A Statistical Framework for Algorithmic Collective Action with Multiple Collectives

Claudio Battiloro[*,1,6] **Pietro Greiner**[*,2,5] **Dario Rancati**[*,3] **Bret Nestor**[4] **Oumaima Amezgar**[5] **Francesca Dominici**[1]

[1]Brown University [2]LawZero [3]Institute of Science and Technology Austria
[4]University of British Columbia [5]Université de Montréal and Mila [6]Harvard University

[*]Equal contribution. Correspondence to: claudio_battiloro@brown.edu

Abstract

As learning systems increasingly shape everyday decisions, Algorithmic Collective Action (ACA), i.e., users coordinating changes to shared data to steer model behavior, offers a complement to regulator-side policy and corporate model design. Real-world collective actions have traditionally been decentralized and fragmented into multiple collectives, despite sharing overarching objectives, with each collective differing in size, strategy, and actionable goals. However, most of the ACA literature focuses on single collective settings. To address this, we propose the first comprehensive statistical framework for ACA with multiple collectives acting on the same system. In particular, we focus on collective action in classification, studying how multiple collectives can influence a classifier's behavior. We provide quantitative statistical bounds on the success of the collectives, considering the role and the interplay of the collectives' sizes and the alignment of their goals. We make such bounds computable by each collective with only partial knowledge of other collectives' sizes and strategies. Finally, we numerically illustrate our framework on simulations inspired by interventions for climate adaptation in smart cities, demonstrating the usefulness of our bounds, and we verify that the bounds hold for classifiers trained on collectively edited synthetic and US Census data.

1 Introduction

AI's rapid expansion is powered by massive training datasets that produce more capable predictive models and enable broader applications. However, as these data-driven systems increasingly shape everyday life, they pose tangible risks, including privacy breaches [24, 21], labor exploitation [26], and biased decisions that exacerbate socio-economic disparities [23].

Responses to these risks span firms, regulators, and users. At the firm level, "Trustworthy AI" programs integrate fairness checks, bias mitigation, privacy audits, and red-teaming across the system lifecycle [1], but they often conflict with objectives tied to model performance and user engagement. At the regulatory level, data-protection regimes such as GDPR [10], PIPEDA [13], and CPRA [32] establish baseline obligations for different stakeholders, although formal adherence alone is insufficient to ensure substantively responsible or equitable outcomes [29, 33]. Finally, at the user level, *Algorithmic Collective Action in machine learning* (ACA) has emerged as a promising framework to coordinate grassroots efforts, by enabling individual users to collectively contribute, withhold, or strategically structure their data to steer model behavior [15, 8].

This work is motivated by the fact that the ACA literature has almost entirely focused on scenarios involving one single collective. However, traditionally, real-world actions have been carried out by *multiple collectives*, whose uneven capacities, temporal misalignment, and varying priorities may result in fragmented, decentralized, and partially overlapping interventions. Different collectives often *vary* widely in size, strategy, and even actionable goals, yet they can still *align* on common overarching objectives (e.g., climate justice or gender equality). Latour's actor-network theory [19] conceptualized why such movements rarely behave as a single unit: what is regarded as a

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

unitary movement is really a web of heterogeneous interactions with no fixed or permanent center. These loose networks manage to act together only through alignment and coordination. Collectives therefore, act coherently only insofar as such coordination and alignment efforts are continuously sustained. An interesting example is the cyberfeminist struggle, which has never been one monolithic front; instead, it consists of distributed efforts ranging from feminist hacker-art collectives to global hashtag campaigns. Haraway's Cyborg Manifesto [14] famously proposed the cyborg as a metaphor for coalition across difference, emphasizing affinity over rigid identity categories. In that spirit, cyberfeminism has taken many decentralized forms around the world while still sharing core aims of challenging patriarchal tech culture and gender inequality. Environmental activism provides another strong example: the climate justice movement is a patchwork of groups employing diverse tactics (from youth-led social media strikes to militant direct-action cells) to achieve various actionable goals (from permit denials to pipeline stoppages and enforcement actions), but they all strive toward the same overarching objective of halting climate breakdown. Here, we aim to reflect and model the need for multiple collectives in ACA.

Related Works. ACA as users steering ML outputs toward a group goal was recently rigorously formalized in [15], building on the Data Leverage framework's view of data as an active instrument rather than a passive input [34]. In [30], the authors proposed a combinatorial model for ACA to study the strategic interaction between drivers and delivery platforms, inspired by the #DeclineNow DoorDash campaign. The work in [4] studied how the firm's chosen learning algorithms affect the success of ACA. Motivated by the growing regulatory focus on privacy and data protection, [31] found that differentially private training can impede the success of algorithmic collective action. The works in [3] and [11] focused on ACA in recommender systems. Simulation-based analyses of platform–worker dynamics were studied in [20]. Complementary directions connect ACA to explainability/recourse [16] and to fairness interventions by minority groups [5]. The authors of [12] proposed an alternative theoretical framework to [15] that empowers collectives via statistical inference, enabling them to learn better ACA strategies and infer the parameters that determine their success. The closest works to ours are [18], [17], and [2]. The authors of [18] introduced a conceptual framework for ACA with multiple collectives, and experiments are performed. However, no theoretical treatment is provided, and most of the focus is on ACA with two collectives. The workshop papers in [17] and [2] generalized the original probabilistic framework from [15] to the (noisy) two collective and the multiple collective settings, respectively. Our work is inspired by the empirical findings of [18] and builds on the framework introduced in [2]. We move substantially beyond prior work by providing, to the best of our knowledge, the *first statistical treatment for Algorithmic Collective Action with multiple collectives*, generalizing the framework from [12].

Contribution. We introduce the first statistical framework for ACA with multiple collectives. We focus on collective action in classification. In particular, we study how M collectives can plant and unplant signals, i.e., bias a classifier to learn or unlearn an association between an altered version of the features x and a chosen, possibly overlapping, multiset of target classes $\{\{y_c\}\}_c$. We measure collective action success using both per-collective and global metrics, where global success can be defined as a summary of the per-collective success measures, such as their average or the minimum. We generalize the statistical framework from [12]. As such, for signal planting, we analyze two regimes: when collectives can act on both features and labels, and when collectives are limited to acting on features only. For each regime, we derive lower bounds on the per-collective success. For signal unplanting, we analyze the case in which collectives employ an adaptive strategy acting on both features and labels, and we again derive a lower bound on the per-collective success. Both for signal planting and unplanting, our results reveal interesting trade-offs driven by the interplay of the sizes of the collectives and by how closely their actionable goals align. Although each collective's bounds depend on other collectives' strategies, we make them individually computable by introducing a hierarchy of approximations, ranging from worst case (the collective knows nothing about others) to best case (the collective knows some aggregated quantity about other collectives' strategies). Finally, we first conceptually illustrate our framework with a use case in the space of interventions for climate adaptation, and we design a synthetic dataset that simulates it, numerically showing the usefulness of and the interesting tradeoffs emerging from our framework. Then, we also verify that the bounds hold for classifiers actually trained on collectively edited synthetic and US Census data. Code is available at `https://github.com/clabat9/multi-collective-aca`.

2 Preliminaries and Setting

We start from the definition of ACA with a single collective as given in [15, 12]. A firm deploys a learning model on data drawn from a certain user population. Within the overall user population,

a collective forms that represents a given share of users (a positive fraction, but not the whole). This collective wants to steer the platform's outcome toward a certain objective. To do so, the collective agrees on one strategy for editing their contributions—potentially changing features, labels, or both—before those contributions reach the platform. The platform then receives a training set that blends unmodified data from the baseline process with the collective's edited data, and it deploys its algorithm on this blended data. The collective's impact is evaluated by a tailored success measure. Given its size, the collective's problem is to choose the editing rule that maximizes this success once the platform has trained on the blended data. Here, we rigorously generalize the statistical treatment of ACA from [12] to a setting with multiple collectives, each acting with a possibly different strategy and actionable goal, but attempting to maximize both per-collective and global measures of success. The measure-theoretic details of the setting are collected in Appendix A.1.

Data space. We define the data space $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ as the product of two finite spaces $\mathcal{X}$ and $\mathcal{Y}$ being the feature and label spaces, respectively.

Population. The population is a finite set Ω of users with a sampling distribution π, and $f_0 : \Omega \to \mathcal{Z}$ assigns a feature–label pair to each user. The population distribution is the push-forward $P_0 := (f_0)_{\#}\pi$, the *pre-edit* distribution over feature–label pairs: for each $z = (x, y) \in \mathcal{Z}$, $P_0(z)$ is the π-fraction of users whose assigned pair is z.

Collectives, Strategies, and Masses. We define an ensemble of M collectives as disjoint subsets $\Omega_1, \ldots, \Omega_M$ of Ω such that $\pi(\Omega_c) = \alpha_c > 0$. We refer to α_c as the mass of the c-th collective and set $\bar{\alpha} := \sum_{c=1}^{M} \alpha_c$. Each collective Ω_c agrees on a (deterministic) strategy $h_c : \mathcal{Z} \to \mathcal{Z}$ from a space of feasible edits. For every collective Ω_c, we make the *neutrality assumption* [15, 12] that the pre-edit distribution within the collective matches the population pre-edit distribution, i.e., $P_0(z) = \mathbb{P}\left[f_0(\omega) = z \,|\, \omega \in \Omega_c\right] \ \forall z \in \mathcal{Z}$. Intuitively, before applying h_c, each collective is a random slice of the same underlying population (membership is independent of unedited examples).

Remark. The neutrality assumption is reasonable in several cases, e.g., when collectives form based on attributes unrelated to data and unknown to the platform. In other cases, collectives may form based on the data they aim to edit; then one can relax the assumption to $\mathrm{TV}\big(P_0,\ \mathbb{P}\left[f_0(\omega) \in \cdot \mid \omega \in \Omega_c\right]\big) \leq \phi_{N,c}$. In what follows we keep neutrality; Appendix F shows that our bounds continue to hold under the relaxed assumption, with explicit bias terms that are linear in $\phi_{N,c}$.

Mixture Distribution. The action of collective Ω_c via h_c maps P_0 to $P_c := (h_c)_{\#}P_0$. Thus, in the infinite-data regime, a firm that trains a learning model on data coming from Ω observes the *post-edit* mixture distribution

$$P \;=\; \big(1 - \bar{\alpha}\big)\, P_0 \;+\; \sum_{c=1}^{M} \alpha_c\, P_c. \tag{1}$$

Training Set and Empirical Distribution. In practice, if a collective action is not performed, a firm would use a training set of N data points $\{Z_i = (X_i, Y_i)\}_{i=1}^{N}$. We model the training set by sampling individuals $S_1, \ldots, S_N$ i.i.d. from π, and setting $Z_i := f_0(S_i)$. Then $\{Z_i\}_{i=1}^{N}$ are i.i.d. from the pre-edit distribution P_0; this covers, e.g., survey-like sampling and platform logging (Appendix A.1). If a collective action is performed, i.e., if the collectives edit these examples according to their strategies h_c, the firm would train its learning model using the training set of the N edited data points. The individuals belonging to collectives that appear in the training set, which we refer to as *sample collectives*, are then captured by the random *index sets*

$$I_{N,c} \;:=\; \big\{\, i \in \{1, \ldots, N\} :\; S_i \in \Omega_c \,\big\}, \quad c = 1, \ldots, M,$$

with sizes $n_{N,c} := |I_{N,c}|$ and empirical masses $\alpha_{N,c} := n_{N,c}/N$, so that $\bar{\alpha}_N := \sum_{c=1}^{M} \alpha_{N,c}$. We define the empirical non-collective mass $\alpha_{N,0} := 1 - \bar{\alpha}_N$, and denote the number of non-collective training points with $n_{N,0} \;:=\; N - \sum_{d=1}^{M} n_{N,d}$. We write $\hat{P}^{\mathrm{pre}}_{N,c}$ and $\hat{P}_{N,c} := (h_c)_{\#}\hat{P}^{\mathrm{pre}}_{N,c}$ for the empirical pre-edit and post-edit distributions of the sample collective $I_{N,c}$, with feature marginals $\hat{P}^{\mathrm{pre},X}_{N,c}$ and $\hat{P}^{X}_{N,c}$, and $\hat{P}_{N,0}$ for the empirical distribution of the non-collective training points (formal definitions in Appendix A.1). Thus the post-edit empirical mixture distribution is

$$\hat{P}_N \;=\; \alpha_{N,0}\hat{P}_{N,0} + \sum_{c:\alpha_{N,c}>0} \alpha_{N,c}\, \hat{P}_{N,c}. \tag{2}$$

As $N \to \infty$, $\hat{P}_N$ converges almost surely to the mixture (1) (Proposition A.1 in Appendix A.1).

We study the case where the firm's learning model is a classifier. In particular, we assume such a classifier is an ε-contamination-suboptimal classifier.

Definition 2.1 (ε-contamination-suboptimal classifier)**.** A classifier $\hat{m} : \mathcal{X} \to \mathcal{Y}$ is ε-contamination-suboptimal on $\mathcal{X}' \subseteq \mathrm{supp}(\hat{P}_N^X)$ under $\hat{P}_N$ if there exists a probability measure $\hat{Q}_N$ on $\mathcal{X} \times \mathcal{Y}$ and a probability measure R_N on $\mathcal{X} \times \mathcal{Y}$ such that

$$\hat{Q}_N = (1-\varepsilon)\hat{P}_N + \varepsilon R_N, \tag{3}$$

and, for every $x \in \mathcal{X}'$ with $\hat{Q}_N^X(x) > 0$,

$$\hat{m}(x) \in \operatorname*{argmax}_{y \in \mathcal{Y}} \hat{Q}_N(y \mid x). \tag{4}$$

Consistently with previous literature [15, 12], we take $\varepsilon < 1/2$. Appendix A.1 relates this contamination model to total-variation suboptimality, under which all results remain valid with $\varepsilon/(1-\varepsilon)$ replaced by 2ε.

Per-collective and Global Success. Given a strategy profile $\{h_c\}_{c=1}^M$, a classifier $\hat{m} : \mathcal{X} \to \mathcal{Y}$ trained on the training set $\{Z_i\}_{i=1}^N$, and a test set of N^{test} data points $\{Z_i^{\text{test}}\}_{i=1}^{N^{\text{test}}}$ again drawn i.i.d. from P_0, each sample collective $I_{N,c}$ is interested in maximizing a per-collective test success metric $\hat{S}_c$, and a global test success metric $\hat{G}(\{\hat{S}_c\}_c)$ being an aggregation of some sort of the per-collective successes; we write $\hat{G}$ for global metrics to stress that they are functions of the per-collective successes $\{\hat{S}_c\}_c$, and reserve $\hat{S}_c$ for the per-collective success. If such aggregation is a $\min$, we obtain an egalitarian metric which takes the per-collective success of the least successful collective as the global measure of success [28], i.e.,

$$\hat{G}_{\min}(\{\hat{S}_c\}_c) := \min_{c:\alpha_{N,c}>0} \hat{S}_c(\alpha_{N,c}, N^{\text{test}}). \tag{5}$$

Alternatively, if we average, we obtain a metric proportional to each collective's mass, i.e.,

$$\hat{G}_{\text{w}}(\{\hat{S}_c\}_c) := \frac{1}{\bar{\alpha}_N} \sum_{c=1}^{M} \alpha_{N,c}\, \hat{S}_c(\alpha_{N,c}, N^{\text{test}}). \tag{6}$$

Similarly to Proposition A.1, in the infinite-data regime, the versions of these success metrics computed under the empirical test measure $\hat{P}_{N^{\text{test}}}$ converge almost surely to their population counterparts computed under P_0.

Remark. In the rest of the paper, we derive bounds for the per-collective success. The bounds on global success can be readily derived by aggregating the per-collective bounds with the chosen summary and adapting union bounds.

Further Setting Details. Following [12], we assume each collective has access only to their own data and not to the data of other collectives or of those who are not part of any collective. However, a collective could have some sort of aggregated knowledge about other collectives (see Sections 3-4). Moreover, we will use Hoeffding's concentration inequality (see Lemma D.1 in [12]) without loss of generality. We denote Hoeffding error terms with $R_\gamma(k) := \sqrt{\frac{\log(1/\gamma)}{2k}}$, for any $\gamma > 0$ and $k \in \mathbb{N}^+$. In the bounds we derive, whenever $R_\gamma(n_{N,0})$ appears, we condition on $n_{N,0} > 0$. Moreover, in the theoretical results below, we assume that the deployed classifier is ε-contamination-suboptimal under the post-edit empirical distribution $\hat{P}_N$ on the relevant planted feature set.

3 Signal Planting

We start by analyzing the setting in which each sample collective $I_{N,c}$ wants the classifier $\hat{m}$ to *learn* an association between an altered version of the features $g_c(x)$ and a chosen target class y_c^*. Formally, let $\hat{P}_{N^{\text{test}}}^X$ denote the empirical distribution of test features (the marginal of $\hat{P}_{N^{\text{test}}}$ on $\mathcal{X}$) and $g_c : \mathcal{X} \to \mathcal{X}$ the feature map used on target-label points induced by h_c, then each sample collective $I_{N,c}$ measures per-collective success as

$$\hat{S}_c(\alpha_{N,c}) = \mathbb{P}_{x \sim \hat{P}_{N^{\text{test}}}^X} \left[\hat{m}(g_c(x)) = y_c^*\right]. \tag{7}$$

We denote with $\mathcal{X}_c^* := g_c(\mathcal{X}) \subseteq \mathcal{X}$ the subset of the feature space representing the image of the eligible set under h_c, i.e., the image of the features alterable by the adoption of h_c. As usually done [15, 12, 31, 2], we consider two cases: the case in which the users can modify both features and labels, referring to the resulting strategies as *feature-label* strategies, and the case in which the users can access both features and labels, but modify only the features, referring to the resulting strategies as *feature-only* strategies.

In the rest of this and the following section, when there is no risk of ambiguity we fix a collective Ω_c, together with its sample counterpart $I_{N,c}$, and assume its empirical mass $\alpha_{N,c} > 0$. Moreover, our results are conditioned on the realization of the sizes of the sample collectives $\{n_{N,0}, n_{N,1}, \ldots, n_{N,M}\}$. For each $x^* \in \mathcal{X}_c^*$, we define the empirical prevalence of the planted feature in collective $I_{N,c}$ and the empirical (planting) *conflict* terms with other collectives as

$$\hat{p}_c^X(x^*) := \hat{P}_{N,c}^X(x^*), \text{ and} \tag{8}$$

$$\Delta_{c\to d}(x^*) := \max_{y' \in \mathcal{Y}\setminus\{y_c^*\}} \big(\hat{P}_{N,d}(x^*, y') - \hat{P}_{N,d}(x^*, y_c^*)\big), \quad d \in \{1, \ldots, M\},\ d \neq c, \tag{9}$$

respectively. Intuitively, $\Delta_{c\to d}(x^*)$ measures how strongly sample collective $I_{N,d}$'s post-edit data associates the planted feature value x^* with some label other than y_c^* compared to y_c^*. Therefore, larger absolute values indicate greater potential for $I_{N,d}$ to *confuse or accrete* (based on its sign) $I_{N,c}$'s planted signal at x^*. To approximate the non-collective conflict using only the pre-edit data of the sample collective $I_{N,c}$, we define

$$\Delta_{c,0}^{(n_{N,c})}(x^*) \ := \ \max_{y' \in \mathcal{Y}\setminus\{y_c^*\}} \big(\hat{P}_{N,c}^{\mathrm{pre}}(x^*, y') - \hat{P}_{N,c}^{\mathrm{pre}}(x^*, y_c^*)\big). \tag{10}$$

Intuitively, $\Delta_{c,0}^{(n_{N,c})}(x^*)$ is a data-driven proxy for how much the non-collective population could counteract the planted signal at x^* by supporting an alternative label over y_c^*, estimated under neutrality.

Feature-label Signal Planting Strategy. For each sample collective $I_{N,c}$, the *feature-label* signal strategy is given by

$$h_c(x, y) = (g_c(x), y_c^*)\,. \tag{11}$$

This strategy forces every edited example to carry the same label y_c^* while mapping its features into the planted set via g_c. It directly aims to create an association between the transformed features and the target class in the training data.

Feature-only Signal Planting Strategy. For each sample collective $I_{N,c}$, assume that $\mathcal{X} \setminus \mathcal{X}_c^* \neq \emptyset$, and fix a point $x_{0,c} \in \mathcal{X} \setminus \mathcal{X}_c^*$. The *feature-only* signal planting strategy is given by

$$h_c(x, y) = \begin{cases} (g_c(x), y)\,, & y = y_c^*, \\ (x_{0,c}, y)\,, & y \neq y_c^*. \end{cases} \tag{12}$$

The intuition here is that, while the features of the target-label examples are modified according to g_c, examples with labels $y \neq y_c^*$ are pushed to a feature value outside the planted signal set. In this setting, the label y is assumed to be accessible or computable, but not modifiable by the collective.

We can now state our main result for this section.

Theorem 3.1 (Signal planting with multiple collectives)**.** *Under the neutrality assumption, assuming that $\hat{m}$ is ε-contamination-suboptimal on $\mathcal{X}_c^*$ under $\hat{P}_N$, and assuming all collectives either employ the feature-label or the feature-only strategy from* (11)-(12)*, respectively, with probability at least $1 - \delta$ over the draw of the training and test samples, with $\delta > 0$, the per-collective success of collective $I_{N,c}$ satisfies*

$$\begin{aligned} \hat{S}_c(\alpha_{N,c}) \ \geq \ & \mathbb{P}_{x\sim\hat{P}_{N,c}^{\mathrm{pre},X}}\Big[\ \alpha_{N,c}\Big(\hat{p}_c^X\big(g_c(x)\big) - 2R_{\tilde{\delta}}(n_{N,c})\Big) - \sum_{d\neq c} \alpha_{N,d}\,\Delta_{c\to d}\big(g_c(x)\big) \\ & - \alpha_{N,0}\,\Delta_{c,0}^{(n_{N,c})}\big(g_c(x)\big) - \ 2\alpha_{N,0}\Big(R_{\tilde{\delta}}(n_{N,0}) + R_{\tilde{\delta}}(n_{N,c})\Big) - \frac{\varepsilon}{1-\varepsilon} > 0\Big] \\ & - R_{\delta/4}(n_{N,c}) - R_{\delta/4}(N^{\mathrm{test}}), \end{aligned} \tag{13}$$

where $K_c \ := \ |\mathcal{X}_c^*|\,|\mathcal{Y}|$ and $\tilde{\delta} \ := \ \frac{\delta}{10K_c}$.

See proof on page 15.

Intuitive Interpretation. The bound in (13) formalizes a race at each planted feature value $x^* = g_c(x)$: collective c succeeds at x^* when the mass it concentrates on (x^*, y_c^*) exceeds, by more than the classifier slack $\varepsilon/(1-\varepsilon)$ and the finite-sample radii, the largest mass that any other data source places on a competing label at the same x^*. The guarantee is the pre-edit probability that c's signal wins this race, and each term identifies one racer: the own signal $\alpha_{N,c}\hat{p}_c^X$, the cross-collective interference $\sum_{d\neq c} \alpha_{N,d}\Delta_{c\to d}$, and the organic resistance of the non-collective population $\alpha_{N,0}\Delta_{c,0}^{(n_{N,c})}$. The

interference terms are sign-indefinite, so the sign of $\Delta_{c\to d}$ classifies every other collective d from the perspective of c, i.e., an *ally* when $\Delta_{c\to d} < 0$ (d's edits support y_c^* at x^*, and its mass adds to c's signal), *neutral* when $\Delta_{c\to d} = 0$ (e.g., d plants feature values outside $\mathcal{X}_c^*$), and an *adversary* when $\Delta_{c\to d} > 0$ (d associates x^* with a competing label). Unlike the single-collective setting, success is thus governed by an alignment-weighted effective mass rather than by size alone. The symmetric regime makes this explicit.

Corollary 3.2 (Symmetric regime)**.** *In the setting of Theorem 3.1, let $M \geq 2$ and $\alpha_{N,d} = \bar{\alpha}_N/M$ for every collective d, and fix $x^* \in \mathcal{X}_c^*$ with $\hat{p}_c^X(x^*) = p$, $\Delta_{c\to d}(x^*) = \Delta$ for all $d \neq c$, and $\Delta_{c,0}^{(n_{N,c})}(x^*) = \Delta_0$. Let $\kappa := (1-\bar{\alpha}_N)\Delta_0 + \frac{\varepsilon}{1-\varepsilon} + \rho_N$, where $\rho_N := \frac{2\bar{\alpha}_N}{M} R_{\tilde{\delta}}(n_{N,c}) + 2(1-\bar{\alpha}_N)\big(R_{\tilde{\delta}}(n_{N,0}) + R_{\tilde{\delta}}(n_{N,c})\big)$ collects the radii. Then the margin in* (13) *is positive at x^* iff*

$$p \;>\; (M-1)\,\Delta + \frac{M}{\bar{\alpha}_N}\,\kappa. \tag{14}$$

In particular, (i) under full alignment ($\Delta = -p$) the condition reads $\bar{\alpha}_N\, p > \kappa$, (ii) under mutual neutrality ($\Delta = 0$) it reads $\frac{\bar{\alpha}_N}{M}\, p > \kappa$, (iii) under full conflict ($\Delta = p$) it fails for every $M \geq 2$ whenever $\kappa \geq 0$, e.g., when $\Delta_0 \geq 0$.

See proof on page 19.

Corollary 3.2 exhibits a trichotomy in the alignment. Under full alignment (all collectives plant the same label at the same features), the collectives pool into a single collective of mass $\bar{\alpha}_N$, and M matters only through the radii. Under mutual neutrality only the individual mass $\bar{\alpha}_N/M$ counts, so fragmentation dilutes success by a factor of M. Under full conflict (equal-size collectives planting distinct labels at shared features), comparable collectives destroy each other's guarantees, consistently with the label-entropy collapse in Figure 2. None of these phenomena exists at $M = 1$.

Impact of Strategies. The lower bound in Theorem 3.1 depends on the collective strategy through the feature map g_c and the resulting empirical measures $\hat{P}_{N,c}$ and $\hat{P}_{N,d}$ that enter $\hat{p}_c^X(\cdot)$ and $\Delta_{c\to d}(\cdot)$. Indeed, the *signal* term $\alpha_{N,c}\hat{p}_c^X(g_c(x))$ increases when the strategy concentrates the collective edited mass on the set $\mathcal{X}_c^*$, while the *overall weighted conflict* $\sum_{d\neq c}\alpha_{N,d}\Delta_{c\to d}(g_c(x))$ reflects how other collectives' edits create competing label associations at the same planted feature values. Under *feature-label* planting, $h_c(x,y) = (g_c(x), y_c^*)$ forces all edited examples to carry y_c^*, typically boosting the signal term at planted values and eliminating within-collective label ambiguity at those values. Under *feature-only* planting, labels are not changed, and only target-label points are mapped via g_c, which can strengthen separation by reducing the presence of non-y_c^* labels inside $\mathcal{X}_c^*$ but may limit how much mass can be moved onto the planted set when $y \neq y_c^*$ points cannot be relabeled.

On the Computability of the Bound. The lower bound in (13) can be evaluated by the sample collective $I_{N,c}$ if it has access to its own pre-edit sample and strategy, and a limited set of global aggregate quantities. Concretely, collective c can compute $\hat{P}_{N,c}^{\mathrm{pre}}$ and its marginal $\hat{P}_{N,c}^{\mathrm{pre},X}$ from $\{(X_i,Y_i) : i \in I_{N,c}\}$, evaluate the proxy non-collective conflict $\Delta_{c,0}^{(n_{N,c})}(x^*)$ for $x^* \in \mathcal{X}_c^*$ using only $\hat{P}_{N,c}^{\mathrm{pre}}$ (under neutrality), and then estimate the outer probability $\mathbb{P}_{x\sim\hat{P}_{N,c}^{\mathrm{pre},X}}[\cdot]$ as the fraction of its own pre-edit features for which the corresponding margin at $g_c(x)$ is positive. The remaining terms are deterministic given the radii, the confidence level δ, and the suboptimality parameter ε. The substantive non-local requirement is the overall weighted conflict $\sum_{d\neq c}\alpha_{N,d}\,\Delta_{c\to d}\big(g_c(x)\big)$; in particular, for each $d \neq c$, it depends on the empirical mass $\alpha_{N,d}$ and the conflict $\Delta_{c\to d}(x^*)$ at the relevant signal features $x^* = g_c(x)$, with $\Delta_{c\to d}(x^*) = \max_{y'\neq y_c^*}\big(\hat{P}_{N,d}(x^*,y') - \hat{P}_{N,d}(x^*,y_c^*)\big)$. Here, we propose different ways to relax these requirements based on how much information about other collectives a collective can access. The best-case scenario is the one in which the collective $I_{N,c}$ can access a single "others" term $\Delta_c^{\mathrm{others}}(x^*)$ that a platform (or trusted aggregator) publishes at the needed signal features without revealing per-collective distributions. Alternatively, the worst case scenario is upper-bounding $\sum_{d\neq c}\alpha_{N,d}\Delta_{c\to d}(x^*)$ by the total collective mass $\bar{\alpha}_N - \alpha_{N,c}$ (or by a known upper bound $\bar{\alpha}_N^{\max} - \alpha_{N,c}$), yielding a generally more conservative bound. Finally, an intermediate scenario is to assume mild symmetry or collective anonymity: while different collectives may adopt different edits, their post-edit behavior at the signal features $x^* = g_c(x)$ is statistically homogeneous in the sense that, from the perspective of collective $I_{N,c}$, another collective induces a conflict level that can be approximated by a proxy being a functional of the $I_{N,c}$'s pre-edit distribution. Under neutrality, $I_{N,c}$ can then estimate $\Delta_{c\to d}(x^*)$ from its own pre-edit sample by relying on such a proxy Φ so that $\Delta_{c\to d}(x^*) \approx \Phi\big(\hat{P}_{N,c}^{\mathrm{pre}}, x^*\big)$, for all $d \neq c$ and relevant $x^* \in \mathcal{X}_c^*$, where Φ can be

derived from a shared model class for strategies (e.g., edits drawn from a known family, or strategies that coincide up to relabeling/noise). Overall, this approach trades stronger structural assumptions (exchangeability/anonymous behavior) for a substantially weaker global information requirement.

On the Consistency of the Bound. In Appendix D, we show how the bound in Theorem 3.1 reduces to the single collective bounds in Theorems 3.3 and 3.5 of [12] when $M = 1$.

4 Signal Unplanting

We now analyze the case where a sample collective $I_{N,c}$ aims to *remove* the association between the planted features $g_c(x)$ and its original target class y_c^*. Accordingly, each sample collective $I_{N,c}$ measures per-collective success as

$$\hat{S}_c(\alpha_{N,c}) \;:=\; \mathbb{P}_{x\sim\hat{P}^X_{N^{\text{test}}}}[\hat{m}(g_c(x)) \neq y_c^*]\,. \tag{15}$$

Throughout this section, we reuse the notation and objects from Section 3 when possible and again employ and generalize the strategies proposed in [12].

Naive Strategy. A simple way to increase (15) is to pick any alternative label $y_c' \in \mathcal{Y} \setminus \{y_c^*\}$ and *plant* the association $(g_c(x), y_c')$ as explained in Section 3. Indeed, for any such y_c', it holds

$$\mathbb{P}_{x\sim\hat{P}^X_{N^{\text{test}}}}[\hat{m}(g_c(x)) \neq y_c^*] \;\geq\; \mathbb{P}_{x\sim\hat{P}^X_{N^{\text{test}}}}[\hat{m}(g_c(x)) = y_c']\,. \tag{16}$$

Thus, the sample collective $I_{N,c}$ may choose y_c' to maximize the right-hand side of (16) and then directly leverage Theorem 3.1 to perform unplanting. However, this reduction forces a single competing label at all planted features.

Adaptive Strategy. An alternative possibility is making the label y_c' depend on the planted feature value $x^* \in \mathcal{X}_c^*$. Intuitively, at each x^*, the collective chooses a label $y_c'(x^*)$ that makes y_c^* least competitive at that feature. To select $y_c'(x^*)$, the collective $I_{N,c}$ selects an estimation sub-collective $E_{N,c} \subseteq I_{N,c}$ of size $0 < n_{e,c} < n_{N,c}$, independently of all data. We denote with $R_{N,c} := I_{N,c} \setminus E_{N,c}$ the remaining fraction of the collective of size $n_{N,c} - n_{e,c}$. Thus, the corresponding pre-edit empirical measures are defined as

$$\hat{P}^{\text{pre},E}_{N,c} := \frac{1}{n_{e,c}} \sum_{i\in E_{N,c}} \delta_{Z_i}, \quad \hat{P}^{\text{pre},R}_{N,c} := \frac{1}{n_{N,c} - n_{e,c}} \sum_{i\in R_{N,c}} \delta_{Z_i}. \tag{17}$$

In this setting, for each $x^* \in \mathcal{X}_c^*$, the feature-dependent selected label is chosen as

$$y_c'(x^*) \;:=\; \operatorname*{arg\,max}_{y\in\mathcal{Y}\setminus\{y_c^*\}} \hat{P}^{\text{pre},E}_{N,c}(x^*, y), \tag{18}$$

The resulting *adaptive unplanting strategy* is the feature–label map

$$h_c(x, y) \;:=\; \big(g_c(x),\, y_c'(g_c(x))\big). \tag{19}$$

In other words, collective $I_{N,c}$ plants the same features as in signal planting, but assigns to each planted feature x^* a competing label $y_c'(x^*) \neq y_c^*$ chosen using $E_{N,c}$.

For each $x^* \in \mathcal{X}_c^*$, we define the empirical (unplanting) conflict terms with other collectives as

$$\Delta^{\text{unp}}_{c\leftarrow d}(x^*) \;:=\; \hat{P}_{N,d}(x^*, y_c^*) - \hat{P}_{N,d}\big(x^*, y_c'(x^*)\big), \quad d \in \{1,\dots,M\},\; d \neq c, \tag{20}$$

and the corresponding non-collective approximate conflict using only data independent of the label-selection step as

$$\Delta^{\text{unp},(n_{N,c}-n_{e,c})}_{c,0}(x^*) \;=\; \hat{P}^{\text{pre},R}_{N,c}(x^*, y_c^*) - \hat{P}^{\text{pre},R}_{N,c}\big(x^*, y_c'(x^*)\big). \tag{21}$$

Intuitively, $\Delta^{\text{unp}}_{c\leftarrow d}(x^*)$ is large when other collectives still support (x^*, y_c^*) more than $(x^*, y_c'(x^*))$, thereby resisting unplanting at x^*. The approximate conflict in (21) plays the analogous role for non-collective users under neutrality. We can now state our main result for this section.

Theorem 4.1 (Signal unplanting with multiple collectives)**.** *Under the neutrality assumption, assuming that $\hat{m}$ is ε-contamination-suboptimal on $\mathcal{X}_c^*$ under $\hat{P}_N$, assuming all collectives employ the adaptive unplanting strategy from* (19) *with estimation splits $\{E_{N,c}\}_c$, then, with probability at least $1-\delta$ over the draw of the training and test samples, with $\delta > 0$, and the random choice of $\{E_{N,c}\}_c$, the per-collective success of $I_{N,c}$ satisfies*

$$\begin{aligned}\hat{S}_c(\alpha_{N,c}) \;\geq\; \mathbb{P}_{x\sim\hat{P}^{\text{pre},R,X}_{N,c}}\Big[\; &\alpha_{N,c}\Big(\hat{p}_c^X\big(g_c(x)\big) - 2R_{\tilde{\delta}}(n_{N,c})\Big) - \sum_{d\neq c} \alpha_{N,d}\,\Delta^{\text{unp}}_{c\leftarrow d}\big(g_c(x)\big) \\ &- \alpha_{N,0}\,\Delta^{\text{unp},(n_{N,c}-n_{e,c})}_{c,0}\big(g_c(x)\big) - 2\alpha_{N,0}\Big(R_{\tilde{\delta}}(n_{N,0}) + R_{\tilde{\delta}}(n_{N,c}-n_{e,c})\Big) - \frac{\varepsilon}{1-\varepsilon} > 0\Big] \\ &- R_{\delta/4}(n_{N,c}-n_{e,c}) - R_{\delta/4}(N^{\text{test}}),\end{aligned} \tag{22}$$

where $K_c := |\mathcal{X}_c^*|\,|\mathcal{Y}|$ and $\tilde{\delta} := \frac{\delta}{10K_c}$.

See proof on page 19.

Intuitive Interpretation. The bound in (22) mirrors (13) with the race reversed: collective c succeeds at x^* when the association it wants to erase, (x^*, y_c^*), is overtaken by the data-driven alternative $(x^*, y_c'(x^*))$ with a margin of $\varepsilon/(1-\varepsilon)$. Accordingly, $\Delta^{\mathrm{unp}}_{c\leftarrow d}$ measures the resistance of collective d, positive when d's data keeps supporting (x^*, y_c^*) against the alternative and negative when d, deliberately or not, helps demote it, while $\alpha_{N,0}\Delta^{\mathrm{unp},(n_{N,c}-n_{e,c})}_{c,0}$ estimates the organic strength of the association in the general population, which is what the collective fights. This reveals an asymmetry between the two goals. In fact, planting can target rare or synthetic feature values whose organic resistance is negligible, whereas unplanting confronts the ambient support for an existing association and pays a cost proportional to the non-collective mass $\alpha_{N,0}$. Corollary 3.2 carries over verbatim with Δ^{unp} in place of Δ and the radii of Theorem 4.1, and yields the same trichotomy: aligned unplanters pool their masses, neutral ones act alone, and collectives protecting associations that others erase generate destructive interference. The estimation split $E_{N,c}$ can be any subset of the collective's sample chosen independently of the data, e.g., uniformly at random as in our experiments. This independence keeps the holdout $R_{N,c}$ i.i.d., as the concentration argument requires. Its size governs a bias–variance trade-off, because enlarging $n_{e,c}$ improves the label-selection step, since $y_c'(x^*)$ is more likely the strongest competitor at x^*, but it shrinks the holdout $n_{N,c} - n_{e,c}$ on which the bound is evaluated and inflates the radii $R_{\tilde{\delta}}(n_{N,c} - n_{e,c})$ and $R_{\delta/4}(n_{N,c} - n_{e,c})$. Since the selection step only ranks the competing labels at each planted value, rather than estimating their probabilities, a sublinear split ($n_{e,c} \ll n_{N,c}$) suffices.

On the Computability of the Bound. The discussion of Section 3 carries over. Collective $I_{N,c}$ selects the alternative labels $y_c'(x^*)$ on its estimation sub-sample as in (18), and evaluates the non-collective conflict and the outer probability on the holdout sub-sample, which is independent of the label selection; the radii now also depend on the split size $n_{e,c}$. The only non-local requirement is again the weighted cross-collective term $\sum_{d\neq c}\alpha_{N,d}\Delta^{\mathrm{unp}}_{c\leftarrow d}(g_c(x))$, with $\Delta^{\mathrm{unp}}_{c\leftarrow d}(x^*) = \hat{P}_{N,d}(x^*, y_c^*) - \hat{P}_{N,d}(x^*, y_c'(x^*))$, so the same hierarchy of approximations applies.

On the Consistency of the Bound. Similarly to Section 3, in Appendix D, we show how the bound in Theorem 4.1 reduces to the single collective bounds in Theorems 3.7 of [12] when $M = 1$.

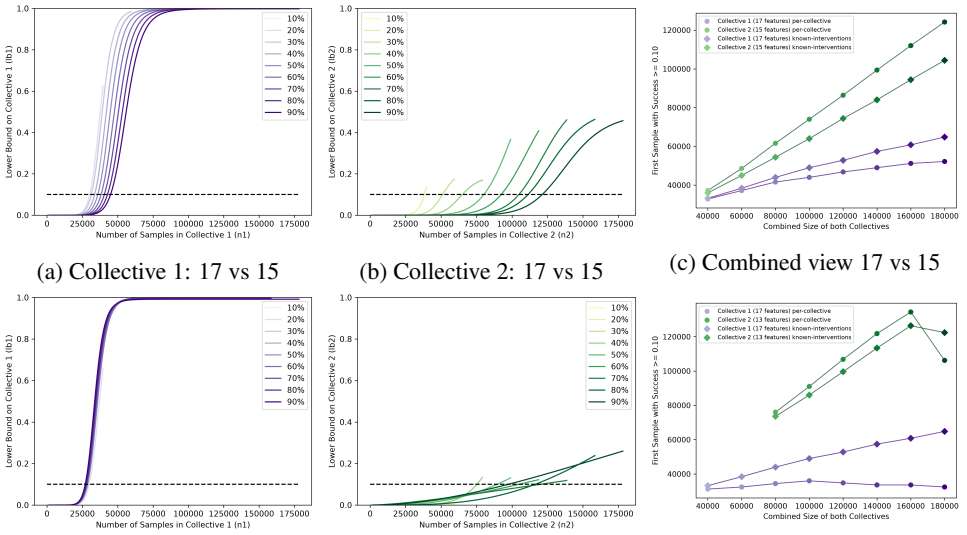

(a) Collective 1: 17 vs 15 (b) Collective 2: 17 vs 15 (c) Combined view 17 vs 15

(d) Collective 1: 17 vs 13 (e) Collective 2: 17 vs 13 (f) Combined view 17 vs 13

Figure 1: Lower bound on success for two feature-label collectives under varying total collective mass, mass split and feature overlap ($N = 2{,}000{,}000$, $N^{\mathrm{test}} = 100{,}000$). (a), (b), (d), (e): per-collective setting, one curve per total mass. (c), (f): smallest collective size reaching a lower bound of 0.1, per-collective and known-interventions settings.

5 Experiments

Our experiments build upon the protocol of [12] and a motivating example in climate adaptation for smart cities (detailed in Appendix B). We simulate a platform collecting neighborhood-level adaptation records, each comprising 18 features (e.g., adaptation strategy type, implementation timeline, funding sources) and one of four label classes: Greening & Shade Expansion, Reflective

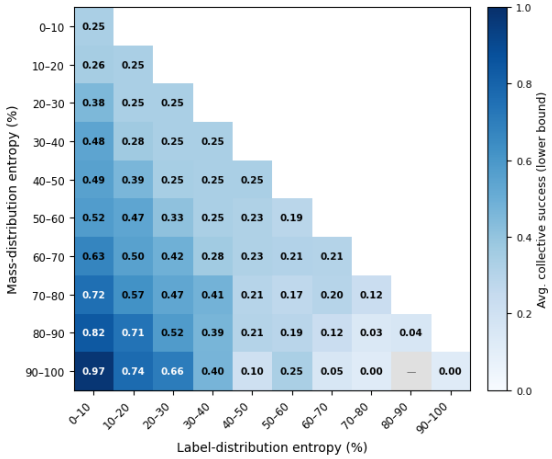

Figure 2: Avg. of the four per-collective success lower bounds, $\hat{G}_{\text{avg}}$, by mass entropy and mass-weighted label entropy ($M = 4$, $N = 500{,}000$, $\bar{\alpha}_N = 0.4$).

Table 1: Trained classifiers: configurations (out of 30) in which the empirical success is at least the bound of Theorem 3.1, and configurations with a non-vacuous bound.

	GBT	LR	Non-vacuous
Synthetic	30	28	22
ACSIncome	30	30	21

Surface Retrofits, Cooling Services & Outreach, and Green Stormwater Infrastructure. Details are presented in Appendix C. We focus on the feature-label regime in the main body and defer feature-only results to Appendix E.

Feature Overlap and Goal Alignment with Two Collectives. We study how goal alignment affects the per-collective lower bounds for $M = 2$, assuming each collective can evaluate the aggregated weighted conflict term (see Section 3). *Collective 1* always plants 17 feature values, while *Collective 2* plants (i) the same 17, (ii) 15, or (iii) 13 of them. Planted labels differ across collectives, so overlap directly induces conflict. We fix $N = 2{,}000{,}000$, $N^{\text{test}} = 100{,}000$, and vary the total collective mass $\bar{\alpha}_N$ from 0.1 to 0.9 in steps of 0.1, sweeping the mass ratio $\alpha_{N,1}/\alpha_{N,2}$ at each levelover 99 equally spaced shares. The *per-collective* and *known-interventions* settings instantiate two different levels of access to the non-local cross-collective term discussed in the computability paragraph of Section 3. In the *per-collective* setting, each collective is assumed to have enough information to evaluate the relevant weighted conflict term using the other collective's actual planted strategy and empirical post-edit distribution. This corresponds to the direct evaluation of the cross-collective contribution in the bound. In the *known-interventions* setting, by contrast, the collective does not know which intervention the other collective implements, but only the set of possible interventions. We therefore replace the exact cross-collective conflict by its average over this intervention set, which here contains the nested patterns of 13 to 17 planted features, each with every target label other than the collective's own (Appendix E). The resulting curves estimate the full-information bound rather than certify it: they are more conservative when the realized strategy conflicts less than the average intervention, and optimistic otherwise. Success curves are smoothed via a Hill curve, and we also report the minimum sample size for 10% success to ease comparisons [27]. Figure 1 reports the 17-vs.-15, (a) and (b), and 17-vs.-13, (d) and (e), regimes. The symmetric 17-vs.-17 case is in Appendix E. A clear trade-off frontier emerges: shifting mass toward one collective raises its bound at the other's expense, while increasing total mass shifts the frontier outward. As overlap decreases, the frontier skews: *Collective 2* loses signal prevalence from fewer planted features, while *Collective 1* benefits from uncontested features that weaken cross-collective conflict. The symmetric case confirms that, absent this asymmetry, performance is driven primarily by the mass split. Finally, from Figure 1 (f), when *Collective 1* plants many more features than *Collective 2* (17 vs. 13), it requires a near-constant number of samples (between 31k and 36k) to obtain a 10% lower bound on success, regardless of the combined size, in the per-collective setting.

Mass Concentration and Goal Alignment with Multiple Collectives. To illustrate how global success responds to overlap beyond two collectives, we now consider $M = 4$ collectives (one per class) sharing the same 18-feature planted pattern under feature-label planting. We sweep $291{,}984$ configurations, fixing $N = 500{,}000$, $N^{\text{test}} = 50{,}000$ and $\bar{\alpha}_N = 0.4$, with per-collective conflict terms. We sample $1{,}985$ mass distributions so as to cover the range of mass entropy uniformly; for each of them we evaluate all $4^4 = 256$ label assignments when the mass entropy is below 0.5, and a subset stratified by label entropy above. Each configuration is summarized by two normalized

entropies: a *mass entropy* measuring how uniform the collective sizes are, and a *mass-weighted label entropy* measuring how concentrated the total collective mass is on the same target label. Both are computed as the normalized Shannon entropy of the relevant distribution. We report the unweighted average of the four per-collective lower bounds, $\hat{G}_{\mathrm{avg}}(\{\hat{S}_c\}_c) := \frac{1}{M}\sum_{c=1}^{M}\hat{S}_c$, a global metric in the sense of Section 2 that gives every collective the same weight regardless of its mass, averaged within 10×10 equal-width entropy bins (Figure 2). Missing entries correspond to bins with no sampled configurations. The main effect is driven by label alignment: in every row, global success is highest when label entropy is low, indicating that concentrating effort on the same target label substantially reduces destructive interference. As label entropy increases, success degrades sharply, with the most severe drop when both entropies are high, i.e., collectives of comparable size targeting different labels strongly conflict at shared planted features, driving the global guarantee near zero. Mass concentration plays a more nuanced role: at intermediate label entropies, higher mass entropy can lift aggregate success once label overlap is partially resolved, but without label alignment, uniform mass alone is insufficient to prevent interference. At low label entropy the average rises from 0.25 for one dominant collective to 0.97 for four comparable ones: aligned collectives pool their masses (Corollary 3.2), but the three small collectives in a concentrated split are individually too small for their own radii, and $\hat{G}_{\mathrm{avg}}$ counts them as failures.

Trained Classifiers. The bounds hold for any ε-contamination-suboptimal classifier. To test them on classifiers that a platform would actually train, two collectives edit the training data (aligned, full-conflict and partial-conflict scenarios, five mass splits at $\bar{\alpha}_N = 0.3$), the platform trains logistic regression (LR) or gradient-boosted trees (GBT), and we compare each collective's empirical success on fresh test data with its bound from Theorem 3.1 ($\delta = 0.05$, $\varepsilon = 0$). We use a synthetic generator with the schema of Appendix C ($N = 2$M) and ACSIncome, built from US Census data [9] $N = 877{,}909$, see Appendix G. GBT never violates the bound (Table 1), which is close to tight where the signal is well identified. LR violates it in two synthetic partial-conflict configurations, where the linear model is misspecified and the suboptimality premise fails.

6 Conclusion

In this work, we introduced the first statistical framework for Algorithmic Collective Action with multiple collectives. We focused on collective action in classification, studying how multiple collectives can bias a classifier to learn an association between an altered version of their features and a chosen, possibly overlapping, set of target classes. We provided quantitative results showcasing the interesting interplay of collectives' sizes and their alignment of values. Finally, we discussed a potential use case for our framework in the space of interventions for climate adaptation, and validated the bounds numerically, exposing multiple tradeoffs. As a final remark, our approach is not naively equivalent to the single-collective approach of [12] instantiated M times. In particular, our results rest on phenomena that only exist for $M \geq 2$. First, the fact that cross-collective conflict terms are sign-indefinite, so every other collective is an ally, neutral, or an adversary, which yields the pooling, dilution, and destructive-interference regimes of Corollary 3.2. Second, folding the other collectives into a worst-case background recovers only the lowest tier of our hierarchy, which our bounds strictly dominate. For instance, it cannot certify aligned collectives that are individually too small but jointly large enough. Third, the finite-sample analysis must decompose the empirical mixture across M random disjoint sample collectives, and the aggregation of per-collective successes into global success is a design dimension with no single-collective analogue. Our framework opens several research directions. It would be interesting to study mixed-objective settings in which some collectives plant signals while others unplant them. One of the limitations of our framework is the initial assumption that collectives are disjoint, and relaxing it is the extension we consider most valuable. Similar to [12], we also assume that the feature space $\mathcal{X}$ is finite. With continuous or high-dimensional features, the union bound over $\mathcal{X}_c^* \times \mathcal{Y}$ must be replaced by uniform convergence over the class of events induced by g_c (e.g., via VC or covering-number arguments), and the conflict terms become functionals that require regularity of P_0. Finally, our framework has an inherently collaborative motivation, from which the described global success metrics stem, but adversarial/game-theoretic systems with multiple antagonist groups can be of interest.

Societal Impact. Coordinated data editing is dual use: the mechanisms that let grassroots collectives steer a classifier can also serve harmful coordination, and the same bounds inform a platform's defensive analysis. The bounds also make power asymmetries quantitative. A collective's mass multiplies the conflict it inflicts, so a small collective's guarantee vanishes once the weighted conflict it faces approaches its own signal mass. The egalitarian metric is designed to surface the worst-served collective, and Figure 2 shows the failure mode of comparable masses with misaligned labels.

Acknowledgments and Disclosure of Funding

CB and FD acknowledge support from the National Institutes of Health Grant 1R01ES037156-01. DR is financed by the project "Building Energy Systems on causal reasoning (BOSS)", funded within the "Technologies and Innovations for the Climate-Neutral City" (TIKS) Programme of the Austrian Research Promotion Agency (FFG). OA is supported by the Social Sciences and Humanities Research Council of Canada (SSHRC) through a Canada Impact+ Research Training Award [award no. 712-2026-1003]. BN declares indirect competing interests in the past 36 months through part-time employment at Modular Insurance Solutions Inc., a course development honorarium from the Harvard Data Science Initiative, and subsequent funding from the UBC Faculty of Medicine AI in Health Network. CB, DR, OA, PG and FD declare no competing interests.

References

[1] Solon Barocas, Moritz Hardt, and Arvind Narayanan. *Fairness and machine learning: Limitations and opportunities*. MIT press, 2023.

[2] Claudio Battiloro, Pietro Greiner, Bret Nestor, Oumaima Amezgar, and Francesca Dominici. Algorithmic collective action with multiple collectives. In *NeurIPS 2025 (Non-archival) Workshop on Algorithmic Collective Action*, 2025.

[3] Joachim Baumann and Celestine Mendler-Dünner. Algorithmic collective action in recommender systems: promoting songs by reordering playlists. *Advances in Neural Information Processing Systems*, 37:119123–119149, 2024.

[4] Omri Ben-Dov, Jake Fawkes, Samira Samadi, and Amartya Sanyal. The role of learning algorithms in collective action. In *International Conference on Machine Learning (ICML)*, volume 235, pages 3443–3461. PMLR, 2024.

[5] Omri Ben-Dov, Samira Samadi, Amartya Sanyal, and Alexandru Tifrea. Fairness for the people, by the people: Minority collective action. In *NeurIPS 2025 (Non-archival) Workshop on Algorithmic Collective Action (ACA)*, 2025.

[6] CAPA Strategies, LLC. Heat mapping – heat watch, 2025.

[7] City of Rotterdam, CIO Office. Rotterdam in transformation: Vision on the digital city 1.0, 2024.

[8] Alicia DeVrio, Motahhare Eslami, and Kenneth Holstein. Building, shifting, & employing power: A taxonomy of responses from below to algorithmic harm. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*, pages 1093–1106, 2024.

[9] Frances Ding, Moritz Hardt, John Miller, and Ludwig Schmidt. Retiring adult: New datasets for fair machine learning. In *Advances in Neural Information Processing Systems*, volume 34, pages 6478–6490, 2021.

[10] European Parliament and Council of the European Union. Regulation (eu) 2016/679 of the european parliament and of the council of 27 april 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing directive 95/46/ec (general data protection regulation). *Official Journal of the European Union*, (L 119):1–88, May 2016.

[11] Ekaterina Fedorova, Madeline Celi Kitch, and Chara Podimata. Altruistic collective action in recommender systems. In *NeurIPS 2025 (Non-archival) Workshop on Algorithmic Collective Action (ACA)*, 2025.

[12] Etienne Gauthier, Francis Bach, and Michael I. Jordan. Statistical collusion by collectives on learning platforms. In *Forty-second International Conference on Machine Learning (ICML)*, 2025.

[13] Government of Canada. Personal information protection and electronic documents act, 2000. S.C. 2000, c. 5.

[14] Donna Haraway. A cyborg manifesto: Science, technology, and socialist-feminism in the late twentieth century. In *The transgender studies reader*, pages 103–118. Routledge, 2013.

[15] Moritz Hardt, Eric Mazumdar, Celestine Mendler-Dünner, and Tijana Zrnic. Algorithmic collective action in machine learning. In *International Conference on Machine Learning (ICML)*, volume 202, pages 12570–12586. PMLR, 2023.

[16] Ayana Hussain, Cole Michael Thacker, Patrick Zhao, Aditya Karan, and Nicholas Vincent. Empowering users together: Connecting algorithmic collective action and explainable ai. In *NeurIPS 2025 (Non-archival) Workshop on Algorithmic Collective Action (ACA)*, 2025.

[17] Aditya Karan, Prabhat Kalle, Nicholas Vincent, and Hari Sundaram. Sync or sink: Bounds on algorithmic collective action with noise and multiple groups. In *NeurIPS 2025 (Non-archival) Workshop on Algorithmic Collective Action (ACA)*, 2025.

[18] Aditya Karan, Nicholas Vincent, Karrie Karahalios, and Hari Sundaram. Algorithmic collective action with two collectives. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, pages 1468–1483, 2025.

[19] Bruno Latour. *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford University Press, Oxford, 2005. ISBN 9780199256044.

[20] Kristina Lewandowska. Workers vs. the algorithm: Simulating collective action in gig-economy platforms. In *NeurIPS 2025 (Non-archival) Workshop on Algorithmic Collective Action (ACA)*, 2025.

[21] Krishna Kanth Nakka, Ahmed Frikha, Ricardo Mendes, Xue Jiang, and Xuebing Zhou. Piiscope: A comprehensive study on training data privacy leakage in pretrained llms. In *Proceedings of the 14th International Joint Conference on Natural Language Processing and the 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*, pages 3731–3765, 2025.

[22] National Oceanic and Atmospheric Administration (NOAA). Mapping campaigns. HEAT.gov, 2025.

[23] NIST. Artificial intelligence risk management framework: Generative artificial intelligence profile. *NIST Trustworthy and Responsible AI Gaithersburg, MD, USA*, 2024.

[24] Paulina Okunytė. Android ai app exposes nearly 2m user images and videos: anyone can watch your videos. Cybernews, February 2026. Published: 19 Feb 2026; last updated: 20 Feb 2026. Accessed: 24 Feb 2026.

[25] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, Jake Vanderplas, Alexandre Passos, David Cournapeau, Matthieu Brucher, Matthieu Perrot, and Édouard Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.

[26] Billy Perrigo. Exclusive: Openai used kenyan workers on less than $2 per hour to make chatgpt less toxic. *TIME*, January 2023. Published Jan 18, 2023. Accessed Feb 24, 2026.

[27] Maithra Raghu, Chiyuan Zhang, Jon Kleinberg, and Samy Bengio. Transfusion: Understanding transfer learning for medical imaging. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL `https://proceedings.neurips.cc/paper_files/paper/2019/file/eb1e78328c46506b46a4ac4a1e378b91-Paper.pdf`.

[28] Shiori Sagawa, Pang Wei Koh, Tatsunori B. Hashimoto, and Percy Liang. Distributionally robust neural networks. In *International Conference on Learning Representations*, 2020. URL `https://openreview.net/forum?id=ryxGuJrFvS`.

[29] Andrew D Selbst, Danah Boyd, Sorelle A Friedler, Suresh Venkatasubramanian, and Janet Vertesi. Fairness and abstraction in sociotechnical systems. In *Proceedings of the conference on fairness, accountability, and transparency*, pages 59–68, 2019.

[30] Dorothee Sigg, Moritz Hardt, and Celestine Mendler-Dünner. Decline now: A combinatorial model for algorithmic collective action. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*, CHI '25. Association for Computing Machinery, 2025. doi: 10.1145/3706598.3713966.

[31] Rushabh Solanki, Meghana Bhange, Ulrich Aïvodji, and Elliot Creager. Crowding out the noise: Algorithmic collective action under differential privacy. In *NeurIPS 2025 (Non-archival) Workshop on Algorithmic Collective Action*, 2025.

[32] State of California. California privacy rights act of 2020 (cpra), 2020. Proposition 24, approved November 3, 2020.

[33] Christine Utz, Martin Degeling, Sascha Fahl, Florian Schaub, and Thorsten Holz. (un) informed consent: Studying gdpr consent notices in the field. In *Proceedings of the 2019 acm sigsac conference on computer and communications security*, pages 973–990, 2019.

[34] Nicholas Vincent, Hanlin Li, Nicole Tilly, Stevie Chancellor, and Brent Hecht. Data leverage: A framework for empowering the public in its relationship with technology companies. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 215–227, 2021.

A Formal Setting and Proofs

A.1 Formal Setting

This section collects the measure-theoretic details of the setting of Section 2. All probability measures on $\mathcal{Z}$ are defined on the discrete σ-algebra $2^{\mathcal{Z}}$ (the power set of $\mathcal{Z}$, of cardinality $2^{|\mathcal{Z}|}$). The population is a probability space $(\Omega, 2^{\Omega}, \pi)$, with Ω finite, and $f_0 : \Omega \to \mathcal{Z}$ is a measurable map, so that $P_0 = (f_0)_{\#}\pi$ is a push-forward measure; each strategy $h_c : \mathcal{Z} \to \mathcal{Z}$ is measurable. If strategies are randomized, replace h_c by a Markov kernel K_c and $(h_c)_{\#}$ by the kernel pushforward $K_c^{\#}$; we use deterministic h_c for notational simplicity.

Sampling model. Sampling individuals i.i.d. from π is general enough to cover different settings, e.g., survey-like sampling, by interpreting π as the survey's sampling design over individuals (e.g., uniform, stratified, or weighted), with f_0 mapping each sampled respondent to their recorded (x, y), and internet/firm sampling, by interpreting π as the platform's exposure/participation process, so that i.i.d. draws from π represent repeated user interactions or users selected by the platform's logging pipeline. In both cases, the observed training data are the (x, y) pairs induced by sampling users via π and then recording their associated outcome through f_0.

Empirical measures. For $n_{N,c} > 0$, the empirical *pre-edit* distribution of collective Ω_c is $\hat{P}^{\mathrm{pre}}_{N,c} := \frac{1}{n_{N,c}} \sum_{i \in I_{N,c}} \delta_{Z_i}$, with feature marginal $\hat{P}^{\mathrm{pre},X}_{N,c}(B) := \frac{1}{n_{N,c}} \sum_{i \in I_{N,c}} \mathbf{1}_{\{X_i \in B\}}$, $B \subseteq \mathcal{X}$. The *post-edit* empirical measure of collective Ω_c is $\hat{P}_{N,c} := (h_c)_{\#}\hat{P}^{\mathrm{pre}}_{N,c}$, with $\hat{P}^{X}_{N,c}$ denoting its marginal on $\mathcal{X}$. For the non-collective users, if $n_{N,0} > 0$, define $\hat{P}_{N,0} := \frac{1}{n_{N,0}} \sum_{i : S_i \notin \cup_{c=1}^{M} \Omega_c} \delta_{Z_i}$. If $n_{N,0} = 0$, we set $\hat{P}_{N,0}$ equal to an arbitrary fixed probability measure on $\mathcal{Z}$; this convention has no effect on $\hat{P}_N$ in (2), because then $\alpha_{N,0} = 0$.

Contamination-suboptimality. The contamination-suboptimality condition above makes explicit the perturbation model underlying the margin argument used in [15, 12]. Namely, the classifier is assumed to be Bayes-optimal, or approximately optimal through the choice of ε, for a distribution $\hat{Q}_N$ obtained by an ε-contamination of the empirical training distribution $\hat{P}_N$. This is slightly stronger and more structured than merely requiring $\mathrm{TV}(\hat{Q}_N, \hat{P}_N) \le \varepsilon$. In fact, under symmetric total variation alone, the threshold $\varepsilon/(1-\varepsilon)$ is not generally implied, and the standard conservative margin threshold is 2ε. Thus, the contamination formulation is the precise source of the $\varepsilon/(1-\varepsilon)$ terms in our bounds. We use this formulation in order to stay close to the perturbation model and constants used in [15, 12], while making explicit the additional structure needed for the $\varepsilon/(1-\varepsilon)$ threshold. All results below remain valid under the looser symmetric total-variation suboptimality model, provided every occurrence of $\varepsilon/(1-\varepsilon)$ in the margin conditions and final bounds is replaced by the conservative threshold 2ε.

As a sanity check, the following proposition shows that the empirical distribution in (2) converges to the distribution in (1) in the infinite-sample regime.

Proposition A.1. *Under the neutrality assumption $P_0(z) = P_0\big(z \mid \omega \in \Omega_c\big)$, it holds $\hat{P}_N \to P$ almost surely. In addition, for each collective Ω_c, it holds $\alpha_{N,c} \to \alpha_c$ and $\hat{P}_{N,c} \to P_c$ almost surely.*

Proof. Fix $z \in \mathcal{Z}$. For $c = 1, \dots, M$ define

$$Y_i^{(c,z)} \;:=\; \mathbf{1}_{\{\, S_i \in \Omega_c,\; h_c(Z_i) = z \,\}}, \qquad Y_i^{(0,z)} \;:=\; \mathbf{1}_{\{\, S_i \notin \cup_{c=1}^{M} \Omega_c,\; Z_i = z \,\}}.$$

Since $(S_i)_{i \ge 1}$ are i.i.d. with law π and $Z_i = f_0(S_i)$, the vectors $\big(Y_i^{(0,z)}, Y_i^{(1,z)}, \dots, Y_i^{(M,z)}\big)$ are i.i.d. By neutrality, $Z_i \mid S_i \in \Omega_c \sim P_0$ for each c, hence

$$\mathbb{E}\, Y_1^{(0,z)} = (1 - \bar{\alpha})\, P_0(z), \qquad \mathbb{E}\, Y_1^{(c,z)} = \alpha_c\, P_c(z) \;\; (c = 1, \dots, M),$$

where $P_c = (h_c)_{\#}P_0$. By the strong law of large numbers,

$$\frac{1}{N} \sum_{i=1}^{N} Y_i^{(0,z)} \underset{\text{a.s.}}{\longrightarrow} (1 - \bar{\alpha})\, P_0(z), \qquad \frac{1}{N} \sum_{i=1}^{N} Y_i^{(c,z)} \underset{\text{a.s.}}{\longrightarrow} \alpha_c\, P_c(z).$$

Noting that

$$\hat{P}_N(z) \;=\; \big(1 - \bar{\alpha}_N\big)\, \hat{P}_{N,0}(z) \;+ \sum_{c : \alpha_{N,c} > 0} \alpha_{N,c}\, \hat{P}_{N,c}(z) \;=\; \frac{1}{N} \sum_{i=1}^{N} \Big(Y_i^{(0,z)} + \sum_{c=1}^{M} Y_i^{(c,z)} \Big),$$

we obtain, for each $z \in \mathcal{Z}$,

$$\hat{P}_N(z) \underset{\text{a.s.}}{\longrightarrow} (1-\bar{\alpha})\, P_0(z) + \sum_{c=1}^{M} \alpha_c\, P_c(z) \;=\; P(z).$$

Since $\mathcal{Z}$ is finite, coordinatewise convergence implies convergence of measures, i.e., $\hat{P}_N \to P$ almost surely. Similarly, $\alpha_{N,c} = N^{-1}\sum_{i=1}^{N} \mathbf{1}_{\{S_i \in \Omega_c\}} \to \alpha_c$ almost surely again by the strong law of large numbers, and, for $\alpha_c > 0$, the convergence $\alpha_{N,c}\hat{P}_{N,c}(z) \to \alpha_c P_c(z)$ for all z together with $\alpha_{N,c} \to \alpha_c$ yields $\hat{P}_{N,c} \to P_c$ almost surely. □

A.2 Proofs

Theorem 3.1 (Signal planting with multiple collectives)**.** *Under the neutrality assumption, assuming that $\hat{m}$ is ε-contamination-suboptimal on $\mathcal{X}_c^*$ under $\hat{P}_N$, and assuming all collectives either employ the feature-label or the feature-only strategy from* (11)-(12)*, respectively, with probability at least $1-\delta$ over the draw of the training and test samples, with $\delta > 0$, the per-collective success of collective $I_{N,c}$ satisfies*

$$\begin{aligned}\hat{S}_c(\alpha_{N,c}) \;\geq\; \mathbb{P}_{x\sim\hat{P}^{\mathrm{pre},X}_{N,c}}\Big[\; & \alpha_{N,c}\Big(\hat{p}_c^{\,X}\big(g_c(x)\big) - 2R_{\tilde{\delta}}(n_{N,c})\Big) - \sum_{d\neq c}\alpha_{N,d}\,\Delta_{c\to d}\big(g_c(x)\big) \\ & - \alpha_{N,0}\,\Delta_{c,0}^{(n_{N,c})}\big(g_c(x)\big) - \; 2\alpha_{N,0}\Big(R_{\tilde{\delta}}(n_{N,0}) + R_{\tilde{\delta}}(n_{N,c})\Big) - \frac{\varepsilon}{1-\varepsilon} > 0\Big] \\ & - R_{\delta/4}(n_{N,c}) - R_{\delta/4}(N^{\mathrm{test}}),\end{aligned} \tag{13}$$

where $K_c := |\mathcal{X}_c^*|\,|\mathcal{Y}|$ and $\tilde{\delta} := \frac{\delta}{10K_c}$.

Proof of Theorem 3.1. The proof follows six steps, here is the outline:

1. We use ε-contamination-suboptimality of $\hat{m}$ to obtain a pointwise margin condition under $\hat{P}_N$ that guarantees correct prediction at a feature value x^* whenever a suitable margin inequality holds.
2. We express this margin in terms of the empirical mixture $\hat{P}_N$ by decomposing contributions from each collective, and define the resulting exact empirical margin $M_c(x^*)$.
3. We lower bound $M_c(x^*)$ by a fixed population-based margin $\overline{M}_c(x^*)$ and then compare $\overline{M}_c(x^*)$ to the corrected empirical margin appearing in the theorem statement.
4. We convert the resulting fixed margin condition into a lower bound on the test success probability under P_0^X, using concentration on the test sample.
5. We replace this population probability by the corresponding empirical probability over the pre-edit feature distribution of collective c, again via Hoeffding's inequality, and then pass pointwise to the corrected empirical margin event.
6. We combine all concentration events via a union bound and obtain the desired lower bound on $\hat{S}_c(\alpha_{N,c})$.

Throughout the proof, we work conditionally on the realized sample sizes $\{n_{N,0}, n_{N,1}, \ldots, n_{N,M}\}$, and all probabilities $\mathbb{P}[\cdot]$ are with respect to the remaining randomness in the training and test draws. We write $R_\gamma(k) = \sqrt{\log(1/\gamma)/(2k)}$.
Step 1. We first need the following simple lemma.

Lemma A.2 (Margin implication for planting)**.** *Assume $\hat{m}$ is ε-contamination-suboptimal on $\mathcal{X}'$ under $\hat{P}_N$. If $x^* \in \mathcal{X}'$ and*

$$\hat{P}_N(x^*, y^*) > \hat{P}_N(x^*, y') + \frac{\varepsilon}{1-\varepsilon} \qquad \forall y' \neq y^*,$$

then $\hat{m}(x^*) = y^*$.

Proof. By ε-contamination-suboptimality, there exists $\hat{Q}_N = (1-\varepsilon)\hat{P}_N + \varepsilon R_N$ such that $\hat{m}(x) \in \operatorname{argmax}_y \hat{Q}_N(y \mid x)$ on $\mathcal{X}'$. Fix $x^* \in \mathcal{X}'$ and $y' \neq y^*$. Then

$$\hat{Q}_N(x^*, y^*) - \hat{Q}_N(x^*, y') = (1-\varepsilon)\big(\hat{P}_N(x^*, y^*) - \hat{P}_N(x^*, y')\big) + \varepsilon\big(R_N(x^*, y^*) - R_N(x^*, y')\big).$$

Since $R_N(x^*, y^*) - R_N(x^*, y') \geq -1$, the assumed margin gives

$$\hat{Q}_N(x^*, y^*) - \hat{Q}_N(x^*, y') > (1-\varepsilon)\frac{\varepsilon}{1-\varepsilon} - \varepsilon = 0.$$

Thus $\hat{Q}_N(x^*, y^*) > \hat{Q}_N(x^*, y')$ for all $y' \neq y^*$. Moreover, the margin condition implies $\hat{P}_N^X(x^*) > 0$, hence $\hat{Q}_N^X(x^*) > 0$. Therefore y^* is the unique maximizer of $\hat{Q}_N(\cdot \mid x^*)$, and $\hat{m}(x^*) = y^*$. □

Since $\hat{m}$ is ε-contamination-suboptimal on $\mathcal{X}_c^*$ under $\hat{P}_N$, the margin implication for planting above gives that, for any $x^* \in \mathcal{X}_c^*$,

$$\forall y' \neq y_c^* :\ \hat{P}_N(x^*, y_c^*) > \hat{P}_N(x^*, y') + \frac{\varepsilon}{1-\varepsilon} \quad \Longrightarrow \quad \hat{m}(x^*) = y_c^*. \tag{23}$$

Step 2. Fix the target collective c and an arbitrary $x^* \in \mathcal{X}_c^*$. By the empirical mixture decomposition,

$$\hat{P}_N = \alpha_{N,0}\,\hat{P}_{N,0} + \sum_{d=1}^{M} \alpha_{N,d}\,\hat{P}_{N,d},$$

we have, for any $y \in \mathcal{Y}$,

$$\hat{P}_N(x^*, y) = \alpha_{N,0}\,\hat{P}_{N,0}(x^*, y) + \sum_{d=1}^{M} \alpha_{N,d}\,\hat{P}_{N,d}(x^*, y). \tag{24}$$

Under either planting strategy, every edited point from collective c with post-edit feature value x^* carries label y_c^*. Hence, for any $x^* \in \mathcal{X}_c^*$ and any $y' \neq y_c^*$,

$$\hat{P}_{N,c}(x^*, y') = 0, \qquad \hat{P}_{N,c}(x^*, y_c^*) = \hat{P}_{N,c}^X(x^*) = \hat{p}_c^X(x^*).$$

Subtracting (24) for (x^*, y') from that for (x^*, y_c^*) and defining

$$\Delta_{c,0}(x^*) \ :=\ \max_{y' \in \mathcal{Y}\setminus\{y_c^*\}} \big(\hat{P}_{N,0}(x^*, y') - \hat{P}_{N,0}(x^*, y_c^*)\big),$$

we obtain for any $y' \neq y_c^*$:

$$\begin{aligned}
\hat{P}_N(x^*, y_c^*) - \hat{P}_N(x^*, y') &= \alpha_{N,c}\,\hat{p}_c^X(x^*) + \sum_{d \neq c} \alpha_{N,d}\big(\hat{P}_{N,d}(x^*, y_c^*) - \hat{P}_{N,d}(x^*, y')\big) \\
&\quad + \alpha_{N,0}\big(\hat{P}_{N,0}(x^*, y_c^*) - \hat{P}_{N,0}(x^*, y')\big) \\
&\geq \alpha_{N,c}\,\hat{p}_c^X(x^*) - \sum_{d \neq c} \alpha_{N,d}\,\Delta_{c\to d}(x^*) - \alpha_{N,0}\,\Delta_{c,0}(x^*).
\end{aligned} \tag{25}$$

Define the exact empirical margin

$$M_c(x^*) \ :=\ \alpha_{N,c}\,\hat{p}_c^X(x^*) - \sum_{d \neq c} \alpha_{N,d}\,\Delta_{c\to d}(x^*) - \alpha_{N,0}\,\Delta_{c,0}(x^*), \qquad x^* \in \mathcal{X}_c^*.$$

Then (25) yields, for all $x^* \in \mathcal{X}_c^*$ and all $y' \neq y_c^*$,

$$\hat{P}_N(x^*, y_c^*) - \hat{P}_N(x^*, y') \ \geq\ M_c(x^*).$$

Combining this with (23) gives

$$x^* \in \mathcal{X}_c^*, \quad M_c(x^*) > \frac{\varepsilon}{1-\varepsilon} \quad \Longrightarrow \quad \hat{m}(x^*) = y_c^*. \tag{26}$$

Step 3. Define the population signal prevalence

$$p_{c,\mathrm{pop}}^X(x^*) \ :=\ P_c^X(x^*), \qquad x^* \in \mathcal{X}_c^*,$$

and the population non-collective conflict

$$\Delta_{c,0}^{\mathrm{pop}}(x^*) \;:=\; \max_{y' \neq y_c^*} \big(P_0(x^*, y') - P_0(x^*, y_c^*)\big).$$

By neutrality, the edited points of collective c are i.i.d. from P_c, the non-collective points are i.i.d. from P_0, and the pre-edit points of collective c are i.i.d. from P_0.

Applying Hoeffding's inequality and a union bound over the scalar one-sided concentration events displayed below, we obtain that, with probability at least $1 - \delta/2$, the following inequalities hold simultaneously. Indeed, the first two families contain $2|\mathcal{X}_c^*|$ scalar events, and the last four families contain $4|\mathcal{X}_c^*||\mathcal{Y}| = 4K_c$ scalar events. Since $|\mathcal{Y}| \geq 2$, the total number of scalar events is at most $5K_c$. With $\tilde{\delta} = \delta/(10K_c)$, the total failure probability is therefore at most $\delta/2$. The inequalities are:

$$\hat{p}_c^X(x^*) \geq p_{c,\mathrm{pop}}^X(x^*) - R_{\tilde{\delta}}(n_{N,c}), \qquad \forall x^* \in \mathcal{X}_c^*, \tag{27}$$

$$\hat{p}_c^X(x^*) \leq p_{c,\mathrm{pop}}^X(x^*) + R_{\tilde{\delta}}(n_{N,c}), \qquad \forall x^* \in \mathcal{X}_c^*, \tag{28}$$

$$\hat{P}_{N,0}(x^*, y) \leq P_0(x^*, y) + R_{\tilde{\delta}}(n_{N,0}), \qquad \forall (x^*, y) \in \mathcal{X}_c^* \times \mathcal{Y}, \tag{29}$$

$$P_0(x^*, y) \leq \hat{P}_{N,0}(x^*, y) + R_{\tilde{\delta}}(n_{N,0}), \qquad \forall (x^*, y) \in \mathcal{X}_c^* \times \mathcal{Y}, \tag{30}$$

$$\hat{P}_{N,c}^{\mathrm{pre}}(x^*, y) \leq P_0(x^*, y) + R_{\tilde{\delta}}(n_{N,c}), \qquad \forall (x^*, y) \in \mathcal{X}_c^* \times \mathcal{Y}, \tag{31}$$

$$P_0(x^*, y) \leq \hat{P}_{N,c}^{\mathrm{pre}}(x^*, y) + R_{\tilde{\delta}}(n_{N,c}), \qquad \forall (x^*, y) \in \mathcal{X}_c^* \times \mathcal{Y}. \tag{32}$$

Using (29)–(30), we obtain, for every $x^* \in \mathcal{X}_c^*$,

$$\Delta_{c,0}(x^*) \;\leq\; \Delta_{c,0}^{\mathrm{pop}}(x^*) + 2R_{\tilde{\delta}}(n_{N,0}). \tag{33}$$

Using (31)–(32), we obtain, for every $x^* \in \mathcal{X}_c^*$,

$$\Delta_{c,0}^{\mathrm{pop}}(x^*) \;\leq\; \Delta_{c,0}^{(n_{N,c})}(x^*) + 2R_{\tilde{\delta}}(n_{N,c}). \tag{34}$$

Define the fixed population-based margin

$$\overline{M}_c(x^*) \;:=\; \alpha_{N,c}\Big(p_{c,\mathrm{pop}}^X(x^*) - R_{\tilde{\delta}}(n_{N,c})\Big) - \sum_{d \neq c} \alpha_{N,d}\, \Delta_{c \to d}(x^*) - \alpha_{N,0}\Big(\Delta_{c,0}^{\mathrm{pop}}(x^*) + 2R_{\tilde{\delta}}(n_{N,0})\Big),$$

and the corrected empirical margin

$$\begin{aligned}\widetilde{M}_c(x^*) \;:=\; \alpha_{N,c}\Big(\hat{p}_c^X(x^*) - 2R_{\tilde{\delta}}(n_{N,c})\Big) - \sum_{d \neq c} \alpha_{N,d}\, \Delta_{c \to d}(x^*) - \alpha_{N,0}\, \Delta_{c,0}^{(n_{N,c})}(x^*) \\ - 2\alpha_{N,0}\Big(R_{\tilde{\delta}}(n_{N,0}) + R_{\tilde{\delta}}(n_{N,c})\Big).\end{aligned}$$

By (27) and (33), we have, for all $x^* \in \mathcal{X}_c^*$,

$$M_c(x^*) \;\geq\; \overline{M}_c(x^*). \tag{35}$$

By (28) and (34), we also have, for all $x^* \in \mathcal{X}_c^*$,

$$\overline{M}_c(x^*) \;\geq\; \widetilde{M}_c(x^*). \tag{36}$$

Combining (26) and (35), we obtain

$$x^* \in \mathcal{X}_c^*, \quad \overline{M}_c(x^*) > \frac{\varepsilon}{1-\varepsilon} \quad \Longrightarrow \quad \hat{m}(x^*) = y_c^*. \tag{37}$$

By (36), this further implies

$$x^* \in \mathcal{X}_c^*, \quad \widetilde{M}_c(x^*) > \frac{\varepsilon}{1-\varepsilon} \quad \Longrightarrow \quad \hat{m}(x^*) = y_c^*. \tag{38}$$

Step 4. Let X^{test} denote a generic feature drawn from the test distribution P_0^X, and write

$$\varphi(x) \;:=\; \mathbf{1}_{\{\hat{m}(g_c(x)) = y_c^*\}}, \qquad x \in \mathcal{X}.$$

By definition,

$$\hat{S}_c(\alpha_{N,c}) = \frac{1}{N^{\mathrm{test}}} \sum_{i=1}^{N^{\mathrm{test}}} \varphi(X_i^{\mathrm{test}}).$$

Applying Hoeffding's inequality to the i.i.d. Bernoulli variables $\varphi(X_i^{\mathrm{test}})$, with failure budget $\delta/4$, we obtain that, with probability at least $1 - \delta/4$,

$$\hat{S}_c(\alpha_{N,c}) \;\geq\; \mathbb{P}_{X\sim P_0^X}\big[\hat{m}(g_c(X)) = y_c^*\big] - R_{\delta/4}(N^{\mathrm{test}}). \tag{39}$$

Since $g_c(X) \in \mathcal{X}_c^*$ for all $X \in \mathcal{X}$, (37) implies

$$\Big\{\overline{M}_c(g_c(X)) > \tfrac{\varepsilon}{1-\varepsilon}\Big\} \subseteq \big\{\hat{m}(g_c(X)) = y_c^*\big\}.$$

Therefore, with probability at least $1 - \delta/4$ conditional on the Step 3 event,

$$\hat{S}_c(\alpha_{N,c}) \;\geq\; \mathbb{P}_{X\sim P_0^X}\Big[\overline{M}_c(g_c(X)) > \tfrac{\varepsilon}{1-\varepsilon}\Big] - R_{\delta/4}(N^{\mathrm{test}}). \tag{40}$$

Step 5. Define the fixed event

$$A_c^{\mathrm{pop}} \;:=\; \Big\{x \in \mathcal{X} : \overline{M}_c(g_c(x)) > \tfrac{\varepsilon}{1-\varepsilon}\Big\}.$$

Conditional on all training data outside collective c, the event A_c^{pop} is fixed. By neutrality and the sampling model, the pre-edit features $\{X_i : i \in I_{N,c}\}$ are i.i.d. with law P_0^X, conditionally on $n_{N,c}$. Applying Hoeffding's inequality to the i.i.d. variables $\mathbf{1}_{\{X_i \in A_c^{\mathrm{pop}}\}}$, $i \in I_{N,c}$, with failure budget $\delta/4$, we obtain that, with probability at least $1 - \delta/4$,

$$\mathbb{P}_{X\sim P_0^X}\big[X \in A_c^{\mathrm{pop}}\big] \;\geq\; \mathbb{P}_{x\sim \hat{P}_{N,c}^{\mathrm{pre},X}}\big[x \in A_c^{\mathrm{pop}}\big] - R_{\delta/4}(n_{N,c}). \tag{41}$$

On the Step 3 event, (36) implies the pointwise inclusion

$$\Big\{\widetilde{M}_c(g_c(x)) > \tfrac{\varepsilon}{1-\varepsilon}\Big\} \subseteq A_c^{\mathrm{pop}}.$$

Hence,

$$\mathbb{P}_{X\sim P_0^X}\Big[\overline{M}_c(g_c(X)) > \tfrac{\varepsilon}{1-\varepsilon}\Big] \;\geq\; \mathbb{P}_{x\sim \hat{P}_{N,c}^{\mathrm{pre},X}}\Big[\widetilde{M}_c(g_c(x)) > \tfrac{\varepsilon}{1-\varepsilon}\Big] - R_{\delta/4}(n_{N,c}). \tag{42}$$

Step 6. By the union bound, the Step 3 concentration event, (39), and (41) hold simultaneously with probability at least $1 - \delta$. Combining (40) and (42), we conclude that, with probability at least $1 - \delta$,

$$\begin{aligned}
\hat{S}_c(\alpha_{N,c}) \geq\; & \mathbb{P}_{x\sim \hat{P}_{N,c}^{\mathrm{pre},X}}\Big[\widetilde{M}_c\big(g_c(x)\big) > \tfrac{\varepsilon}{1-\varepsilon}\Big] - R_{\delta/4}(n_{N,c}) - R_{\delta/4}(N^{\mathrm{test}}) \\
=\; & \mathbb{P}_{x\sim \hat{P}_{N,c}^{\mathrm{pre},X}}\Big[\; \alpha_{N,c}\Big(\hat{p}_c^X\big(g_c(x)\big) - 2R_{\tilde{\delta}}(n_{N,c})\Big) - \sum_{d\neq c} \alpha_{N,d}\, \Delta_{c\to d}\big(g_c(x)\big) \\
& - \alpha_{N,0}\, \Delta_{c,0}^{(n_{N,c})}\big(g_c(x)\big) - \; 2\alpha_{N,0}\Big(R_{\tilde{\delta}}(n_{N,0}) + R_{\tilde{\delta}}(n_{N,c})\Big) - \frac{\varepsilon}{1-\varepsilon} > 0\Big] \\
& - R_{\delta/4}(n_{N,c}) - R_{\delta/4}(N^{\mathrm{test}}),
\end{aligned}$$

which is exactly (13). This completes the proof. □

Corollary 3.2 (Symmetric regime)**.** *In the setting of Theorem 3.1, let $M \geq 2$ and $\alpha_{N,d} = \bar{\alpha}_N/M$ for every collective d, and fix $x^* \in \mathcal{X}_c^*$ with $\hat{p}_c^X(x^*) = p$, $\Delta_{c\to d}(x^*) = \Delta$ for all $d \neq c$, and $\Delta_{c,0}^{(n_{N,c})}(x^*) = \Delta_0$. Let $\kappa := (1-\bar{\alpha}_N)\Delta_0 + \frac{\varepsilon}{1-\varepsilon} + \rho_N$, where $\rho_N := \frac{2\bar{\alpha}_N}{M} R_{\tilde{\delta}}(n_{N,c}) + 2(1-\bar{\alpha}_N)\big(R_{\tilde{\delta}}(n_{N,0}) + R_{\tilde{\delta}}(n_{N,c})\big)$ collects the radii. Then the margin in* (13) *is positive at x^* iff*

$$p \;>\; (M-1)\,\Delta + \frac{M}{\bar{\alpha}_N}\,\kappa. \tag{14}$$

In particular, (i) under full alignment ($\Delta = -p$) the condition reads $\bar{\alpha}_N\, p > \kappa$, (ii) under mutual neutrality ($\Delta = 0$) it reads $\frac{\bar{\alpha}_N}{M}\, p > \kappa$, (iii) under full conflict ($\Delta = p$) it fails for every $M \geq 2$ whenever $\kappa \geq 0$, e.g., when $\Delta_0 \geq 0$.

Proof of Corollary 3.2. With $\alpha_{N,c} = \bar{\alpha}_N/M$ for every collective and $\alpha_{N,0} = 1 - \bar{\alpha}_N$, the margin of (13) at x^* is

$$\frac{\bar{\alpha}_N}{M}\big(p - 2R_{\tilde{\delta}}(n_{N,c})\big) - (M-1)\frac{\bar{\alpha}_N}{M}\,\Delta - (1-\bar{\alpha}_N)\Delta_0 \\ - 2(1-\bar{\alpha}_N)\big(R_{\tilde{\delta}}(n_{N,0}) + R_{\tilde{\delta}}(n_{N,c})\big) - \frac{\varepsilon}{1-\varepsilon} = \frac{\bar{\alpha}_N}{M}\big(p - (M-1)\Delta\big) - \kappa,$$

which is positive if and only if (14) holds, since $\bar{\alpha}_N/M > 0$. Substituting $\Delta = -p$ and $\Delta = 0$ gives (i) and (ii). For $\Delta = p$ the margin equals $\frac{\bar{\alpha}_N}{M}(2-M)\,p - \kappa \le -\kappa \le 0$, because $M \ge 2$ and $p \ge 0$; this gives (iii). □

Theorem 4.1 (Signal unplanting with multiple collectives)**.** *Under the neutrality assumption, assuming that $\hat{m}$ is ε-contamination-suboptimal on $\mathcal{X}_c^*$ under $\hat{P}_N$, assuming all collectives employ the adaptive unplanting strategy from* (19) *with estimation splits $\{E_{N,c}\}_c$, then, with probability at least $1-\delta$ over the draw of the training and test samples, with $\delta > 0$, and the random choice of $\{E_{N,c}\}_c$, the per-collective success of $I_{N,c}$ satisfies*

$$\hat{S}_c(\alpha_{N,c}) \;\ge\; \mathbb{P}_{x\sim\hat{P}^{\mathrm{pre},R,X}_{N,c}}\Big[\,\alpha_{N,c}\Big(\hat{p}_c^X\big(g_c(x)\big) - 2R_{\tilde{\delta}}(n_{N,c})\Big) - \sum_{d\neq c}\alpha_{N,d}\,\Delta^{\mathrm{unp}}_{c\leftarrow d}\big(g_c(x)\big) \\ - \alpha_{N,0}\,\Delta^{\mathrm{unp},(n_{N,c}-n_{e,c})}_{c,0}\big(g_c(x)\big) - 2\alpha_{N,0}\Big(R_{\tilde{\delta}}(n_{N,0}) + R_{\tilde{\delta}}(n_{N,c}-n_{e,c})\Big) - \frac{\varepsilon}{1-\varepsilon} > 0\Big] \\ - R_{\delta/4}(n_{N,c}-n_{e,c}) - R_{\delta/4}(N^{\mathrm{test}}), \tag{22}$$

where $K_c := |\mathcal{X}_c^*|\,|\mathcal{Y}|$ and $\tilde{\delta} := \frac{\delta}{10K_c}$.

Proof of Theorem 4.1. The proof follows closely the same six-step structure as Theorem 3.1. We work conditionally on the realized sample sizes

$$\{n_{N,0}, n_{N,1}, \ldots, n_{N,M}\}.$$

The split $E_{N,c}$ is chosen independently of the data. For concentration bounds involving the data-dependent labels $y'_c(\cdot)$, we further condition on the estimation sigma-field

$$\mathcal{F}_{E,c} := \sigma(E_{N,c}, \{Z_i : i \in E_{N,c}\})\,.$$

Under this conditioning, the map $y'_c(\cdot)$ is fixed, while the holdout pre-edit points in $R_{N,c}$ remain i.i.d. from P_0 and independent of the label-selection step. The concentration bound for the post-edit feature marginal $\hat{p}_c^X$ is justified separately, before conditioning on the observed estimation data, because the feature part of the adaptive strategy is the fixed map g_c and does not depend on $y'_c(\cdot)$.

Step 1. We first prove a simple margin implication.

Lemma A.3 (Contamination margin implication: unplanting)**.** *Assume $\hat{m}$ is ε-contamination-suboptimal on $\mathcal{X}'$ under $\hat{P}_N$. If $x^* \in \mathcal{X}'$ and there exists $y' \neq y^*$ such that*

$$\hat{P}_N(x^*, y') > \hat{P}_N(x^*, y^*) + \frac{\varepsilon}{1-\varepsilon},$$

then

$$\hat{m}(x^*) \neq y^*.$$

Proof. By ε-contamination-suboptimality, there exist probability measures $\hat{Q}_N$ and R_N on $\mathcal{X}\times\mathcal{Y}$ such that

$$\hat{Q}_N = (1-\varepsilon)\hat{P}_N + \varepsilon R_N,$$

and such that, for every $x \in \mathcal{X}'$ with $\hat{Q}_N^X(x) > 0$,

$$\hat{m}(x) \in \operatorname*{argmax}_{y\in\mathcal{Y}} \hat{Q}_N(y \mid x).$$

Fix $x^* \in \mathcal{X}'$ and suppose that, for some $y' \neq y^*$,

$$\hat{P}_N(x^*, y') > \hat{P}_N(x^*, y^*) + \frac{\varepsilon}{1-\varepsilon}.$$

Then

$$\hat{P}_N(x^*, y') - \hat{P}_N(x^*, y^*) > \frac{\varepsilon}{1-\varepsilon}.$$

Using the contamination representation, we have

$$\begin{aligned}\hat{Q}_N(x^*, y') - \hat{Q}_N(x^*, y^*) = (1-\varepsilon)\big(\hat{P}_N(x^*, y') - \hat{P}_N(x^*, y^*)\big) \\ + \varepsilon\big(R_N(x^*, y') - R_N(x^*, y^*)\big).\end{aligned}$$

Since R_N is a probability measure,

$$R_N(x^*, y') - R_N(x^*, y^*) \geq -1.$$

Therefore,

$$\begin{aligned}\hat{Q}_N(x^*, y') - \hat{Q}_N(x^*, y^*) &> (1-\varepsilon)\frac{\varepsilon}{1-\varepsilon} - \varepsilon \\ &= 0.\end{aligned}$$

Hence

$$\hat{Q}_N(x^*, y') > \hat{Q}_N(x^*, y^*).$$

Moreover, the strict inequality

$$\hat{P}_N(x^*, y') > \hat{P}_N(x^*, y^*) + \frac{\varepsilon}{1-\varepsilon}$$

implies $\hat{P}_N^X(x^*) > 0$. Since $\varepsilon < 1$, we also have

$$\hat{Q}_N^X(x^*) = (1-\varepsilon)\hat{P}_N^X(x^*) + \varepsilon R_N^X(x^*) > 0.$$

Thus the conditional probabilities $\hat{Q}_N(y \mid x^*)$ are well defined. Because

$$\hat{Q}_N(x^*, y') > \hat{Q}_N(x^*, y^*),$$

and both labels share the same marginal denominator $\hat{Q}_N^X(x^*)$, we get

$$\hat{Q}_N(y' \mid x^*) > \hat{Q}_N(y^* \mid x^*).$$

Therefore y^* cannot be an argmax of $\hat{Q}_N(\cdot \mid x^*)$. Since $\hat{m}(x^*)$ is chosen among the argmax labels under $\hat{Q}_N$, it follows that

$$\hat{m}(x^*) \neq y^*.$$

□

Since $\hat{m}$ is ε-contamination-suboptimal on $\mathcal{X}_c^*$ under $\hat{P}_N$, the margin implication above gives that, for any $x^* \in \mathcal{X}_c^*$,

$$\hat{P}_N(x^*, y_c'(x^*)) > \hat{P}_N(x^*, y_c^*) + \frac{\varepsilon}{1-\varepsilon} \quad \Longrightarrow \quad \hat{m}(x^*) \neq y_c^*. \tag{43}$$

Step 2. Fix $x^* \in \mathcal{X}_c^*$. Using the empirical mixture decomposition of $\hat{P}_N$ and the adaptive feature–label unplanting property

$$\hat{P}_{N,c}(x^*, y_c^*) = 0, \qquad \hat{P}_{N,c}(x^*, y_c'(x^*)) = \hat{p}_c^X(x^*),$$

we obtain

$$\begin{aligned}&\hat{P}_N(x^*, y_c'(x^*)) - \hat{P}_N(x^*, y_c^*) \\ &\quad = \alpha_{N,c}\hat{p}_c^X(x^*) + \sum_{d \neq c} \alpha_{N,d}\Big(\hat{P}_{N,d}(x^*, y_c'(x^*)) - \hat{P}_{N,d}(x^*, y_c^*)\Big) \\ &\qquad + \alpha_{N,0}\Big(\hat{P}_{N,0}(x^*, y_c'(x^*)) - \hat{P}_{N,0}(x^*, y_c^*)\Big) \\ &\quad = \alpha_{N,c}\hat{p}_c^X(x^*) - \sum_{d \neq c} \alpha_{N,d}\Delta_{c \leftarrow d}^{\mathrm{unp}}(x^*) - \alpha_{N,0}\Delta_{c,0}^{\mathrm{unp}}(x^*),\end{aligned}$$

where

$$\Delta_{c,0}^{\mathrm{unp}}(x^*) := \hat{P}_{N,0}(x^*, y_c^*) - \hat{P}_{N,0}(x^*, y_c'(x^*)).$$

Define the exact empirical unplanting margin

$$M_c^{\mathrm{unp}}(x^*) := \alpha_{N,c}\hat{p}_c^X(x^*) - \sum_{d\neq c}\alpha_{N,d}\Delta_{c\leftarrow d}^{\mathrm{unp}}(x^*) - \alpha_{N,0}\Delta_{c,0}^{\mathrm{unp}}(x^*).$$

Then, by (43),

$$x^* \in \mathcal{X}_c^*, \quad M_c^{\mathrm{unp}}(x^*) > \frac{\varepsilon}{1-\varepsilon} \quad \Longrightarrow \quad \hat{m}(x^*) \neq y_c^*. \tag{44}$$

Step 3. Define the population signal prevalence

$$p_{c,\mathrm{pop}}^X(x^*) := (g_c)_{\#}P_0^X(x^*) = P_0^X\left(g_c^{-1}(\{x^*\})\right), \qquad x^* \in \mathcal{X}_c^*.$$

Since the feature part of the adaptive unplanting strategy is the fixed map g_c, the empirical post-edit feature marginal $\hat{p}_c^X$ is the empirical marginal of the fixed pushforward of the collective pre-edit sample under g_c. Hence, conditionally on $n_{N,c}$, Hoeffding's inequality gives the two-sided feature-marginal concentration bounds

$$\hat{p}_c^X(x^*) \geq p_{c,\mathrm{pop}}^X(x^*) - R_{\tilde{\delta}}(n_{N,c}), \qquad \forall x^* \in \mathcal{X}_c^*, \tag{45}$$

$$\hat{p}_c^X(x^*) \leq p_{c,\mathrm{pop}}^X(x^*) + R_{\tilde{\delta}}(n_{N,c}), \qquad \forall x^* \in \mathcal{X}_c^*. \tag{46}$$

Next, condition on the estimation sigma-field $\mathcal{F}_{E,c}$. Under this conditioning, the map $y_c'(\cdot)$ is fixed, and the holdout pre-edit points in $R_{N,c}$ are i.i.d. from P_0. Define the population non-collective unplanting conflict

$$\Delta_{c,0}^{\mathrm{unp,pop}}(x^*) := P_0(x^*, y_c^*) - P_0(x^*, y_c'(x^*)).$$

Together with the two feature-marginal concentration families (45)–(46), Hoeffding's inequality and a union bound over the scalar one-sided events displayed below imply that, with probability at least $1-\delta/2$, all the concentration inequalities in this step hold simultaneously:

$$\hat{P}_{N,0}(x^*, y) \leq P_0(x^*, y) + R_{\tilde{\delta}}(n_{N,0}), \qquad \forall (x^*, y) \in \mathcal{X}_c^* \times \mathcal{Y}, \tag{47}$$

$$P_0(x^*, y) \leq \hat{P}_{N,0}(x^*, y) + R_{\tilde{\delta}}(n_{N,0}), \qquad \forall (x^*, y) \in \mathcal{X}_c^* \times \mathcal{Y}, \tag{48}$$

$$\hat{P}_{N,c}^{\mathrm{pre},R}(x^*, y) \leq P_0(x^*, y) + R_{\tilde{\delta}}(n_{N,c} - n_{e,c}), \qquad \forall (x^*, y) \in \mathcal{X}_c^* \times \mathcal{Y}, \tag{49}$$

$$P_0(x^*, y) \leq \hat{P}_{N,c}^{\mathrm{pre},R}(x^*, y) + R_{\tilde{\delta}}(n_{N,c} - n_{e,c}), \qquad \forall (x^*, y) \in \mathcal{X}_c^* \times \mathcal{Y}. \tag{50}$$

Indeed, the two feature-marginal families contain $2|\mathcal{X}_c^*|$ scalar events and the four joint-mass families contain $4|\mathcal{X}_c^*||\mathcal{Y}| = 4K_c$ scalar events. Since $|\mathcal{Y}| \geq 2$, the total number of scalar events is at most $5K_c$. With $\tilde{\delta} = \delta/(10K_c)$, the total failure probability is therefore at most $\delta/2$.
Using (47)–(48), we obtain, for every $x^* \in \mathcal{X}_c^*$,

$$\Delta_{c,0}^{\mathrm{unp}}(x^*) \leq \Delta_{c,0}^{\mathrm{unp,pop}}(x^*) + 2R_{\tilde{\delta}}(n_{N,0}). \tag{51}$$

Using (49)–(50), we obtain, for every $x^* \in \mathcal{X}_c^*$,

$$\Delta_{c,0}^{\mathrm{unp,pop}}(x^*) \leq \Delta_{c,0}^{\mathrm{unp},(n_{N,c}-n_{e,c})}(x^*) + 2R_{\tilde{\delta}}(n_{N,c} - n_{e,c}). \tag{52}$$

Define the fixed population-based unplanting margin

$$\overline{M}_c^{\mathrm{unp}}(x^*) := \alpha_{N,c}\Big(p_{c,\mathrm{pop}}^X(x^*) - R_{\tilde{\delta}}(n_{N,c})\Big) - \sum_{d\neq c}\alpha_{N,d}\Delta_{c\leftarrow d}^{\mathrm{unp}}(x^*) - \alpha_{N,0}\Big(\Delta_{c,0}^{\mathrm{unp,pop}}(x^*) + 2R_{\tilde{\delta}}(n_{N,0})\Big),$$

and the corrected empirical unplanting margin

$$\widetilde{M}_c^{\mathrm{unp}}(x^*) := \alpha_{N,c}\Big(\hat{p}_c^X(x^*) - 2R_{\tilde{\delta}}(n_{N,c})\Big) - \sum_{d\neq c}\alpha_{N,d}\Delta_{c\leftarrow d}^{\mathrm{unp}}(x^*) - \alpha_{N,0}\Delta_{c,0}^{\mathrm{unp},(n_{N,c}-n_{e,c})}(x^*) - 2\alpha_{N,0}\Big(R_{\tilde{\delta}}(n_{N,0}) + R_{\tilde{\delta}}(n_{N,c} - n_{e,c})\Big).$$

On the concentration event above, for all $x^* \in \mathcal{X}_c^*$,

$$M_c^{\mathrm{unp}}(x^*) \geq \overline{M}_c^{\mathrm{unp}}(x^*), \tag{53}$$

and, by (46) and (52),

$$\overline{M}_c^{\mathrm{unp}}(x^*) \geq \widetilde{M}_c^{\mathrm{unp}}(x^*). \tag{54}$$

Combining (44) and (53), we obtain

$$x^* \in \mathcal{X}_c^*, \quad \overline{M}_c^{\mathrm{unp}}(x^*) > \frac{\varepsilon}{1-\varepsilon} \quad \Longrightarrow \quad \hat{m}(x^*) \neq y_c^*. \tag{55}$$

By (54), this further implies

$$x^* \in \mathcal{X}_c^*, \quad \widetilde{M}_c^{\mathrm{unp}}(x^*) > \frac{\varepsilon}{1-\varepsilon} \quad \Longrightarrow \quad \hat{m}(x^*) \neq y_c^*. \tag{56}$$

Step 4. Let

$$\varphi(x) := \mathbf{1}_{\{\hat{m}(g_c(x)) \neq y_c^*\}}, \qquad x \in \mathcal{X}.$$

By definition,

$$\hat{S}_c(\alpha_{N,c}) = \frac{1}{N^{\mathrm{test}}} \sum_{i=1}^{N^{\mathrm{test}}} \varphi(X_i^{\mathrm{test}}).$$

Conditionally on the training sample, the variables $\varphi(X_i^{\mathrm{test}})$ are i.i.d. Bernoulli random variables. Applying Hoeffding's inequality with failure budget $\delta/4$, we obtain that, with probability at least $1 - \delta/4$,

$$\hat{S}_c(\alpha_{N,c}) \geq \mathbb{P}_{X \sim P_0^X}\left[\hat{m}(g_c(X)) \neq y_c^*\right] - R_{\delta/4}(N^{\mathrm{test}}). \tag{57}$$

Since $g_c(X) \in \mathcal{X}_c^*$ for all $X \in \mathcal{X}$, (55) implies

$$\left\{\overline{M}_c^{\mathrm{unp}}(g_c(X)) > \frac{\varepsilon}{1-\varepsilon}\right\} \subseteq \left\{\hat{m}(g_c(X)) \neq y_c^*\right\}.$$

Therefore, on the Step 3 concentration event,

$$\hat{S}_c(\alpha_{N,c}) \geq \mathbb{P}_{X \sim P_0^X}\left[\overline{M}_c^{\mathrm{unp}}(g_c(X)) > \frac{\varepsilon}{1-\varepsilon}\right] - R_{\delta/4}(N^{\mathrm{test}}). \tag{58}$$

Step 5. Define the fixed event

$$A_c^{\mathrm{unp,pop}} := \left\{x \in \mathcal{X} : \overline{M}_c^{\mathrm{unp}}(g_c(x)) > \frac{\varepsilon}{1-\varepsilon}\right\}.$$

Conditional on $\mathcal{F}_{E,c}$ and on all training data outside the holdout split $R_{N,c}$, the event $A_c^{\mathrm{unp,pop}}$ is fixed. By neutrality and the sampling model, the pre-edit features $\{X_i : i \in R_{N,c}\}$ are i.i.d. with law P_0^X, conditionally on $n_{N,c} - n_{e,c}$. Applying Hoeffding's inequality to the i.i.d. variables $\mathbf{1}_{\{X_i \in A_c^{\mathrm{unp,pop}}\}}$, $i \in R_{N,c}$, with failure budget $\delta/4$, we obtain that, with probability at least $1 - \delta/4$,

$$\mathbb{P}_{X \sim P_0^X}\left[X \in A_c^{\mathrm{unp,pop}}\right] \geq \mathbb{P}_{x \sim \hat{P}_{N,c}^{\mathrm{pre},R,X}}\left[x \in A_c^{\mathrm{unp,pop}}\right] - R_{\delta/4}(n_{N,c} - n_{e,c}). \tag{59}$$

On the Step 3 concentration event, (54) implies the pointwise inclusion

$$\left\{\widetilde{M}_c^{\mathrm{unp}}(g_c(x)) > \frac{\varepsilon}{1-\varepsilon}\right\} \subseteq A_c^{\mathrm{unp,pop}}.$$

Hence,

$$\mathbb{P}_{X \sim P_0^X}\left[\overline{M}_c^{\mathrm{unp}}(g_c(X)) > \frac{\varepsilon}{1-\varepsilon}\right] \geq \mathbb{P}_{x \sim \hat{P}_{N,c}^{\mathrm{pre},R,X}}\left[\widetilde{M}_c^{\mathrm{unp}}(g_c(x)) > \frac{\varepsilon}{1-\varepsilon}\right] - R_{\delta/4}(n_{N,c} - n_{e,c}). \tag{60}$$

Step 6. By the union bound, the Step 3 concentration event, (57), and (59) hold simultaneously with probability at least $1-\delta$. Combining (58) and (60), we conclude that, with probability at least $1-\delta$,

$$\begin{aligned}\hat{S}_c(\alpha_{N,c}) &\geq \mathbb{P}_{x\sim\hat{P}^{\mathrm{pre},R,X}_{N,c}}\left[\widetilde{M}^{\mathrm{unp}}_c\big(g_c(x)\big) > \frac{\varepsilon}{1-\varepsilon}\right] - R_{\delta/4}(n_{N,c}-n_{e,c}) - R_{\delta/4}(N^{\mathrm{test}}) \\ &= \mathbb{P}_{x\sim\hat{P}^{\mathrm{pre},R,X}_{N,c}}\Big[\,\alpha_{N,c}\Big(\hat{p}^X_c\big(g_c(x)\big) - 2R_{\tilde{\delta}}(n_{N,c})\Big) - \sum_{d\neq c}\alpha_{N,d}\,\Delta^{\mathrm{unp}}_{c\leftarrow d}\big(g_c(x)\big) \\ &\qquad - \alpha_{N,0}\,\Delta^{\mathrm{unp},(n_{N,c}-n_{e,c})}_{c,0}\big(g_c(x)\big) - 2\alpha_{N,0}\Big(R_{\tilde{\delta}}(n_{N,0}) + R_{\tilde{\delta}}(n_{N,c}-n_{e,c})\Big) \\ &\qquad - \frac{\varepsilon}{1-\varepsilon} > 0\Big] - R_{\delta/4}(n_{N,c}-n_{e,c}) - R_{\delta/4}(N^{\mathrm{test}}),\end{aligned}$$

which is exactly (22). This completes the proof. □

B A Climate Adaptation Use Case

In this section, we provide a real-world use case for ACA with multiple collectives in the space of interventions for climate adaptation. Consider a national adaptation intake platform used by a city to prioritize interventions at the neighborhood level. The platform stores standardized neighborhood-level adaptation records to identify which interventions are most needed. These records combine official sources (e.g., inspections, infrastructure inventories, incident reports) with community inputs (e.g., forum discussions, neighborhood association reports), and are stored in a standardized format with a fixed set of categorical fields, so that each neighborhood is summarized in the same structured way. The model recommends one intervention per neighborhood from a finite menu $\mathcal{Y}$. The training corpus is large as it contains millions of neighborhood planning units (e.g., census tracts/block groups) pooled across many municipalities, while the decision remains neighborhood-level. Assume there are M collectives, being grassroots associations. Each collective is run by residents from different neighborhoods that face similar conditions (e.g., recurrent flooding, extreme heat exposure, limited cooling access) and therefore has a nuanced, ground-level understanding of local needs. The collective seeks to coordinately change the adaptation records corresponding to its member neighborhoods (and the shared community inputs that feed into those records) to steer the classifier toward the intervention $y^*_c \in \mathcal{Y}$ it has judged to be most necessary for those neighborhoods. Each collective thus deploys an editing strategy h_c to achieve its actionable goal. For example, in a feature-label regime, a collective with residents from waterfront neighborhoods documents chronic flooding with household infiltration tests and evidence collected by interviewing community members, then updates the neighborhoods' records so the categorical fields consistently reflect that risk profile and explicitly requests rain gardens/bioswales in the intervention request field. Similarly, collectives would likely need to resort to feature-only strategies if the city's intake pipeline does not allow community label edits, i.e., if the intervention field is locked and labeled by staff/administrative codes. Global success would then emphasize equity (lifting the worst-served neighborhood with $\hat{G}_{\min}$) or scale (a mass-weighted average with $\hat{G}_{\mathrm{w}}$). Even if this example is structured in a slightly simplistic way, there is precedent for community-reported concerns to be ingested into decision-making tools. NOAA and CAPA-led "Heat Watch" campaigns have produced heat maps in $120+$ communities and feed them into planning and public-health practice [22, 6]. European cities are deploying digital twins for climate-resilience analysis and participatory planning, e.g., [7].

Information Regimes as Protocols. The information regimes of Section 3 correspond to concrete protocols in this use case. (i) In the best case, the platform or a regulator acts as a trusted aggregator and publishes the aggregate conflict $\sum_{d\neq c}\alpha_{N,d}\Delta_{c\to d}(x^*)$ at the planted feature values, in the spirit of transparency reporting, without revealing any per-collective distribution. (ii) Without a trusted party, the aggregate conflict is a sum of bounded per-collective empirical averages, so the collectives can compute it jointly by secure aggregation, optionally adding differentially private noise; since the bound already absorbs additive perturbations of the conflict term through its radii, we expect calibrated noise to enter as one more additive term of the same form, degrading the guarantee gracefully rather than invalidating it. (iii) Aligned collectives need only exchange their planted sets and target labels, not their data, to determine the sign of $\Delta_{c\to d}$ and whether mass pooling applies (Corollary 3.2); in the example above, two waterfront associations requesting rain gardens for overlapping record profiles can establish that they are allies by comparing the fields they plan to edit.

C Simulated Neighborhood Adaptation-Records Dataset

We build on the codebase of [12] and generate a dataset of $\mathbf{3{,}000{,}000}$ instances. Each instance represents a neighborhood-level adaptation record for a *planning unit* (e.g., a census tract/block

group), characterized by multiple categorical features that capture key aspects of local conditions, planned measures, governance, and operational capacity. Conceptually, the base dataset is generated by drawing instances *i.i.d.* from a population distribution over planning units.

The features included in this synthetic dataset are as follows:

- Adaptation Strategy Type: the type of adaptation strategy, categorized as *Infrastructure Upgrade*, *Nature-Based Solution*, *Policy/Regulation*, *Community Outreach*, *Emergency Preparedness*, *Technology Deployment*, *Land-Use Planning*, or *Public Health Services*.
- Primary Energy Source for Adaptation Measures: the primary energy source used by the deployed measures, which can be *Renewable (Solar)*, *Renewable (Wind)*, *Grid Electricity*, or *None*.
- Implementation Timeline: the timeline for implementation, classified as *Short-term (<1 yr)*, *Medium-term (1–5 yr)*, or *Long-term (>5 yr)*.
- Geographic Scope: the geographic scope of the measure, identified as *Neighborhood*, *City-wide*, or *Regional*.
- Resilience Rating: the resilience rating of the neighborhood profile, rated as *Very Low*, *Low*, *Moderate*, *High*, or *Very High*.
- Construction Material for Adaptation Projects: the primary construction material, which can be *Concrete*, *Steel*, or *Timber*.
- Public Awareness / Communication Platform: the level of the communication platform, ranging from *None*, *Basic Alerts (SMS/email)*, *Interactive Dashboard*, or *Integrated Smart-city Platform*.
- Funding Source / Jurisdiction: the source/jurisdiction of funding, that we will denote by *F1*, *F2*, *F3*, *F4*, and *F5*.
- Project Maintenance Horizon: the maintenance horizon, available in options of *1 year*, *3 years*, *5 years*, or *10 years*.
- Number of Access Points to Critical Infrastructure: the number of access points, which can be *0 (sealed)*, *1–2*, or *3–4*.
- Capacity of Adaptation Facility: the capacity of an adaptation facility, with options for *<50*, *50–200*, *200–500*, or *500+*.
- Cooling Provision: indicates the type of cooling provision (*None* or *Mechanical (AC)*).
- Decision-Support / Early-Warning System: the level of decision-support, which can be *None*, *Basic (threshold alerts)*, or *Advanced (predictive modeling)*.
- Infrastructure Surface Type: the surface type, categorized as *Standard Pavement*, *Permeable Pavement*, or *Heat-Reflective Coating*.
- Green Roof / Solar Installation: indicates whether a roof retrofit is present (*None* or *Retrofit Present*).
- Public Notification System: the public notification system, which can be *None*, *Sirens*, *Public Address (PA) System*, or *Multilingual Broadcast*.
- Automated Operation / Smart Controls: indicates whether smart controls are present (*None* or *Fully automated (AI-driven)*).
- Interoperability with External Sensors / IoT: indicates whether IoT interoperability is present (*No* or *Full*).

Additionally, each adaptation record is assigned an Intervention Recommendation label based on a scoring system that considers various factors such as resilience rating, geographic scope, cooling provision, maintenance horizon, and others. The possible outcomes are classified into four categories: Greening & Shade Expansion, Reflective Surface Retrofits, Cooling Services & Outreach, and Green Stormwater Infrastructure.

In our experiments, we then generate separate datasets for the collectives sampled without replacement from this base dataset. Unless otherwise specified, we choose $N = 2{,}000{,}000$ for the training set and $N^{\text{test}} = 100{,}000$ for the test set; the two-collective feature-label experiments, the four-collective heatmap, the unplanting experiments and the single-collective ablations state their own sizes. In all the experiments, we set $\delta = 0.05$ and $\varepsilon = 0$.

D Reduction to the single-collective bounds of [12]

Relation to the single-collective bounds of [12]. We show that our multi-collective guarantees specialize to the corresponding single-collective guarantees of [12] when $M = 1$, for both *signal planting* and *signal unplanting*. Throughout this appendix we condition on the realized sample sizes, as in the main text.

Signal Planting. Set $M = 1$ and $c = 1$ in Theorem 3.1. Write

$$n := n_{N,1}, \qquad \alpha := \alpha_{N,1} = \frac{n}{N}, \qquad \alpha_{N,0} = 1 - \alpha = \frac{N-n}{N},$$

and denote $g := g_1$, $y^* := y_1^*$, and $\mathcal{X}^* := \mathcal{X}_1^* = g(\mathcal{X})$. Since there are no other collectives, the cross-collective term vanishes and (13) specializes to

$$\begin{aligned} \hat{S}(\alpha) \;\geq\; & \mathbb{P}_{x\sim\hat{P}_n^{\mathrm{pre},X}}\Big[\alpha\Big(\hat{p}^X\big(g(x)\big) - 2R_{\tilde{\delta}}(n)\Big) - (1-\alpha)\,\Delta_0^{(n)}\big(g(x)\big) \\ & \qquad - 2(1-\alpha)\Big(R_{\tilde{\delta}}(N-n) + R_{\tilde{\delta}}(n)\Big) - \frac{\varepsilon}{1-\varepsilon} > 0\Big] \\ & - \; R_{\delta/4}(n) - R_{\delta/4}(N^{\mathrm{test}}), \end{aligned} \tag{61}$$

where, consistently with Section 2, $\hat{P}_n^{\mathrm{pre}}$ is the empirical pre-edit distribution on the collective sample, $\hat{p}^X(x^*) := \hat{P}_n^X(x^*)$ is the empirical post-edit feature marginal under the chosen planting strategy, and

$$\Delta_0^{(n)}(x^*) := \max_{y'\in\mathcal{Y}\setminus\{y^*\}} \big(\hat{P}_n^{\mathrm{pre}}(x^*, y') - \hat{P}_n^{\mathrm{pre}}(x^*, y^*)\big). \tag{62}$$

Here

$$\tilde{\delta} = \frac{\delta}{10|\mathcal{X}^*|\,|\mathcal{Y}|}.$$

We now match (61) with [12, Theorems 3.3 and 3.5].

Identification of empirical objects. Let $D^{(n)}$ denote the pre-edit collective dataset and $\tilde{D}^{(n)}$ its post-edit version. Under our notation,

$$\hat{P}_{D^{(n)}} \equiv \hat{P}_n^{\mathrm{pre}}, \qquad \hat{P}_{\tilde{D}^{(n)}} \equiv \hat{P}_n, \qquad \hat{P}^X_{\tilde{D}^{(n)}} \equiv \hat{P}_n^X.$$

Under feature–label planting, $h(x, y) = (g(x), y^*)$, so for any $x^* \in \mathcal{X}^*$ we have $\hat{P}_{\tilde{D}^{(n)}}(x^*, y') = 0$ for all $y' \neq y^*$ and hence

$$\hat{P}^X_{\tilde{D}^{(n)}}(x^*) = \hat{P}_{\tilde{D}^{(n)}}(x^*, y^*).$$

Therefore the signal term in (61) can be read either as a feature-marginal prevalence, as we do, or as a joint mass on (x^*, y^*). Moreover, the non-collective conflict term (62) coincides with the corresponding empirical conflict quantity in [12, Theorem 3.3] after identifying x^* with the planted feature value.

Outer probability / change of variables. Our outer probability is written as $\mathbb{P}_{x\sim\hat{P}_n^{\mathrm{pre},X}}[\cdot]$ with the margin evaluated at $g(x)$. Since

$$(g)_{\#}\hat{P}_n^{\mathrm{pre},X} = \hat{P}_n^X$$

by construction, for any measurable Ψ on $\mathcal{X}^*$ we have

$$\mathbb{P}_{x\sim\hat{P}_n^{\mathrm{pre},X}}\big[\Psi(g(x))\big] = \mathbb{P}_{x^*\sim\hat{P}_n^X}\big[\Psi(x^*)\big].$$

Thus (61) can equivalently be written with the outer probability taken over post-edit signal features, matching the presentation in [12].

Comparison to [12]. With the identifications above, (61) has the same functional structure as [12, Theorem 3.3] (feature–label planting) and [12, Theorem 3.5] (feature-only planting): an outer empirical probability over collective features, a positive term proportional to n/N times the post-edit prevalence of the planted feature, a negative term proportional to $(N-n)/N$ times an empirical conflict term computed from the collective's pre-edit data, a suboptimality offset $\varepsilon/(1-\varepsilon)$, and explicit Hoeffding radii. The remaining differences are only in the bookkeeping of constants, through the particular split of δ and the resulting Hoeffding radii.

Signal Unplanting. Set $M = 1$ and $c = 1$ in Theorem 4.1. Keep the shorthand

$$n := n_{N,1}, \qquad \alpha := \frac{n}{N}, \qquad \alpha_{N,0} = 1 - \alpha = \frac{N-n}{N},$$

and let $n_e := n_{e,1}$ denote the size of the estimation sub-collective. Write again $g := g_1$, $y^* := y_1^*$, and $\mathcal{X}^* := g(\mathcal{X})$. Let $E_N \subseteq I_{N,1}$ be the estimation index set of size n_e and $R_N := I_{N,1} \setminus E_N$ its complement. Define the associated empirical measures

$$\hat{P}_n^{\mathrm{pre},E} := \frac{1}{n_e} \sum_{i \in E_N} \delta_{Z_i}, \qquad \hat{P}_n^{\mathrm{pre},R} := \frac{1}{n - n_e} \sum_{i \in R_N} \delta_{Z_i}.$$

For each $x^* \in \mathcal{X}^*$, define the data-driven alternative label

$$y'(x^*) \;:=\; \arg \max_{y \in \mathcal{Y} \setminus \{y^*\}} \hat{P}_n^{\mathrm{pre},E}(x^*, y), \tag{63}$$

with deterministic tie-breaking, and consider the adaptive unplanting rule

$$h(x, y) = (g(x), y'(g(x))).$$

In the single-collective case, the cross-collective term disappears and (22) specializes to

$$\begin{aligned} \hat{S}(\alpha) \;\geq\; \mathbb{P}_{x \sim \hat{P}_n^{\mathrm{pre},R,X}} \Big[& \alpha\Big(\hat{p}^X\big(g(x)\big) - 2R_{\tilde{\delta}}(n)\Big) - (1-\alpha)\,\Delta_0^{\mathrm{unp},(n-n_e)}\big(g(x)\big) \\ & - 2(1-\alpha)\Big(R_{\tilde{\delta}}(N-n) + R_{\tilde{\delta}}(n-n_e)\Big) - \frac{\varepsilon}{1-\varepsilon} > 0 \Big] \\ - \;\; & R_{\delta/4}(n-n_e) - R_{\delta/4}(N^{\mathrm{test}}), \end{aligned} \tag{64}$$

where the approximate non-collective conflict term is computed from the independent holdout split R_N:

$$\Delta_0^{\mathrm{unp},(n-n_e)}(x^*) := \hat{P}_n^{\mathrm{pre},R}(x^*, y^*) - \hat{P}_n^{\mathrm{pre},R}\big(x^*, y'(x^*)\big). \tag{65}$$

Match to [12]. Under the identification

$$D^{(n_e)} \leftrightarrow \text{estimation split } E_N,\; D^{(n-n_e)} \leftrightarrow \text{holdout split } R_N,\; \tilde{D}^{(n)} \leftrightarrow \text{post-edit collective dataset},$$

the label-selection rule (63) coincides with the adaptive unplanting label-selection rule in [12] up to notation, and the conflict term (65) matches the corresponding holdout-based empirical conflict term in [12, Theorem 3.7]. The outer probability in (64) is taken over the holdout empirical feature distribution, because the same holdout split is used in the proof to estimate the non-collective conflict independently of the label-selection step. Equivalently, since the post-edit holdout feature distribution is the pushforward of $\hat{P}_n^{\mathrm{pre},R,X}$ under g, one may rewrite the outer probability over the corresponding post-edit holdout signal features.

Therefore, (64) specializes the multi-collective adaptive unplanting guarantee to the single-collective setting and has the same functional structure as [12, Theorem 3.7]: a positive signal-prevalence term, a holdout-based empirical conflict term for the non-collective population, the suboptimality offset $\varepsilon/(1-\varepsilon)$, and explicit concentration radii. The remaining differences are again only in constant bookkeeping through the chosen confidence split.

E Additional Experimental Results

E.1 Two Collective Experiments

The two-collective feature-label bounds of Figure 1 and Figures 3–8 are computed with an implementation of Theorem 3.1 that follows its statement exactly, available in our code, with $N = 2{,}000{,}000$, $N^{\mathrm{test}} = 100{,}000$, total collective mass $\bar{\alpha}_N \in \{0.1, 0.2, \ldots, 0.9\}$ (combined collective size from 200,000 to 1,800,000) and 99 equally spaced shares of Collective 1 at each mass. In the known-interventions setting, each collective knows a pool of possible interventions: the nested patterns that plant 13, 14, 15, 16 or 17 of the 17 features, each with every target label other than its own, for 15 interventions. It estimates the post-edit distribution of each intervention by applying it to its own pre-edit sample, which is valid under neutrality, averages the resulting conflicts over the pool, and weights the average by the total mass of the other collectives.

We report additional experiments in the 2-collective feature overlap setting described in Section 5 in both the feature-label and the feature-only settings. In the feature-label setting we observe a qualitative behavior consistent with the trend observed in the main text: reallocating mass toward

one collective improves its lower bound while reducing the other's, and increasing the total planted mass shifts the attainable guarantees upward. The symmetric 17-vs.-17 case is driven primarily by the mass split (Figures 3 & 6). Reducing the number of planted features for *Collective 2* introduces the asymmetry observed in Figure 1: As *Collective 2*'s planted features decrease, *Collective 2* loses signal prevalence, whereas *Collective 1* benefits from features that are no longer directly contested (Figures 4, 5, 7, & 8).

In the 17 vs. 13 case, under the *per-collective* setting, we observe that collective 1 is insensitive to the mass of collective 2 (Figure 5). The quantity of features prevails over the mass proportions of each collective. This is not the case in the *known-interventions* setting (Figure 8), where the averaged conflict does not reflect how little the realized 13-feature pattern contests Collective 1's planted values, and Collective 1's threshold grows with the combined size, from 33,200 to 64,800 samples. More generally, the known-interventions estimate is more conservative than the per-collective bound when the actual opponent conflicts less than the average intervention in the pool (Collective 1 in the 17-vs.-15 and 17-vs.-13 cases), and optimistic when it conflicts more (both collectives in the 17-vs.-17 case and Collective 2 in the 17-vs.-15 case). Like the *per-collective* setting, the maximum success that Collective 2 can reach decreases as Collective 2 plants fewer features (0.99, 0.42 and 0.33 for 17, 15 and 13 features).

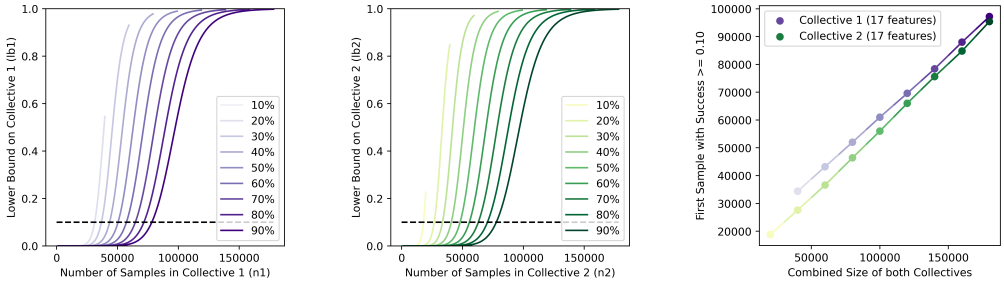

(a) Collective 1 lower bound on success with a partial conflict with 17 vs. 17 features and a per-collective strategy.

(b) Collective 2 lower bound on success with a partial conflict with 17 vs. 17 features and a per-collective strategy.

(c) The number of samples required to exceed a 0.1 lower bound on success in a 17 vs. 17 collective scenario.

Figure 3: The bound on success for two collectives under varying population sizes and collective sizes.

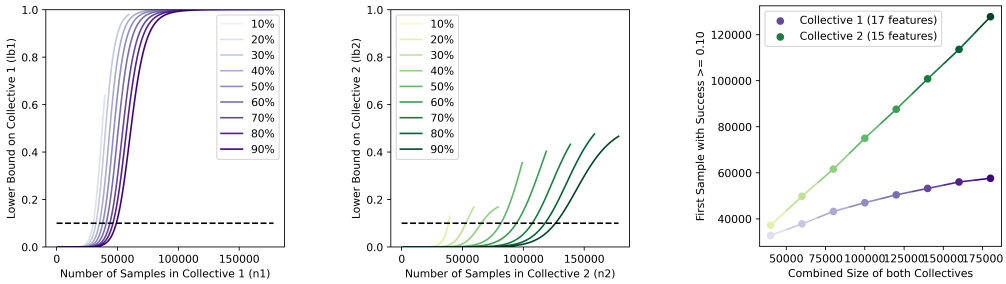

(a) Collective 1 lower bound on success with a partial conflict with 17 vs. 15 features and a per-collective strategy.

(b) Collective 2 lower bound on success with a partial conflict with 17 vs. 15 features and a per-collective strategy.

(c) The number of samples required to exceed a 0.1 lower bound on success in a 17 vs. 15 collective scenario.

Figure 4: The bound on success for two collectives under varying population sizes and collective sizes.

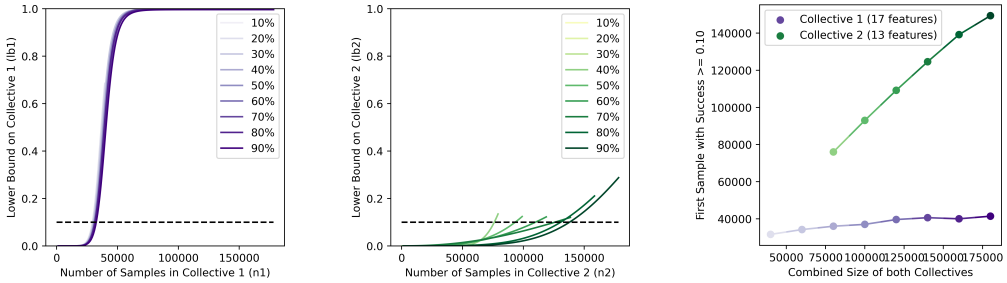

(a) Collective 1 lower bound on success with a partial conflict with 17 vs. 13 features and a per-collective strategy.

(b) Collective 2 lower bound on success with a partial conflict with 17 vs. 13 features and a per-collective strategy.

(c) The number of samples required to exceed a 0.1 lower bound on success in a 17 vs. 13 collective scenario.

Figure 5: The bound on success for two collectives under varying population sizes and collective sizes.

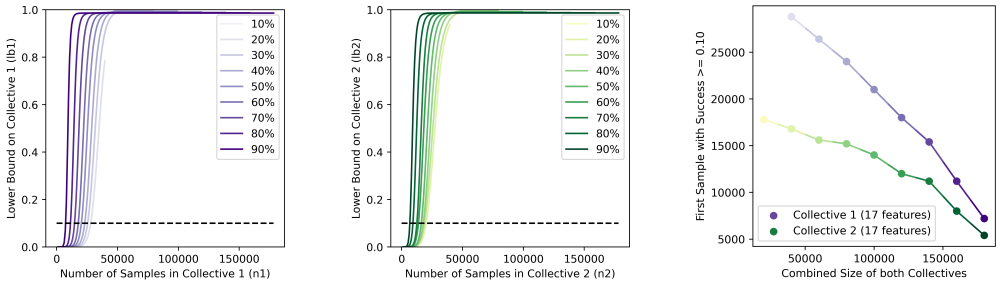

(a) Collective 1 lower bound on success with a partial conflict with 17 vs. 17 features and a known-interventions strategy.

(b) Collective 2 lower bound on success with a partial conflict with 17 vs. 17 features and a known-interventions strategy.

(c) The number of samples required to exceed a 0.1 lower bound on success in a 17 vs. 17 collective scenario.

Figure 6: The bound on success for two collectives under varying population sizes and collective sizes.

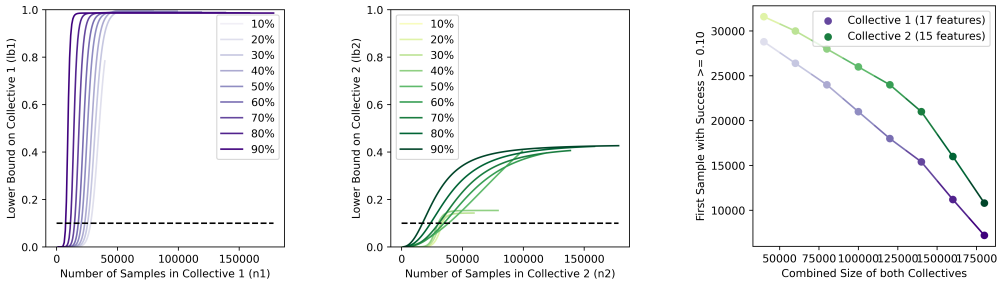

(a) Collective 1 lower bound on success with a partial conflict with 17 vs. 15 features and a known-interventions strategy.

(b) Collective 2 lower bound on success with a partial conflict with 17 vs. 15 features and a known-interventions strategy.

(c) The number of samples required to exceed a 0.1 lower bound on success in a 17 vs. 15 collective scenario.

Figure 7: The bound on success for two collectives under varying population sizes and collective sizes.

For feature-only planting, we report the per-collective setting (Figures 9, 10, 11, 12, & 13). As the overlap decreases from 17-vs.-17 to 17-vs.-13, the conflict between collectives weakens for the collective retaining more planted features, while the collective with fewer planted features requires more favorable mass allocation to achieve comparable lower bounds.

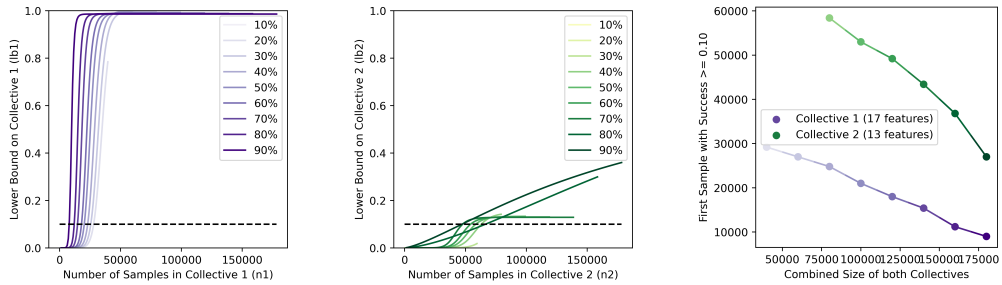

(a) Collective 1 lower bound on success with a partial conflict with 17 vs. 13 features and a known-interventions strategy.

(b) Collective 2 lower bound on success with a partial conflict with 17 vs. 13 features and a known-interventions strategy.

(c) The number of samples required to exceed a 0.1 lower bound on success in a 17 vs. 13 collective scenario.

Figure 8: The bound on success for two collectives under varying population sizes and collective sizes.

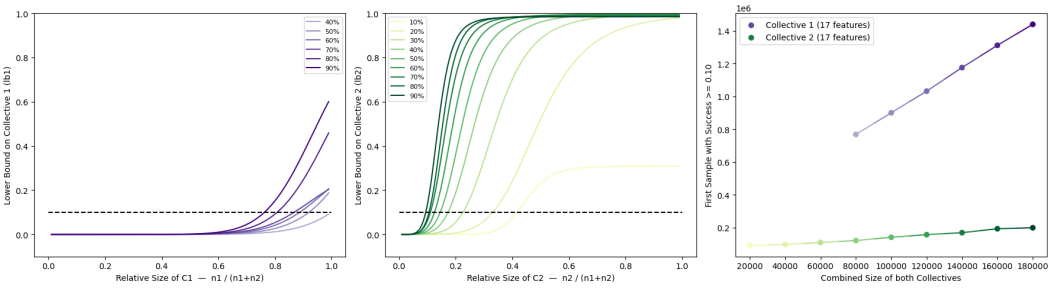

Figure 9: The bound on feature-only planting success for two collectives, 17 vs 17 features, under varying population sizes and collective sizes.

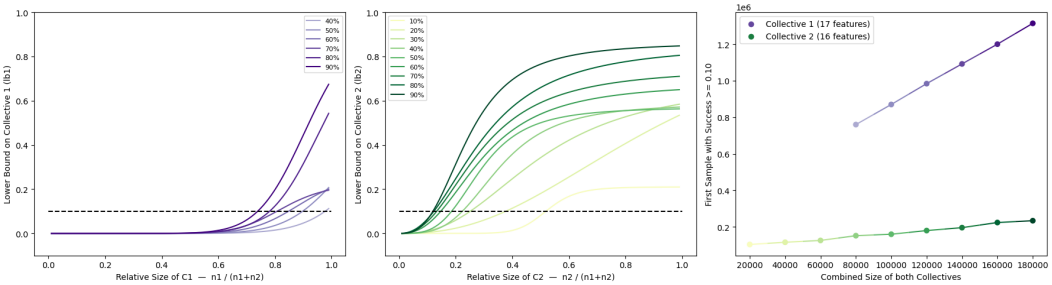

Figure 10: The bound on feature-only planting success for two collectives, 17 vs 16 features, under varying population sizes and collective sizes.

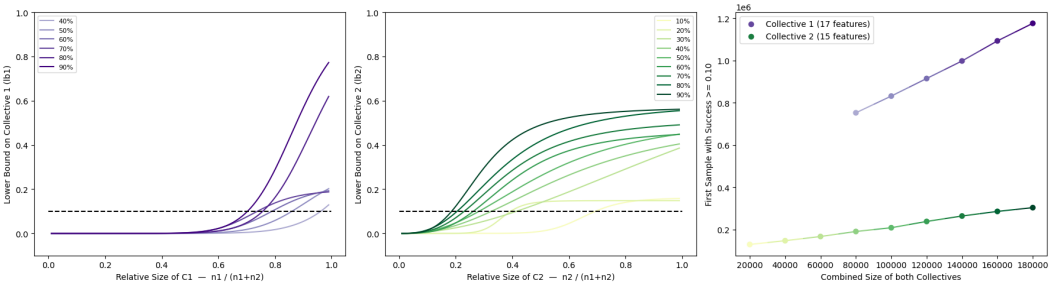

Figure 11: The bound on feature-only planting success for two collectives, 17 vs 15 features, under varying population sizes and collective sizes.

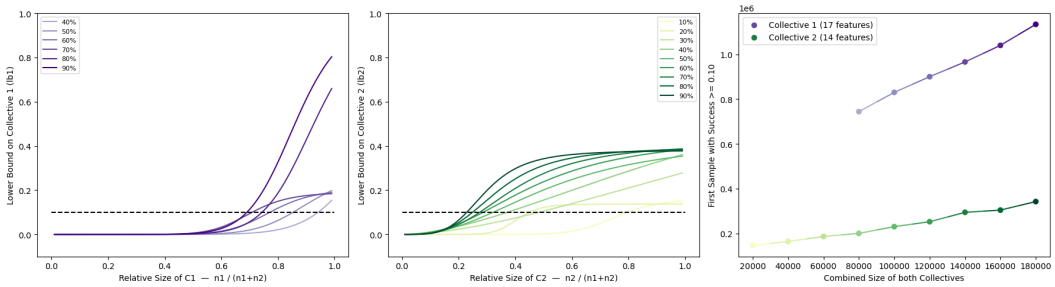

Figure 12: The bound on feature-only planting success for two collectives, 17 vs 14 features, under varying population sizes and collective sizes.

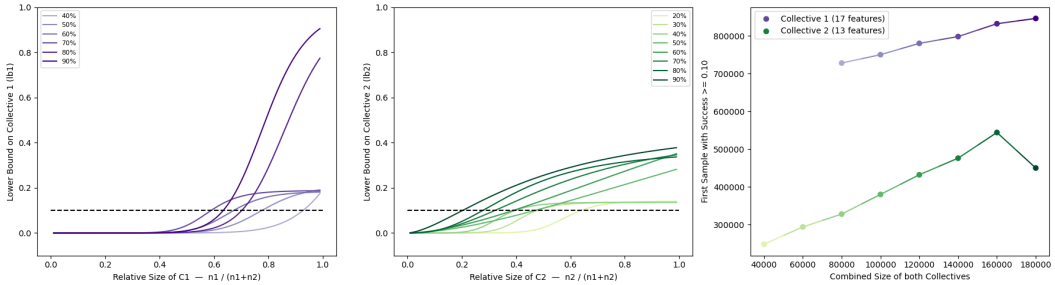

Figure 13: The bound on feature-only planting success for two collectives, 17 vs 13 features, under varying population sizes and collective sizes.

E.2 Single-Collective Ablations

We reproduce the single-collective lower bound numbers of [12] by running our multi-collective implementation with a single collective ($M = 1$). We use the same 17-feature signal g, $N^{\text{test}} = 100{,}000$, $\delta = 0.05$ and $\varepsilon = 0$ as the original work, and recreate the following experiments:

- Feature-label planting of the label *Poor* with total dataset size $N \in \{500\text{K}, 1\text{M}, 2\text{M}\}$. The original paper sweeps the collective size from 1% to 30% of N; we evaluate the bound on the active window 6% to 16% in steps of 0.5%, outside of which the bound is flat at 0 or at its saturation value. We report the results in Figure 14;
- Feature-only planting with a fixed dataset size of $N = 1\text{M}$, the transformation x_0 of the original work, and the label y^* varying between *Poor*, *Average* and *Excellent*. The collective size spans the active window of each label: 5% to 20% for *Excellent*, 55% to 99.5% for *Poor* and 60% to 99.5% for *Average*, in steps of 0.5%. We report the results in Figure 15;
- Signal unplanting of the label *Good* for a dataset of size $N = 1\text{M}$ with the adaptive strategy and an estimation split of $n_e = 2000$, for collective sizes from 0.5% to 15% in steps of 0.5%. We report the results in Figure 16.

For the two planting experiments, the $M = 1$ bounds coincide with the single-collective numbers of [12]: on every collective size evaluated by both, the bound values agree to within 0.002, with the exception of a handful of points located exactly at a jump of the step-like bound, where the transition occurs one grid step (0.5% of N) earlier or later. The only systematic difference is that our implementation clips the bound at 0, where the original reports values of about -0.01. This validates the method proposed here as a strict generalization of the single-collective approach.

E.3 Signal Unplanting with Two Collectives

We complement the planting experiments of Section 5 with two collectives that unplant. Collective 1 maps 17 features through g_1 and removes their association with *Greening & Shade Expansion*; Collective 2 maps a nested subset of 15 of those features through g_2 and targets *Green Stormwater Infrastructure*. We set $N = 200{,}000$, $N^{\text{test}} = 50{,}000$, $\delta = 0.05$ and $\varepsilon = 0$, and use per-collective conflict terms. The total collective mass is $\bar{\alpha}_N \in \{0.90, 0.75, 0.60\}$, of which Collective 1 holds a share s_1 and Collective 2 the rest. We sweep s_1 over 50 equally spaced values in $[0.01, 0.99]$ and report five representative shares. Table 2 reports the lower bound of Theorem 4.1 for the adaptive

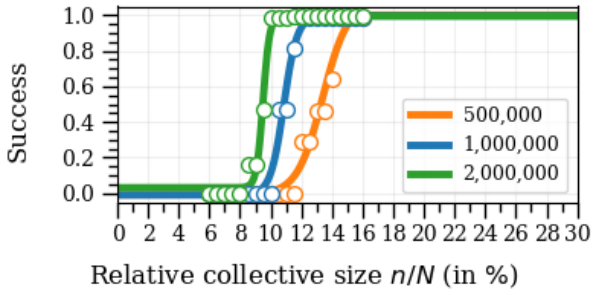

Figure 14: Re-creating lower bound numbers for the feature-label case from the Single-collective reference paper

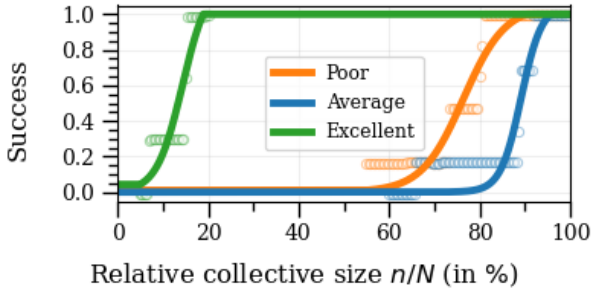

Figure 15: Re-creating lower bound numbers for the feature-only case from the Single-collective reference paper

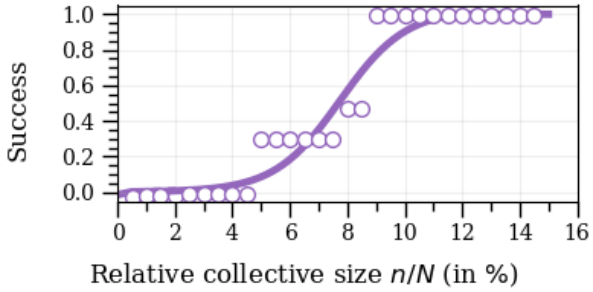

Figure 16: Re-creating lower bound numbers for the signal-unplanting case from the Single-collective reference paper

strategy, with $n_{e,c} = 2000$ (Appendix E.4), and for the naive strategy of Section 4; we clip small negative bound values to 0.

Neither strategy dominates. For Collective 1, the adaptive bound is at least as large as the naive one in 11 of the 15 cells, and in the remaining four the naive bound is larger by about 0.17. For Collective 2, whose map is nested in Collective 1's so that all of its planted values are contested, both bounds are low: the adaptive bound falls below 0.05 in 9 of the 15 cells, and the naive bound is larger in 8 of them. For Collective 1, the adaptive bound is never vacuous and, at the shares reported in the table, increases with s_1 at every mass level, whereas the naive bound is vacuous in 6 of the 30 cells and is not monotone (Collective 1 at $\bar{\alpha}_N = 0.90$: 0.978, 0 and 0.983 for $s_1 = 0.11$, 0.21 and 0.31). Once Collective 1 holds half of the collective mass, its adaptive bound exceeds 0.98 at every mass level.

E.4 Size of the Estimation Split

The adaptive strategy selects the alternative labels $y'_c(x^*)$ on the estimation split $E_{N,c}$ and evaluates the bound on the holdout $R_{N,c}$. We repeat the experiment of Appendix E.3 at $\bar{\alpha}_N = 0.90$ with $n_{e,c} \in \{100, 2000, 10000\}$ (Table 3). The bounds for $n_{e,c} = 2000$ and $n_{e,c} = 10000$ agree to within 0.02 in every cell, so the label selection has converged at $n_{e,c} = 2000$. With $n_{e,c} = 100$ the selected labels are noisy, and the bound, which is valid for the strategy it certifies, differs by more than 0.1 in four of the ten cells (e.g., Collective 2: 0.256, 0.226 and 0.142, against 0.042, 0.015 and 0.013 at $n_{e,c} = 2000$). Since a larger $n_{e,c}$ removes data from the holdout and inflates the radii, $n_{e,c} = 2000$ is the smallest converged choice; it is also the value used by Gauthier et al. [12] for a single collective.

Table 2: Signal unplanting with two collectives: lower bounds of Theorem 4.1 for the adaptive strategy ($n_{e,c} = 2000$) and the naive strategy, by Collective 1's share s_1 of the total collective mass $\bar{\alpha}_N$. A bound of 0 is vacuous.

	Adaptive, $\bar{\alpha}_N =$			Naive, $\bar{\alpha}_N =$		
s_1	0.90	0.75	0.60	0.90	0.75	0.60
Collective 1 (17 features)						
0.11	0.812	0.642	0.290	0.978	0.000	0.000
0.21	0.815	0.815	0.814	0.000	0.981	0.288
0.31	0.987	0.816	0.815	0.983	0.983	0.981
0.51	0.988	0.988	0.987	0.985	0.984	0.984
0.71	0.989	0.989	0.988	0.986	0.986	0.985
Collective 2 (15 features)						
0.11	0.422	0.423	0.144	0.418	0.364	0.137
0.21	0.200	0.145	0.144	0.193	0.364	0.362
0.31	0.042	0.011	0.012	0.364	0.362	0.138
0.51	0.015	0.012	0.013	0.363	0.000	0.137
0.71	0.013	0.014	0.041	0.000	0.134	0.000

Table 3: Size $n_{e,c}$ of the estimation split: adaptive unplanting bounds at total collective mass $\bar{\alpha}_N = 0.90$.

	Collective 1, $n_{e,c} =$			Collective 2, $n_{e,c} =$		
s_1	100	2000	10000	100	2000	10000
0.11	0.811	0.812	0.812	0.422	0.422	0.422
0.21	0.986	0.815	0.815	0.200	0.200	0.200
0.31	0.987	0.987	0.986	0.256	0.042	0.042
0.51	0.988	0.988	0.988	0.226	0.015	0.012
0.71	0.989	0.989	0.989	0.142	0.013	0.014

F Collectives that Are Not Neutral

Neutrality requires each collective to be a random slice of the population before editing, which fails for collectives that form around shared characteristics. Here we relax it to closeness in total variation. Let $P_{0,c} := \mathbb{P}[f_0(\omega) \in \cdot \mid \omega \in \Omega_c]$ be the pre-edit law of collective c and $P_{0,\mathrm{nc}} := \mathbb{P}[f_0(\omega) \in \cdot \mid \omega \notin \cup_{d=1}^{M} \Omega_d]$ that of the non-collective users, and assume

$$\mathrm{TV}(P_0, P_{0,c}) \;\le\; \phi_{N,c}, \qquad c = 1, \dots, M. \tag{66}$$

The bias of the non-collective population is not a new primitive: the mixture identity $P_0 = (1-\bar{\alpha})P_{0,\mathrm{nc}} + \sum_c \alpha_c P_{0,c}$ gives $(1-\bar{\alpha})(P_0 - P_{0,\mathrm{nc}}) = \sum_c \alpha_c (P_{0,c} - P_0)$, hence $\phi_{N,0} := \mathrm{TV}(P_0, P_{0,\mathrm{nc}}) \le \sum_c \alpha_c \phi_{N,c}/(1-\bar{\alpha})$. Neutrality is the case $\phi_{N,c} = \phi_{N,0} = 0$.

Proposition F.1. *Under the assumptions of Theorem 3.1, with neutrality replaced by* (66)*, with probability at least* $1-\delta$*,*

$$\hat{S}_c(\alpha_{N,c}) \;\ge\; \mathbb{P}_{x \sim \hat{P}^{\mathrm{pre},X}_{N,c}} \Big[m_c\big(g_c(x)\big) - 2\alpha_{N,0}\big(\phi_{N,c} + \phi_{N,0}\big) > 0 \Big] \\ - R_{\delta/4}(n_{N,c}) - R_{\delta/4}(N^{\mathrm{test}}) - \phi_{N,c},$$

where $m_c(x^*)$ *denotes the margin inside the probability in* (13)*. The analogous statement holds for Theorem 4.1, with the holdout sample* $R_{N,c}$ *of size* $n_{N,c} - n_{e,c}$ *in place of the sample of collective* c*.*

Proof. Neutrality enters the proof of Theorem 3.1 in two places. In Step 3 it centers the non-collective sample and the pre-edit sample of collective c at P_0, which yields (29)–(32) and, from them, the conflict comparisons (33) and (34). In Step 5 it makes the pre-edit features of collective c i.i.d. with law P_0^X, which yields (41). The other steps do not use it: the margin lemma and the mixture decomposition are distribution-free, the cross-collective conflicts $\Delta_{c \to d}$ are empirical quantities of the other collectives' post-edit samples, and the concentration of the signal term holds verbatim once $p^X_{c,\mathrm{pop}}$ is taken to be the feature law of the edited points of collective c, $((h_c)_{\#} P_{0,c})^X$, which does not appear in the final bound.

Without neutrality, Hoeffding's inequality centers the non-collective sample at $P_{0,\mathrm{nc}}$ and the pre-edit sample of collective c at $P_{0,c}$. Since $|Q(B) - Q'(B)| \le \mathrm{TV}(Q, Q')$ for every event B, (29)–(30) hold with $R_{\tilde{\delta}}(n_{N,0}) + \phi_{N,0}$, and (31)–(32) with $R_{\tilde{\delta}}(n_{N,c}) + \phi_{N,c}$, in place of the radii. A conflict term is a maximum of differences of two such probabilities, so (33) and (34) gain the additive terms $2\phi_{N,0}$ and $2\phi_{N,c}$, respectively. Both comparisons enter the margin multiplied by $\alpha_{N,0}$, which subtracts $2\alpha_{N,0}(\phi_{N,c} + \phi_{N,0})$ from it. In Step 5, Hoeffding's inequality applies to the law $P^X_{0,c}$ of the pre-edit features of collective c, and $\mathrm{TV}(P^X_0, P^X_{0,c}) \le \mathrm{TV}(P_0, P_{0,c}) \le \phi_{N,c}$ gives $P^X_0(A^{\mathrm{pop}}_c) \ge P^X_{0,c}(A^{\mathrm{pop}}_c) - \phi_{N,c}$, which adds $-\phi_{N,c}$ outside the probability. The concentration events, and hence the failure probability, are unchanged. In the proof of Theorem 4.1, neutrality enters in the same two places, through the holdout sample. □

Three consequences follow. First, the degradation is linear in $(\phi_{N,c}, \phi_{N,0})$, with no threshold effect, and Theorems 3.1 and 4.1 are recovered at $\phi = 0$. Second, a non-representative collective pays twice: its data are a biased mirror of the non-collective population, which costs $2\alpha_{N,0}\phi_{N,c}$ inside the margin, and its features are a biased proxy for where the test mass lies, which costs $\phi_{N,c}$ outside it. Third, the penalty inside the margin comes from estimating the non-collective conflict through the population law, so it does not depend on the information regime for the cross-collective term: even in the best case of Section 3, where an aggregator publishes the cross-collective term, the non-collective conflict is still estimated from the collective's own sample.

G Validity for Trained Classifiers

Theorem 3.1 holds for any classifier that is ε-contamination-suboptimal under the post-edit distribution. Here we test it on classifiers that a platform would actually train on collectively edited data; Table 1 in Section 5 summarizes the results.

Protocol. Two collectives perform feature-label planting with total mass $\bar{\alpha}_N = 0.3$, split as $\alpha_{N,1} = x\,\bar{\alpha}_N$ and $\alpha_{N,2} = (1-x)\,\bar{\alpha}_N$ for $x \in \{0.2, 0.35, 0.5, 0.65, 0.8\}$; each training point joins a collective independently, as in the sampling model of Section 2. In the *aligned* scenario both collectives plant the same feature values with the same target label; under *full conflict* they plant the same feature values with different labels; under *partial conflict* Collective 2 plants a subset of Collective 1's feature values with a different label. The platform trains multinomial logistic regression (LR) on one-hot features (L2 penalty with $C = 1$, up to 500 L-BFGS iterations) or histogram gradient-boosted trees (GBT) with native categorical features and default hyperparameters, both from scikit-learn [25]. The empirical success of collective c is the fraction of $N^{\mathrm{test}} = 100{,}000$ fresh test points x with $\hat{m}(g_c(x)) = y^*_c$, and its bound is Theorem 3.1 in the per-collective regime with $\delta = 0.05$ and $\varepsilon = 0$. Setting $\varepsilon = 0$ gives the most demanding test: the bound is largest at $\varepsilon = 0$ and decreases with ε, so validity at $\varepsilon = 0$ implies validity for every $\varepsilon > 0$. All draws use a single fixed seed.

Data. *Synthetic:* a generator written for these experiments, with the 18 features and 4 labels of Appendix C. It differs from the generator behind Figures 1 and 2: features are drawn independently from fixed non-uniform marginals, and labels from an additive per-feature score plus Gumbel noise. We use $N = 2{,}000{,}000$. Collective 1 plants 17 of the 18 features (all but cooling provision); under partial conflict Collective 2 plants 15 of them. *ACSIncome* [9]: US Census American Community Survey microdata (2018, one-year) for the 12 most populous states, with the binary label "income above \$50K" (positive rate 0.381) and 10 features, discretized to keep $\mathcal{X}$ finite: age, class of worker, education, marital status, occupation group, place of birth, relationship, weekly hours, sex and race. We use $N = 877{,}309$ training and $100{,}000$ test points. Collective 1 plants 9 of the 10 features (all but sex) with target "above \$50K"; Collective 2 plants the same 9 features (8 under partial conflict, also leaving marital status free), with the same target when aligned and "at most \$50K" otherwise.

Results. Table 4 reports every configuration and Figure 17 plots them. GBT never violates the bound, on either dataset, and the bound is close to tight where the signal is well identified: under synthetic partial conflict, Collective 2 is certified 0.780 against an empirical success of 0.786. LR satisfies the bound on ACSIncome but violates it twice on the synthetic data, for Collective 2 under partial conflict at $x = 0.5$ and $x = 0.65$. There, Collective 1's signal is a conjunction of feature values that Collective 2's nested pattern contests with a different label, which a model linear in one-hot features cannot represent; LR is then far from Bayes-optimal and the premise of Theorem 3.1 fails. Removing the regularization ($C = \infty$) leaves LR's empirical success unchanged in all 30 configurations, so the violations are not a regularization artifact. The phenomena predicted by the theory appear on both datasets. Aligned collectives pool their masses and are both certified at about 0.99 for every split (Corollary 3.2(i)). Under full conflict the larger collective is certified and the smaller one is not, and

at the even split neither is (Corollary 3.2(iii)). Under partial conflict Collective 1 wins the contested values once its mass is large enough ($x \geq 0.35$ on the synthetic data, $x \geq 0.5$ on ACSIncome), while Collective 2 keeps its uncontested planted values. On ACSIncome, Collective 2's bound under partial conflict (0.31 against 0.46 at $x = 0.5$ and $x = 0.65$) is limited by finite-sample concentration over its 10 planted values, some of which carry too little mass to be certified at this sample size. A run takes about 9 minutes (synthetic) and 2 minutes (ACSIncome) on an Apple M3 Pro CPU; the code is available at `https://github.com/clabat9/multi-collective-aca`.

Table 4: Trained classifiers: Theorem 3.1 lower bound and empirical success of gradient-boosted trees (GBT) and logistic regression (LR), for each scenario, collective c and mass split x. A bound of 0 is vacuous. Violations of the bound are in bold.

			Synthetic, $N = 2{,}000{,}000$			ACSIncome, $N = 877{,}909$		
Scenario	c	x	Bound	GBT	LR	Bound	GBT	LR
Aligned	1	0.20	0.991	1.000	1.000	0.989	1.000	1.000
	1	0.35	0.992	1.000	1.000	0.990	1.000	1.000
	1	0.50	0.993	1.000	1.000	0.991	1.000	1.000
	1	0.65	0.993	1.000	1.000	0.992	1.000	1.000
	1	0.80	0.993	1.000	1.000	0.992	1.000	1.000
	2	0.20	0.993	1.000	1.000	0.992	1.000	1.000
	2	0.35	0.993	1.000	1.000	0.992	1.000	1.000
	2	0.50	0.993	1.000	1.000	0.991	1.000	1.000
	2	0.65	0.992	1.000	1.000	0.990	1.000	1.000
	2	0.80	0.991	1.000	1.000	0.989	1.000	1.000
Full conflict	1	0.20	0.000	0.000	0.000	0.000	0.000	0.000
	1	0.35	0.000	0.000	0.000	0.000	0.000	0.000
	1	0.50	0.000	1.000	0.517	0.000	0.520	0.520
	1	0.65	0.993	1.000	1.000	0.992	1.000	1.000
	1	0.80	0.993	1.000	1.000	0.992	1.000	1.000
	2	0.20	0.993	1.000	1.000	0.992	1.000	1.000
	2	0.35	0.993	1.000	1.000	0.992	1.000	1.000
	2	0.50	0.000	0.000	0.483	0.000	0.480	0.480
	2	0.65	0.000	0.000	0.000	0.000	0.000	0.000
	2	0.80	0.000	0.000	0.000	0.000	0.000	0.000
Partial conflict	1	0.20	0.000	1.000	0.000	0.000	0.000	0.000
	1	0.35	0.992	1.000	1.000	0.000	0.480	0.000
	1	0.50	0.993	1.000	1.000	0.991	1.000	1.000
	1	0.65	0.993	1.000	1.000	0.992	1.000	1.000
	1	0.80	0.993	1.000	1.000	0.992	1.000	1.000
	2	0.20	0.780	0.786	1.000	0.911	1.000	1.000
	2	0.35	0.780	0.786	0.786	0.309	0.756	1.000
	2	0.50	0.779	0.786	**0.694**	0.310	0.460	0.460
	2	0.65	0.692	0.786	**0.607**	0.309	0.460	0.460
	2	0.80	0.000	0.786	0.438	0.000	0.460	0.327

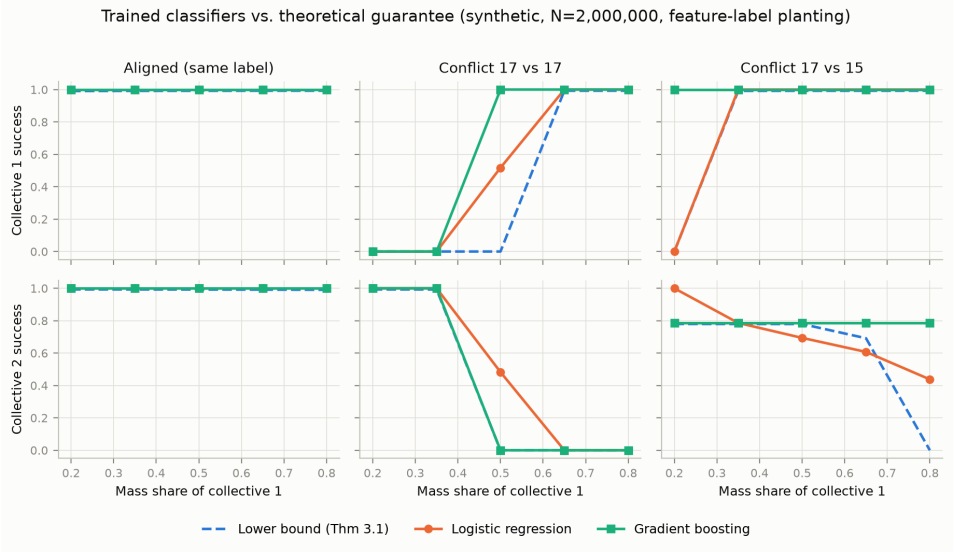

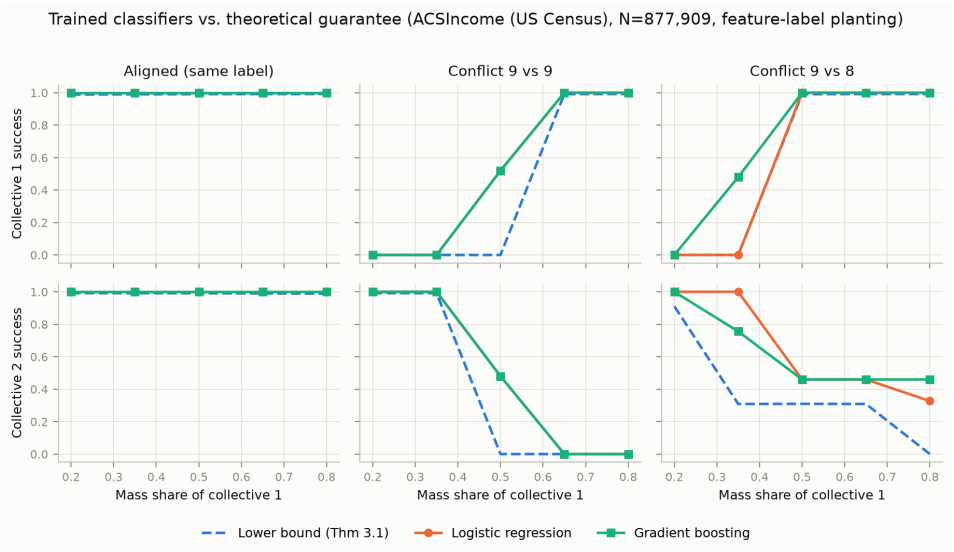

Figure 17: Trained classifiers: lower bound of Theorem 3.1 (dashed) and empirical success of logistic regression and gradient-boosted trees, by Collective 1's share of the total collective mass. Top: synthetic data. Bottom: ACSIncome.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: The abstract and introduction state the main contributions: (i) a statistical framework for ACA with multiple collectives, (ii) lower bounds on per-collective success for signal planting (Theorem 3.1, with Corollary 3.2) and unplanting (Theorem 4.1), with a computability hierarchy, and (iii) numerical illustrations on a synthetic climate-adaptation dataset, together with a check of the bounds for classifiers trained on synthetic and US Census data. Limitations, including finite feature spaces and disjoint collectives, are discussed in Section 6, and neutrality in Section 2 and Appendix F.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: Section 6 discusses the disjointness assumption on collectives and what relaxing it requires, the finite feature space, and the absence of adversarial or game-theoretic modeling; its Societal Impact paragraph discusses dual use. The neutrality assumption is discussed in Section 2 and relaxed in Appendix F, and the role of the suboptimality premise is illustrated by the trained-classifier experiments of Appendix G.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

   Justification: Theorems 3.1 and 4.1 state all assumptions (neutrality, ε-contamination-suboptimality, strategy class). Complete proofs are in Appendix A.2, each following a six-step structure outlined at its start; the supporting margin lemmas are stated and proved within the proofs. Corollary 3.2 and Propositions A.1 and F.1 are also proved, and Appendix D verifies consistency with the single-collective results of [12].

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: Appendix C describes the synthetic dataset (3,000,000 instances, 18 categorical features, 4 label classes), whose generator is included in our code. Section 5 and Appendices E and G specify N, N^{test}, δ, ε, the mass sweeps, the overlap regimes and the conflict settings, including the intervention pool of the known-interventions setting, and, for the trained classifiers, the data, models and hyperparameters.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes]

   Justification: Code to reproduce all experiments is available at `https://github.com/clabat9/multi-collective-aca`, with a pinned conda environment and instructions for each figure and table. The synthetic dataset is generated by a seeded script in the repository, and the ACSIncome data are downloaded from the US Census Bureau with the folktables package [9].

6. **Experimental setting/details**

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

Answer: [Yes]

Justification: Section 5 specifies the key parameters: $N = 200{,}000$, $N^{\text{test}} = 50{,}000$ (two collectives) or $20{,}000$ (four collectives), $\delta = 0.05$, $\varepsilon = 0$, the total-mass levels, the mass-ratio sweeps, the overlap regimes (17 vs. 17/15/13) and the conflict settings. Appendix E gives the intervention pool of the known-interventions setting and the unplanting setup, Appendix G the data, models and hyperparameters of the trained classifiers, and Appendix C the features and labels of the synthetic dataset.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [No]

   Justification: The bound experiments evaluate deterministic functions of the samples, and the confidence level $\delta = 0.05$ is part of each guarantee. The trained-classifier experiments use a single random seed, and we report every configuration rather than averages, so no error bars are given.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: All experiments run on a single laptop CPU (Apple M3 Pro, 18 GB of memory). The trained-classifier experiments take about 9 minutes (synthetic) and 2 minutes (ACSIncome); with the released code, the bound sweeps behind Figures 1 and 2 take under an hour in total.

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes]

   Justification: The research involves theoretical analysis, synthetic data and public, de-identified US Census microdata (ACS PUMS) accessed through folktables. No human subjects, private data, or ethically sensitive applications are involved.

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [Yes]

    Justification: The paper is motivated by positive impacts, such as empowering grassroots collectives in climate adaptation and enabling coordinated action against harmful classifiers (Section 1). The Societal Impact paragraph of Section 6 discusses the dual use of coordinated data editing, which could also serve harmful coordination and which the same bounds help platforms analyze defensively, and how the bounds quantify power asymmetries between collectives.

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

    Answer: [N/A]

    Justification: We release code for computing the bounds and running the experiments of the paper; no pretrained models, scraped datasets, or other high-risk artifacts are released.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We credit [12], whose public code (which states no license) is the starting point of our bound experiments and dataset generator. Our repository credits it explicitly. ACSIncome [9] is built from public US Census Bureau ACS PUMS data, used under the Census Bureau's terms through the folktables package (MIT license), and models are trained with scikit-learn [25] (BSD-3-Clause license).

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes]

Justification: We release our code under the MIT license, with documentation of the environment, of each experiment, and of the parts adapted from [12]. We release no new datasets: the synthetic data are generated by our code, and the Census data are downloaded from the source.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [N/A]

Justification: The paper does not involve crowdsourcing or research with human subjects; besides synthetic data, it uses only public, de-identified Census microdata.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [N/A]

Justification: No human subjects research is conducted in this paper.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

Answer: [N/A]

Justification: The core methodology does not involve LLMs as an important, original, or non-standard component. LLMs were used for editing, as assistance in proving theorems, and as coding assistants for the experiments.