Counterfactually Fair Dynamic Assignment: A Case Study on Policing

Tasfia Mashiat, Xavier Gitiaux, Huzefa Rangwala, Sanmay Das

George Mason University

{tmashiat, xgitiaux, rangwala, sanmay} @gmu.edu

Motivation

- Relative levels of police force deployment to different areas can unfairly target minoritized groups and lead to problematic feedback loops [1].
- We combine theoretical modeling of fair police deployments across geographic areas with a novel dataset of actual deployments over time across beats in Chicago.
- Our methods enable novel kinds of analysis of real-world data and potentially more fair responsive dynamic deployment algorithms.

Allocation Policies

- ***Optimal.*** A benevolent planner that knows the structure of the causal model in Figure 2(a).
- ***PropFair.*** This method estimates parameters and criminal activity Z from the data and then derives an allocation of the police force from our causal model.
- ***PropArrest and PropCrime.*** allocates police resources proportionally to past arrest and past concerning crime rates.
- ***PropPolyaUrnAR and PropPolyaUrnCR.*** two baseline allocation policies based on the Polya urn model proposed by Ensign et al. [1].

Empirical Experiment

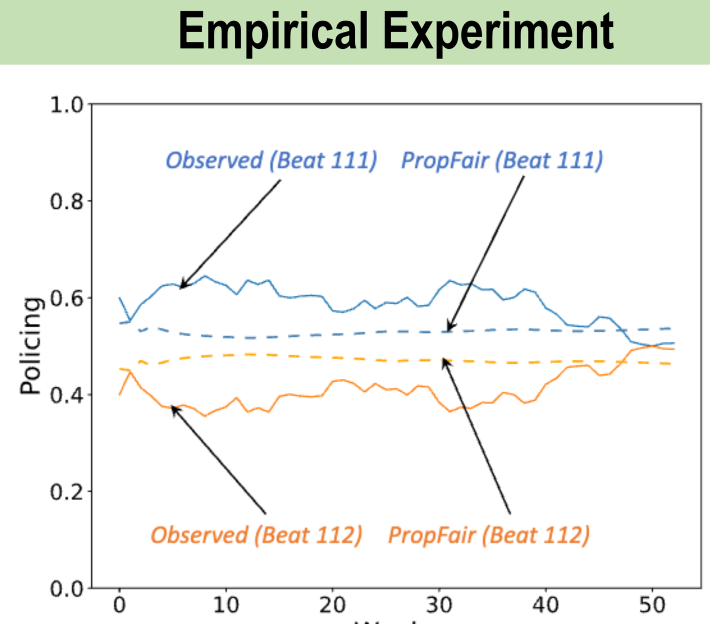

(a) (111,112)

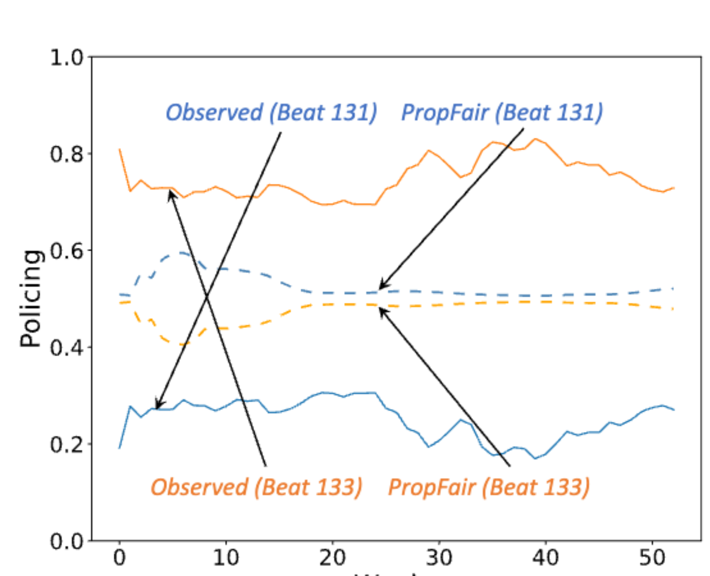

(b) (131,133)

Figure 3. Observed and predicted assignment of police officers as estimated by our method. Police Beats (111,112) and (131,133) are neighboring beats of District 1 of Chicago police department. Y axis shows the proportion of policing in the pair of beats. Beat 133 seems over-policed (observed) and does indeed have a higher fraction of the minority population.

References

1. Ensign. D.. Friedler. S.A.. Neville. S.. Scheidegger, C. and Venkatasubramanian. S.. 2018. Januarv. Runawav feedback loops in predictive policing. In Conterence on Fairness, Accountability and Transparency (pp. 160-171).
2. Kusner, M.J., Lottus, J., Russell, C. and Silva, R., 201/. Countertactual Fairness. In Advances In Neural Information Processing Systems (pp. 4066-4076).

Contributions

- ***A Novel Real-world Dataset:***

  We construct a novel real-world dataset by merging three different sources:
  1. Population demographics from the American Community Survey.
  2. Publicly available data on crimes, arrests, and stops from the website of the Chicago Police Department.
  3. Police deployment levels obtained using a Freedom of Information Act (FOIA) request from the police department.

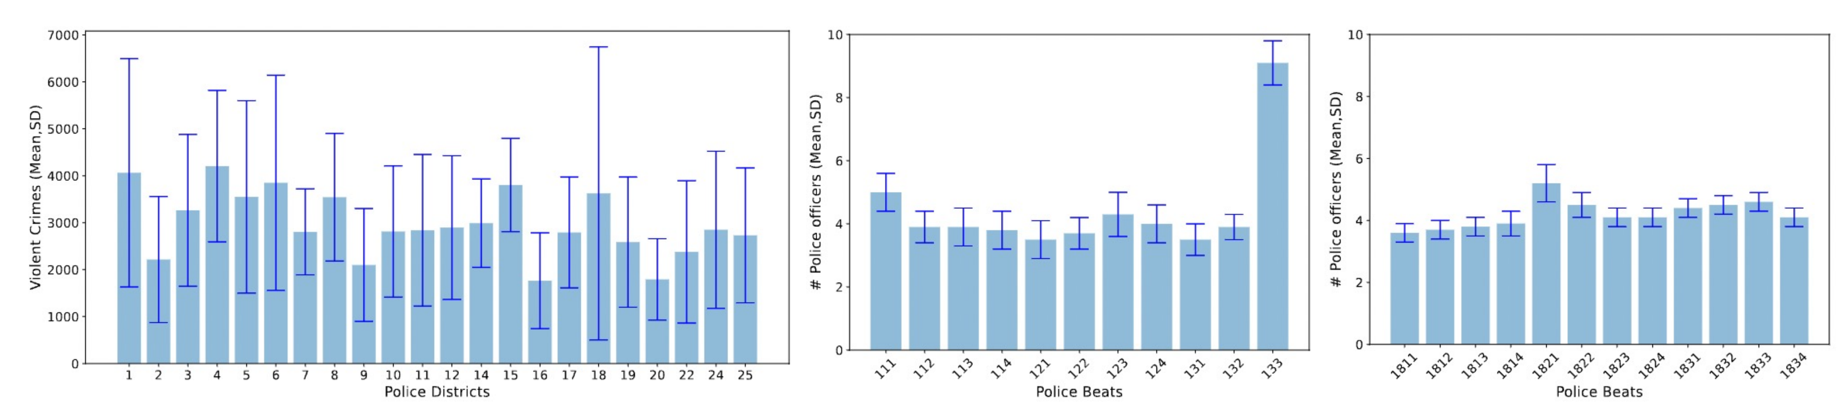

(a) The mean and std of Violent crimes (b) District 1 (c) District 18

Figure 1. (a) The mean and standard deviation of observed Violent crimes in each district. Weekly assignment of police officers in Districts 1 (b) and 18 (c). X-axis represents the police beats and Y-axis represents the number of police officers assigned each week.

- ***Causal Model for Predictive Policing:***

  Our approach to predictive policing is to construct a causal model that could reasonably describe the underlying data generating process and to use this causal model to estimate an optimal policy allocation that is not affected by the biases introduced by the data collection.

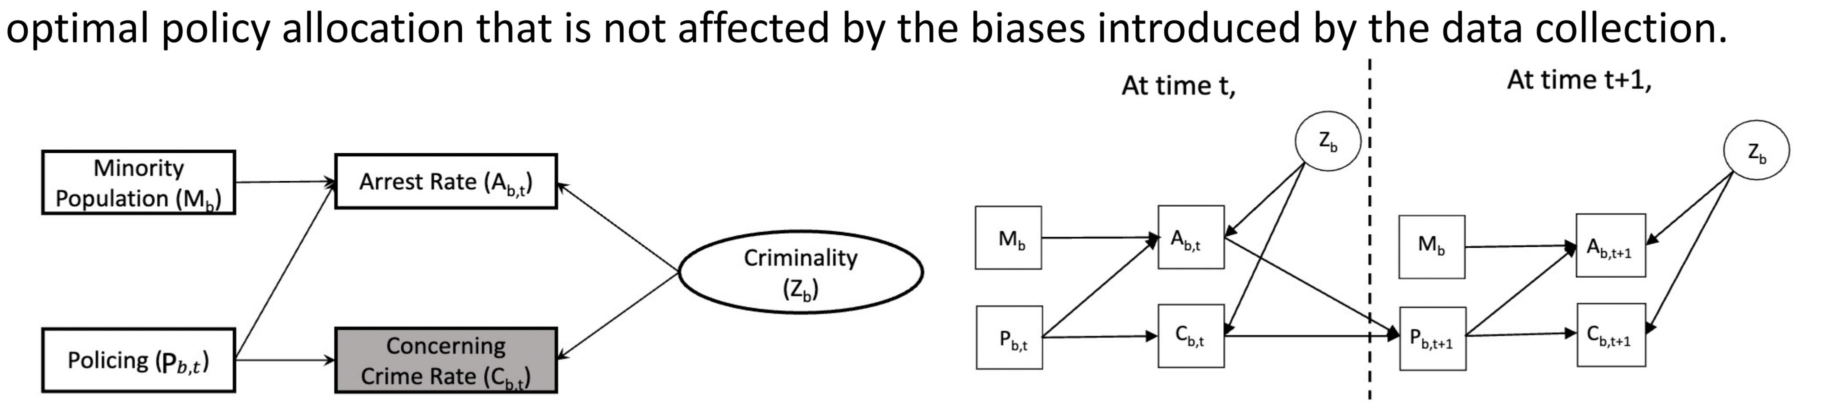

(a) (b)

Figure 2. (a) A causal model for policing and crime. Sensitive attributes: Policing, Minority Population; Observed features: Arrest Rate; Outcome Variable: Concerning Crime Rate; Latent Variable: Criminality. (b) In our causal model, there is no direct arrow between ($A_{b,t}$, $C_{b,t}$) and ($A_{b,t+1}$, $C_{b,t+1}$). We assume that time t variables influence time $t+1$ variables only through the observable $P_{b,t+1}$.

We parameterize the structural causal model:

$$Z_b \sim \mathcal{N}(0,1)$$
$$C_{b,t} = \alpha_d^C + \gamma_t^C + (1 - \beta_d^C \theta(P_{b,t}))Z_b + \epsilon_{b,t}^C$$
$$A_{b,t} = \alpha_d^A + \gamma_t^A + \beta_d^A \theta(P_{b,t})Z_b + \delta_d M_b + \epsilon_{b,t}^A$$

For district d and time t we allocate the total police force $\overline{P_d}$ in *B* beats to minimize the total criminal activities in the district:

$$\min_{P_1,\ldots,P_B} \sum_b^B C_{b,t} \text{ s. t } \sum_b^B P_{b,t} \leq \overline{P_d} \text{ and } \forall b,t, P_{b,t} > 0.$$

Simulated Experiment

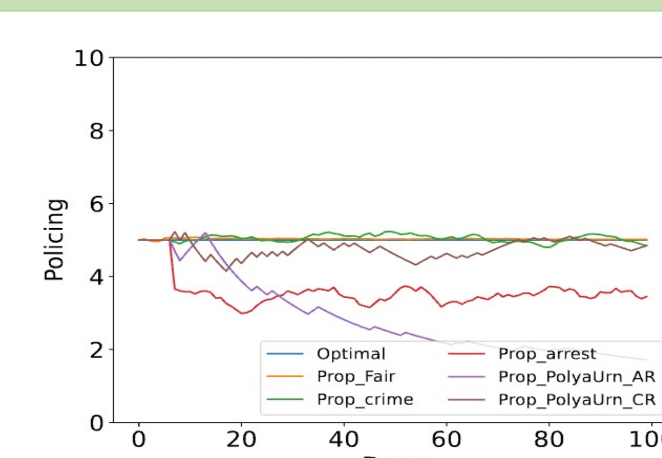

(a) Case 1: Beats differ by percentages of minority population ($Z_{b_1} = 0.5$, $Z_{b_2} = 0.5$, $M_{b_1} = 0.2$, $M_{b_2} = 0.8$).

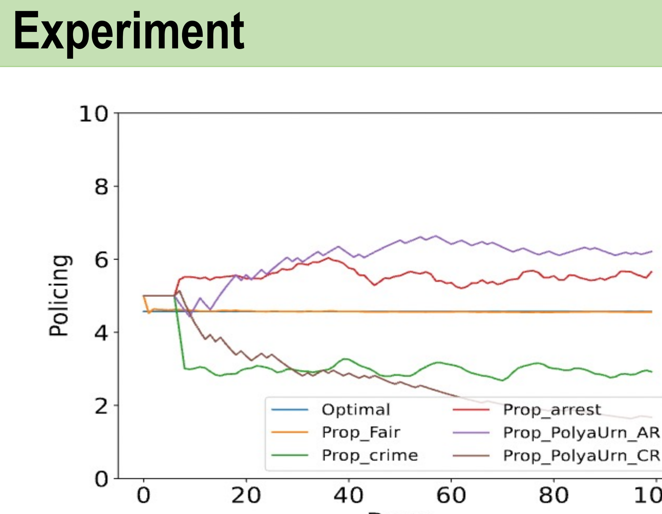

(b) Case 2: Beats differ by level of criminal activity ($Z_{b_1} = 0.3$, $Z_{b_2} = 0.7$, $M_{b_1} = 0.8$, $M_{b_2} = 0.2$).

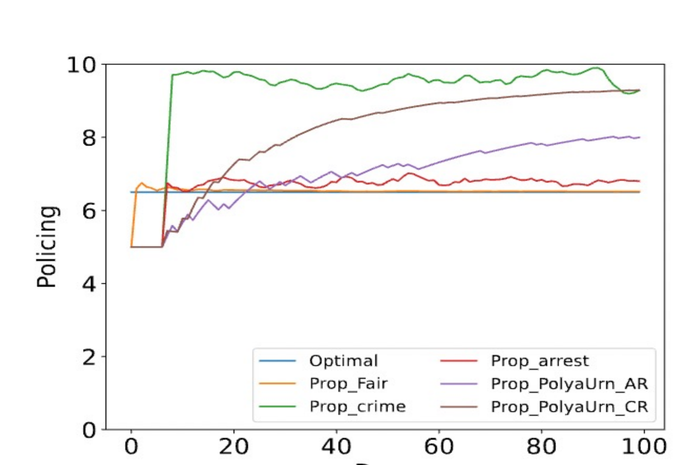

(c) Case 3: Beats differ by level of criminal activity ($Z_{b_1} = 0.95$, $Z_{b_2} = 0.05$, $M_{b_1} = 0.8$, $M_{b_2} = 0.8$).

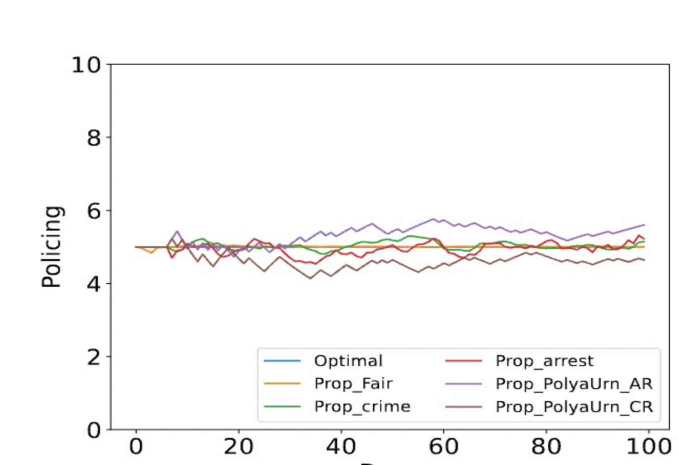

(d) Case 4: Beats are the same ($Z_{b_1} = 0.5$, $Z_{b_2} = 0.5$, $M_{b_1} = 0.25$, $M_{b_2} = 0.25$).

Figure. 4. Comparative analysis of baseline policies (PropArrest, PropCrime, PropPolyaUrnAR, PropPolyaUrnCR), optimal policy (Benevolent planner's policy) and proposed policy (PropFair). The graphs show the fraction of police allocated to Beat B_1 (the remaining fraction is allocated to B_2).