Uncertainty and the Social Planner's Problem: Why Sample Complexity Matters

Cyrus Cousins
cyrus_cousins@brown.edu
Brown University Department of Computer Science
Providence, Rhode Island, USA

ABSTRACT

Welfare measures overall utility across a population, whereas malfare measures overall disutility, and the social planner's problem can be cast either as maximizing the former or minimizing the latter. We show novel bounds on the expectations and tail probabilities of estimators of welfare, malfare, and regret of per-group (dis)utility values, where estimates are made from a finite sample drawn from each group. In particular, we consider *estimating* these quantities for individual functions (e.g., allocations or classifiers) with standard probabilistic bounds, and optimizing and *bounding generalization error* over *hypothesis classes* (i.e., we quantify overfitting) using Rademacher averages. We then study algorithmic fairness through the lens of sample complexity, finding that because marginalized or minority groups are often understudied, and fewer data are therefore available, the social planner is more likely to overfit to these groups, thus even models that *seem fair in training* can be *systematically biased* against such groups. We argue that this effect can be mitigated by ensuring sufficient sample sizes for each group, and our sample complexity analysis characterizes these sample sizes. Motivated by these conclusions, we present *progressive sampling* algorithms to efficiently optimize various fairness objectives.

ACM Reference Format:
Cyrus Cousins. 2022. Uncertainty and the Social Planner's Problem: Why Sample Complexity Matters. In *2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22), June 21–24, 2022, Seoul, Republic of Korea.* ACM, New York, NY, USA, 12 pages. https://doi.org/10.1145/3531146.3533243

1 INTRODUCTION

Machine learning systems in settings like facial recognition [10, 11, 13] and medicine [3, 4, 12] exhibit *differential accuracy* across race and other protected groups. This can lead to discrimination: for example, facial recognition in policing yields disproportionate false-arrest rates [24], and machine learning in medicine can lead to inequity of health outcomes [18], both of which exacerbate existing structural inequalities impacting minority groups. In recent years, researchers have proposed welfare-centric fair learning models, which constrain or optimize welfare [15, 22, 26, 28, 38, 44, 45] or malfare [1, 14, 20, 30, 33, 43] to promote *fair learning* across *all groups*, as well as regret-based methods [7, 39], which similarly promote fairness by minimizing the *maximum dissatisfaction* of any group, relative to their preferred outcome, i.e., how much excess risk is incurred, or utility is lost, to any group by *compromising* on a *shared model.*

We study sampling and learning problems in the optimization of welfare, malfare, and regret objectives. In particular, our setting subsumes the *minimax fair learning* [1, 20, 30, 33, 43] and *fair-PAC learning* [14] settings, by considering arbitrary malfare or welfare functions, as well as the *multi-group agnostic PAC learning* [7, 39] setting, by considering *arbitrary malfare functions* — rather than just the maximum — of per-group regret values. This extension naturally and smoothly interpolates between minimizing *utilitarian* (i.e., weighted average) and *egalitarian* (i.e., maximum) malfare of risk or regret. Crucially, this allows for fine-grained control over the desired fairness concept, and mitigates the minority rule issues of minimax methods, while remaining axiomatically grounded in *cardinal welfare theory.* We bound the generalization error of optimizing welfare, malfare, and regret objectives, and find that while the power-mean malfare is always easy to estimate, due to Lipschitz-continuity (as studied by Cousins [14]), our learning algorithms work for any malfare, welfare, or regret objective that is *continuous* and *monotonic* in per-group (dis)utility values.

We then study algorithmic fairness through the lens of sample complexity, finding that because marginalized or minority groups are often understudied, and fewer data are therefore available, the social planner is more likely to *overfit* to these groups. Consequently, even models that *seem fair* in training can be *systematically biased* against such groups. Section 3 shows that this effect can be mitigated with sufficient per-group sample sizes, and §4 presents *progressive sampling* methods, which dynamically sample until a near-optimal model (w.r.t. some fairness objective) is learned.

Our bounds leverage the specific character of the objective at hand; for example, utilitarian welfare is sensitive to the *average* confidence radius across groups, whereas egalitarian welfare is more sensitive to the confidence radii of disadvantaged (i.e., low-utility or high-risk) groups. Furthermore, our progressive sampling methods are tailored to three realistic models of data generation: in the *joint sampling* model, each sample contains a piece of information for every group, in the *mixture sampling* model, samples are annotated with (sets of) group labels, and in the *conditional sampling* model, we are allowed to choose from which groups to sample. While our settings and modelling assumptions are practically motivated, this is a highly theoretical paper, and all novel results are meticulously proven in §6.

FAccT '22, June 21–24, 2022, Seoul, Republic of Korea

ACM ISBN 978-1-4503-9352-2/22/06…$15.00
https://doi.org/10.1145/3531146.3533243

2 LEARNING FRAMEWORK AND OBJECTIVES

In this section, we introduce the functional form of the objects and random spaces that we operate over, and we define our learning objectives. In particular, §2.1 presents the welfare, malfare, and regret objectives, which compile per-group sentiment values into a *cardinal objective value* that can be optimized and analyzed, then §2.2 reifies this abstract mathematics with three realistic models of data-collection, each of which requires its own statistical treatment to efficiently learn from data, i.e., to optimize and bound objectives, while minimizing the cost of obtaining said data.

We henceforth assume a *supervised learning setting*, where $\mathcal{X}$ is the domain and $\mathcal{Y}$ is the codomain. We also assume either a *loss function* [1] $\ell(\cdot,\cdot): \mathcal{Y}' \times \mathcal{Y} \to \mathbb{R}_{0+}$ or a *utility function* $\mathrm{u}(\cdot,\cdot): \mathcal{Y}' \times \mathcal{Y} \to \mathbb{R}_{0+}$, which map predictions and labels onto *negatively connoted* loss or disutility, or *positively connoted* gain or utility, generically termed a *sentiment function* $\mathrm{s}(\cdot,\cdot): \mathcal{Y}' \times \mathcal{Y} \to \mathbb{R}_{0+}$. In most supervised learning settings, a single probability distribution $\mathcal{D}$ over $\mathcal{X} \times \mathcal{Y}$ suffices, but we assume a set $\mathcal{Z}$ of g groups, and we model the experiences and conditions of each group as its own distribution, i.e., we have $\mathcal{D}_1, \dots, \mathcal{D}_g$. For convenience, we often compose the sentiment function with a predictor or model $h(\cdot): \mathcal{X} \to \mathcal{Y}'$, taking $(\mathrm{s} \circ h)(x, y) \doteq \mathrm{s}(h(x), y)$, thus we quantify model performance for group i as $\mathbb{E}_{\mathcal{D}_i}[\mathrm{s} \circ h]$.

2.1 Fair Learning with Malfare, Welfare, and Regret Objectives

Here we define the welfare, malfare, and regret objectives. While the details differ, each of these is a function of the expected utility or loss (generically sentiment) of some $h: \mathcal{X} \to \mathcal{Y}$ for each of the g groups, and we are interested in selecting the model or hypothesis h from some *hypothesis class* $\mathcal{H} \subseteq \mathcal{X} \to \mathcal{Y}'$ that optimizes the given objective.

Malfare and Welfare. A welfare function $\mathrm{W}(\mathcal{S}; \boldsymbol{w})$ measures overall positive utility $\mathcal{S}$ across a population weighted by $\boldsymbol{w}$, whereas a malfare function $\mathrm{M}(\mathcal{S}; \boldsymbol{w})$ measures overall disutility $\mathcal{S}$, and generically, we say an aggregator function $\mathrm{M}(\mathcal{S}; \boldsymbol{w})$ measures overall sentiment $\mathcal{S}$. The prototypical example is the *utilitarian* (or *Benthamite*) aggregate, defined as $\mathrm{M}_1(\mathcal{S}; \boldsymbol{w}) \doteq \mathcal{S} \cdot \boldsymbol{w}$, which simply averages sentiment across the population (e.g., welfare as per-capita income, or malfare as per-capita medical expenditure), and the second-fiddle is the *egalitarian* (or *Rawlsian*) welfare $\mathrm{W}_{-\infty}(\mathcal{S}; \boldsymbol{w})$ (minimum) or malfare $\mathrm{M}_{\infty}(\mathcal{S}; \boldsymbol{w})$ (maximum), which summarizes a population's sentiment as that of its most disadvantaged member. We assume throughout that $\boldsymbol{w} \in (0,1)^g$ is a *probability vector*, thus $\|\boldsymbol{w}\|_1 = 1$, and $\mathcal{S} \in \mathbb{R}^g_{0+}$ is *nonnegative*. *Ab initio*, our first objective is, via the social planner's problem, to maximize welfare [22, 28, 38, 44], or by extension (e.g., in chores manna or harm allocation [25, 29], or in machine learning [1, 14]) to minimize malfare, i.e., we seek to approximate

$$h^\star \doteq \operatorname*{argmin}_{h \in \mathcal{H}} \mathrm{M}\left(i \mapsto \mathop{\mathbb{E}}_{\mathcal{D}_i}[\ell \circ h]; \boldsymbol{w}\right) ,$$

$$\text{or} \quad h^\star \doteq \operatorname*{argmax}_{h \in \mathcal{H}} \mathrm{W}\left(i \mapsto \mathop{\mathbb{E}}_{\mathcal{D}_i}[\mathrm{u} \circ h]; \boldsymbol{w}\right) . \tag{1}$$

Intuitively, the utilitarian case seeks to optimize overall or average sentiment, whereas the egalitarian case instead seeks to lift up the most disadvantaged group, and thus promote equality, perhaps at the expense of overall (total) utility.

Notably, welfare maximization generalizes *utility maximization* to multiple groups, and malfare minimization likewise generalizes *risk minimization*, and the well-studied *minimax fair-learning* framework arises as the special-case of *egalitarian malfare minimization.* In general, we assume only *monotonicity* and *continuity* of aggregator functions; however, there are a set of relatively standard axioms that, when taken together, restricts the class of interest to the *power-mean family* [14]. This is convenient, as all power-mean malfare functions are Lipschitz-continuous, which in §4 leads to stronger estimation guarantees and more efficient sampling algorithms than ε-δ limit-continuity.

Definition 2.1 (Axioms of Cardinal Welfare and Malfare)**.** Suppose an aggregator function $\mathrm{M}(\mathcal{S}; \boldsymbol{w})$. For each item, assume (if necessary) that the axiom applies for all $\mathcal{S}, \mathcal{S}' \in \mathbb{R}^g_{0+}$, scalars $\alpha, \beta \in \mathbb{R}_{0+}$, and probability vectors $\boldsymbol{w} \in (0,1)^g$.

(1) *Strict Monotonicity*: $\mathcal{S}' \neq \mathbf{0} \implies \mathrm{M}(\mathcal{S}; \boldsymbol{w}) < \mathrm{M}(\mathcal{S} + \mathcal{S}'; \boldsymbol{w})$.
(2) *Weighted Symmetry*: Suppose $g' \in \mathbb{Z}_+$, $\mathcal{S}' \in \mathbb{R}^{g'}_{0+}$, and probability vector $\boldsymbol{w}' \in (0,1)^{g'}$, such that for all $u \in \mathbb{R}_{0+}$, it holds that $\sum_{i \text{ s.t. } \mathcal{S}_i = u} \boldsymbol{w}_i = \sum_{i \text{ s.t. } \mathcal{S}'_i = u} \boldsymbol{w}'_i$. Then $\mathrm{M}(\mathcal{S}; \boldsymbol{w}) = \mathrm{M}(\mathcal{S}'; \boldsymbol{w}')$.
(3) *Continuity*: $\mathrm{M}(\mathcal{S}; \boldsymbol{w})$ is a continuous function (in the standard ε-δ limit-continuity sense) in both $\mathcal{S}$ and $\boldsymbol{w}$.
(4) *Independence of Unconcerned Agents* : $\mathrm{M}(\langle \mathcal{S}_{1:g-1}, \alpha \rangle; \boldsymbol{w}) \leq \mathrm{M}(\langle \mathcal{S}'_{1:g-1}, \alpha \rangle; \boldsymbol{w}) \implies \mathrm{M}(\langle \mathcal{S}_{1:g-1}, \beta \rangle; \boldsymbol{w}) \leq \mathrm{M}(\langle \mathcal{S}'_{1:g-1}, \beta \rangle; \boldsymbol{w})$.
(5) *Multiplicative Linearity*: $\mathrm{M}(\alpha\mathcal{S}; \boldsymbol{w}) = \alpha\mathrm{M}(\mathcal{S}; \boldsymbol{w})$.
(6) *Unit Scale*: $\mathrm{M}(\mathbf{1}; \boldsymbol{w}) = \mathrm{M}(\langle 1, \dots, 1 \rangle; \boldsymbol{w}) = 1$.
(7) *Pigou-Dalton Transfer Principle*: Suppose $\mu = \boldsymbol{w} \cdot \mathcal{S} = \boldsymbol{w} \cdot \mathcal{S}'$, and for all $i \in \mathcal{Z}$: $|\mu - \mathcal{S}'_i| \leq |\mu - \mathcal{S}_i|$. Then for utility and welfare, $\mathrm{W}(\mathcal{S}'; \boldsymbol{w}) \geq \mathrm{W}(\mathcal{S}; \boldsymbol{w})$, and for disutility and malfare, $\mathrm{M}(\mathcal{S}'; \boldsymbol{w}) \leq \mathrm{M}(\mathcal{S}; \boldsymbol{w})$.

Axioms 1–4 are essentially the standard *axioms of cardinal welfare* [37, 40], modified to include the weights $\boldsymbol{w}$, and omitting any of them leads to rather perverse aggregator functions. Axiom 5 (multiplicative linearity) strengthens the traditional *independence of common scale* axiom, and ensures that the *units* of welfare or malfare must match those of sentiment, and axiom 6 (unit scale) merely specifies a multiplicative constant. Finally, axiom 7, the *Pigou-Dalton transfer principle* [17, 36], characterizes *fairness* in the sense of *equitable redistribution* of utility (welfare) or disutility (malfare).

Theorem 2.2 (Aggregator Function Properties [14, theorems 2.4 and 2.5])**.** Suppose aggregator function $\mathrm{M}(\mathcal{S}; \boldsymbol{w})$, and assume arbitrary sentiment vector $\mathcal{S} \in \mathbb{R}^g_{0+}$ and probability vector $\boldsymbol{w} \in (0,1)^g$. The following then hold.

[1] Often $\mathcal{Y}' = \mathcal{Y}$, such as in standard classification and regression settings, but this is not universally the case. For instance, in probabilistic classification or regression (i.e., conditional density estimation), $\mathcal{Y}'$ is a space of *distributions over* $\mathcal{Y}$, and in *interval estimation*, $\mathcal{Y}'$ is a space of *sets over* $\mathcal{Y}$.

(1) *Power-Mean Factorization*: Axioms 1–6 imply $\exists p \in \mathbb{R}$ s.t.

$$\mathrm{M}(\mathcal{S};\boldsymbol{w}) = \mathrm{M}_p(\mathcal{S};\boldsymbol{w}) \doteq f_p^{-1}\left(\sum_{i=1}^{g} \boldsymbol{w}_i f_p(\mathcal{S}_i)\right)$$
$$=_{p\neq 0} \sqrt[p]{\sum_{i=1}^{g} \boldsymbol{w}_i \mathcal{S}_i^p} \ , \quad \text{with} \begin{cases} p = 0 & f_0(x) \doteq \ln(x) \\ p \neq 0 & f_p(x) \doteq \mathrm{sgn}(p) x^p \end{cases} .$$

(2) *Fair Welfare and Malfare*: Axioms 1–7 imply $p \in (-\infty, 1]$ for welfare and $p \in [1, \infty)$ for malfare.
(3) *Lipschitz-Continuity*: For all $p \geq 1$, it holds that $|\mathrm{M}_p(\mathcal{S};\boldsymbol{w}) - \mathrm{M}_p(\mathcal{S}';\boldsymbol{w})| \leq \mathrm{M}_p(|\mathcal{S} - \mathcal{S}'|;\boldsymbol{w}) \leq \max_{i\in\mathcal{Z}} |\mathcal{S}_i - \mathcal{S}'_i|$.

In closing, we note that utilitarian philosophy is often criticized for permitting great inequality by ignoring the needs of smaller or less visible groups, whereas egalitarian philosophy is criticized for ignoring the masses in favor of outliers and disadvantaged groups, and its inherent susceptibility to minority rule. Concretely, utilitarian aggregates only weakly satisfy the Pigou-Dalton principle, thus do not incentivize equitable redistribution , and egalitarian aggregates satisfy only weakmonotonicity, thus only incentivize gains in the *most disadvantaged* group(s). Power-means provide a spectrum of intermediaries, so exactly how tradeoffs should be made may depend on the application, as well as the culturosocietal values of the social planner. They are also *statistically convenient*, as many of our estimation guarantees hold in terms of generic Lipschitz-continuity assumptions, and thus apply to any power-mean malfare function.

Malfare of Regret. Regret measures the *relative dissatisfaction* of group i with some $h \in \mathcal{H}$, relative to their preferred $\boldsymbol{h}_i^\star \in \mathcal{H}$. We define the (per-group) preferred outcome $\boldsymbol{h}_i^\star$ as the model group i would select for themselves, i.e.,

$$\boldsymbol{h}_i^\star \doteq \operatorname*{argmin}_{h\in\mathcal{H}} \mathop{\mathbb{E}}_{\mathcal{D}_i}[\ell \circ h] \quad \text{or} \quad \boldsymbol{h}_i^\star \doteq \operatorname*{argmax}_{h\in\mathcal{H}} \mathop{\mathbb{E}}_{\mathcal{D}_i}[\mathrm{u} \circ h] \ , \tag{2}$$

for loss or utility, respectively, and we let $\mathcal{S}_i^\star$ denote the *optimal expected sentiment* for group i, i.e., $\mathcal{S}_i^\star \doteq \mathbb{E}_{\mathcal{D}_i}[\mathrm{s} \circ \boldsymbol{h}_i^\star]$. We now formally define the *regret* of group i on some outcome or model $h \in \mathcal{H}$ as

$$\mathrm{Reg}_i(h) \doteq \mathop{\mathbb{E}}_{\mathcal{D}_i}[\ell \circ h] - \mathcal{S}_i^\star \ , \quad \mathrm{Reg}_i(h) \doteq \mathcal{S}_i^\star - \mathop{\mathbb{E}}_{\mathcal{D}_i}[\mathrm{u} \circ h] \ ,$$
$$\text{or generically,} \quad \mathrm{Reg}_i(h) \doteq |\mathop{\mathbb{E}}_{\mathcal{D}_i}[\mathrm{s} \circ h] - \mathcal{S}_i^\star| \ . \tag{3}$$

Intuitively $\mathrm{Reg}_i(h)$ is nonnegative, and it quantifies the amount by which group i prefers their optimal $\boldsymbol{h}_i^\star$ to h.

Several authors [7, 39] minimize the worst-case (over groups) *regret* of the selected $\hat{h}$, and the statistical and computational questions that arise are studied under the umbrella of "multi-group agnostic PAC learning." We generalize this notion, optimizing not just *worst-case* (i.e., egalitarian), but *arbitrary malfare functions*, of per-group regret values, which allows for greater flexibility and resistance to the usual issues of egalitarian malfare. In particular, we seek

$$h^\star \doteq \operatorname*{argmin}_{h\in\mathcal{H}} \mathrm{M}\!\!\mathrm{M}\left(i \mapsto \mathrm{Reg}_i(h);\boldsymbol{w}\right)$$
$$= \operatorname*{argmin}_{h\in\mathcal{H}} \mathrm{M}\!\!\mathrm{M}\left(i \mapsto |\mathop{\mathbb{E}}_{\mathcal{D}_i}[\mathrm{s} \circ h] - \mathcal{S}_i^\star|;\boldsymbol{w}\right) \ . \tag{4}$$

Curiously, since we seek to measure overall regret, and regret is a nonnegative quantity with negative connotation, we always summarize it with a malfare function $\mathrm{M}\!\!\mathrm{M}(\cdot;\boldsymbol{w})$, even when we began with a *utility function*. Intuitively, this is because we can never hope to select a shared function $\hat{h}$ that group i prefers to $\boldsymbol{h}_i^\star$, thus *excess dissatisfaction* is always positive in both the loss and utility cases. In some sense, the malfare of regret thus measures the *price of sharing* in a society, as the shared model $\hat{h}$ is naturally compared [23] to letting each group select their own model $\hat{h}_i$.

Previous work summarizes regret across groups by taking the *largest regret* amongst them. This is analogous to game-theoretic regret (i.e., the maximum over agents of *utility differences* at *adjacent profiles*), but even there, *any malfare function* could reasonably aggregate per-group regret values. We argue that considering only *egalitarian regret* may act as an enforcer of the status quo, if one group is particularly happy with their $\boldsymbol{h}_i^\star$ and is thus aggrieved by any compromise — perhaps best summarized by the adage, "To those accustomed to privilege, equality feels like oppression." We mitigate this issue by summarizing regret with a power-mean malfare function $\mathrm{M}\!\!\mathrm{M}_p(\cdot;\boldsymbol{w})$, instead of the egalitarian malfare, in order to lessen the impact of the most aggrieved group. In particular, this class smoothly and nonlinearly interpolates between the worst-case (egalitarian) $\mathrm{M}\!\!\mathrm{M}_\infty(\cdot;\boldsymbol{w})$ regret and the utilitarian $\mathrm{M}_1(\cdot;\boldsymbol{w})$ welfare or malfare.

Fascinatingly, we find that utilitarian regret minimization reduces to utilitarian malfare or welfare optimization, as all terms involving per-group optimal sentiment can be factored into an additive constant from these objectives; observe

$$\mathrm{M}\!\!\mathrm{M}_1\left(i \mapsto \mathrm{Reg}_i(h);\boldsymbol{w}\right) = \mathrm{M}\!\!\mathrm{M}_1\left(i \mapsto |\mathop{\mathbb{E}}_{\mathcal{D}_i}[\mathrm{s} \circ h] - \mathcal{S}_i^\star|;\boldsymbol{w}\right)$$
$$= \left|\left(\sum_{i=1}^{g} \boldsymbol{w}_i \mathop{\mathbb{E}}_{\mathcal{D}_i}[\mathrm{s} \circ h]\right) - \boldsymbol{w} \cdot \mathcal{S}^\star\right|$$
$$= \begin{cases} \mathrm{s} = \ell & \mathrm{M}\!\!\mathrm{M}_1\left(\mathbb{E}_{\mathcal{D}_i}[\ell \circ h];\boldsymbol{w}\right) - \boldsymbol{w} \cdot \mathcal{S}^\star \\ \mathrm{s} = \mathrm{u} & \boldsymbol{w} \cdot \mathcal{S}^\star - \mathrm{W}_1\left(\mathbb{E}_{\mathcal{D}_i}[\mathrm{u} \circ h];\boldsymbol{w}\right) \end{cases} , \tag{5}$$

namely $\mathcal{S}^\star$ appears only in the additive constant $\boldsymbol{w} \cdot \mathcal{S}^\star$, which is independent of h. From this perspective, we conclude that while the utilitarian regret is not particularly interesting, the power-mean malfare of regret *interpolates between* minimizing largest regret, with its minority rule issues, and optimizing utilitarian welfare or malfare.

2.2 Three Sampling Models for Populations with Multiple Groups

In order to study efficient sampling, we must first quantify the difficulty or cost of a sampling-based estimation routine, which requires a *sampling model*. Within a single-group population, methods like *i.i.d. sampling*, *importance sampling*, or *sampling without replacement* are near-ubiquitous, and all can measure cost as *sample size* $m \in \mathbb{Z}_+$; however, in group-sensitive settings, we must consider how samples from different groups are obtained, and what the cost of collecting these samples is. In the context of this work, we don't argue for a one-size-fits-all solution, but rather we discuss three sampling models, and show that they fit key applications in the computer science domain and beyond.

(1) *Joint Sampling*: Each i.i.d. sample contains a piece of information for each of the g groups, with arbitrary dependencies *between groups*. For example, per-group representatives could be shown a shared $x \in \mathcal{X}$ and asked for their feedback, which

would then be used to establish some $\mathcal{Y}_i$ for each group i. Thus each sample is in the space $\mathcal{X} \times \mathcal{Y}^g$ if the $\mathcal{X}$ components are shared between groups, or more generally in $(\mathcal{X} \times \mathcal{Y})^g$. This setting also arises in *multi-objective reinforcement learning* [15, 44], as well as various bandit problems and empirical game theoretic analysis [48], where each query of an *action* or *strategy profile* yields a sample of the utility values of each *player*, *agent*, or *group*.

(2) *Mixture Sampling*: For each sample, the data are only relevant to a nonempty *subset of groups* $z \in 2^{\mathcal{Z}}$, thus samples are elements of $\mathcal{X} \times \mathcal{Y} \times 2^{\mathcal{Z}}$. This generality is useful for studying concepts like *intersectionalism* and *multicalibration* [39], where individuals may belong to multiple groups, but the case of *mutually exclusive* groups (i.e., each sample belongs to exactly one group) is also convenient [23]. This model naturally represents a mixed population being sampled i.i.d., where the group identities of the sample are left up to chance (i.e., roughly proportional to group frequencies), and is thus the most appropriate model for learning from [14] existing datasets with *group identity features* [21].

(3) *Conditional Sampling*: Here we *actively choose* from which group to sample, in contrast to the mixture sampling model, where we simply cast our net and "get what we get." In particular, we sample i.i.d. $(\mathcal{X}, \mathcal{Y})$ pairs *conditioned on* some group $z \in \mathcal{Z}$, thus we may select sample sizes $\boldsymbol{m}_{1:g} \in \mathbb{Z}_+^g$ and draw a sample $(\boldsymbol{x}, \boldsymbol{y}) \in (\mathcal{X} \times \mathcal{Y})^{\boldsymbol{m}_1} \times \cdots \times (\mathcal{X} \times \mathcal{Y})^{\boldsymbol{m}_g}$. This is a natural model in *active sampling* [1] and *scientific inquiry* settings, where initial results guide further study and resource expenditure, and similar conditional sampling structure arises in *stratified sampling* settings.

In mixture sampling, we generally assume *unit cost* $C = 1$ per sample, and in joint sampling, we assume *constant cost* $C > 1$ per joint-sample, as it is more expensive to set up a properly controlled joint sampling distribution. On the other hand, in conditional sampling, some groups may be more difficult or costly to study than others, so we assume a *cost model* $\boldsymbol{C}_{1:g} \in \mathbb{R}_+^g$, where $\boldsymbol{C}_i$ is the per-sample cost for group i, thus the total cost of a sample with per-group sizes $\boldsymbol{m}_{1:g}$ is $\boldsymbol{m} \cdot \boldsymbol{C}$. Note that the extra control of the conditional sampling model is extremely convenient and powerful, however it is generally more expensive than mixture sampling. These costs are entirely application dependent, so we take no stance on which is preferable, and rather focus on developing efficient learning algorithms under each sampling model.

3 STATISTICAL ANALYSIS AND ESTIMATION GUARANTEES

In this section, we discuss the statistics of estimating malfare and welfare functions. In particular, we assume a set $\mathcal{Z}$ of g groups, and we want to estimate the malfare, welfare, or regret of per-group expected loss or utility of some h, i.e.,

$$\hat{\mathrm{M}} \approx \mathrm{M}\big(i \mapsto \mathop{\mathbb{E}}_{\mathcal{D}_i}[\mathrm{s} \circ h]; \boldsymbol{w}\big) \ , \quad \text{or} \quad \hat{\mathrm{M}} \approx \mathrm{M}\big(i \mapsto \mathrm{Reg}_i(h); \boldsymbol{w}\big) \ ,$$

where $\mathcal{D}_{1:g}$ are distributions over $\mathcal{X} \times \mathcal{Y}$, and $\mathrm{M}(\cdot; \boldsymbol{w})$ generically represents some aggregator function. Estimating the expected loss or utility of one group is a well-studied sampling problem, but generalizing to the welfare, malfare, or regret of multiple groups introduces some subtleties. We start by noting that while the empirical mean is an unbiased estimator of expected utility or loss *of a single group*, in general there is no unbiased estimator of welfare or malfare (essentially due to their nonlinear nature, much like with the *standard deviation*). Thus rather than unbiased estimators, we seek *additive error* (AE) bounds of the form $\mathbb{P}\big(|\mathrm{M} - \hat{\mathrm{M}}| \leq \varepsilon\big) \geq 1-\delta$, where ε is the *confidence radius* (a.k.a. the *margin of error*), and δ is the *failure probability* (or, by alternative convention, $1 - \delta$ is the *level of confidence*).

In machine learning, we optimize over a *hypothesis class* $\mathcal{H} \subseteq \mathcal{X} \to \mathcal{Y}'$, thus we seek some sample-dependent $\hat{h} \in \mathcal{H}$ with true objective value within ε of that of the optimal $h^\star \in \mathcal{H}$. At times, we are also interested in related statistics, like the objective values of $\hat{h}$ and $h^\star$, and in general, tools to bound the deviations between the empirical and true objective values for any $h \in \mathcal{H}$ are sufficient to bound these quantities. The rest of this section pursues such bounds, assuming a fixed *failure probability* δ and *sample size* $\boldsymbol{m}_i$ for each group $i \in \mathcal{Z}$. In particular, section 3.1 reviews known results for uniformly estimating expectations across $\mathcal{H}$, section 3.2 builds upon these results to uniformly estimate malfare, welfare, and regret values, and section 3.3 then studies how varying per-group sentiment values and confidence radii impacts these bounds, and quantifies the incremental value of sampling from each group as a function of these quantities.

3.1 Uniform Convergence Bounds for Mean Estimation

In this work, the common functional form of our additive error (AE) bounds is *data dependent uniform convergence*, vectorized to operate over samples from multiple groups, rather than on a single-group sample. Occasionally, we are interested in the scalar form $\mathrm{AES}(m, \delta, \boldsymbol{x}, \boldsymbol{y}) : \mathbb{Z}_+ \times (0,1) \times \mathcal{X}^m \times \mathcal{Y}^m \to \mathbb{R}_{0+}$, which operates on a single group, but unless otherwise stated, we refer to the vector bound $\mathrm{AEV}(\boldsymbol{m}, \delta, \boldsymbol{x}, \boldsymbol{y}) : \mathbb{Z}_+^g \times (0,1) \times (\mathcal{X}^{\boldsymbol{m}_1} \times \cdots \times \mathcal{X}^{\boldsymbol{m}_g}) \times (\mathcal{Y}^{\boldsymbol{m}_1} \times \cdots \times \mathcal{Y}^{\boldsymbol{m}_g}) \to \mathbb{R}_{0+}$. In particular, given a sample $(\boldsymbol{x}, \boldsymbol{y}) \sim \mathcal{D}_1^{\boldsymbol{m}_1} \times \cdots \times \mathcal{D}_g^{\boldsymbol{m}_g}$, we require a *function*[2] $\mathrm{AEV}(\dots)$ such that

$$\hat{\boldsymbol{\varepsilon}} \leftarrow \mathrm{AEV}(\boldsymbol{m}, \delta, \boldsymbol{x}, \boldsymbol{y}) \implies \mathop{\mathbb{P}}_{\boldsymbol{x}, \boldsymbol{y}, \hat{\boldsymbol{\varepsilon}}}\left(\max_{i \in \mathcal{Z}} \sup_{h \in \mathcal{H}} \big| \mathop{\hat{\mathbb{E}}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\mathrm{s} \circ h] - \mathop{\mathbb{E}}_{\mathcal{D}_i}[\mathrm{s} \circ h]\big| - \hat{\boldsymbol{\varepsilon}}_i > 0\right) < \delta. \quad (6)$$

Section 3.2 explores how $\mathrm{AEV}(\dots)$ can be used to bound malfare, welfare, and regret, and the remainder of this subsection is dedicated to showing non-trivial bounds of this form for machine learning applications. All of our AE bounds assume *bounded sentiment range* $r \doteq \sup_{y' \in \mathcal{Y}', y \in \mathcal{Y}} \mathrm{s}(y', y)$, but this can usually be relaxed if we instead assume a *moment condition*, e.g., each $\mathrm{s} \circ h$ is sub-exponential, sub-gamma, sub-Poisson, or sub-Gaussian [8].

Data-dependent uniform convergence bounds, i.e., those of the form $\mathrm{AES}(m, \delta, \boldsymbol{x}, \boldsymbol{y})$, are invaluable for studying a population about which little is known. Such bounds require data to evaluate, thus we

[2] Going forward, we present only scalar bounds, but it is to be understood that given additive error *scalar bound* $\mathrm{AES}(\dots)$ and a finite group count g, we may construct the additive error *vector bound* $\mathrm{AEV}(\boldsymbol{m}, \delta, \boldsymbol{x}, \boldsymbol{y}) \leftarrow \big\langle \mathrm{AES}(\boldsymbol{m}_1, \frac{\delta}{g}, \boldsymbol{x}_{1,:}, \boldsymbol{y}_{1,:}), \dots, \mathrm{AES}(\boldsymbol{m}_g, \frac{\delta}{g}, \boldsymbol{x}_{g,:}, \boldsymbol{y}_{g,:})\big\rangle$ via the union bound.

cannot determine *a priori* how much data will be required to meet a given confidence radius. This contrasts distribution-free bounds, which must have worst-case dependence on the distribution, and take the form $\mathrm{AES}(m,\delta) \leq \sup_{(\boldsymbol{x},\boldsymbol{y})\in(\mathcal{X}\times\mathcal{Y})^m} \mathrm{AES}(m,\delta,\boldsymbol{x},\boldsymbol{y})$. In section 4, when constructing schedules for progressive sampling, we often assume knowledge of $\mathrm{AES}(m,\delta)$, but this is usually possible via this worst-case RHS bound. We first present simple bounds for *bounded finite hypothesis classes*, which depend on the *sentiment range* r, hypothesis class size $|\mathcal{H}|$, variances $\mathbb{V}[\cdot]$, and empirical variances $\hat{\mathbb{V}}[\cdot]$.

Theorem 3.1 (Uniform Convergence for Bounded Finite Hypothesis Classes)**.** theorem We may bound the distribution-free $\mathrm{AES}(m,\delta)$, the distribution-dependent $\mathrm{AES}_{\mathcal{D}}(m,\delta)$, and the data-dependent $\mathrm{AES}(m,\delta,\boldsymbol{x},\boldsymbol{y})$ scalar additive error as

(1) $\boldsymbol{\varepsilon} \leftarrow \sqrt{\frac{r^2 \ln \frac{2|\mathcal{H}|}{\delta}}{2m}}$ [27] ;

(2) $\boldsymbol{\varepsilon} \leftarrow \frac{r \ln \frac{2|\mathcal{H}|}{\delta}}{3m} + \sup_{h\in\mathcal{H}} \sqrt{\frac{2\,\mathbb{V}_{\mathcal{D}}[\mathrm{s}\circ h] \ln \frac{2|\mathcal{H}|}{\delta}}{m}}$ [6] ; &

(3) $\hat{\boldsymbol{\varepsilon}} \leftarrow \frac{7r \ln \frac{2|\mathcal{H}|+1}{\delta}}{3(m-1)} + \sup_{h\in\mathcal{H}} \sqrt{\frac{2\,\hat{\mathbb{V}}_{\boldsymbol{x},\boldsymbol{y}}[\mathrm{s}\circ h] \ln \frac{2|\mathcal{H}|+1}{\delta}}{(m-1)}}$ [16] .

Note that supremum variances and empirical variances are properties of the distribution and sample, respectively. Dependence on variance is necessary (similar terms appear in mean-estimation lower-bounds [19, 32]), however the $\ln|\mathcal{H}|$ union bound terms are loose, and the bounds are vacuous for infinite (continuous) $\mathcal{H}$. We now state results using Rademacher averages [5, 42] that tolerate infinite $\mathcal{H}$, while preserving the variance-dependence of item 2.

Theorem 3.2 (Uniform Convergence with Rademacher Averages)**.** theorem Suppose hypothesis class $\mathcal{H}$ and sentiment function $\mathrm{s}(\cdot,\cdot)$, take $(\boldsymbol{x},\boldsymbol{y}) \sim \mathcal{D}^m$ and $\boldsymbol{\sigma} \sim \mathcal{U}^m(\pm 1)$, i.e., $\boldsymbol{\sigma}$ is uniformly distributed on $(\pm 1)^m$, and define the *Rademacher average* $\mathcal{R}_m(\mathrm{s}\circ\mathcal{H},\mathcal{D})$ and *Bousquet variance proxy* $\mathcal{V}_m(\mathrm{s}\circ\mathcal{H},\mathcal{D})$ [see 9] as

$$\mathcal{R}_m(\mathrm{s}\circ\mathcal{H},\mathcal{D}) \doteq \mathop{\mathbb{E}}_{\boldsymbol{x},\boldsymbol{y},\boldsymbol{\sigma}}\left[\sup_{h\in\mathcal{H}} \left|\frac{1}{m}\sum_{i=1}^{m} \mathrm{s}\circ h(\boldsymbol{x}_i)\boldsymbol{\sigma}_i\right|\right] \ ,$$
$$\mathcal{V}_m(\mathrm{s}\circ\mathcal{H},\mathcal{D}) \doteq \sup_{h\in\mathcal{H}} \mathop{\mathbb{V}}_{\mathcal{D}}[\mathrm{s}\circ h] + 4r\mathcal{R}_m(\mathrm{s}\circ\mathcal{H},\mathcal{D}) \ . \tag{7}$$

We may then bound $\mathrm{AES}_{\mathcal{D}}(m,\delta)$ as $\boldsymbol{\varepsilon} \leftarrow 2\mathcal{R}_m(\mathrm{s}\circ\mathcal{H},\mathcal{D}) + \frac{r\ln\frac{1}{\delta}}{3m} + \sqrt{\frac{2\mathcal{V}_m(\mathrm{s}\circ\mathcal{H},\mathcal{D})\ln\frac{1}{\delta}}{m}}$.

Data-dependent analogues of theorem 3.2 are possible using *empirical Rademacher averages* and *variances* at no asymptotic cost [16]. In the worst case, theorem 3.2 performs comparably to theorem 3.1 item 2, however it improves when *correlations* exist between elements of $\mathcal{H}$, because the *effective size* of $\mathcal{H}$ is smaller for the purposes of realizing the supremum in the Rademacher average, see (7). The abstract inequalities of theorem 3.2 are quite opaque, so we now provide concrete bounds on the Rademacher averages of practical infinite hypothesis classes. The below results hold *for any distribution* $\mathcal{D}$, and are thus distribution-free, although similar distribution-dependent or data-dependent bounds are possible.

Property 3.3 (Practical Bounds on Rademacher Averages)**.** (1) Suppose $\mathcal{H}$ has Vapnik-Chervonenkis (VC) dimension d [46, 47], and $\ell(\hat{y},y) \doteq 1 - \mathbb{1}_y(\hat{y})$ is the 0-1 loss. Then for some absolute constant c, $\mathcal{R}_m(\ell\circ\mathcal{H},\mathcal{D}) \leq \sqrt{\frac{cd}{m}}$, which implies bounds for linear classifiers, bounded-depth decision trees [31], and many classes of neural network [2].

(2) Suppose $\mathcal{X} \doteq \{\vec{x} \in \mathbb{R}^\infty \mid \|\vec{x}\|_2 \leq R\}$ is the R-radius $\mathcal{L}_2$ ball in $\mathbb{R}^\infty$, $\mathcal{H} \doteq \{\vec{x} \mapsto \vec{w}\cdot\vec{x} \mid \|\vec{w}\|_2 \leq \gamma\}$ is a γ-regularized linear class, $\mathcal{Y} \doteq [-R\gamma, R\gamma]$, and $\ell(\cdot,\cdot)$ is a λ-Lipschitz loss function s.t. $\ell(y,y) = 0$. Then $r \leq 2\lambda R\gamma$ and $\mathcal{R}_m(\ell\circ\mathcal{H},\mathcal{D}) \leq \frac{2\lambda R\gamma}{\sqrt{m}}$. This implies bounds for (kernelized) SVM, generalized linear models [35], and bounded linear regression.

3.2 From Mean Estimation to Welfare, Malfare, and Regret Bounds

We now adapt the AE bounds of section 3.1 on expectations to bound malfare, welfare, and regret in terms of empirical estimates thereof. In particular, the strategy for each is to combine tail bounds for mean-estimation with the *monotonicity axiom* (definition 2.1 item 1) to bound the tails and expectations of our desiderata. We use the uniform convergence bounds of section 3.1 to bound the error of these estimates, thus we need only propagate this uncertainty through the appropriate aggregator functions. In general, aggregator functions are nonlinear, and optimizing over $\mathcal{H}$ results in estimation bias, thus the plug-in estimator is biased, however, we still obtain *tail bounds* on our objectives via AEV(. . .). Because the plug-in estimator is biased, we also consider various LCB-and-UCB-style estimates, which when optimized yield safer function choices and partially control for overfitting. Finally, in some cases, integrating over worst-case uncertainty from the tail bounds of AEV(. . .) yields convenient bounds on the expectation (and thus the bias) of the plug-in estimator.

Welfare and Malfare. Due to the lack of an unbiased estimator for welfare and malfare, we study the simple plug-in estimator $\hat{\mathrm{M}}$, as employed by [16], and introduce a pair of lower and upper estimators $(\hat{\mathrm{M}}^{\downarrow}, \hat{\mathrm{M}}^{\uparrow})$. In particular, we take

$$\hat{\mathrm{M}} \doteq \underbrace{\mathrm{M}\left(i \mapsto \hat{\mathop{\mathbb{E}}}_{\boldsymbol{x}_{i,:},\boldsymbol{y}_{i,:}}[\mathrm{s}\circ h]; \boldsymbol{w}\right)}_{\text{Plug-In Estimate}} \ ,$$
$$\hat{\mathrm{M}}^{\downarrow} \doteq \underbrace{\mathrm{M}\left(i \mapsto 0 \vee \hat{\mathop{\mathbb{E}}}_{\boldsymbol{x}_{i,:},\boldsymbol{y}_{i,:}}[\mathrm{s}\circ h] - \hat{\boldsymbol{\varepsilon}}_i; \boldsymbol{w}\right)}_{\text{LCB Estimate}} \ ,$$
$$\&\ \hat{\mathrm{M}}^{\uparrow} \doteq \underbrace{\mathrm{M}\left(i \mapsto r \wedge \hat{\mathop{\mathbb{E}}}_{\boldsymbol{x}_{i,:},\boldsymbol{y}_{i,:}}[\mathrm{s}\circ h] + \hat{\boldsymbol{\varepsilon}}_i; \boldsymbol{w}\right)}_{\text{UCB Estimate}} \ , \tag{8}$$

where $\vee$ and $\wedge$ are the (minimum precedence) infix binary max and min operators. By monotonicity (axiom 1), it holds that $\hat{\mathrm{M}}^{\downarrow} \leq \hat{\mathrm{M}} \leq \hat{\mathrm{M}}^{\uparrow}$. The *lower* and *upper confidence bound* estimates are convenient, both to show high probability bounds, and to sandwich the plug-in estimator, which we use to bound its bias. We first show tail bounds for the estimation of welfare and malfare in terms of their plug-in, LCB, and UCB estimates, and we then bound the bias of $\hat{\mathrm{M}}$.

Theorem 3.4 (Welfare and Malfare Tail Bounds)**.** theorem Suppose sentiment function $\mathrm{s}(\cdot,\cdot): \mathcal{Y}' \times \mathcal{Y} \to \mathbb{R}_{0+}$, per-group probability distributions $\mathcal{D}_{1:g}$, sample size vector $\boldsymbol{m} \in \mathbb{Z}_+^g$, samples $(\boldsymbol{x}, \boldsymbol{y}) \sim \mathcal{D}_1^{\boldsymbol{m}_1} \times \cdots \times \mathcal{D}_g^{\boldsymbol{m}_g}$, failure probability $\delta \in (0,1)$, and additive error bound $\mathrm{AEV}(\ldots)$, and let $\hat{\boldsymbol{\varepsilon}} \leftarrow \mathrm{AEV}(\boldsymbol{m}, \delta, \boldsymbol{x}, \boldsymbol{y})$. Then for all $h \in \mathcal{H}$ and all monotonic aggregator functions $\mathrm{M}(\cdot; \boldsymbol{w})$, it holds with probability at least $1-\delta$ over $\boldsymbol{x}, \boldsymbol{y}$, and $\hat{\boldsymbol{\varepsilon}}$ that

$$\underbrace{\mathrm{M}\left(i \mapsto 0 \vee \mathop{\mathbb{E}}_{\mathcal{D}_i}[\mathrm{s} \circ h] - \hat{\boldsymbol{\varepsilon}}_i; \boldsymbol{w}\right)}_{\text{True LB}} \leq \underbrace{\mathrm{M}\left(i \mapsto \mathop{\hat{\mathbb{E}}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\mathrm{s} \circ h]; \boldsymbol{w}\right)}_{\text{Plug-In Estimate } \hat{\mathrm{M}}} \leq \underbrace{\mathrm{M}\left(i \mapsto r \wedge \mathop{\mathbb{E}}_{\mathcal{D}_i}[\mathrm{s} \circ h] + \hat{\boldsymbol{\varepsilon}}_i; \boldsymbol{w}\right)}_{\text{True UB}} , \tag{9}$$

$$\& \underbrace{\mathrm{M}\left(i \mapsto 0 \vee \mathop{\hat{\mathbb{E}}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\mathrm{s} \circ h] - \hat{\boldsymbol{\varepsilon}}_i; \boldsymbol{w}\right)}_{\text{LCB Estimate } \hat{\mathrm{M}}^{\downarrow}} \leq \underbrace{\mathrm{M}\left(i \mapsto \mathop{\mathbb{E}}_{\mathcal{D}_i}[\mathrm{s} \circ h]; \boldsymbol{w}\right)}_{\text{True Aggregate M}} \leq \underbrace{\mathrm{M}\left(i \mapsto r \wedge \mathop{\hat{\mathbb{E}}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\mathrm{s} \circ h] + \hat{\boldsymbol{\varepsilon}}_i; \boldsymbol{w}\right)}_{\text{UCB Estimate } \hat{\mathrm{M}}^{\uparrow}} , \tag{10}$$

thus if $\mathrm{M}(\cdot; \boldsymbol{w})$ is λ-Lipschitz-continuous w.r.t. some norm $\|\cdot\|_{\mathrm{M}}$, we have

$$|\underbrace{\mathrm{M}\left(i \mapsto \mathop{\hat{\mathbb{E}}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\mathrm{s} \circ h]; \boldsymbol{w}\right)}_{\text{Plug-In Estimate}} - \underbrace{\mathrm{M}\left(i \mapsto \mathop{\mathbb{E}}_{\mathcal{D}_i}[\mathrm{s} \circ h]; \boldsymbol{w}\right)}_{\text{True Aggregate}}| \leq \lambda \|\hat{\boldsymbol{\varepsilon}}\|_{\mathrm{M}} . \tag{11}$$

From (10), we see that minimizing $\hat{\mathrm{M}}^{\uparrow}$ (or maximizing $\hat{\mathrm{M}}^{\downarrow}$) is in some sense a *safe choice*, as w.h.p. we can bound the true aggregate value in terms of the UCB or LCB. This idea is reminiscent of the *sample variance penalization* algorithm of [34], wherein ERM is supplanted by minimizing an *upper-bound* on risk; in that case with variance-dependent bounds, but here the bound depends on the structure of the malfare or welfare objective at hand. It should also be noted that while the final Lipschitz form (11) is concise and convenient for all Lipschitz-continuous aggregator functions (e.g., all $p \geq 1$ power-mean malfare functions, see theorem 2.2 item 3), it can be quite loose. For example, under $\pm$ uncertainty intervals, the egalitarian welfare $\mathrm{W}_{-\infty}(\langle 4\pm1, 9\pm8\rangle; \boldsymbol{w}) = \min(4\pm1, 9\pm8)$ must be on the interval 3 ± 2, despite (11) giving confidence radius 8. Thus while (11) is convenient for intuition and analysis, when possible (9) or (10) should be favored.

Theorem 3.5 (Welfare and Malfare Expectation Bounds)**.** theorem Suppose as in theorem 3.4, and assume also that $\mathrm{AEV}(\boldsymbol{m}, \delta, \boldsymbol{x}, \boldsymbol{y}) = \mathrm{AEV}(\boldsymbol{m}, \delta)$ is a *distribution-free* or *distribution-dependent* (but *not data-dependent*) bound. Then

$$|\,\mathrm{M} - \mathbb{E}[\hat{\mathrm{M}}]| \leq \mathbb{E}\left[\,|\mathrm{M} - \hat{\mathrm{M}}|\right] \leq \lambda \int_0^1 \|\mathrm{AEV}(\boldsymbol{m}, \delta)\|_{\mathrm{M}} \,\mathrm{d}\delta \ .$$

The above theorems give general recipes for bounding tails and expectations, so for demonstrative purposes, we instantiate them with theorem 3.1 for malfare estimation. Similar bounds can be derived for learning with theorem 3.2.

Theorem 3.6 (Bernstein-Type Malfare Bounds)**.** corollary-thmwmestexpbern

Suppose as in theorem 3.1, and also per-group sample size m (i.e., $\boldsymbol{m} = \langle m, \ldots, m\rangle$) and $p \geq 1$ power-mean malfare function $\mathrm{M}_p(\cdot; \boldsymbol{w})$. Now, let variance proxy v be defined in three cases as $v \doteq \mathrm{M}_{1/2}(\boldsymbol{v}; \boldsymbol{w}) = \left(\sum_{i=1}^g \boldsymbol{w}_i \sqrt{\boldsymbol{v}_i}\right)^2$ for $p = 1$, $v \doteq \boldsymbol{w} \cdot \boldsymbol{v}$ for $p \in (1,2]$, or $v \doteq \|\boldsymbol{v}\|_\infty$ for $p > 2$. Then for all $\delta \in (0,1)$, we have

(1) $\mathbb{P}\left(|\mathrm{M} - \hat{\mathrm{M}}| \geq \frac{r \ln \frac{g}{\delta}}{3m} + \sqrt{\frac{2v \ln \frac{g}{\delta}}{m}}\right) \leq \delta$;
(2) $\mathbb{E}[|\mathrm{M} - \hat{\mathrm{M}}|] \leq \frac{r \ln(2eg)}{3m} + \sqrt{\frac{2v \ln(2eg)}{m}}$; &
(3) $\mathrm{M} \leq \mathbb{E}[\hat{\mathrm{M}}] \leq \mathrm{M} + \frac{r \ln(eg)}{3m} + \sqrt{\frac{2v \ln(eg)}{m}}$.

Estimating the Malfare of Regret. Regret is difficult to bound, as it depends both on the expected sentiment of the selected $\hat{h}$, and also on $\mathcal{H}$ through the (unknown) per-group optimal sentiments $\mathcal{S}^\star_{1:g}$. We thus introduce the estimators

$$\hat{\mathcal{S}}_i \doteq \inf_{h \in \mathcal{H}} \mathop{\hat{\mathbb{E}}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\ell \circ h] \ , \quad \text{or} \quad \hat{\mathcal{S}}_i \doteq \sup_{h \in \mathcal{H}} \mathop{\hat{\mathbb{E}}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\mathrm{u} \circ h] \ , \tag{12}$$

for loss or utility, respectively, cf. (2). By analogy with (3), the plug-in estimator for the regret malfare minimizer is then

$$\hat{h} \doteq \operatorname*{argmin}_{h \in \mathcal{H}} \mathrm{M}\left(i \mapsto |\mathop{\hat{\mathbb{E}}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\mathrm{s} \circ h] - \hat{\mathcal{S}}_i|; \boldsymbol{w}\right) \ . \tag{13}$$

The following theorem bounds the difference between the true and empirical malfare of regret.

Theorem 3.7 (Regret Estimation Bounds)**.** theorem Suppose sentiment function $\mathrm{s}(\cdot,\cdot): \mathcal{Y}' \times \mathcal{Y} \to \mathbb{R}_{0+}$, per-group probability distributions $\mathcal{D}_{1:g}$, sample size vector $\boldsymbol{m} \in \mathbb{Z}_+^g$, samples $(\boldsymbol{x}, \boldsymbol{y}) \sim \mathcal{D}_1^{\boldsymbol{m}_1} \times \cdots \times \mathcal{D}_g^{\boldsymbol{m}_g}$, failure probability $\delta \in (0,1)$, and additive error bound $\mathrm{AEV}(\ldots)$, and let $\hat{\boldsymbol{\varepsilon}} \leftarrow \mathrm{AEV}(\boldsymbol{m}, \delta, \boldsymbol{x}, \boldsymbol{y})$. Then for all $h \in \mathcal{H}$ and all monotonic malfare functions $\mathrm{M}(\cdot; \boldsymbol{w})$, it holds with probability at least $1-\delta$ over $\boldsymbol{x}, \boldsymbol{y}$, and $\hat{\boldsymbol{\varepsilon}}$ that

$$\underbrace{\mathrm{M}\left(i \mapsto 0 \vee \left|\mathop{\mathbb{E}}_{\mathcal{D}_i}[\mathrm{s} \circ h] - \mathcal{S}^\star_i\right| - 2\hat{\boldsymbol{\varepsilon}}_i; \boldsymbol{w}\right)}_{\text{True Regret Malfare LB}} \leq \underbrace{\mathrm{M}\left(i \mapsto \left|\mathop{\hat{\mathbb{E}}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\mathrm{s} \circ h] - \hat{\mathcal{S}}_i\right|; \boldsymbol{w}\right)}_{\text{Plug-In Regret Malfare}} \leq \underbrace{\mathrm{M}\left(i \mapsto r \wedge \left|\mathop{\mathbb{E}}_{\mathcal{D}_i}[\mathrm{s} \circ h] - \mathcal{S}^\star_i\right| + 2\hat{\boldsymbol{\varepsilon}}_i; \boldsymbol{w}\right)}_{\text{True Regret Malfare UB}} , \tag{14}$$

$$\& \underbrace{\mathrm{M}\left(i \mapsto 0 \vee \left|\mathop{\hat{\mathbb{E}}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\mathrm{s} \circ h] - \hat{\mathcal{S}}_i\right| - 2\hat{\boldsymbol{\varepsilon}}_i; \boldsymbol{w}\right)}_{\text{LCB Estimate}} \leq \underbrace{\mathrm{M}\left(i \mapsto \left|\mathop{\mathbb{E}}_{\mathcal{D}_i}[\mathrm{s} \circ h] - \mathcal{S}^\star_i\right|; \boldsymbol{w}\right)}_{\text{True Regret Malfare}} \leq \underbrace{\mathrm{M}\left(i \mapsto r \wedge \left|\mathop{\hat{\mathbb{E}}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\mathrm{s} \circ h] - \hat{\mathcal{S}}_i\right| + 2\hat{\boldsymbol{\varepsilon}}_i; \boldsymbol{w}\right)}_{\text{UCB Estimate}} , \tag{15}$$

thus if $\mathrm{M}(\cdot; \boldsymbol{w})$ is λ-Lipschitz-continuous w.r.t. some norm $\|\cdot\|_{\mathrm{M}}$, we have

$$\left| \underbrace{\mathrm{M}\left(i \mapsto \left| \hat{\mathbb{E}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\mathrm{s} \circ h] - \hat{\mathcal{S}}_i \right|; \boldsymbol{w}\right)}_{\text{Plug-In Regret Malfare}} - \underbrace{\mathrm{M}\left(i \mapsto \left| \mathbb{E}_{\mathcal{D}_i}[\mathrm{s} \circ h] - \mathcal{S}^\star_i \right|; \boldsymbol{w}\right)}_{\text{True Regret Malfare}} \right| \leq 2\lambda \|\hat{\varepsilon}\|_{\mathrm{M}} . \quad (16)$$

Note that similar bounds on the *expectation* of the regret plug-in estimator can be shown along the lines of theorem 3.6, *mutatis mutandis* for regret. Note also that theorem 3.7 matches theorem 3.4 up to a 2-factor attached to the confidence radius, thus in some sense regret is "about twice as difficult" to estimate as malfare or welfare.

3.3 Information Asymmetry and Where Best to Sample

An intuitive notion of fairness would suggest that we should draw equally-sized samples for each group, or perhaps samples proportional to population frequencies. If the goal is to optimize or bound welfare, malfare, or regret, such intuitive notions should be rejected, as they are critically flawed. We now discuss the ways in which samples drawn from one group or another may be more or less valuable to for the purposes of estimating or optimizing these objectives.

As a brief thought experiment, suppose we want to estimate the egalitarian welfare of a population consisting of two groups. Suppose also that their utilities are similarly difficult to estimate, and their expected utilities are $\langle 1, 10 \rangle$. In such a setting, nearly all sampling effort should be invested in estimating the utility of group one, as once group two is estimated to within ±9 AE, there is *no further benefit* to improving their estimate. Thus the optimal sampling strategy depends on the true expected utilities, the difficulties of estimating utilities for each group, and the objective in question, and *in no way* resembles the naïve uniform or proportional so-called "fair sampling strategies" described above. We argue that such naïve strategies are dangerous, as they introduce subtle biases and fairness issues, but the rationale for alternative sampling strategies is only apparent through the lens of sample complexity.

We now ask the questions, "Given a sample, what do we need to obtain sharper bounds?" and "How much will bounds improve with a larger sample?" We begin with a soft discussion as to why samples from different groups may contribute more or less information to an estimate, which we measure as the improvement to tail bounds that additional samples may yield. In particular, for malfare, we discuss the improvement to upper bounds, but the entire discussion can be directly translated to welfare and lower bounds in the usual manner. We then quantify these factors mathematically, and we develop these ideas further in section 4.2, where they are used to adaptively choose from which group to sample.

(1) *Variable estimation difficulty or overfit potential*: Often it is inherently more difficult to give bounds on the expected sentiment for some groups than for others. This can be due to differences in variances (see theorem 3.1) or in uniform convergence bounds (see theorem 3.2), and in general, occurs when $\hat{\boldsymbol{\varepsilon}} \leftarrow \mathrm{AEV}(\dots)$ has $\hat{\boldsymbol{\varepsilon}}_i \ll \hat{\boldsymbol{\varepsilon}}_j$, even while $\boldsymbol{m}_i \approx \boldsymbol{m}_j$.
(2) *Variable task difficulty*: Some groups may be inherently easier or harder to satisfy than others; e.g., regression and classification problems are generally easier for groups with labels that are more homogeneous, and regret varies with the optimal expected sentiment $\mathcal{S}^\star_i$. This is crucial, because most malfare and welfare functions are more sensitive[3] to high-risk or low-utility groups, thus the ease of satisfying a group effects their impact on malfare and welfare values.
(3) *Aggregator function interactions*: Complicated interactions also occur through the malfare or welfare function. When learning over $\mathcal{H}$, the set of near-optimal functions is more relevant than those that are clearly bad choices overall, and groups that tend to be mutually satisfied (i.e., are correlated) are less impactful to the overall objective. Weight values in malfare or welfare functions may also differ between groups, and higher-weighted groups are usually more impactful.

Quantifying the Incremental Value of Sampling. We measure the impact of sampling by asking the question, "What is the incremental value of a single sample drawn for some group?" In particular, we quantify the value of the sample as the *reduction in uncertainty*, as measured by the infimum UCB (over $\mathcal{H}$), and although this is inherently a discrete question, we approximate the answer for the power-mean malfare with tools from calculus of infinitesimals.

Note that all such analysis is necessarily heuristic, as we fundamentally cannot answer this question without more information: it is precisely because we are trying to *estimate unknown means* that we can't know how the samples we draw will impact the *empirical means*. For now, we heuristically assume that our estimated expectations are reasonably accurate, and consider what will happen as tail bounds sharpen with additional samples. The strategy we thus employ is to make a reasonable guess as to how sampling might impact the UCB by assuming that the empirical mean will not be strongly affected, and all confidence intervals over m samples will contract at a $\Theta\sqrt{\frac{1}{m}}$ rate.

Property 3.8 (Incremental Gain of Sampling)**.** Suppose power-mean malfare $\mathrm{M}_p(\cdot; \boldsymbol{w})$, sample $(\boldsymbol{x}, \boldsymbol{y})$ with group sample sizes $\boldsymbol{m}_{1:g}$, and let $\boldsymbol{x}', \boldsymbol{y}'$ extend $\boldsymbol{x}, \boldsymbol{y}$ to sample sizes $\boldsymbol{m}'$, where $\boldsymbol{m}' = \boldsymbol{m} + \mathbb{1}_i$, i.e., group i has one additional sample. Now, let $\hat{\boldsymbol{\varepsilon}} \leftarrow \mathrm{AEV}(\boldsymbol{m}, \delta, \boldsymbol{x}, \boldsymbol{y})$ and $\tilde{\boldsymbol{\varepsilon}} \leftarrow \mathrm{AEV}(\boldsymbol{m}', \delta, \boldsymbol{x}', \boldsymbol{y}')$, and take $\hat{h} \doteq \operatorname{argmin}_{h \in \mathcal{H}} \mathrm{M}_p\left(i \mapsto \hat{\mathbb{E}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\ell \circ h] + \hat{\boldsymbol{\varepsilon}}_i; \boldsymbol{w}\right)$, $\hat{\mathrm{M}} \doteq \mathrm{M}_p\left(i \mapsto \hat{\mathbb{E}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\ell \circ \hat{h}]; \boldsymbol{w}\right)$, $\hat{\mathrm{M}}^\uparrow \doteq \mathrm{M}_p\left(i \mapsto \hat{\mathbb{E}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\ell \circ \hat{h}] + \hat{\boldsymbol{\varepsilon}}_i; \boldsymbol{w}\right)$, and $\tilde{\mathrm{M}}^\uparrow \doteq \inf_{h \in \mathcal{H}} \mathrm{M}_p\left(i \mapsto \hat{\mathbb{E}}_{\boldsymbol{x}'_{i,:}, \boldsymbol{y}'_{i,:}}[\ell \circ h] + \tilde{\boldsymbol{\varepsilon}}_i; \boldsymbol{w}\right)$.

[3] In particular, this holds for all $p \neq 1$ power means, and is axiomatically justified by the *Pigou-Dalton transfer principle* (definition 2.1 item 7).

Then the *incremental impact* of sampling from group i on the UCB is approximately

$$\hat{\text{M}}^{\uparrow} - \tilde{\text{M}}^{\uparrow} \approx \frac{\hat{\boldsymbol{\varepsilon}}_i \boldsymbol{w}_i}{2\boldsymbol{m}_i + \frac{3}{2}} \left(\frac{\hat{\mathbb{E}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\ell \circ \hat{h}] + \hat{\boldsymbol{\varepsilon}}_i}{\hat{\text{M}}^{\uparrow}} \right)^{p-1}$$
$$\approx \frac{\hat{\boldsymbol{\varepsilon}}_i \boldsymbol{w}_i}{2\boldsymbol{m}_i} \left(\frac{\hat{\mathbb{E}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\ell \circ \hat{h}]}{\hat{\text{M}}} \right)^{p-1} . \quad (17)$$

Observe that (17) characterizes the knowledge gain of sampling from group i. This gain is *proportional to* the current bound radius $\hat{\boldsymbol{\varepsilon}}_i$, the group weight $\boldsymbol{w}_i$, and the $(p-1)$th power of the ratio of the UCB risk of group i to the UCB malfare, i.e., $(\hat{\mathbb{E}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\ell \circ \hat{h}] + \hat{\boldsymbol{\varepsilon}}_i / \hat{\text{M}}^{\uparrow})^{p-1}$, and *inversely proportional* to the amount of effort $\boldsymbol{m}_i$ already put forth into studying group i. These terms line up with the soft arguments at the top of section 3.3 as to where sampling should occur, but it is only via precisely studying sample complexity and estimation error that we gain quantifiable mathematical insight. In particular, the weight term $\boldsymbol{w}_i$ appears directly, and $\frac{\hat{\boldsymbol{\varepsilon}}_i}{\boldsymbol{m}_i}$ captures both the difficulty of estimating this group, and also the diminishing incremental improvement produced by further sampling. The ratio between the risk of group i and the malfare then captures how important group i is relative to the other groups, and this term being raised to the $(p-1)$th power nonlinearly adjusts its impact; higher p saturate high-risk groups, tending towards egalitarianism, whereas in the $p = 1$ (utilitarian) case, this term is 1. Finally, for optimization problems, the dependence on $\hat{h}$ captures other dependencies; namely the behavior of $\text{M}(\cdot; \boldsymbol{w})$ *near the optimal* $h \in \mathcal{H}$ is what matters.

This analysis parallels concerns in *stratified sampling*, wherein subpopulations are sampled individually, generally to produce an improved mean estimator. In particular, we suggest a form of *disproportionate allocation*, i.e., per-group sample sizes are not necessarily proportional to their population frequencies. Rather than simply considering *variances* to estimate *means*, we holistically consider the objective and uncertainty over various quantities, thus our sample-size selection-strategy is a variant of the *minimax sampling ratio* [41] method. Chen et al. [12] also suggest disproportionate allocation in fair learning, albeit only for bounding differences of per-group statistics. Similar concerns also arise in *optimizing* minimax-fair models, wherein Abernethy et al. [1] present an algorithm that takes gradient steps to improve a model for the highest-risk group, though it is unclear whether such methods generalize beyond the egalitarian case.

4 PROGRESSIVE AND ACTIVE SAMPLING ALGORITHMS

Section 3 considers *fixed sample sizes* $\boldsymbol{m}_{1:g}$ and *failure probabilities* δ, and bounds the *confidence radius* ε. In this section, we want a fixed ε-δ AE guarantee, but we are willing to let an algorithm select the sample size m (or per-group sample sizes $\boldsymbol{m}_{1:g}$). In particular, due to the cost of sampling and processing data, we want our algorithm to minimize m (or cost measured as some function of $\boldsymbol{m}$), while constraining ε and δ to user-supplied levels. Some cases are simpler than others; the *joint sampling model* yields a standard progressive sampling method with a fixed sampling schedule, and the method under *mixture sampling* is similar, except a subtle conditioning argument allows us to use variably-sized per-group sample sizes based on the order groups are sampled in. For the *conditional sampling model*, we develop an *active sampling* approach, which makes cost-sensitive decisions as to which group to sample at each iteration. More details on sampling schedules and other aspects of our progressive sampling algorithms are given in section 7.

We can't simply draw samples one-by-one, compute bounds using $\hat{\boldsymbol{\varepsilon}} \leftarrow \text{AEV}(\boldsymbol{m}, \delta, \boldsymbol{x}, \boldsymbol{y})$ after each sample, and terminate when a sufficiently sharp bound is available, because the possibility of *early termination* leads to the *multiple comparisons problem*, wherein by chance the desired confidence radius is met at some timestep. Progressive sampling algorithms correct for this by establishing a *sampling schedule* $\boldsymbol{s}$ and *failure probability schedule* $\boldsymbol{\delta}$, which usually dictate that, at timestep t, we take a tail-bound with $\delta = \boldsymbol{\delta}_t$ and sample size $\boldsymbol{s}_t$, while ensuring that all bounds hold simultaneously (by union bound) with probability at least $1 - \delta$. Due to this union bound, it is *inefficient* to take bounds after drawing every sample. Furthermore, for technical reasons, we henceforth assume a few mild regularity conditions:

1. The *sampling schedule* $\boldsymbol{s} \in \mathbb{Z}_+^\infty$ is a strictly monotonically increasing sequence, i.e., for all $t \in \mathbb{Z}_+$, $\boldsymbol{s}_t \leq \boldsymbol{s}_{t+1}$;
2. The *failure probability schedule* $\boldsymbol{\delta} \in [0, 1)^\infty$ is a sequence that sums to some $\delta \in (0, 1)$, i.e., $\sum_{i=1}^\infty \boldsymbol{\delta}_t = \|\boldsymbol{\delta}\|_1 = \delta$; &
3. The *distribution-free bound* $\sup_{\boldsymbol{x}, \boldsymbol{y}} \|\text{AEV}(\boldsymbol{m}, \delta, \boldsymbol{x}, \boldsymbol{y})\|$ is monotonically decreasing in $\boldsymbol{m}_{1:g}$ and δ for any norm $\|\cdot\|$.

In order to prove that a progressive sampling algorithm produces a (probabilistically) correct answer, it is crucial to show that it does not loop indefinitely. We now introduce ε-convergent schedules, which require all sentiment values to eventually be ε-δ estimated w.r.t. some norm $\|\cdot\|_\text{M}$, yielding welfare, malfare, or regret bounds via theorems 3.4 and 3.7.

Definition 4.1 (ε-Uniformly-Convergent Schedule)**.** For any $\varepsilon \geq 0$, a sampling schedule $\boldsymbol{s}$ and failure probability schedule $\boldsymbol{\delta}$ are ε-uniformly-convergent w.r.t. AEV(. . .) and some norm $\|\cdot\|_\text{M}$ if

$$\inf_{t \in \mathbb{Z}_+} \sup_{(\boldsymbol{x}, \boldsymbol{y}) \in (\mathcal{X} \times \mathcal{Y})^{\boldsymbol{s}_t \times g}} \|\text{AEV}(\langle \boldsymbol{s}_t, \ldots, \boldsymbol{s}_t \rangle, \boldsymbol{\delta}_t, \boldsymbol{x}, \boldsymbol{y})\|_\text{M} \leq \varepsilon \ . \quad (18)$$

Intuitively, definition 4.1 captures the idea that no matter how unlucky we are with the sampled $\boldsymbol{x}, \boldsymbol{y}$, if AEV(. . .) bounds tails once for each timestep t of the schedule, with per-group samples of size at least $\boldsymbol{s}_t$ and failure probability $\boldsymbol{\delta}_t$, then at some point an ε-estimate of the objective will be produced. Note that neither data-dependent AEV(. . .) bounds on sentiment values, nor sufficient per-group error radii to estimate the objective, are known *a priori*, thus it is not always possible to select a *sufficient* static sample size, however, definition 4.1 is *more flexible*, as it requires *only the existence* of a (possibly unknown) sufficient sample size. Even when a sufficient sample size is known, unless it is also *necessary*, progressive sampling is usually more sample-efficient, often terminating closer to the *necessary sample size*.

With this definition in hand, we now construct finite ε-, and infinite 0-, uniformly-convergent schedules. In the context of this work (see theorems 4.5 and 4.6), the finite schedule can be employed with a Lipschitz-continuous objective and an *a priori* known distribution-free bound on AEV(. . .), and when the objective is continuous but not Lipschitz-continuous, or the class $\mathcal{H}$ is uniformly-convergent at an unknown rate, the infinite schedule can still be used. Both are based on geometrically-increasing sample sizes, which are efficient because they never "overshoot" any sample size by more than a constant factor, yet they cover an exponentially large range of sample sizes in a linear number of timesteps.

Definition 4.2 (Geometric-Uniform Schedule)**.** Suppose *optimistic size* $\alpha \geq 1$, *common ratio* $\beta > 1$, and *schedule length* $T \in \mathbb{Z}_+$. The geometric-uniform schedule then takes (geometric) $\boldsymbol{s}_t \doteq \lceil \alpha\beta^t \rceil$ and (uniform) $\boldsymbol{\delta}_t \doteq \frac{\delta}{T}\mathbb{1}_{1,\ldots,T}(t)$.

Definition 4.3 (Double-Geometric Schedule)**.** Suppose *optimistic size* $\alpha > 0$ and *common ratio* $\beta > 1$. The double-geometric schedule then takes (geometric) $\boldsymbol{s}_t \doteq \lceil \alpha\beta^t \rceil$ and (geometric) $\boldsymbol{\delta}_t \doteq \frac{\delta(\beta-1)}{\beta^t}$.

Lemma 4.4 (Sufficient Conditions for Uniformly-Convergent Geometric Schedules)**.** Suppose as in definition 4.2, and assume

$$\sup_{(\boldsymbol{x},\boldsymbol{y})\in(\mathcal{X}\times\mathcal{Y})^{s_T\times g}} \|\mathrm{AEV}\left(\langle s_T,\ldots,s_T\rangle, \tfrac{\delta}{T}, \boldsymbol{x}, \boldsymbol{y}\right)\|_{\mathrm{M}} \leq \varepsilon \ . \tag{19}$$

Then the geometric-uniform schedule $(\boldsymbol{s},\ \boldsymbol{\delta})$ is ε-uniformly-convergent. Furthermore, suppose as in definition 4.3, $\alpha \geq \frac{1}{\delta}$, and

$$\lim_{m\to\infty} \sup_{(\boldsymbol{x},\boldsymbol{y})\in(\mathcal{X}\times\mathcal{Y})^{m\times g}} \|\mathrm{AEV}\left(\langle m,\ldots,m\rangle, \tfrac{\beta-1}{\beta(m+1)}, \boldsymbol{x}, \boldsymbol{y}\right)\|_{\mathrm{M}} = 0 \ . \tag{20}$$

Then the double-geometric schedule $(\boldsymbol{s},\ \boldsymbol{\delta})$ is 0-uniformly-convergent.

The initial and final sample sizes of the geometric-uniform schedule are $\boldsymbol{s}_1 = \lceil \alpha\beta \rceil$ and $\boldsymbol{s}_T = \lceil \alpha\beta^T \rceil$, and often one can set s_1/β and $\boldsymbol{s}_T$ to *minimal sufficient* and *maximal necessary* sample sizes (as a function of T, the objective, and other parameters). To maximize *statistical efficiency* while controlling the value of β, we may select the minimal T such that $\lceil \log_\beta \frac{s_T}{s_1} \rceil = T$.[4] In particular, assuming a λ-Lipschitz objective, the Hoeffding (item 1) and empirical Bernstein (item 3) bounds of theorem 3.1 imply ε-uniformly convergent schedules via (19) of length $T \in \Theta(\log \frac{\lambda r}{\varepsilon})$. For the double-geometric schedule, we may similarly set s_1/β to a minimal sufficient sample size, and here there is no T parameter (the schedule is infinite), thus we may simply select β as desired. This yields 0-uniformly convergent schedules, since each of the bounds of theorem 3.1 satisfy (20), as do those of theorem 3.2, so long as $\lim_{m\to\infty}\max_{i\in\mathcal{Z}}\mathcal{R}_m(\mathrm{s}\circ\mathcal{H},\mathcal{D}_i) = 0$.

Both of the above schedule types are *efficient*, in the sense that for the smallest (per-group) static sample size $m^\star$ at which we obtain the bound $\boldsymbol{\varepsilon}^\star$, some $\hat{m} \leq \lceil \beta m^\star \rceil$ is contained in the schedule, and the bound $\hat{\boldsymbol{\varepsilon}} \leftarrow \mathrm{AEV}(\langle \hat{m},\ldots,\hat{m}\rangle,\ldots)$ exceeds $\boldsymbol{\varepsilon}^\star$ only because it uses a smaller δ value. In particular, assuming all bounds are $\Theta\sqrt{u}$ for $u \doteq \ln\frac{g}{\delta}$, we have for each group i that $\frac{\varepsilon_i^\star}{\hat{\varepsilon}_i} \in \Theta\sqrt{\frac{u}{\log(T)+u}}$ and $\frac{\varepsilon_i^\star}{\hat{\varepsilon}_i} \in \Theta\sqrt{\frac{u}{\log(m^\star)+u}}$ for the geometric-uniform and double-geometric schedules, respectively. Note also that $\log(T) \in \Theta(\log\log\frac{r\lambda}{\varepsilon})$, whereas $\log(m^\star) \in \mathrm{O}(\log\frac{r\lambda u}{\varepsilon})$, thus the geometric-uniform schedule is preferable, unless $m^\star$ is *exponentially smaller* than the above bound, e.g., if $\lambda = \infty$, or if a nonlinear objective is more stable to perturbations of each $\mathcal{S}_i$ about its optimum than the Lipschitz constant λ would indicate.

4.1 The Linear Progressive Sampling Algorithm

The core of *linear progressive sampling* (algorithm 1) is quite simple. At timestep $t = 1$, we *guess* that a sample of size $\boldsymbol{s}_1$ for all groups will be sufficient to ε-δ optimize the objective, we draw at least such a sample (7 for joint sampling, or lines 9–12 for mixture sampling), compute tail bounds (line 14), then determine the UCB-optimal $\hat{h}$ (line 16). If our bounds indicate that $\hat{h}$ is provably near-optimal, algorithm 1 terminates, otherwise, our guess was incorrect, so we increment t, draw at least $\boldsymbol{s}_t$ samples (per-group), and repeat. The basic principle is quite flexible, so algorithm 1 can maximize welfare or minimize malfare of risk or regret via the LinearPSLoss(...) and LinearPSUtility(...) routines.

Theorem 4.5 shows that algorithm 1 learns an optimal $h \in \mathcal{H}$ to within user-specified ε-δ AE. We require only monotonicity (axiom 1) and continuity (axiom 3) of $\mathrm{M}(\cdot;\boldsymbol{w})$, though the power-mean malfare family is convenient, as Lipschitz-continuity (thm. 2.2 item 3) permits efficient ε-uniformly-convergent schedules (def. 4.2). NB this result generalizes to welfare objectives, *mutatis mutandis* (flipping infima and suprema), via the negation reduction of lines 25–30.

Theorem 4.5 (Linear PS Guarantees)**.** Suppose $(\hat{h},\hat{\mu},\hat{\varepsilon},\mathrm{M}^{\star\downarrow}) \leftarrow$ LinearPSLoss($\mathcal{H}$, $\ell(\cdot,\cdot)$, $\mathcal{D}$, AEV(...), $\boldsymbol{s}$, $\boldsymbol{\delta}$, ε, $\mathrm{M}(\cdot;\boldsymbol{w})$, Reg), $\mathrm{M}(\mathcal{S};\boldsymbol{w})$ is continuous and monotonic in $\mathcal{S}$ with (possibly infinite) Lipschitz constant λ_{M} w.r.t. $\|\cdot\|_{\mathrm{M}}$, and the schedules $(\boldsymbol{s},\boldsymbol{\delta})$ are $\frac{\varepsilon}{\lambda_{\mathrm{M}}(1+\mathbb{1}_{\mathrm{Reg}})}$-uniformly-convergent w.r.t. AEV(...) and $\|\cdot\|_{\mathrm{M}}$. Now take μ to be the true objective value of $\hat{h}$ and $\mu^\star$ to be the true objective value of the optimal $h^\star$, i.e., if Reg = False, take $\mu \doteq \mathrm{M}(i \mapsto \mathbb{E}_{\mathcal{D}_i}[\ell\circ\hat{h}];\boldsymbol{w})$ and $\mu^\star \doteq \inf_{h\in\mathcal{H}}\mathrm{M}(i \mapsto \mathbb{E}_{\mathcal{D}_i}[\ell\circ h];\boldsymbol{w})$, or if Reg = True, take (see section 3) $\mu \doteq \mathrm{M}(i \mapsto \mathrm{Reg}_i(\hat{h});\boldsymbol{w})$ and $\mu^\star \doteq \inf_{h\in\mathcal{H}} \mathrm{M}\left(i \mapsto \mathrm{Reg}_i(h);\boldsymbol{w}\right)$. Then, with probability at least $1-\delta$, the output $(\hat{h},\hat{\mu},\hat{\varepsilon},\mathrm{M}^{\star\downarrow})$ obeys

(1) $|\hat{\mu}-\mu| \leq \hat{\varepsilon} \leq \varepsilon$; &
(2) $\mathrm{M}^{\star\downarrow} \leq \mu^\star \leq \mu \leq \hat{\mu}+\hat{\varepsilon} \leq \mathrm{M}^{\star\downarrow}+2\varepsilon$.

4.2 The Braided Progressive Sampling Algorithm

Under the joint and mixture sampling models (algorithm 1), progress is *linear* (i.e., *sequential*, as no decisions are made except when to terminate); we begin with (at least) $\boldsymbol{s}_1$ samples per group, and advance until we reach a sufficient sample size to terminate with the desired guarantee. For the conditional sampling model, we present *braided progressive sampling* (algorithm 2), which is *actively making decisions*, thus linear analysis is not applicable. At each *iteration* (line 6) of algorithm 2, a group index $i \in \mathcal{Z}$ is chosen (line 16) to optimize an estimate of knowledge-gain via logic similar to that of section 3.3 (due to space limitations, the details are deferred to section 7.2), and group i is sampled for one additional *timestep* (line 17), i.e., the sample associated with group i is extended from size $\boldsymbol{s}_{\boldsymbol{t}_i}$ to $\boldsymbol{s}_{1+\boldsymbol{t}_i}$, where $\boldsymbol{t}_i$ denotes the current timestep for group i. The remainder of algorithm 2 is essentially the same as algorithm 1; after sampling, we optimize (line 9) a UCB-optimal $\hat{h}$, bound the

[4]The base-β logarithm arises intuitively, as the *number of times* the sample size must increase by a factor β to reach $\boldsymbol{s}_T$ from $\boldsymbol{s}_1$.

Algorithm 1 Fair Learning with Linear Progressive Sampling under the Joint and Mixture Sampling Models

1: **procedure** LinearPSLoss($\mathcal{H}, \ell(\cdot,\cdot), \mathcal{D}, \mathrm{AEV}(\dots), \boldsymbol{s}, \boldsymbol{\delta}, \varepsilon, \mathrm{M}(\cdot;\boldsymbol{w}), \mathrm{Reg}) \to (\hat{h}, \hat{\mu}, \hat{\varepsilon}, \mathrm{M}^{\star\downarrow})$
2: **input**: Hypothesis class $\mathcal{H} \subseteq \mathcal{X} \to \mathcal{Y}'$, loss function $\ell(\cdot,\cdot) : \mathcal{Y}' \times \mathcal{Y} \to [0, c]$, joint or mixture distribution $\mathcal{D}$, additive error vector bound $\mathrm{AEV}(\boldsymbol{m}, \delta, \boldsymbol{x}, \boldsymbol{y})$, schedule $\boldsymbol{s} \in \mathbb{Z}_+^\infty$ and $\boldsymbol{\delta} \in [0,1)^\infty$, confidence radius ε, weighted malfare $\mathrm{M}(\cdot;\boldsymbol{w})$, and Boolean Reg
3: **output**: Empirically UCB-optimal $\hat{h}$, empirical malfare estimate $\hat{\mu}$, confidence radius $\hat{\varepsilon}$, and lower bound on minimal malfare $\mathrm{M}^{\star\downarrow}$
4: $\boldsymbol{m}_{1:g} \leftarrow \boldsymbol{0}$; $\boldsymbol{x}_{1:g} \leftarrow \langle\langle\rangle, \dots, \langle\rangle\rangle$; $\boldsymbol{y}_{1:g} \leftarrow \langle\langle\rangle, \dots, \langle\rangle\rangle$ ▷ Initialize per-group sample counts, empty per-group sample lists
5: **for** $t \in 1, 2, \dots$ **do** ▷ Progressive sampling timesteps
6: **if** $\mathcal{D}$ is joint sampler **then**
7: $(\boldsymbol{x}_{1:g, s_{t-1}+1:s_t}, \boldsymbol{y}_{1:g, s_{t-1}+1:s_t}) \sim \mathcal{D}^{s_t - s_{t-1}}$; $\forall i \in \mathcal{Z}: \boldsymbol{m}_i \leftarrow s_t$ ▷ Sample from joint distribution (assume $s_0 = 0$)
8: **else if** $\mathcal{D}$ is mixture sampler **then**
9: **while** $\min_i \boldsymbol{m}_i < s_t$ **do**
10: $(x, y, z) \sim \mathcal{D}$ ▷ Draw $\mathcal{X} \times \mathcal{Y} \times 2^{\mathcal{Z}}$ triplet (domain, codomain, groups)
11: $\forall i \in z: \boldsymbol{m}_i \leftarrow \boldsymbol{m}_i + 1$; $(\boldsymbol{x}_{i,\boldsymbol{m}_i}, \boldsymbol{y}_{i,\boldsymbol{m}_i}) \leftarrow (x, y)$ ▷ Increment counts and store samples for each group i associated with (x, y)
12: **end while**
13: **end if**
14: $\hat{\boldsymbol{\varepsilon}}_{1:g} \leftarrow (1 + \mathbb{1}_{\mathrm{Reg}})\, \mathrm{AEV}(\boldsymbol{m}, \boldsymbol{\delta}_t, \boldsymbol{x}, \boldsymbol{y})$ ▷ Bound additive error of per-group supremum deviations (w.h.p.)
15: $\forall i \in \mathcal{Z}: \hat{\mathcal{S}}_i \leftarrow \big(\inf_{h \in \mathcal{H}} \hat{\mathbb{E}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\ell \circ h]\big)$ **if** Reg **else** 0 ▷ Set regret baseline of per-group minimal empirical risks (or 0 if ¬Reg)
16: $\hat{h} \leftarrow \operatorname{argmin}_{h \in \mathcal{H}} \mathrm{M}\big(i \mapsto c \wedge \hat{\mathbb{E}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\ell \circ h] - \hat{\mathcal{S}}_i + \hat{\boldsymbol{\varepsilon}}_i; \boldsymbol{w}\big)$ ▷ Compute UCB-optimal $\hat{h}$
17: $\mathrm{M}^{\star\downarrow} \leftarrow \inf_{h \in \mathcal{H}} \mathrm{M}\big(i \mapsto 0 \vee \hat{\mathbb{E}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\ell \circ h] - \hat{\mathcal{S}}_i - \hat{\boldsymbol{\varepsilon}}_i; \boldsymbol{w}\big)$ ▷ Lower-bound optimal $\mathrm{M}^\star$
18: $(\hat{\mathrm{M}}^\downarrow, \hat{\mathrm{M}}^\uparrow) \leftarrow \Big(\mathrm{M}\big(i \mapsto 0 \vee \hat{\mathbb{E}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\ell \circ \hat{h}] - \hat{\mathcal{S}}_i - \hat{\boldsymbol{\varepsilon}}_i; \boldsymbol{w}\big), \mathrm{M}\big(i \mapsto c \wedge \hat{\mathbb{E}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}}[\ell \circ \hat{h}] - \hat{\mathcal{S}}_i + \hat{\boldsymbol{\varepsilon}}_i; \boldsymbol{w}\big)\Big)$ ▷ LCB and UCB on $\hat{h}$ (regret) malfare
19: **if** $\hat{\mathrm{M}}^\uparrow \leq \mathrm{M}^{\star\downarrow} + 2\varepsilon$ **then** ▷ Check if desired error guarantee is met (termination condition)
20: $(\hat{\mu}, \hat{\varepsilon}) \leftarrow \Big(\frac{1}{2}(\hat{\mathrm{M}}^\downarrow + \hat{\mathrm{M}}^\uparrow), \frac{1}{2}(\hat{\mathrm{M}}^\uparrow - \hat{\mathrm{M}}^\downarrow)\Big)$ ▷ Symmetric estimate $\hat{\mu}$ and confidence radius $\hat{\varepsilon}$ of (regret) malfare of $\hat{h}$
21: **return** $(\hat{h}, \hat{\mu}, \hat{\varepsilon}, \mathrm{M}^{\star\downarrow})$ ▷ Return UCB-optimal $\hat{h}$, $\hat{\varepsilon}$-estimate of $\mathrm{M}(\cdot;\boldsymbol{w})$, and lower-bound on optimal malfare $\mathrm{M}^{\star\downarrow}$
22: **end if**
23: **end for**
24: **end procedure**
25: **procedure** LinearPSUtility($\mathcal{H}, \mathrm{u}(\cdot,\cdot), \mathcal{D}, \mathrm{AEV}(\dots), \boldsymbol{s}, \boldsymbol{\delta}, \varepsilon, \mathrm{M}(\cdot;\boldsymbol{w}), \mathrm{Reg}) \to (\hat{h}, \hat{\mu}, \hat{\varepsilon}, \mathrm{M}^{\star\uparrow})$
26: **input**: Utility function $\mathrm{u}(\cdot,\cdot) : \mathcal{Y}' \times \mathcal{Y} \to [0, c]$, weighted aggregator function $\mathrm{M}(\cdot;\boldsymbol{w})$ (malfare if Reg, otherwise welfare), see line 2
27: **output**: Empirically LCB-optimal $\hat{h}$, empirical welfare $\hat{\mu}$, confidence radius $\hat{\varepsilon}$, and UB on maximal welfare $\mathrm{M}^{\star\uparrow}$ (*or similar for regret*)
28: $(\hat{h}, \hat{\mu}, \hat{\varepsilon}, \mathrm{M}^{\star\uparrow}) \leftarrow$ LinearPSLoss($\mathcal{H}, c - \mathrm{u}(\cdot,\cdot), \mathcal{D}, \mathrm{AEV}(\dots), \boldsymbol{s}, \boldsymbol{\delta}, \varepsilon, (2\mathbb{1}_{\mathrm{Reg}} - 1)\mathrm{M}(\mathcal{S}_i \mapsto c - \mathcal{S}_i; \boldsymbol{w}), \mathrm{Reg})$ ▷ Negate to flip inf and sup
29: **return** $(\hat{h}, r - \hat{\mu}, \hat{\varepsilon}, (2\mathbb{1}_{\mathrm{Reg}} - 1)\mathrm{M}^{\star\uparrow})$
30: **end procedure**

objective (lines 10–11), and terminate if the user supplied guarantee is met, otherwise we continue.

There is thus a *lattice* of possible sample size vectors $\boldsymbol{m}$. To avoid a union bound over this (exponentially large) lattice, we analyze the method as a *braid*, in that g progressive sampling sequences are *concurrently active*, and at each iteration we select some group i, and advance the schedule by one timestep for only group i. Consequently, we must use (line 17) the additive error *scalar bound* $\hat{\boldsymbol{\varepsilon}}_i \leftarrow \mathrm{AES}\big(\boldsymbol{m}_i, \frac{\boldsymbol{\delta}_{t_i}}{g}, \boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}\big)$, i.e., we operate on *one group at a time*, rather than over all groups as in the linear algorithm (algorithm 1 line 14). Similar analysis is employed for *multi-armed bandits*, where a union bound is taken over all timesteps and each arm being sampled. We now show correctness of algorithm 2.

Theorem 4.6 (Braided PS Guarantees)**.** Suppose $(\hat{h}, \hat{\mu}, \hat{\varepsilon}, \mathrm{M}^{\star\downarrow}) \leftarrow$ BraidedPSLoss($\mathcal{H}$, $\ell(\cdot, \cdot)$, $\mathcal{D}$, AES($\dots$), $\boldsymbol{s}$, $\boldsymbol{\delta}$, ε, $\mathrm{M}(\cdot; \boldsymbol{w})$, Reg), $\mathrm{M}(\mathcal{S}; \boldsymbol{w})$ is continuous and strictly monotonic in $\mathcal{S}$ with (possibly infinite) Lipschitz constant λ_{M} w.r.t. $\|\cdot\|_{\mathrm{M}}$, and the schedules $(\boldsymbol{s}, \boldsymbol{\delta})$ are $\frac{\varepsilon}{\lambda_{\mathrm{M}}(1+\mathbb{1}_{\mathrm{Reg}})}$-uniformly-convergent w.r.t. $\|\cdot\|_{\mathrm{M}}$ and the *additive error vector bound* $\mathrm{AEV}(\boldsymbol{m}, \delta, \boldsymbol{x}, \boldsymbol{y}) \leftarrow \big\langle \mathrm{AES}(\boldsymbol{m}_1, \frac{\delta}{g}, \boldsymbol{x}_1, \boldsymbol{y}_1), \dots, \mathrm{AES}(\boldsymbol{m}_g, \frac{\delta}{g}, \boldsymbol{x}_g, \boldsymbol{y}_g)\big\rangle$. Now take μ to be the true objective value of $\hat{h}$ and $\mu^\star$ to be the true objective value of the optimal $h^\star$ (see theorem 4.5). Then, with probability at least $1 - \delta$, we have

(1) $|\hat{\mu} - \mu| \leq \hat{\varepsilon} \leq \varepsilon$; &
(2) $\mathrm{M}^{\star\downarrow} \leq \mu^\star \leq \mu \leq \hat{\mu} + \hat{\varepsilon} \leq \mathrm{M}^{\star\downarrow} + 2\varepsilon$.

5 CONCLUSION

This work generalizes existing theories of fair machine learning, with welfare, malfare, and regret objectives, thus subsuming the *minimax fair learning* [1, 20, 30, 33, 43], *multi-group agnostic PAC learning* [7, 39], and *fair-PAC learning* [14] settings, while enjoying rigorous statistical learning guarantees and the axiomatization of cardinal welfare theory. In particular, we bound the *generalization error* and *sample complexity* of UCB-optimal models, either given a fixed sample, or to meet a user-supplied ε-δ optimality guarantee via progressive sampling. Our bounds leverage the specific character of the objective at hand, and our progressive sampling methods are tailored to three realistic models of data generation. We stress that while training UCB-optimal models is analytically convenient, there is also an important fairness impact to this decision, as fair malfare functions (e.g., egalitarian) place strong emphasis on the most disadvantaged groups, which are often understudied minority groups. Cousins [14] notes that optimizing *empirical malfare* $\hat{\mathrm{M}}$

Algorithm 2 Fair Learning with Braided Progressive Sampling under the Conditional Sampling Model

procedure $\textsc{BraidedPSLoss}(\mathcal{H}, \ell(\cdot,\cdot), \mathcal{D}_{1:g}, \boldsymbol{C}_{1:g}, \mathrm{AES}(\dots), \boldsymbol{s}, \boldsymbol{\delta}, \varepsilon, \mathrm{M}(\cdot;\boldsymbol{w}), \textsc{Reg}) \rightarrow (\hat{h}, \hat{\mu}, \hat{\varepsilon}, \mathrm{M}^{\star\downarrow})$

  input: Hypothesis class $\mathcal{H}$, loss function $\ell(\cdot,\cdot)$, per-group distributions $\mathcal{D}_{1:g}$, cost model $\boldsymbol{C}_{1:g} \in \mathbb{R}_+^g$, additive error scalar bound $\mathrm{AES}(m, \delta, \boldsymbol{x}, \boldsymbol{y})$, schedule $\boldsymbol{s} \in \mathbb{Z}_+^\infty$ and $\boldsymbol{\delta} \in [0,1)^\infty$, confidence radius ε, weighted malfare $\mathrm{M}(\cdot;\boldsymbol{w})$, and Boolean $\textsc{Reg}$

  output: Empirically UCB-optimal $\hat{h}$, empirical malfare estimate $\hat{\mu}$, confidence radius $\hat{\varepsilon}$, and lower bound on minimal malfare $\mathrm{M}^{\star\downarrow}$

  $\boldsymbol{t}_{1:g} \leftarrow \mathbf{1}$ ▹ Initialize per-group timestep indices

  $\forall i \in \mathcal{Z}: \ (\boldsymbol{x}_{i,1:s_1}, \boldsymbol{y}_{i,1:s_1}) \sim \mathcal{D}_i^{s_1}; \hat{\boldsymbol{\varepsilon}}_i \leftarrow (1 + \mathbb{1}_{\textsc{Reg}})\, \mathrm{AES}(\boldsymbol{s}_1, \frac{\delta_1}{g}, \boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:})$ ▹ Draw initial sample for all groups & bound error

  loop ▹ Loop over braided algorithm iterations

    $\forall i, j \in \mathcal{Z}, t \in \mathbb{Z}_+ : \ \tilde{\boldsymbol{\varepsilon}}_{j,t}^{(i)} \leftarrow \left(\hat{\boldsymbol{\varepsilon}}_j \ \textbf{if} \ i \neq j \ \textbf{else} \ \hat{\boldsymbol{\varepsilon}}_j \sqrt{\frac{s_{t_j} \ln \frac{g}{\delta_{t_j}}}{s_{t+t_j} \ln \frac{g}{\delta_{t+t_j}}}} \right)$ ▹ Estimate of $\hat{\boldsymbol{\varepsilon}}_j$ after sampling group i for t more iterations

    $\forall i \in \mathcal{Z}: \ \hat{\mathcal{S}}_i \leftarrow \left(\inf_{h \in \mathcal{H}} \hat{\mathbb{E}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}} [\ell \circ h] \right) \ \textbf{if} \ \textsc{Reg} \ \textbf{else} \ 0$ ▹ Set regret baseline of per-group minimal empirical risks (or 0 if $\neg\textsc{Reg}$)

    $\hat{h} \leftarrow \operatorname*{argmin}_{h \in \mathcal{H}} \mathrm{M}\big(i \mapsto c \wedge \hat{\mathbb{E}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}} [\ell \circ h] - \hat{\mathcal{S}}_i + \hat{\boldsymbol{\varepsilon}}_i; \boldsymbol{w}\big)$ ▹ Compute UCB-optimal $\hat{h}$

    $\mathrm{M}^{\star\downarrow} \leftarrow \inf_{h \in \mathcal{H}} \mathrm{M}\big(i \mapsto 0 \vee \hat{\mathbb{E}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}} [\ell \circ h] - \hat{\mathcal{S}}_i - \hat{\boldsymbol{\varepsilon}}_i; \boldsymbol{w}\big)$ ▹ Lower-bound optimal $\mathrm{M}^\star$

    $(\hat{\mathrm{M}}^\downarrow, \hat{\mathrm{M}}^\uparrow) \leftarrow \Big(\mathrm{M}\big(i \mapsto 0 \vee \hat{\mathbb{E}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}} [\ell \circ \hat{h}] - \hat{\mathcal{S}}_i - \hat{\boldsymbol{\varepsilon}}_i; \boldsymbol{w}\big), \mathrm{M}\big(i \mapsto c \wedge \hat{\mathbb{E}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}} [\ell \circ \hat{h}] - \hat{\mathcal{S}}_i + \hat{\boldsymbol{\varepsilon}}_i; \boldsymbol{w}\big) \Big)$ ▹ LCB and UCB on $\hat{h}$ (regret) malfare

    if $\hat{\mathrm{M}}^\uparrow \leq \mathrm{M}^{\star\downarrow} + 2\varepsilon$ **then** ▹ Check if desired error guarantee is met (termination condition)

      $(\hat{\mu}, \hat{\varepsilon}) \leftarrow \left(\frac{1}{2}(\hat{\mathrm{M}}^\downarrow + \hat{\mathrm{M}}^\uparrow), \frac{1}{2}(\hat{\mathrm{M}}^\uparrow - \hat{\mathrm{M}}^\downarrow) \right)$ ▹ Symmetric estimate of $\hat{\mu}$ of malfare or regret of $\hat{h}$

      return $(\hat{h}, \hat{\mu}, \hat{\varepsilon}, \mathrm{M}^{\star\downarrow})$

    end if

    $i \leftarrow \operatorname*{argmax}_{i \in \mathcal{Z}} \sup_{t \in \mathbb{Z}_+} \underbrace{\frac{1}{C_i(\boldsymbol{s}_{t+t_i} - \boldsymbol{s}_{t_i})}}_{\textsc{Reciprocal Cost}} \underbrace{\left(\hat{\mathrm{M}}^\uparrow - \mathrm{M}\Big(j \mapsto c \wedge \hat{\mathbb{E}}_{\boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:}} [\ell \circ h] - \hat{\mathcal{S}}_i + \tilde{\boldsymbol{\varepsilon}}_{j,t}^{(i)}; \boldsymbol{w} \Big) \right)}_{\textsc{Estimated (Regret) Malfare Improvement}}$ ▹ Maximize improvement:cost ratio

    $(\boldsymbol{x}_{i,1+s_{t_i}:s_{1+t_i}}, \boldsymbol{y}_{i,1+s_{t_i}:s_{1+t_i}}) \sim \mathcal{D}_i^{s_{1+t_i} - s_{t_i}}; \boldsymbol{t}_i \leftarrow \boldsymbol{t}_i + 1; \hat{\boldsymbol{\varepsilon}}_i \leftarrow (1 + \mathbb{1}_{\textsc{Reg}})\, \mathrm{AES}(\boldsymbol{s}_{t_i}, \frac{\delta_{t_i}}{g}, \boldsymbol{x}_{i,:}, \boldsymbol{y}_{i,:})$ ▹ Sample group i & bound error

  end loop

end procedure

overfits to small numbers of sampled minorities, however we argue that training UCB-optimal models (i.e., optimizing $\hat{\mathrm{M}}^\uparrow$) factors *uncertainty* into training, so that the needs of understudied groups (i.e., those with large $\hat{\boldsymbol{\varepsilon}}_i$ values) are better addressed.

Our active learning setting under the conditional sampling model is philosophically intriguing, as we find that optimally investing sampling effort under uncertainty is challenging, depends on the objective at hand, and has important fairness impact. In section 3.3, we see that a host of factors involving the objective, function class $\mathcal{H}$, and per-group distributions $\mathcal{D}_{1:g}$ all interact to determine the sharpness of welfare, malfare, and regret bounds, and property 3.8 quantifies the incremental UCB improvement of sampling each group. This analysis answers questions raised by Chen et al. [12] as to how sampling-error impacts fairness, and generalizes the analysis of Shekhar et al. [43] from the egalitarian special-case to arbitrary power-mean malfare functions. Algorithm 2 then incorporates these ideas into an *active sampling algorithm*, which dynamically select groups to sample based on projected UCB improvement. Notably, algorithm 1 does use uniform sample sizes under the joint sampling model, and uses whatever is available under the mixture sampling model, as these are natural choices under these sampling models. In contrast, under the conditional sampling model, algorithm 2 is able to make more intelligent decisions as to where to allocate sampling effort.

We thus conclude that (welfare-centric) fairness, statistical uncertainty, and sample complexity analysis are tightly intertwined, and must all be considered to best allocate resources in service of the social planner. We are hopeful that this analysis and algorithmic study will lead to a greater emphasis on sample complexity and finite sample analysis for the social planner's problem, which is traditionally analyzed in terms of the asymptotic Bayesian methods of classical economics. In particular, we are hopeful that this analysis emphasizes and mathematically supports the call for greater visibility of minority groups and the importance of incorporating diverse data into (fair) machine learning systems.

ACKNOWLEDGMENTS

This work was supported in part by NSF award RI-1813444 and DARPA/AFRL grant FA8750.

REFERENCES

[1] Jacob Abernethy, Pranjal Awasthi, Matthäus Kleindessner, Jamie Morgenstern, Chris Russell, and Jie Zhang. 2020. Active Sampling for Min-Max Fairness. *arXiv preprint arXiv:2006.06879* (2020).

[2] Martin Anthony and Peter L Bartlett. 2009. *Neural network learning: Theoretical foundations*. Cambridge University Press.

[3] Ahmed Ashraf, Shehroz Khan, Nikhil Bhagwat, Mallar Chakravarty, and Babak Taati. 2018. Learning to unlearn: Building immunity to dataset bias in medical imaging studies. *Machine Learning for Health (ML4H) Workshop at Advances at Neural Information Processing Systems 31* (2018).

[4] Kristine Bærøe, Torbjørn Gundersen, Edmund Henden, and Kjetil Rommetveit. 2022. Can medical algorithms be fair? Three ethical quandaries and one dilemma. *BMJ Health & Care Informatics* 29, 1 (2022).

[5] Peter L Bartlett and Shahar Mendelson. 2002. Rademacher and Gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research* 3, Nov (2002), 463–482.

[6] George Bennett. 1962. Probability inequalities for the sum of independent random variables. *J. Amer. Statist. Assoc.* 57, 297 (1962), 33–45.

[7] Avrim Blum and Thodoris Lykouris. 2020. Advancing Subgroup Fairness via Sleeping Experts. In *Innovations in Theoretical Computer Science Conference (ITCS)*, Vol. 11.

[8] Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. 2013. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press.

[9] Olivier Bousquet. 2002. A Bennett concentration inequality and its application to suprema of empirical processes. *Comptes Rendus Mathematique* 334, 6 (2002), 495–500.

[10] Joy Buolamwini and Timnit Gebru. 2018. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency*. PMLR, 77–91.

[11] Jacqueline G Cavazos, P Jonathon Phillips, Carlos D Castillo, and Alice J O'Toole. 2020. Accuracy comparison across face recognition algorithms: Where are we on measuring race bias? *IEEE Transactions on Biometrics, Behavior, and Identity Science* (2020).
[12] Irene Chen, Fredrik D Johansson, and David Sontag. 2018. Why is my classifier discriminatory? *Advances in Neural Information Processing Systems* 31 (2018).
[13] Cynthia M Cook, John J Howard, Yevgeniy B Sirotin, Jerry L Tipton, and Arun R Vemury. 2019. Demographic effects in facial recognition and their dependence on image acquisition: An evaluation of eleven commercial systems. *IEEE Transactions on Biometrics, Behavior, and Identity Science* 1, 1 (2019), 32–41.
[14] Cyrus Cousins. 2021. An Axiomatic Theory of Provably-Fair Welfare-Centric Machine Learning. In *Advances in Neural Information Processing Systems*.
[15] Cyrus Cousins, Kavosh Asadi, and Michael L. Littman. 2022. Fair E3: Efficient Welfare-Centric Fair Reinforcement Learning. In *5th Multidisciplinary Conference on Reinforcement Learning and Decision Making*. RLDM.
[16] Cyrus Cousins and Matteo Riondato. 2020. Sharp uniform convergence bounds through empirical centralization. *Advances in Neural Information Processing Systems* 33 (2020).
[17] Hugh Dalton. 1920. The measurement of the inequality of incomes. *The Economic Journal* 30, 119 (1920), 348–361.
[18] Matthew DeCamp and Charlotta Lindvall. 2020. Latent bias and the implementation of artificial intelligence in medicine. *Journal of the American Medical Informatics Association* 27, 12 (2020), 2020–2023.
[19] Luc Devroye, Matthieu Lerasle, Gabor Lugosi, Roberto I Oliveira, et al. 2016. Sub-Gaussian mean estimators. *The Annals of Statistics* 44, 6 (2016), 2695–2725.
[20] Emily Diana, Wesley Gill, Michael Kearns, Krishnaram Kenthapadi, and Aaron Roth. 2021. Minimax group fairness: Algorithms and experiments. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*. 66–76.
[21] Frances Ding, Moritz Hardt, John Miller, and Ludwig Schmidt. 2021. Retiring adult: New datasets for fair machine learning. *Advances in Neural Information Processing Systems* 34 (2021).
[22] Virginie Do and Nicolas Usunier. 2022. Optimizing generalized Gini indices for fairness in rankings. In *45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2022)*.
[23] Cynthia Dwork, Nicole Immorlica, Adam Tauman Kalai, and Mark D. M. Leiserson. 2018. Decoupled Classifiers for Group-Fair and Efficient Machine Learning. In *Conference on Fairness, Accountability and Transparency, FAT 2018, 23–24 February 2018, New York, NY, USA (Proceedings of Machine Learning Research, Vol. 81)*. PMLR, 119–133.
[24] Clare Garvie, Alvaro Bedoya, and Jonathan Frankle. 2016. The Perpetual Line-Up. Unregulated Police Face Recognition in America. https://www.perpetuallineup.org
[25] Hoda Heidari, Solon Barocas, Jon M Kleinberg, and Karen Levy. 2021. On Modeling Human Perceptions of Allocation Policies with Uncertain Outcomes. In *EC '21: The 22nd ACM Conference on Economics and Computation, Budapest, Hungary, July 18–23, 2021*. ACM, 589–609.
[26] Hoda Heidari, Claudio Ferrari, Krishna Gummadi, and Andreas Krause. 2018. Fairness behind a veil of ignorance: A welfare analysis for automated decision making. In *Advances in Neural Information Processing Systems*. 1265–1276.
[27] Wassily Hoeffding. 1963. Probability Inequalities for Sums of Bounded Random Variables. *J. Amer. Statist. Assoc.* 58, 301 (1963), 13–30.
[28] Lily Hu and Yiling Chen. 2020. Fair classification and social welfare. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*. 535–545.
[29] Rucha Kulkarni, Ruta Mehta, and Setareh Taki. 2021. Indivisible Mixed Manna: On the Computability of MMS + PO Allocations. In *Proceedings of the 22nd ACM Conference on Economics and Computation*. 683–684.
[30] Preethi Lahoti, Alex Beutel, Jilin Chen, Kang Lee, Flavien Prost, Nithum Thain, Xuezhi Wang, and Ed Chi. 2020. Fairness without demographics through adversarially reweighted learning. *Advances in neural information processing systems* 33 (2020), 728–740.
[31] Jean-Samuel Leboeuf, Frédéric LeBlanc, and Mario Marchand. 2020. Decision trees as partitioning machines to characterize their generalization properties. *Advances in Neural Information Processing Systems* 33 (2020).
[32] Gábor Lugosi and Shahar Mendelson. 2019. Mean estimation and regression under heavy-tailed distributions: A survey. *Foundations of Computational Mathematics* 19, 5 (2019), 1145–1190.
[33] Natalia Martinez, Martin Bertran, and Guillermo Sapiro. 2020. Minimax pareto fairness: A multi objective perspective. In *International Conference on Machine Learning*. PMLR, 6755–6764.
[34] Andreas Maurer and Massimiliano Pontil. 2009. Empirical Bernstein Bounds and Sample-Variance Penalization. In *International Conference on Computational Learning Theory*. Springer.
[35] John Ashworth Nelder and Robert WM Wedderburn. 1972. Generalized linear models. *Journal of the Royal Statistical Society: Series A (General)* 135, 3 (1972), 370–384.
[36] Arthur Cecil Pigou. 1912. *Wealth and welfare*. Macmillan and Company, limited.
[37] Kevin WS Roberts. 1980. Interpersonal comparability and social choice theory. *The Review of Economic Studies* (1980), 421–439.
[38] Esther Rolf, Max Simchowitz, Sarah Dean, Lydia T Liu, Daniel Björkegren, Moritz Hardt, and Joshua Blumenstock. 2020. Balancing Competing Objectives with Noisy Data: Score-Based Classifiers for Welfare-Aware Machine Learning. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13–18 July 2020, Virtual Event (Proceedings of Machine Learning Research, Vol. 119)*. PMLR, 8158–8168.
[39] Guy N Rothblum and Gal Yona. 2021. Multi-group Agnostic PAC Learnability. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18–24 July 2021, Virtual Event (Proceedings of Machine Learning Research, Vol. 139)*. PMLR, 9107–9115.
[40] Amartya Sen. 1977. On weights and measures: Informational constraints in social welfare analysis. *Econometrica: Journal of the Econometric Society* (1977), 1539–1572.
[41] Mohammad Shahrokh Esfahani and Edward R Dougherty. 2014. Effect of separate sampling on classification accuracy. *Bioinformatics* 30, 2 (2014), 242–250.
[42] Shai Shalev-Shwartz and Shai Ben-David. 2014. *Understanding machine learning: From theory to algorithms*. Cambridge University Press.
[43] Shubhanshu Shekhar, Greg Fields, Mohammad Ghavamzadeh, and Tara Javidi. 2021. Adaptive Sampling for Minimax Fair Classification. *Advances in Neural Information Processing Systems* 34 (2021).
[44] Umer Siddique, Paul Weng, and Matthieu Zimmer. 2020. Learning Fair Policies in Multi-Objective (Deep) Reinforcement Learning with Average and Discounted Rewards. In *International Conference on Machine Learning*. PMLR, 8905–8915.
[45] Till Speicher, Hoda Heidari, Nina Grgić-Hlača, Krishna P Gummadi, Adish Singla, Adrian Weller, and Muhammad Bilal Zafar. 2018. A unified approach to quantifying algorithmic unfairness: Measuring individual & group unfairness via inequality indices. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 2239–2248.
[46] Vladimir Vapnik and Aleksei Chervonenkis. 1968. The uniform convergence of frequencies of the appearance of events to their probabilities. In *Doklady Akademii Nauk*, Vol. 181. Russian Academy of Sciences, 781–783.
[47] Vladimir Vapnik and Aleksei Chervonenkis. 1971. On the uniform convergence of relative frequencies of events to their probabilities. *Theory of Probability and its Applications* 16, 2 (1971), 264–281.
[48] Enrique Areyan Viqueira, Cyrus Cousins, and Amy Greenwald. 2020. Improved Algorithms for Learning Equilibria in Simulation-Based Games. In *Proceedings of the 19th International Conference on Autonomous Agents and MultiAgent Systems*. 79–87.