ML for SS: Policy prediction

Dr. Richard M. Crowley

rcrowley@smu.edu.sg

https://rmc.link/

Overview

Papers

> Kleinberg, Mullainathan and Obermeyer 2015 AER

- A simple explanation of *why* ML is more appropriate for policy problems

> Athey and Imbens 2017 JEP

- Provides a broad overview of causality using traditional econometrics approaches and machine learning approaches

> Chalfin et al. 2016 AER

- A couple simple examples of policy problems approached using ML

A conceptual dive into a policy prediction problem

The problem

> In the criminal justice system, for instance, judges have to decide whether to detain or release arrestees as they await adjudication of their case—a decision that depends on a prediction about the arrestee's probability of committing a crime.
>
> — **Kleinberg, Mullainathan and Obermeyer 2015 AER**

Kleinberg et al. 2018 QJE

> Attempts to predict the likelihood of a released offender comitting a crime

Some notes:

1. Main data is from NYC, NY
    - The problem judges are supposed to solve is whether to release someone before a trial; in New York, the only consideration is *flight risk*
        - Not likelihood to commit a crime
2. Data used in the model includes case characteristics, prior criminal record, age
    - Excludes race, ethnicity, and gender
        - These are correlated with judge decisions, however
3. Apply a GBM + cross validation

Difficulties of related algorithms in practice: COMPAS

> “Machine Bias” by ProPublica

Prediction Fails Differently for Black Defendants

	WHITE	AFRICAN AMERICAN
Labeled Higher Risk, But Didn't Re-Offend	23.5%	44.9%
Labeled Lower Risk, Yet Did Re-Offend	47.7%	28.0%

Overall, Northpointe's assessment tool correctly predicts recidivism 61 percent of the time. But blacks are almost twice as likely as whites to be labeled a higher risk but not actually re-offend. It makes the opposite mistake among whites: They are much more likely than blacks to be labeled lower risk but go on to commit other crimes. (Source: ProPublica analysis of data from Broward County, Fla.)

- A more technical write-up is available at this link
- The code (R) is available on Github

The difficulty is in the econometrics

Black defendants

	Re-offend	Did not
Low risk	0.1439394	0.2678571
High risk	0.3704004	0.2178030

White defendants

	Re-offend	Did not
Low risk	0.1878566	0.4641402
High risk	0.2057865	0.1422168

	black	white
FPR|1-Specificity	0.4484680	0.2345430
TPR|Sensitivity	0.7201473	0.5227743
Precision|PPV	0.6297148	0.5913349
Accuracy	0.6382576	0.6699267

- Propublica's analysis shows a difference in False Positive Rate, $\frac{FP}{FP+TN} \Big| \mathrm{Race}$
- COMPAS is optimized using PPV, $\frac{TP}{TP+FP} \Big| \mathrm{Race}$

> What is optimal to optimize here?

- Note: Observed recidivism rates are different by race – this is why the above statistics don't agree!

How this is addressed in Kleinberg et al. 2018 QJE

In addition to an “efficient” algorithm (minimize recidivism), consider an algorithm that, at a given threshold…

1. Guarantees no more black individuals are jailed than a judge would jail
2. Matches judges jailing rates of black and hispanic defendants
3. Requires equal rates of jailing across all races (white, black, hispanic)
4. Requires no race to be jailed at rates above option #3 or judges’ rates

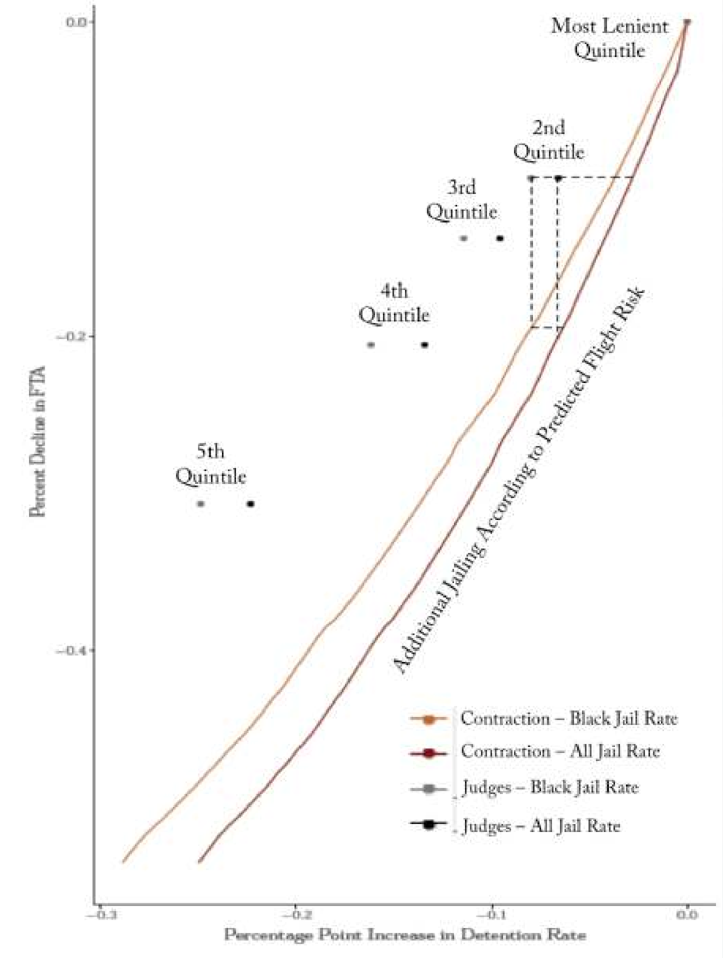

Figure A.5: Effect of Contraction on Racial Bias

All 4 algorithms decreased *fail to appear* outcomes when trained on them.

Conclusion

Wrap-up

> Policy prediction problems are practical in economics

- Anything where what matters is the outcome, not the cause

> Policy prediction can be done using tools we have already covered in this course

- LASSO, Random Forest, GBM, etc.

Packages used for these slides

Python

N/A

R

- kableExtra
- knitr
- quarto
- reticulate
- revealjs

References

- Athey, Susan, and Guido W. Imbens. "The state of applied econometrics: Causality and policy evaluation." Journal of Economic Perspectives 31, no. 2 (2017): 3-32.
- Chalfin, Aaron, Oren Danieli, Andrew Hillis, Zubin Jelveh, Michael Luca, Jens Ludwig, and Sendhil Mullainathan. "Productivity and selection of human capital with machine learning." American Economic Review 106, no. 5 (2016): 124-27.
- Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan. "Human decisions and machine predictions." The quarterly journal of economics 133, no. 1 (2018): 237-293.
- Kleinberg, Jon, Jens Ludwig, Sendhil Mullainathan, and Ziad Obermeyer. "Prediction policy problems." American Economic Review 105, no. 5 (2015): 491-95.