The Effects of Competition and Regulation on Error Inequality in Data-Driven Markets

Hadi Elzayn
University of Pennsylvania
hads@sas.upenn.edu

Ben Fish
Microsoft Research
Benjamin.Fish@microsoft.com

Motivation

Fairness is a major concern in machine learning models. Empirically, minority groups see worse performance in machine learning models

- GenderShades, Predictive Inequity, Speech recognition, Facial recognition [see 1-4]

Data-Driven Markets are increasingly common. We focus on settings where consumers and users *both* prefer more accurate models

- Speech recognition, Search, Matching
- **Not** Loans, Insurance, Predictive policing

In market settings, does monopoly result in unequal error rates? If so, does *competition* mitigate inequality? If not, can we use regulation? Use learning theory + economics to answer these questions

Framework

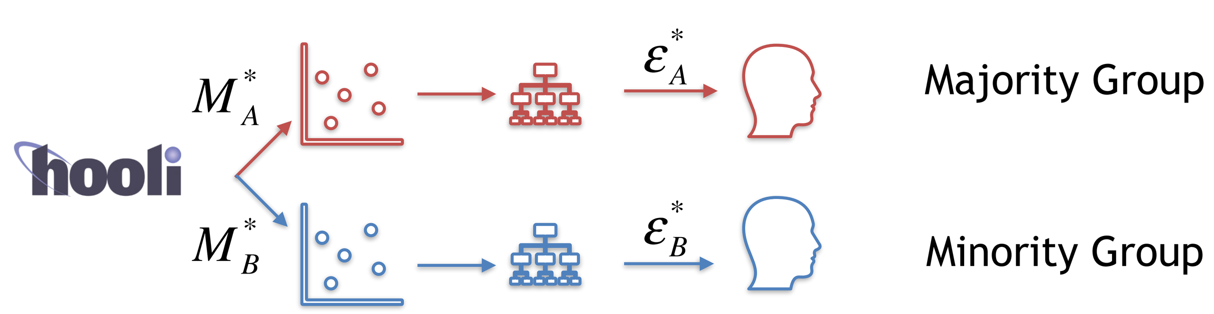

$\mathcal{G}$ user groups. Firm i buys M_{gi} i.i.d. datapoints from $\mathcal{D}_g$, then learn separate models for each group with worst-case excess error rates ε_{gi}. Users choose among firms/outside options depending on ε_{gi} (and demand type). Firms gain revenue from demand weighted by market shares μ_g. Firms optimize profits with respect M_{gi}. *Minority* groups have less market power and higher data cost (see below).

Learning Theory & Cost Structure

$$\varepsilon(h) - \min_{h' \in \mathcal{H}} \varepsilon(h') \leq K \cdot \sqrt[q]{\frac{d_{\mathcal{H}} + \log 1/\delta}{M}}$$

Probably Approximately Correct (PAC) framework - firms choose hypothesis from fixed class; empirical risk minimization produces small error with high probability.

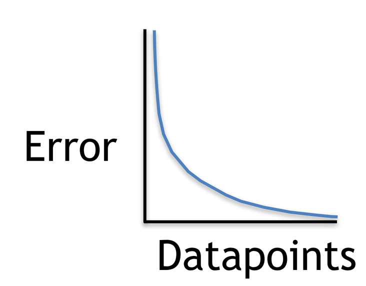

Firms pay fixed cost of Φ_g and γ_g per datapoint. Learning rate q links M_g and ε_g. Assumption: minority group has higher fixed cost ($\Phi_g > \Phi_{g'}$) and/or per-datapoint cost ($\gamma_g > \gamma_{g'}$).

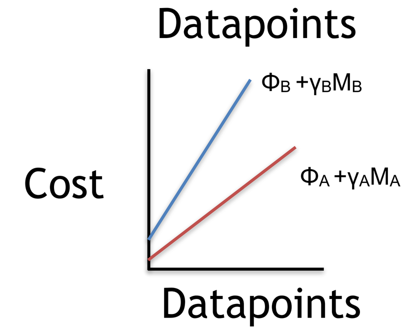

Revenue

Firms' revenue is measured in market share, weighted by market sizes μ_g. We consider various demands, including linear, proportional split, or Bertrand-like.

β_g, ρ_g captures demand *elasticity* or market competitiveness. Assumption: minority group has less elastic ($\beta_g < \beta_{g'}$) or less competitive ($\rho_g < \rho_{g'}$) demand. Also assume that minority market size is smaller, i.e. $\mu_g < \mu_{g'}$.

Demand	Formal Definition
Multilinear	$D_{gi}(\varepsilon_{gi}) = \alpha_g - \beta_g \varepsilon_{gi} + \lambda_{gj}$
Proportional	$D_{gi} = \varepsilon_{gj}^{\wedge}(\rho_g) / (\varepsilon_{gi}^{\wedge}(\rho_g) + \varepsilon_{gj}^{\wedge}(\rho_g))$
Bertrand-like	$D_{gi} = 1$ iff $\varepsilon_{gi} = \min(\varepsilon_{gj}, \varepsilon_{gj})$ up to some tolerance ζ_g

Linear Demand

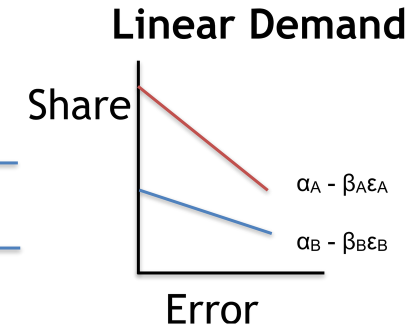

Monopolist

Consider a monopolist facing linear demand (Maj. Group A, Red; Min. Group B, Blue). The profit-maximizing first-order conditions with respect to data investment requires marginal revenue equal marginal cost.

$$\pi_g(M_g) = \mu_g D_g(M_g) - C(M_g)$$

FOC: $$\mu_g \frac{\partial D_g}{\partial M_g} = \frac{\partial C_g}{\partial M_g}$$

Theorem in a picture: Optimal choice of data investment for Minority group (M_B^*) is smaller than for Majority group (M_A^*). Since error is linked via:

$$\varepsilon_g = \mathcal{O}(M_g^{1/q})$$

This implies that $\varepsilon_A^* < \varepsilon_B^*$.

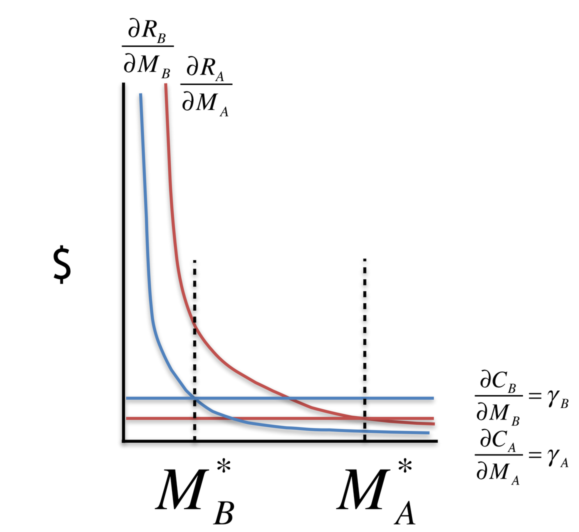

THEOREM 1 (MONOPOLY INEQUALITY). *Suppose a monopolist with learning rate q faces linear demand. Then in any interior optimum, for every pair of groups g and g′, the error inequality is given by:*

$$\frac{\varepsilon_g^*}{\varepsilon_{g'}^*} = \left(\frac{\mu_{g'} \beta_{g'} \gamma_g}{\mu_g \beta_g \gamma_{g'}} \right)^{\frac{1}{q+1}}$$

Competition

Can competition mitigate this inequality, relative to monopoly setting? We look for Nash equilibria under several models of competition.

Answer: No, under all but the most extreme models of competition.

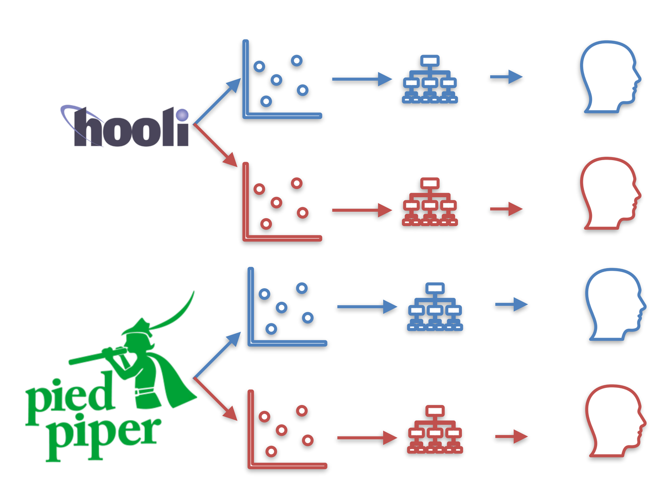

Demand	Improvement?	Weak Assumptions?
Multilinear	✗	✓
Proportional	✗	✓
Bertrand-like	✓	✗

THEOREM 3 (INEQUALITY UNDER PROPORTIONAL DEMAND). Suppose two firms with learning rate q compete under proportional demand. Then in any interior equilibrium, error inequality is given by:

$$\frac{\varepsilon_{gi}^*}{\varepsilon_{g'i}^*} = \left(\frac{\rho_{g'} \mu_{g'}}{\rho_g \mu_g} \right)^{\frac{1}{q}} \frac{f(\gamma_{gi}, \gamma_{gj}, q)}{f(\gamma_{g'i}, \gamma g' j, q)}$$

Regulation

Consider a regulator with ability to constrain the monopolist either to set approximately equal error rates across groups (1+χ ratio) or absolute maximum error χ for each group.

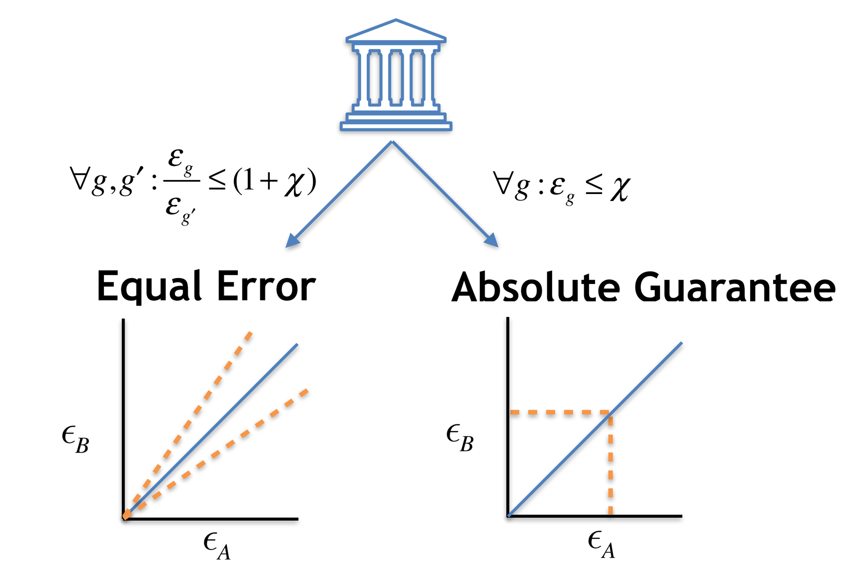

Either type of regulation *can* improve the situation, with different tradeoffs. *Equal error* imposes a *Price of Fairness* (regulated ε_A > monopolist's choice of ε_A); absolute guarantees can only be imposed up to a point. Both types of regulation will affect the monopolist's profit.

	Equal Error	Absolute Guarantee
Majority Group Price of Fairness (POF)	$\mathrm{PoF}_{1+\chi} \leq \left(1 + \frac{\gamma_B}{\gamma_A} \frac{1}{1+\chi}\right)^{\frac{1}{q+1}}$	None
Monopolist Profit Price of Fairness (MonPOF)	$\lim_{\mu_B \to \infty} \mathrm{MonPoF}_{1+\chi} = 1$ $\lim_{\mu_B \to 0} \mathrm{MonPoF}_{1+\chi} = K > 1$	$\lim_{\mu_B \to \infty} \mathrm{MonPoF}_{1+\chi} = 1$ $\lim_{\mu_B \to 0} \mathrm{MonPoF}_{1+\chi} = K > 1$
Limitations?	None	Smallest χ solves: $\chi^{q+1} + B\chi^q - C = 0$

Implications

Economic incentives create error inequality in data-driven markets
Adding **competition does not mitigate** this incentive under reasonable models of competition
Regulators can improve minority welfare by adding constraints
Not all regulation is created equal – equal error guarantees impose a **price of fairness** on the majority, while absolute error guarantees do not. On the other hand, absolute guarantees do not necessarily prevent relative error inequality across groups
Firms also pay a quantifiable **price of fairness** that disappears as the minority group becomes large in absolute size even if still relatively small
Society must choose when, where, and how to regulate error inequality; **should not expect the market to take care of itself**
These results are relatively **robust** to reasonable modeling choices, but do not capture all possible scenarios.
Plenty of room for future work!

References

[1] Buolamwini, Joy, and Timnit Gebru. "Gender shades: Intersectional accuracy disparities in commercial gender classification." *Conference on fairness, accountability and transparency*. 2018.
[2] Wilson, Benjamin, Judy Hoffman, and Jamie Morgenstern. "Predictive inequity in object detection." *arXiv preprint arXiv:1902.11097* (2019).
[3] Blodgett, Su Lin, and Brendan O'Connor. "Racial disparity in natural language processing: A case study of social media african-american english." *arXiv preprint arXiv: 1707.00061* (2017).
[4] Obermeyer, Ziad, and Sendhil Mullainathan. "Dissecting Racial Bias in an Algorithm that Guides Health Decisions for 70 Million People." *Proceedings of the Conference on Fairness, Accountability, and Transparency*. ACM, 2019.
[5] Kearns, Michael J., Umesh Virkumar Vazirani, and Umesh Vazirani. *An introduction to computational learning theory*. MIT press, 1994.
[7] Tirole, Jean. *The theory of industrial organization*. MIT press, 1988.