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

Hadi Elzayn[*]
University of Pennsylvania
hads@sas.upenn.edu

Benjamin Fish
Microsoft Research
benjamin.fish@microsoft.com

Abstract

Recent work has documented instances of *unfairness* in deployed machine learning models, and significant researcher effort has been dedicated to creating algorithms that intrinsically consider fairness. In this work, we highlight another source of unfairness: market forces that drive differential investment in the data pipeline for different groups. We develop a high-level framework, based on insights from learning theory and industrial organization, to study this phenomenon. First, we show that our model predicts unfairness in a monopoly setting. Then, we show that under all but the most extreme models, competition does not eliminate this tendency, and may even exacerbate it. Finally, we consider two avenues for regulating a machine-learning driven monopolist - relative error inequality and absolute error-bounds - and quantify the price of fairness. Our results suggest that mitigating fairness concerns may require *policy*-driven solutions, not only technological ones.

1 Introduction

As machine learning has become more integrated into products, markets, and decision-making throughout society, researchers, practitioners, and activists have identified many instances of *unfairness* in predictions or decisions made or influenced by machine-learned models [5, 9, 18, 19, 22], including in extremely high stakes settings like object recognition systems for autonomous vehicles [22] or healthcare treatment assignment [19]. With the hope of mitigating unfairness, researchers have engaged in empirical and theoretical investigations to understand the reasons behind unfairness including historical bias [17] and selection bias [14] in training data; feedback loops [17]; sample size disparity [7]; the inability to fit group-specific models for legal, ethical, or practical reasons [15]; and use of the wrong loss function [19]. Researchers have developed many innovative technical solutions to these problems (see, e.g. [1, 3, 6, 8, 10, 11, 23]), yet unfairness in practice persists.

We highlight a simple but important point: while technical solutions to unfairness are certainly necessary, mitigating unfairness in practice may require tackling the *economic* incentives driving unfairness. In this paper, we are interested in data-driven markets with disjoint consumer groups and the error inequality between these groups. Our framework is built on learning theory and industrial organization; while stylized, it captures salient market features and elides previously described sources of unfairness. We first show that a monopolist will invest less in data collection (and thus model accuracy) for minority groups, where minority status is defined by a group's market power and cost of data collection. We then show that this phenomenon is not mitigated by introducing a competing firm. Finally, we analyze the effect of regulatory constraints on the monopolist. We conclude by discussing the policy implications of these results.

[*]Work completed as part of an internship at Microsoft Research.

AI for Social Good workshop at NeurIPS (2019), Vancouver, Canada.

2 Modeling Framework

At a high level, our models consist of firms and consumers. Firms choose how much data to purchase (modeled as a continuous quantity for simplicity); they use this data to create a machine learning model, which is then used to serve consumers. Consumers are split into non-overlapping groups representing their market segment; they choose which firm, if any, to use based on the performance of the models for their own group. Because users do not know the realized error of firms' models, they use the *worst-case* excess error guarantees implied by the firms' data choices as their only proxy for performance. Firms' profits are their market share less their costs invested in data collection. We assume that firms have unlimited budget and that the data sources are group-specific and potentially infinite; we also assume that firms make no distributional similarity assumptions across groups, and so must train models separately rather than engage in transfer learning. These assumptions, taken together, eliminate many causes of unfairness previously highlighted in the literature.

Using $-i$ to denote Firm i's competitors (if any), we can write profit maximization problem as:

$$\max_{M_i} \pi_i(M_i, M_{-i}) = \max_{M_i} \sum_g \mu_g D_{gi}(\epsilon_{gi}(M_{gi}), \epsilon_{g,-i}(M_{g,-i})) - C(M_i), \tag{1}$$

potentially subject to constraints imposed by regulation. Here, M_i is the vector of choices of data for each group; $\epsilon_{gi}(M_{gi})$ is the implied worst-case error guarantee for group g under Firm i's model; μ_g captures the total market size of Group g; D_{gi} is the share of group g that Firm i captures; and $C(M_i)$ is the total cost Firm i must pay for its data choices. If the firm has no competitors, profit maximization is a simple optimization problem; otherwise, the firm must take into account competitors' actions. In the multi-firm case, we look for *Nash equilibria* (see e.g. [12] for more on game theory). In either case, D_{gi} will play a role in the form of the optimal solutions, but our results are qualitatively similar for all but the most unrealistic models of demand we study.

For simplicity, we model the link between data purchase and error rates using the PAC model of learning [16]. Firms can sample (potentially infinitely many) independently drawn feature-label pairs from fixed distributions $\mathcal{D}_g$; we assume these $\mathcal{D}_g$ are the relevant distributions consumers from each group care about. Firms have access to a hypothesis class $\mathcal{H}$; each hypothesis has its own *risk* $R(h) = E_{x,y\sim\mathcal{D}_g}[h(x) \neq y]$. We allow both the *realizable* setting, when there is some $h \in \mathcal{H}$ with risk zero, as well as the *agnostic* setting; PAC-learning theory tells us that in the agnostic setting, if a hypothesis class has finite *Vapnis-Chervonenkis dimension* (VC), then with probability $1 - \delta$, empirical risk minimization will output a hypothesis h such that:

$$R(h) - \min_{h' \in \mathcal{H}} R(h') \leq K\sqrt{\frac{d_{\mathcal{H}} + \log(1/\delta)}{M}},$$

where d_H is the VC dimension of $\mathcal{H}$ and K is a universal constant. (See [20] for more on empirical risk minimization, VC dimension, and the various kinds of PAC learning.) A similar statement holds without the square root in the realizable setting, and other settings can be modeled with different learning rates. Moreover, these bounds are tight in the worst case. Thus, we can link firms' group-specific worst-case excess error ϵ_{gi} to purchased data points as $\epsilon_{gi} = C_{gi}/M_{gi}^{1/q}$, where q captures the learning rate and C_{gi} is a group-specific constant. Then, assuming that firms pay some fixed startup cost for each group ϕ_{gi} and linear (possibly differing marginal costs) c_{gi} per datapoint, we can write the write firms' costs as a function of error: $C(\epsilon_{gi}) = \gamma_{gi}/\epsilon_{gi}^q$, where γ_{gi} captures the constants.

3 Monopoly

We start with the case where there is one firm in the market and demand is linear. That is:

$$D_g(\epsilon_g) = \alpha_g - \beta_g \epsilon_g,$$

where $0 < \beta \leq \alpha \leq 1$. We can also consider *piecewise linear* demand (capped at 1 and floored at 0) that does not have such restrictions; this does not qualitatively affect the results, so we omit it here. Now, under linear demand, we can write the firm's problem as

$$\max_{\epsilon} \left[\sum_{g \in \mathcal{G}} \mu_g \left(\alpha_g - \beta_g \epsilon_g\right) - \sum_g \left(\phi_g + \gamma_g/\epsilon_g^q\right) \right]. \tag{2}$$

Simple calculus and the fact that profit is concave yield Lemma 1. We defer proof of this and all other omitted proofs to the appendix.

Lemma 1. *At any* interior *optimal solution ϵ^M, linear demand and learning rate q imply that $\forall g \in \mathcal{G}$:*

$$\epsilon_g^M = \left(\frac{q\gamma_g}{\mu_g \beta_g}\right)^{\frac{1}{q+1}}. \tag{3}$$

In this paper, we will only be concerned with *interior* optimal solutions and equilibria, since solutions in which firms choose to decline from investing entirely in a given group will obviously engender inequality. Equation 3 intuitively says that a group's worst-case excess error rate will vary proportionally with the cost of acquiring data for that group and *inversely* with its market power (that is, the product of $\mu_g \beta_g$). Simply taking the ratio of errors leads us to our first inequality result, Theorem 1:

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:*

$$\frac{\epsilon_g^M}{\epsilon_{g'}^M} = \left(\frac{\mu_{g'} \beta_{g'} \gamma_g}{\mu_g \beta_g \gamma_{g'}}\right)^{\frac{1}{q+1}}.$$

Theorem 1 demonstrates that under a monopoly, groups that have relatively smaller market power and higher cost of data acquisition will suffer relatively greater error rates. Of course, monopoly favors the monopolist at the expense of consumers, and the canonical remedy is more competition. Can competition also mitigate between-group error inequality? The next section explores this possiblity.

4 Competition

Suppose instead that two firms compete in the market. Several natural models of demand, along a spectrum of rationality, could capture consumer behavior. Here we consider a *proportional split* split demand, similar to the well-known Tullock contest [21]. We discuss other models in the appendix.

With two firms, proportionally split demand with competition exponent ρ_g in each group is given by:

$$D_{gi}(\epsilon_{gi}, \epsilon_{gj}) = 1 - \frac{\epsilon_{gi}^{\rho_g}}{\epsilon_{gi}^{\rho_g} + \epsilon_{gj}^{\rho_g}} = \frac{\epsilon_{gj}^{\rho_g}}{\epsilon_{gi}^{\rho_g} + \epsilon_{gj}^{\rho_g}},$$

and the firm's problem becomes

$$\max_{\epsilon_i} \left[\sum_{g \in \mathcal{G}} \mu_g \frac{\epsilon_{gj}^*}{\epsilon_{gi} + \epsilon_{gj}^*} - \sum_{g \in \mathcal{G}} \left(\phi_{gi} + \frac{\gamma_{gi}}{\epsilon_{gi}^q}\right)\right], \tag{4}$$

where ϵ_{gj}^* is its opponent's (equilibrium) choice of error for group g. Lemma 2 describes equilibrium.

Lemma 2. *Suppose two firms compete for proportional demand with parameters q and ρ. There exists a range of market and cost structure parameters for which $(\epsilon_{gi}^*, \epsilon_{gj}^*)$ is an equilibrium, where*

$$\epsilon_{gi}^* = \left(\frac{q\gamma_{gi}}{\rho_g \mu_g}\right)^{\frac{1}{q}} \frac{(\gamma_{gi}^q + \gamma_{gj}^q)^{\frac{2}{q}}}{\gamma_{gi}^q \gamma_{gj}^q} = \left(\frac{q}{\rho_g \mu_g}\right)^{\frac{1}{q}} \frac{(\gamma_{gi}^q + \gamma_{gj}^q)^{\frac{2}{q}}}{\gamma_{gi}^{1-\frac{1}{q}} \gamma_{gj}^q}.$$

For a (nonempty) subset of this parameter range, this equilibrium is unique.

The parameter range referenced of Lemma 2 ensures that it is not a best response for firms to choose to invest no data collection at all in any group; outside this range, there are equilibria in which firms do not invest in data collection for some groups, which will clearly lead to inequality. For brevity, we do not precisely specify these conditions, but they simply require that firms' cost structures are not *too* different and that the fixed costs are not *too* large relative to the total possible gain.

Theorem 2 (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:*

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

where $F(\gamma_{gi}, \gamma_{gj}, \gamma_{g'i}, \gamma_{g'j}, q) = \frac{(\gamma_{gi}^q + \gamma_{gj}^q)^{\frac{2}{q}}}{\gamma_{gi}^{1-\frac{1}{q}} \gamma_{g'j}^q} \Big/ \frac{(\gamma_{g'i}^q + \gamma_{g'j}^q)^{\frac{2}{q}}}{\gamma_{g'i}^{1-\frac{1}{q}} \gamma_{g'j}^q}$.

Theorem 2 may look complicated, but the qualitative interpretation is simple: like in Theorem 1, the ratio of error rates between groups is inversely related to the ratio between their market powers (where the competition exponent in each market plays the previous role played by elasticity). In fact, the dependence under this model is worse (since the exponent is larger).

5 Regulation

In this section, we consider regulating a monopolist with the intent of improving the outcome for a minority. For clarity of exposition, we assume that there are two groups, A and B, with $\mu_A\beta_A \geq \mu_B\beta_B$ and $\gamma_A \leq \gamma_B$. We consider two types of regulation: the first is the constraint that $\forall\, g, g', \epsilon_g/\epsilon_{g'} \leq (1+\chi)$, which we call the *approximately equal error* constraint; the second requires that $\forall g, \epsilon_g \leq \chi$, which we call the *absolute error guarantees*. The firm's problem then becomes that in Equation 2 but subject to the imposed constraints. We begin with the following useful Lemma:

Lemma 3 (Saturation)**.** *If the regulator applies equal or absolute error constraints, and the error constraints change the firm's behavior, then the constraints will be saturated for one or both groups.*

Formally, Lemma 3 says that in the equal error rate case, the monopolist will set $\epsilon_B^R = (1+\chi)\epsilon_A^R$, and under absolute error guarantees, will set ϵ_B^R to χ (and possibly ϵ_A^R to χ as well). We can then substitute these into the firm's problem to find the new optimal solutions ϵ_A^R and ϵ_B^R, and obtain:

Theorem 3 (Price of Fairness)**.** *Under both type of regulations, the error rate of the minority group improves. That is, $\epsilon_B^R \leq \epsilon_B^M$. Under equal error constraints, $\epsilon_A^R \geq \epsilon_A^M$. But, for a fixed ratio $\frac{\mu_B}{\mu_A}$, a simple upper bound on $\epsilon_A^R/\epsilon_A^M$ is:*

$$\frac{\epsilon_A^R}{\epsilon_A^M} \leq \left(1 + \frac{\gamma_B}{\gamma_A}\frac{1}{1+\chi}\right)^{\frac{1}{q+1}}$$

Under absolute error guarantees, by contrast, $\epsilon_A^R \leq \epsilon_A^M$ – the minority pays no 'price of fairness'.

Society may wish to maintain some profits as an incentive to innovate; Theorem 4 shows the limiting behavior of the ratio of unconstrained (π^M) and regulated π^R monopoly profits. It also says that there is a minimum error value the regulator can guarantee (without making the firm unprofitable).

Theorem 4 (Firm Price of Fairness)**.** *Under both types of regulation, the monopolist pays a price of fairness via lost profit; however, for any constant r and regulation-specific constants C, we can write:*

$$\lim_{\substack{\mu_B,\mu_A\to\infty \\ \mu_A = r\mu_B}} \frac{\pi^M}{\pi^R} = 1 \qquad\qquad \lim_{\mu_B\to 0} \frac{\pi^M}{\pi^R} = C > 1$$

Under absolute error guarantees, there exists some minimum error ϵ_0 the monopolist can guarantee.

Taken together, Theorems 3 and 4 elucidate the (quantifiable) tradeoffs that may influence whether and how society chooses to regulate for fairness. We highlight that that while these quantities are important inputs to this choice, they may ultimately be trumped by ethical, practical, or other considerations.

6 Discussion

In this work, we identify economic incentives leading to unfairness in data-driven markets. At a high level, we show that monopolists are incentivized to invest less in minority groups (as measured by market size, elasticity, and data costs) because they are less profitable; that competition does not mitigate this incentive towards inequality; and that judicious regulation *can* improve outcomes, potentially at a cost in terms of profits or, depending on the regulation, error rates for the majority.

We view this paper as highlighting an important and understudied point of view, but certainly not as the last word. We made many choices that situate our models in particular contexts; for example, the assumption that firms and users benefit from improved accuracy does not capture many settings that currently are or will soon be urgent domains of adjudicating fairness concerns - machine learning in loans, insurance, and facial recognition systems are obvious cases, but the potential landscape, and consequent scope for unfairness, is vast. We hope that future work will explore other settings and further clarify the possibility - and perhaps necessity- of leveraging policy tools in addition to algorithmic solutions in order to combat unfairness in machine learning.

References

[1] Alekh Agarwal et al. "A Reductions Approach to Fair Classification". In: (2017).

[2] Rajiv D. Banker, Inder Khosla, and Kingshuk K. Sinha. "Quality and Competition". In: *Manage. Sci.* 44.9 (Sept. 1998), pp. 1179–1192. ISSN: 0025-1909. DOI: 10.1287/mnsc.44.9.1179.

[3] R. Berk et al. "A Convex Framework for Fair Regression". In: *FAT-ML Workshop* (2017).

[4] Steven T. Berry. "Estimating Discrete-Choice Models of Product Differentiation". In: *The RAND Journal of Economics* 25.2 (1994), pp. 242–262. ISSN: 07416261.

[5] Su Lin Blodgett and Brendan O'Connor. "Racial disparity in natural language processing: A case study of social media african-american english". In: *arXiv preprint arXiv:1707.00061* (2017).

[6] Flavio Calmon et al. "Optimized pre-processing for discrimination prevention". In: *Advances in Neural Information Processing Systems*. 2017, pp. 3992–4001.

[7] Irene Chen, Fredrik D Johansson, and David Sontag. "Why is my classifier discriminatory?" In: *Advances in Neural Information Processing Systems*. 2018, pp. 3539–3550.

[8] Sam Corbett-Davies et al. "Algorithmic decision making and the cost of fairness". In: *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM. 2017, pp. 797–806.

[9] Michael D Ekstrand et al. "All the cool kids, how do they fit in?: Popularity and demographic biases in recommender evaluation and effectiveness". In: *Conference on Fairness, Accountability and Transparency*. 2018, pp. 172–186.

[10] Hadi Elzayn et al. "Fair algorithms for learning in allocation problems". In: *Proceedings of the Conference on Fairness, Accountability, and Transparency*. ACM. 2019, pp. 170–179.

[11] Danielle Ensign et al. "Runaway feedback loops in predictive policing". In: *arXiv preprint arXiv:1706.09847* (2017).

[12] D. Fudenberg et al. *Game Theory*. Mit Press. MIT Press, 1991. ISBN: 9780262061414.

[13] Kenneth L Judd. "Cournot versus Bertrand: A dynamic resolution". In: (1996).

[14] Nathan Kallus and Angela Zhou. "Residual unfairness in fair machine learning from prejudiced data". In: *arXiv preprint arXiv:1806.02887* (2018).

[15] M. Kearns and A. Roth. *The Ethical Algorithm: The Science of Socially Aware Algorithm Design*. Oxford University Press, Incorporated, 2019. ISBN: 9780190948207.

[16] M.J. Kearns et al. *An Introduction to Computational Learning Theory*. Mit Press. MIT Press, 1994. ISBN: 9780262111935.

[17] Kristian Lum and William Isaac. "To predict and serve?" In: *Significance* 13.5 (2016), pp. 14–19.

[18] Rishabh Mehrotra et al. "Auditing Search Engines for Differential Satisfaction Across Demographics". In: *Proceedings of the 26th International Conference on World Wide Web Companion*. WWW '17 Companion. Perth, Australia: International World Wide Web Conferences Steering Committee, 2017, pp. 626–633. ISBN: 978-1-4503-4914-7. DOI: 10.1145/3041021.3054197.

[19] Ziad Obermeyer and Sendhil Mullainathan. "Dissecting Racial Bias in an Algorithm that Guides Health Decisions for 70 Million People". In: *Proceedings of the Conference on Fairness, Accountability, and Transparency*. ACM. 2019, pp. 89–89.

[20] S. Shalev-Shwartz and S. Ben-David. *Understanding Machine Learning: From Theory to Algorithms*. Understanding Machine Learning: From Theory to Algorithms. Cambridge University Press, 2014. ISBN: 9781107057135.

[21] Gordon Tullock. "Efficient rent seeking". In: *Efficient Rent-Seeking*. Springer, 2001, pp. 3–16.

[22] Benjamin Wilson, Judy Hoffman, and Jamie Morgenstern. "Predictive Inequity in Object Detection". In: *CoRR* abs/1902.11097 (2019). arXiv: 1902.11097.

[23] Blake Woodworth et al. "Learning non-discriminatory predictors". In: *arXiv preprint arXiv:1702.06081* (2017).

A Appendix

A.1 Additional Models of Competition

We briefly discuss two other models of competition. We call the first *multilinear* demand; it is a simple generalization of linear demand in which a firm's demand diminishes with its own error rate and rises with the error rate of its competitor. Formally, we write $D_{gi}(\epsilon_{gi}, \epsilon_{gj}) = \alpha_{gi} - \beta_{gi}\epsilon_{gi} + \lambda_g \epsilon_{gj}$,where $0 < \lambda_g < \beta_{gi} \leq \alpha_{gi}$ and $\alpha_{gi} + \lambda_g \leq 1$. This model is analogous to multilinear demand in price as in [2]. Conceptually, this model makes sense if consumers as a whole view firms' products as *imperfect substitutes* [4]; for example, they may have some brand loyalty, so that some consumers will use the firm with higher error (up to a point). Under this model, firms have dominant strategies (that is, their optimal strategies don't actually depend on what the other firm does), that work out to acting as if they are 'mini-monopolists' of their own market segments. In concordance with this formal similarity, firms make the same error and inequality choices as a monopolist facing the same linear demand - assuming, of course, that both firms make positive profit. If this assumption doesn't hold, then one or both firm may not invest in some group at all, but again, that trivially will engender unfairness.

A *different* conclusion will be reached if we assume that the firm with lower error captures *all* the market. This is analogous to *Bertrand* price competition [13] and, like in Bertrand competition, there can be no equilibrium, as there is alway an incentive to slightly under cut the other firm. We can remedy this by assuming some ζ_g *tolerance* - when error levels drop below ζ_g, consumers split their allegiance as if the firms produce the same error. In this case, there is a unique interior equilibrium in which both firms set ϵ_g to ζ_g; then inequality will be given by $\frac{\epsilon_g}{\epsilon_{g'}} = \frac{\zeta_g}{\zeta_{g'}}$. That is, in this model, market power no longer affects inequality! However, this model, like the Bertrand model that it inherits its properties from, should be considered extremely unrealistic in practice; it requires that all consumers can perfectly distinguish, and uniquely care about, arbitrarily small error differences (above ζ_g). Moreover, the conclusions of the Bertrand model - duopolies pricing as if they were price-takers in a perfectly competitive market - clashes with observations of real-world duopolies; the analogous conclusion of the Bertrand-like model in data-driven markets is similarly dubious.

A.2 Deferred Proofs

Proof of Lemma 1. Recall that

$$\pi(\epsilon) = \sum_{g \in \mathcal{G}} \mu_g \left(\alpha_g - \beta_g \epsilon_g\right) - \sum_{g \in \mathcal{G}} \left(\phi_g + \frac{\gamma_g}{\epsilon_g^q}\right).$$

Now, we notice that this profit function is separable into the sum of profits from each market. Differentiating with respect to ϵ_g separately and setting to zero,we arrive at the first-order conditions:

$$\frac{\partial \pi}{\partial \epsilon_g} = -\mu_g \beta_g + \frac{q\gamma_g}{\epsilon_g^{q+1}} = 0.$$

Solving this equation yields ϵ_g^*. Profit is concave on its domain since $\frac{\partial}{\partial \epsilon_g^2}\pi = -(q)(q-1)\gamma_g \epsilon_g^{-q-2} < 0$, hence ϵ_g^* is a maximizer of π_g. If the profit at ϵ_g^* is larger than $\epsilon_g^* = 1$, then ϵ_g^* will be the global optimum, as $\lim_{\epsilon_g \to 0} \pi = -\infty$.

□

Notice that if, for all g, $\pi_g\left(\epsilon_g^*\right) > 0$, $\pi_g\left(\epsilon_g^*\right) > \pi_g(1)$, and $\epsilon_g^* < 1$, then the interior optimum exists and is unique.

Proof of Lemma 2. Under the proportional split model of demand, each firm's profit depends not only on its own action, but also that of the other firm. Again, this calls for a game theoretic notion of solution. We look for a pure strategy *Nash Equilibrium*. Recall that in an equilibrium, both firms must be best-responding and have no incentive to deviate. That is, both strategies must be *best-responses to each other*.

To find an equilibrium, we first find the best-response of Firm i, *given* the choices of Firm j. Fixing ϵ_j, the profit of Firm i given the choice of ϵ is as follows:

$$\pi(\epsilon_i, \epsilon_j) = \sum_{g \in \mathcal{G}} \left[\mu_g \frac{\epsilon_{gj}^{\rho_g}}{\epsilon_{gi}^{\rho_g} + \epsilon_{gj}^{\rho_g}}\right] - \sum_{g \in \mathcal{G}} \phi_{gi} + \frac{\gamma_{gi}}{\epsilon_{gi}^q}$$

Taking the derivative:

$$\frac{\partial \pi}{\partial \epsilon_{gi}} = -\mu_g \epsilon_{gj}^{\rho_g} (\epsilon_{gi}^{\rho_g} + \epsilon_{gj}^{\rho_g})^{-2} \left(\rho_g \epsilon_{gi}^{\rho_g - 1}\right) + \frac{q\gamma_{gi}}{\epsilon_{gi}^{q+1}}$$

Setting to zero yields the first-order condition:

$$\frac{q\gamma_{gi}}{\epsilon_{gi}^{q+1}} = \frac{\rho_g\mu_g\epsilon_{gi}^{\rho_g-1}\epsilon_{gj}^{\rho_g}}{(\epsilon_{gi}^{\rho_g}+\epsilon_{gj}^{\rho_g})^2} \implies \frac{\rho_g\mu_g\epsilon_{gi}^{\rho_g+q}\epsilon_{gj}^{\rho_g}}{q\gamma_{gi}} = (\epsilon_{gi}^{\rho_g}+\epsilon_{gj}^{\rho_g})^2$$

Symmetric logic will apply to Firm j. Bu then using the fact that the first order condition must hold simultaneously for both firms, we have that in equilibrium:

$$\frac{\rho_g\mu_g\epsilon_{gi}^{\rho_g+q}\epsilon_{gj}^{\rho_g}}{q\gamma_{gi}} = \frac{\rho_g\mu_g\epsilon_{gj}^{\rho_g+q}\epsilon_{gi}^{\rho_g}}{q\gamma_{gj}} \implies \epsilon_{gj} = \epsilon_{gi}\left(\frac{\gamma_{gj}}{\gamma_{gi}}\right)^{\frac{1}{q}}$$

That is, equilibrium requires a specific relationship between firms' choices. Substituting this back in, we have:

$$\frac{\rho_g\mu_g\epsilon_{gi}^{\rho_g+q}\epsilon_{gi}^{\rho_g}\left(\frac{\gamma_{gj}}{\gamma_{gi}}\right)^{\frac{\rho_g}{q}}}{q\gamma_{gi}} = \left(\epsilon_{gi}^{\rho_g}+\epsilon_{gi}^{\rho_g}\left(\frac{\gamma_{gj}}{\gamma_{gi}}\right)^{\frac{\rho_g}{q}}\right)^2 = \epsilon_{gi}^{2\rho}\left(1+\left(\frac{\gamma_{gj}}{\gamma_{gi}}\right)^{\frac{\rho_g}{q}}\right)^2$$

Solving gives that

$$\epsilon_{gi} = \left[\frac{q\gamma_{gi}}{\rho_g\mu_g}\left(\frac{\gamma_{gi}}{\gamma_{gj}}\right)^{\frac{\rho_g}{q}}\left(1+\left(\frac{\gamma_{gj}}{\gamma_{gi}}\right)^{\frac{\rho_g}{q}}\right)^2\right]^{\frac{1}{q}} = \left[\frac{q}{\rho_g\mu_g}\frac{\left(\gamma_{gi}^{\frac{\rho_g}{q}}+\gamma_{gj}^{\frac{\rho_g}{q}}\right)^2}{\gamma_{gi}^{\frac{\rho_g}{q}}\gamma_{gj}^{\frac{\rho_g}{q}}}\right]^{\frac{1}{q}} \tag{5}$$

Equation 5 gives the critical point, but we must compare its profits with endpoints. The profit can be written as:

$$\pi_{gi}(\epsilon_{gi}^*,\epsilon_{gj}^*) = \frac{\left(\frac{q}{\rho_g\mu_g}\right)^{\frac{1}{q}}\frac{(\gamma_{gi}^q+\gamma_{gj}^q)^{\frac{2}{q}}}{\gamma_{gi}^{1-\frac{1}{q}}\gamma_{gj}^q}}{\left(\frac{q}{\rho_g\mu_g}\right)^{\frac{1}{q}}\frac{(\gamma_{gi}^q+\gamma_{gj}^q)^{\frac{2}{q}}}{\gamma_{gj}^{1-\frac{1}{q}}\gamma_{gi}^q}+\left(\frac{q}{\rho_g\mu_g}\right)^{\frac{1}{q}}\frac{(\gamma_{gi}^q+\gamma_{gj}^q)^{\frac{2}{q}}}{\gamma_{gi}^{1-\frac{1}{q}}\gamma_{gj}^q}} - \frac{\gamma_{gi}}{\epsilon_{gi}^q} = \frac{1}{1+\left[\frac{\gamma_{gj}^{1-\rho_g/q-q}}{\gamma_{gi}^{1-\rho_g/q-q}}\right]^{\rho_g}} - \frac{\gamma_{gi}}{\epsilon_{gi}^q}$$

Substituting back in ϵ_{gi}^q, it is:

$$\frac{1}{1+\left[\frac{\gamma_{gj}^{1-\rho_g/q-q}}{\gamma_{gi}^{1-\rho_g/q-q}}\right]^{\rho_g}} - \frac{\gamma_{gi}\rho_g\mu_g\gamma_{gi}^{\frac{\rho_g}{q}}\gamma_{gj}^{\frac{\rho_g}{q}}}{q\left(\gamma_{gi}^{\frac{\rho_g}{q}}+\gamma_{gj}^{\frac{\rho_g}{q}}\right)^2}$$

For this interior equilibrium to hold, it must be that $\pi_{gi}^*(\epsilon_{gi}^*,\epsilon_{gj}^*) \geq \pi_{gi}^*(\epsilon',\epsilon_{gj}^*)$ for all other choices of ϵ_{gj}^*. Note that $\pi_{gi,\epsilon_{gj}^*}(\epsilon)$ is continuous away from 0. Moreover, for small enough ϵ_0, $\pi_{gi,\epsilon_{gj}^*}(\epsilon) < 0$, since the market size is bounded by costs can be come arbitrarily negative. Hence, we can consider the maximizing this function on the compact set $[\epsilon_0, 1]$. Since $\epsilon_{gi,\epsilon_{gj}^*}(\epsilon)$ is continuous on this set, and ϵ_{gi}^* satisfies the first-order condition, the only possible maxima of this function are ϵ_0 or 1. At ϵ_0, the firm is making zero profits, so any choice with positive profits eliminates it. At $\epsilon = 1$, the firm can also choose to not invest anything in data (and receive the same revenue but no data costs), so the condition that makes $\pi_{gi,\epsilon_{gj}^*}(\epsilon_{gi}^*) > \pi_{gi,\epsilon_{gj}^*}(1)$ will be sufficient to make this an equilibrium.

This condition holds if

$$\frac{1}{1+\left[\frac{\gamma_{gj}^{1-\rho_g/q-q}}{\gamma_{gi}^{1-\rho_g/q-q}}\right]^{\rho_g}} - \frac{\gamma_{gi}\rho_g\mu_g\gamma_{gi}^{\frac{\rho_g}{q}}\gamma_{gj}^{\frac{\rho_g}{q}}}{q\left(\gamma_{gi}^{\frac{\rho_g}{q}}+\gamma_{gj}^{\frac{\rho_g}{q}}\right)^2} \geq \pi_{gi}(1,\epsilon_{gj^*}) \tag{6}$$

We call Inequality 6 the *nondeviation condition*. We can write:

$$\pi_{gi}(1,\epsilon_{gj^*}) = \frac{\left(\frac{q}{\rho_g\mu_g}\right)^{\frac{1}{q}}\frac{(\gamma_{gi}^q+\gamma_{gj}^q)^{\frac{2}{q}}}{\gamma_{gi}^{1-\frac{1}{q}}\gamma_{gj}^q}}{1+\left(\frac{q}{\rho_g\mu_g}\right)^{\frac{1}{q}}\frac{(\gamma_{gi}^q+\gamma_{gj}^q)^{\frac{2}{q}}}{\gamma_{gi}^{1-\frac{1}{q}}\gamma_{gj}^q}} = \frac{1}{1+\left(\frac{q}{\rho_g\mu_g}\right)^{-\frac{1}{q}}\frac{(\gamma_{gi}^q+\gamma_{gj}^q)^{\frac{-2}{q}}}{\gamma_{gi}^{\frac{1}{q}-1}\gamma_{gj}^{-q}}},$$

so Inequality 6 asks that

$$\frac{1}{1+\left[\frac{\gamma_{gj}^{1-\rho_g/q-q}}{\gamma_{gi}^{1-\rho_g/q-q}}\right]^{\rho_g}} - \frac{\gamma_{gi}\rho_g\mu_g\gamma_{gi}^{\frac{\rho_g}{q}}\gamma_{gj}^{\frac{\rho_g}{q}}}{q\left(\gamma_{gi}^{\frac{\rho_g}{q}}+\gamma_{gj}^{\frac{\rho_g}{q}}\right)^2} \geq \frac{1}{1+\left(\frac{q}{\rho_g\mu_g}\right)^{-\frac{1}{q}}\frac{(\gamma_{gi}^q+\gamma_{gj}^q)^{\frac{-2}{q}}}{\gamma_{gi}^{\frac{1}{q}-1}\gamma_{gj}^{-q}}}$$

We have shown that if the nondeviation condition holds for each group and each firm, then $(\epsilon_{gi}^*, \epsilon_{gj}^*)$ is a Nash Equilibrium in pure strategies under proportionally split demand with competition exponent ρ_g in each group and learning rate q. If a further conditions holds, namely that there exists a preferred strategy to non-investment if the opponent does not invest, then the equilibrium is unique. Call this the *investment* condition: there exists $\epsilon \in (0, 1)$ such that:

$$\frac{1}{\epsilon_g^\rho+1} - \frac{\gamma_{gi}}{\epsilon^q} > \frac{\mu_g}{2} \iff \epsilon^q - \gamma_{gi}(\epsilon^{\rho_g}+1) \geq \frac{\mu_g\left((\epsilon^{\rho_g}+1)\epsilon^q\right)}{2}. \tag{7}$$

Equivalently, we need to ensure that:

$$\epsilon^q - \gamma_{gi}(\epsilon^{\rho_g}+1) - \frac{\mu_g\left((\epsilon^{\rho_g}+1)\epsilon^q\right)}{2} \geq 0 \tag{8}$$

has a solution in $(0, 1)$. This will not always be the case, of course; when it is not, there is an equilibrium in which both firms prefer not to invest in collecting data from one group at all, which certainly implies inequality.

□

Proof of Lemma 3. First, consider absolute equality guarantees. First, notice that Lemma 1 gives that if $\mu_A\beta_A \geq \mu_B\beta_B$ and $\gamma_A \leq \gamma_B$, then $\epsilon_B^M > \epsilon_A^M(1+\chi)$. Fix any choice ϵ_A^R. We claim that if $\epsilon_B^R < (1+\chi)\epsilon_A^R$, the pair $(\epsilon_A^R, \epsilon_B^R)$ cannot be a constrained profit maximizer. There are two cases. Either $\epsilon_B^R < \epsilon_A^R(1+\chi) < \epsilon_B^M$, or $\epsilon_B^R < \epsilon_A^R(1+\chi)$ but $\epsilon_B^M < \epsilon_A^R(1+\chi)$. The latter case is trivial - by the separability of the profit function, the feasible pair $(\epsilon_A^R, \epsilon_B^M)$ would increase profits. Now consider the case that $\epsilon_B^R < \epsilon_A^R(1+\chi) < \epsilon_B^M$. Let $\tilde{\epsilon} = (1-\alpha)(\epsilon_A^R, \epsilon_B^R) + \alpha(\epsilon_A^R, \epsilon_B^M)$. The profit function is concave, so Jensen's inequality gives that for $\alpha \in [0, 1]$:

$$\pi(\tilde{\epsilon}) > (1-\alpha)\pi(\epsilon_A^R, \epsilon_B^R) + \alpha(\epsilon_A^R, \epsilon_B^M) \geq \pi(\epsilon_A^R, \epsilon_B^R)$$

and by choosing $\alpha < \frac{\epsilon_B^M - (1+\chi)\epsilon_A^R}{\epsilon_B^M - \epsilon_B^R}$, $\tilde{\epsilon}$ is indeed feasible.

Now consider absolute error constraints. First, note that the absolute error constraints are separable, so that the firm's problem is additively separable. That is, it can consider the populations separately. So consider population g. If $\epsilon_g^M \leq \chi$, then the unconstrained optimal is less than the regulatory bound, so the regulation won't change the behavior of the firm. Thus if $\epsilon_g^R \neq \epsilon_g^M$, it must be that $\epsilon_G^M > \chi$. Since ϵ_g^R must be feasible, we must have that $\epsilon_g^R \leq \chi$. Now suppose that $\epsilon_g^R < \chi$.

Let $\tilde{\epsilon}_g = (1-\alpha)\epsilon_g^R + \alpha\epsilon_g^M$. Since π_g is concave, Jensen's inequality implies that

$$\pi_g(\tilde{\epsilon}_g) = \pi_g((1-\alpha)\epsilon_g^R + \alpha\epsilon_g^M) \geq (1-\alpha)\pi_g(\epsilon_g^R) + \alpha\pi_g(\epsilon_g^M) \geq \pi_g(\epsilon_g^R)$$

so moving to $\tilde{\epsilon}_g$ can only improve the profits. $\tilde{\epsilon}_g$ will be feasible if

$$\tilde{\epsilon}_g \leq \chi \iff (1-\alpha)\epsilon_g^R + \alpha\epsilon_g^M \leq \chi$$

So choosing $\alpha < \frac{\chi - \epsilon_g^R}{\epsilon_g^M - \epsilon_g^R}$ will suffice. (Notice that $\frac{\chi - \epsilon_g^R}{\epsilon_g^M - \epsilon_g^R} \in [0, 1]$, since $\epsilon_g^M > \chi \implies \epsilon_g^M - \epsilon_g^R > \chi - \epsilon_g^R$, and both numerator and denominator are positive by assumption.)

But that means that if $\epsilon_g^R < \chi$, we can find another point with strictly more profit. Thus, ϵ_g^R was not optimal.

□

Proof of Theorem 3. First, we consider the equal error guarantee. Lemma 3 means that the regulated monopolist's problem is equivalent to substituting $(1+\chi)\epsilon_A$ for ϵ_B in the original formulation. The solution to this new formulation gives that $\epsilon_A^R = \frac{q(\gamma_A + \gamma_B/(1+\chi)^q)}{\mu_A\beta_A + \mu_B\beta_B(1+\chi)}$.

$$\left(\frac{\epsilon_A^R}{\epsilon_A^M}\right)^{q+1} = \frac{q\left(\gamma_A + \gamma_B/(1+\chi)^q\right)/\left(\mu_A\beta_A + \mu_B\beta_B(1+\chi)\right)}{q\gamma_A/(\mu_A\beta_A)} = \frac{\mu_A\beta_A\gamma_A + \mu_A\beta_A\gamma_B/(1+\chi)^q}{\mu_A\beta_A\gamma_A + \mu_B\beta_B\gamma_A(1+\chi)}$$

Now using the fact that $\frac{\epsilon_A^R}{\epsilon_A^M} \geq 1 \iff \left(\frac{\epsilon_A^R}{\epsilon_A^M}\right)^{q+1} \geq 1$ Now using the elementary fact that for positive x, y, z, $(x+y)/(x+z) \geq 1 \iff y \geq z$, we can see that

$$\left(\frac{\epsilon_A^R}{\epsilon_A^M}\right)^{q+1} \geq 1 \iff \mu_A\beta_A\gamma_B/(1+\chi)^q \geq \mu_B\beta_B\gamma_A(1+\chi) \iff \frac{\gamma_B}{\mu_B\beta_B} \geq \frac{\gamma_A}{\mu_A\beta_A}(1+\chi)^{q+1}$$

But recalling that the monopolist's optimal solution is $\epsilon_g^M = (\frac{q\gamma_g}{\mu_g\beta_g})^{\frac{1}{q+1}}$, we can rewrite the previous inequality:

$$\epsilon_B^{M\,q+1} \geq \epsilon_A^{M\,q+1}(1+\chi)^{q+1} \iff \epsilon_B^M \geq \epsilon_A^M(1+\chi)$$

which holds under our assumptions on market power and cost. A similar calculation will give that $\epsilon_B^R \leq \epsilon_B^M$ under the same conditions.

An upper bound on $\frac{\epsilon_A^R}{\epsilon_A^M}$ follows by diving ϵ_A^M and ϵ_A^R and dropping the term $\frac{\mu_A\beta_A}{\mu_A\beta_A+\mu_B\beta_B} < 1$ from the product.

The claim for the absolute error guarantee follows directly from Lemma 3 – $\epsilon_A^R \leq \chi$, so if $\epsilon_A^M > \chi$ and the constraint is saturated, then $\epsilon_A^R = \chi < \epsilon_A^M$.

□

Proof of Theorem 4. Under equal error constraints, Lemma 3 implies that $\epsilon_B^R = \epsilon_A^R(1+\chi)$. Then the regulated monopolist's problem is equivalent to the problem of a monopolist facing a single population with demand $\mu_A\alpha + \mu_B\alpha + B - \mu_A\beta_A\epsilon_A - \mu_B\beta_B\epsilon_B(1+\chi)$, with costs $\phi_A + \phi_B - \gamma_A/\epsilon_A^q - \gamma_B/(\epsilon_A^q(1+\chi)^q)$. Plugging in the optimal solution for the monopolist's problem from Lemma 1 and rearranging shows that the optimal *profit* is given by $pi^*(\epsilon^*) = \sum_{g\in\mathcal{G}} \mu_g\alpha_g - (\mu_g\beta_g)^{\frac{q}{q+1}}\gamma_g^{\frac{1}{q+1}}Q$, where Q is a constant related to q. Thus, we can write

$$\text{MonPoF}_{1+\chi} = \frac{\mu_A\alpha_A + r\mu_A\alpha_B - Q\mu_A^{\frac{q}{q+1}}\left[\beta_A^{\frac{q}{q+1}}\gamma_A^{\frac{1}{q+1}} + r^{\frac{q}{q+1}}\beta_B^{\frac{q}{q+1}}\gamma_B^{\frac{1}{q+1}}\right]}{\mu_A\alpha_A + r\mu_A\alpha_B - Q\mu_A^{\frac{q}{q+1}}\left[(\beta_A + r\beta_B(1+\chi))^{\frac{q}{q+1}}(\gamma_A + \gamma_B\frac{1}{(1+\chi)}^q)^{\frac{1}{q+1}}\right]}$$

Factoring out μ_A and using the fact that if $\mu_B \to \infty$, $\mu_A \to \infty$, we can see that this quantity tends to 1.

For the second claim, we can simply substitute in $\mu_B = 0$ and factor out $\mu_A\alpha_A$ to get

$$\lim_{\mu_B\to 0}\text{MonPoF} = \frac{\mu_A\alpha_A[1 - Q\left(\frac{\gamma_A}{\mu_A\alpha_A}\right)^{\frac{1}{q+1}}]}{\mu_A\alpha_A[1 - Q\left(\frac{\gamma_A+\gamma_B/(1+\chi)^{\frac{1}{q+1}}}{\mu_A\alpha_A}\right)^{\frac{1}{q+1}}]} = \frac{1 - Q/q\epsilon_A^M}{1 - Q/q\epsilon_A^M\left(1 + \frac{\gamma_B}{\gamma_A}\frac{1}{(1+\chi)^{\frac{1}{q+1}}}\right)^{\frac{1}{q+1}}}$$

and the latter term is a constant above 1.

For the absolute error guarantees: Notice that profit and constraints are separable across groups, so the monopolist can optimize separately. Thus, the monopolist can optimize each group separately. By the saturation lemma, we can thus have three possibilities: neither constraint binds, the constraint on the minority group binds, or the constraint on both groups bind (the constraint will not bind on the majority group without binding on the minority group since $\epsilon_A^M \leq \epsilon_B^M$). If we write π_g for the profit accrued solely from group g, then, we can write that ϵ^R/ϵ^M will be 1 in the first case, $\frac{\pi_A^M+\pi_B^M}{\pi_A^M+\pi_B(\chi)}$ in the second case, or $\frac{\pi_A^M+\pi_B^M}{\pi_A(\chi)+\pi_B^\chi}$.

Now, suppose that $\mu_B, \mu_A \to \infty$ at a constant ratio. Lemma 1 implies that $\epsilon_A^M, \epsilon_B^M \to 0$. Hence, as the market size grows, eventually ϵ_B^M and ϵ_A^M will be less than χ, so that $\epsilon_B^R = \epsilon_B^M$ and $\epsilon_A^R = \epsilon_A^M$. Then in the limit, case 1 obtains, so $\pi^M/\pi^R \to 1$. For $\mu_B \to 0$, we will eventually be in either Case 2 or Case 3. Note that as $\mu_B \to 0$, ϵ_B^M will eventually be larger than χ, so the limit will be obtained at either Case 2 or Case 3. Notice that as $\mu_B \to 0$, $\pi_B \to 0$; moreover, for small enough μ_B, the optimal choice for the unconstrained monopolist eventually becomes to set error as high as possible ($\epsilon_B^M = 1$). Thus, in Case 2, we can write the π^M/π^R as:

$$\lim_{\mu_B\to\infty}\frac{\pi^M}{\pi^R} = \frac{\pi_A(\epsilon_A^M) - \gamma_B - \phi_B}{\pi_A(\epsilon_A)^M - \gamma_B/\chi^q - \phi_B} \geq 1$$

Alternatively, if Case 3 obtains, then we can write:

$$\lim_{\mu_B \to \infty} \text{MonPoF}_\chi = \frac{\pi_A(\chi) - \gamma_B - \phi_B}{\pi_A(\chi) - \gamma_B/\chi^q - \phi_B} \geq 1$$

Finally, for the claim that there is some minimum error the monopolist can guarantee, First, by the saturation lemma, when $\chi < \epsilon_A^M$, then the optimal choice of the monopolist is $(\epsilon_A^R, \epsilon_B^R) = (\chi, \chi)$. then the profit can be written as a function of χ:

$$\pi(\chi) = \mu_A(\alpha_A - \beta_A\chi) + \mu_B(\alpha_B - \beta_B\chi) - \phi_A - \phi_B - \frac{\gamma_A}{\chi^q} - \frac{\gamma_B}{\chi^q} \tag{9}$$

Since π is concave and the global optimum is at $(\epsilon_A^M, \epsilon_B^M)$, we know that decreasing χ will decrease eventually decrease the attainable objective value (otherwise, the global maximum would not have been attained at $(\epsilon_A^M, \epsilon_B^M)$). The point at which objective value becomes 0 is the solution to $0 = \mu_A(\alpha_A - \beta_A\chi) + \mu_B(\alpha_B - \beta_B\chi) - \phi_A - \phi_B - \frac{\gamma_A}{\chi^q} - \frac{\gamma_B}{\chi^q}$. Multiplying by χ^q gives a polynomial equation whose solution can be found via the quadratic or cubic formulas in the realizable or agnostic settings, respectively, or numerically for other settings. If the solution ϵ_0 is in $[0, 1]$, then the above reasoning implies that the polynomial is negative for $\chi < \epsilon_0$. □