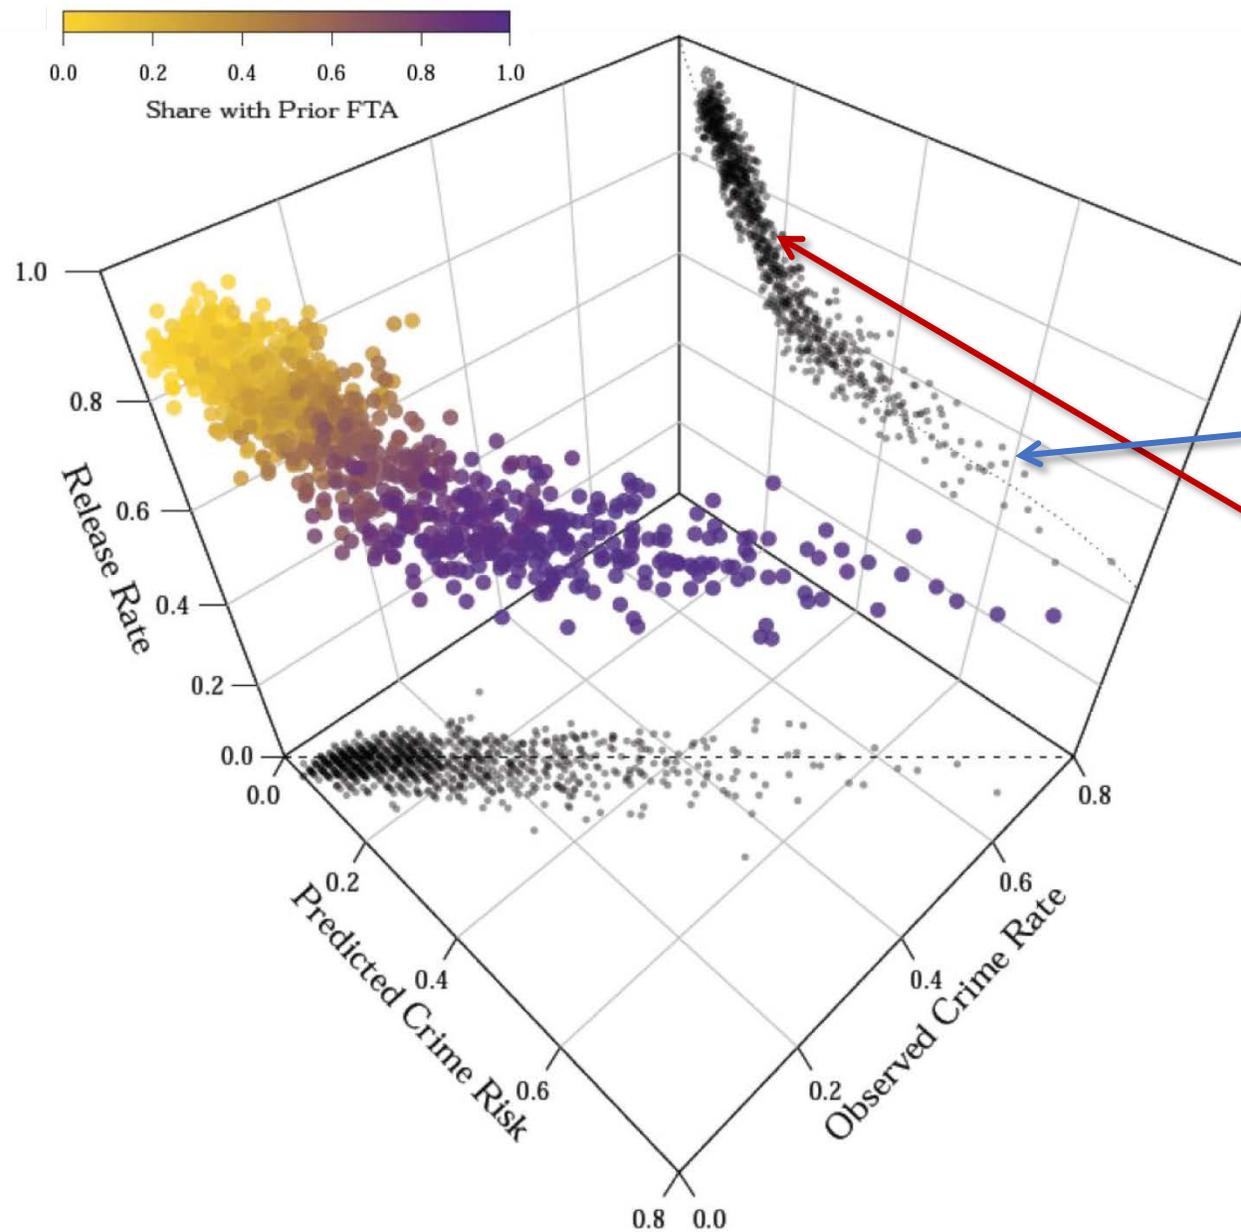
Predicted crime risk is highly correlated with observed crime rates (for defendants who are released)

# Judges' Release Decisions vs. Machine Predictions and Crime Risk



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

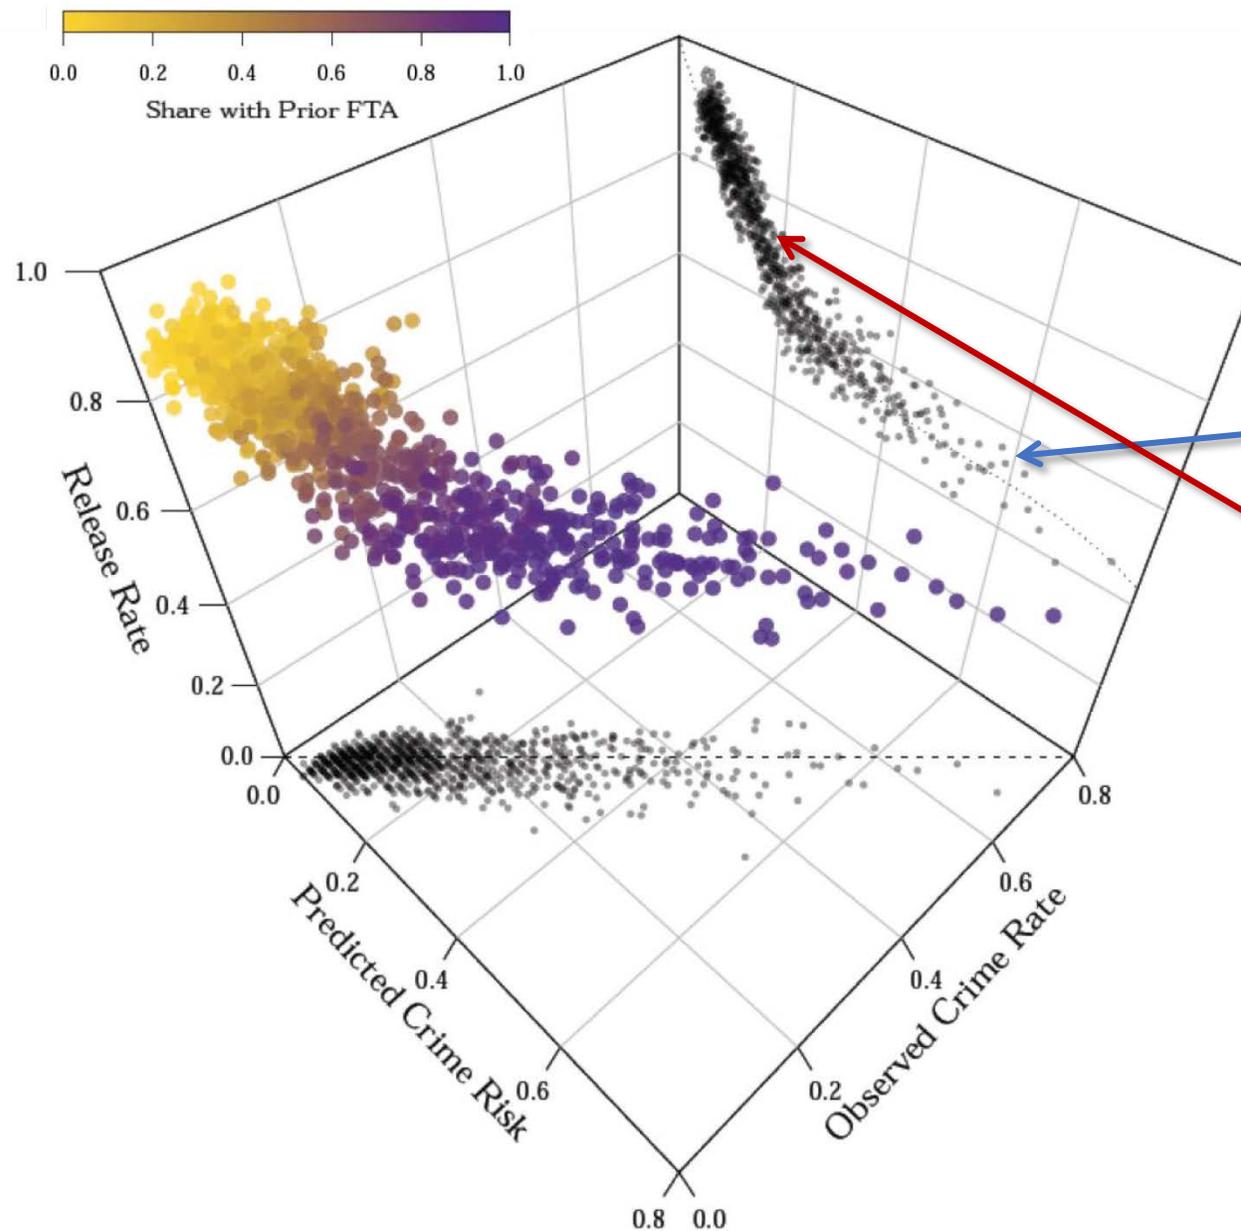
# Judges' Release Decisions vs. Machine Predictions and Crime Risk



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

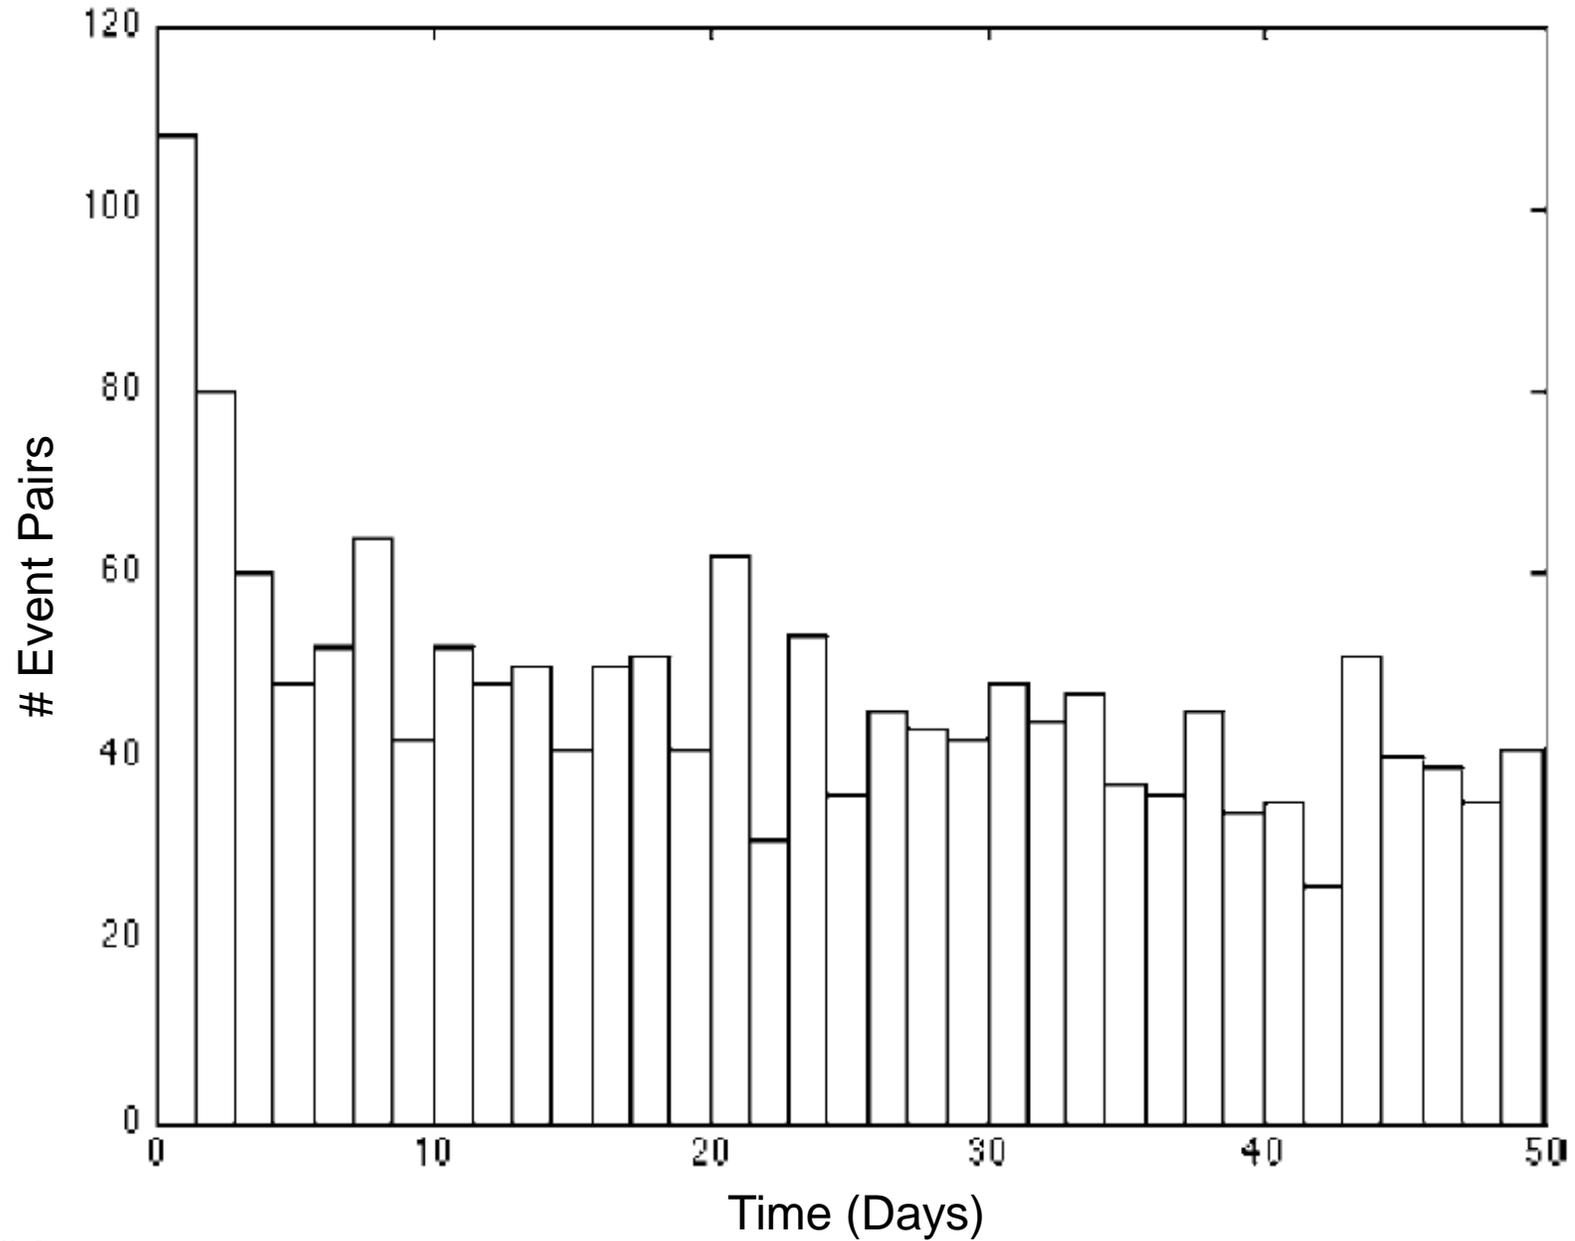
# Comparing Machine Predictions to Human Predictions

- 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”)

# 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?
- Discuss these issues and practical application of predictive analytics in guest lecture on Thursday
  - Lynn Overmann and Catie Bialick (Arnold Ventures) and Peter Koutoujian (Middlesex County Sherrif)