
Unravelling in Collaborative Learning

Aymeric Capitaine¹

Etienne Boursier²

Antoine Scheid¹

Eric Moulines¹

Michael I. Jordan³

El-Mahdi El-Mhamdi¹

Alain Durmus¹

¹ Centre de Mathématiques Appliquées – CNRS – École polytechnique – Palaiseau, 91120, France

² INRIA Saclay, Université Paris Saclay, LMO - Orsay, 91400, France

³ Inria, Ecole Normale Supérieure, PSL Research University - Paris, 75, France

Abstract

Collaborative learning offers a promising avenue for leveraging decentralized data. However, collaboration in groups of strategic learners is not a given. In this work, we consider strategic agents who wish to train a model together but have sampling distributions of different quality. The collaboration is organized by a benevolent aggregator who gathers samples so as to maximize total welfare, but is unaware of data quality. This setting allows us to shed light on the deleterious effect of *adverse selection* in collaborative learning. More precisely, we demonstrate that when data quality indices are private, the coalition may undergo a phenomenon known as *unravelling*, wherein it shrinks up to the point that it becomes empty or solely comprised of the worst agent. We show how this issue can be addressed without making use of external transfers, by proposing a novel method inspired by probabilistic verification. This approach makes the grand coalition a Nash equilibrium with high probability despite information asymmetry, thereby breaking unravelling.

1 Introduction

Collaborative learning is a framework in which multiple agents share their data and computational resources to address a common learning task [Blum et al., 2017, Kairouz et al., 2021]. A significant challenge arises when the quality of these distributions is unknown centrally and agents are strategic. Indeed, participants may be tempted to withhold or misrepresent the quality of their data to gain a competitive advantage. These strategic behaviors and their consequences have been studied extensively in the literature on information economics [Mas-Colell et al., 1995, Laffont and Martimort, 2001]. In particular, information asymmetry is known to result in *adverse selection*, whereby low-quality goods end up dominating the market. In the current paper we study collaborative learning from the perspective of information economics.

A vivid illustration of adverse selection is found in Akerlof’s seminal work on the market for lemons [second-hand cars of low quality, Akerlof, 1970]. Because buyers cannot properly assess the quality of cars on the second-hand market, their inclination to pay decreases. As a consequence, sellers with high-quality cars withdraw from the market, since the proposed price falls below their reservation price. This in turn lowers the buyers’ expectation regarding the average quality of cars on the market, so their willingness to pay decreases even more, which *de facto* crowds out additional cars. The

market may therefore enter a death spiral, up to the point where only low-quality cars are exchanged in any competitive Nash equilibrium. This phenomenon is known as *unravelling*. The insurance market serves as another poignant example of this effect [Rothschild and Stiglitz, 1976, Einav and Finkelstein, 2011, Hendren, 2013]. In insurance, information asymmetry arises due to the fact that insureds possess private knowledge about their individual risk profiles, which insurers lack. Individuals with higher risk are more inclined to purchase policy, while those with lower risk opt out. Consequently, insurers are left with a pool of policyholders skewed towards higher risk, leading to increased premiums to cover potential losses. This, in turn, prompts low-risk individuals to exit the market, exacerbating adverse selection further—a cycle reminiscent of the unravelling described by Akerlof.

We study the problem of whether collaborative learning could also fall victim to unravelling. We consider strategic agents who have access to sampling distributions of varying quality and wish to jointly train a model. They delegate the training process to a central authority who collects samples so as to maximize total welfare. We ask whether adverse selection can arise when data quality is private information. In particular, the presence of a low-quality data owner may harm the model, prompting high-quality owner to leave the collaboration and train a model entirely on their own. Their departure would decrease the average data quality even more, and create a vicious circle. In the worst case, the coalition of learners would reduce to the lowest data quality owner alone. This question is of prime importance from a practical point of view, because unravelling could jeopardize the long-run stability of collaborative models deployed at large scale.

Our contribution is threefold:

1. We provide a rigorous framework for analyzing collaborative learning with strategic agents having data distributions of varying quality. On the one hand, we leverage tools from domain adaptation to capture a notion of data quality formally. On the other hand, we model collaboration as a principal-agent problem, where the principal is an aggregator in charge of collecting samples so as to maximize social welfare. This setup allows us to derive the benchmark welfare-maximizing collaboration scheme when data quality is public information.
2. We show that when data quality is private, a naive aggregation strategy which consists in asking agents to declare their quality type and applying the optimal scheme results in a complete unravelling. More precisely, the set of agents willing to collaborate is either empty or made of the lowest-quality data owner alone at any pure Nash equilibrium.
3. We present solutions to unravelling. When transfers are allowed, the VCG mechanism suffices to re-establish optimality. When transfers are not possible, we leverage probabilistic verification techniques to design a mechanism which breaks unravelling. More precisely, we ensure that the optimal, grand coalition ranks with high probability among the Nash equilibria of the game induced by our mechanism. We demonstrate how to implement our mechanism practically in the setting of classification.

Related work. The issue of information asymmetry in machine learning has been an area of recent activity. Several learning settings have been considered, including bandits [Wei et al., 2024], linear regression [Donahue and Kleinberg, 2021a,b], classification [Blum et al., 2021] and empirical risk minimization [Dorner et al., 2023, Liu et al., 2023] in a federated context [Tu et al., 2022].

Most of these studies focus on the sub-problem of *moral hazard*, where agents take actions that are unobserved by others. This situation usually results in under-provision of effort and inefficiency at the collective scale [Laffont and Martimort, 2001]. This issue appears naturally in federated learning, because model updates are performed locally. For instance, Karimireddy et al. [2022] show that heterogeneity in sampling costs results in total free-riding without a proper incentive scheme. Huang et al. [2022] show that under-provision of data points in federated learning arises from privacy concerns. Yan et al. [2023] consider a federated classification game where agents can reduce the noise in their data distribution but incur a costly effort to do so. In the same vein, Huang et al. [2023] study the case of agents who are interested in different models, and may skew their sampling measure accordingly. Saig et al. [2023] and Ananthakrishnan et al. [2023] study hidden actions when a principal delegates a predictive task to another agent, and show that thresholds or linear contracts are able to approximate the optimal contracts.

Adverse selection is another type of information asymmetry, where preferences of agents are unobserved rather than actions. This issue also naturally arises in collaborative learning, because the data distributions from which agents sample or about which they care may not be public. While

data heterogeneity is a widely explored topic in federated learning (see for instance [Gao et al. 2022](#) for a general survey, and [Fu et al. 2023](#) for the specific problem of *client selection*), the strategic aspect has been rarely considered. Most studies doing so focus on hidden sampling costs (see, e.g., [Karimireddy et al. 2022](#), or [Wei et al. 2024](#) in a bandit context), but few address the fundamental problem of distribution shift. [Ananthkrishnan et al. \[2023\]](#) mentions the issue of adverse selection when a principal delegates a predictive task to an agent, and provide qualitative insights about the optimal contract. However, they consider a single agent and leave aside the question of participation. Finally, [Werner et al. \[2024\]](#) and [Tsoy and Konstantinov \[2024\]](#) study the question of the stability of collaborative learning between competing agents. However, their analysis relies on ad-hoc market structures rather than information asymmetry per se. As such, our work is the first to demonstrate the effect of imperfect information on the sustainability of collaborative learning.

Organization. Section 2 presents our model and assumptions. Section 3 studies a full information benchmark, which allows us to derive the welfare-maximizing contribution scheme. In Section 4, we turn to the more realistic case where data quality is private information. We first show that in this case, a naive aggregation method leads to total unravelling. Second, we introduce a mechanism which breaks unravelling by inducing a game where the grand coalition is a pure strategy Nash equilibrium with high probability.

2 Model

Statistical framework. Let (X, \mathcal{X}) and (Y, \mathcal{Y}) be two measurable spaces and denote by \mathcal{P} a family of probability measures on $(X \times Y, \mathcal{X} \otimes \mathcal{Y})$. We consider $[J] = \{1, \dots, J\}$ agents who aim to perform a prediction task associated with a hypothesis class $\mathcal{G} \subset \{g : X \rightarrow Y\}$, a loss function $\ell : Y \times Y$ and a probability measure $P_0 \in \mathcal{P}$. Each agent seeks to minimize $g \in \mathcal{G} \mapsto \mathcal{R}_{P_0}(g)$ where for any probability distribution $P \in \mathcal{P}$, $\mathcal{R}_P(g) = \int \ell(g(x), y) dP(x, y)$ is the risk associated to $g \in \mathcal{G}$ with respect to P .

We leverage tools and results from statistical learning theory. Denote by $\{(X_i, Y_i)\}_{i \geq 1}$ the canonical process on $X \times Y$ and denote by \mathbb{P}_P and \mathbb{E}_P the canonical probability and expectation under which $\{(X_i, Y_i)\}_{i \geq 1}$ are i.i.d. random variables with distribution $P \in \mathcal{P}$. With this notation, we can introduce our assumptions on \mathcal{G} and \mathcal{P} .

H1. For $\delta \in (0, 1)$, there exist $\alpha_\delta > 0$, $\beta > 0$ and $\gamma > 0$ such that for any distribution $P \in \mathcal{P}$, hypothesis $g \in \mathcal{G}$ and $n \geq 1$

$$\mathbb{P}_P \left(\left| \mathcal{R}_P(g) - \frac{1}{n} \sum_{i=1}^n \ell(g(X_i), Y_i) \right| \leq \frac{\alpha_\delta}{(1+n)^\gamma} + \beta \right) \geq 1 - \delta.$$

This assumption covers a wide range of situations. For instance in the classification case where $Y = \{-1, 1\}$ and $\ell : (y, \tilde{y}) \mapsto \mathbb{1}_{\{y\tilde{y} \leq 0\}}$, $\alpha_\delta = \sqrt{\ln(1/\delta)}/2$, $\beta = 2\text{RAD}(\mathcal{G})$ and $\gamma = 1/2$ where $\text{RAD}(\mathcal{G})$ is the empirical Rademacher complexity of \mathcal{G} [[Bousquet et al., 2003](#)]. Similarly, the Bayesian PAC approach in the linear regression context with bounded loss leads to $\alpha_\delta = \text{KL}(\rho \|\pi) + \ln(1/\delta)$, $\beta = \|\ell\|_\infty^2/8$ and $\gamma = 1$ where ρ and π are any distributions on \mathcal{G} [[Shalaeva et al., 2019](#), Corollary 4].

For ease of notation, we let $\mathcal{R}_j(g)$ serve as a shorthand for $\mathcal{R}_{P_j}(g)$. It is moreover assumed that agent $j \in [J]$ cannot directly sample from P_0 , but has instead access to a distribution $P_j \in \mathcal{P}$ which deviates from P_0 according to the \mathcal{G} -divergence:

H2. For any $j \in [J]$, $P_j \in \mathcal{P}$ has finite \mathcal{G} -divergence: $\theta_j = \sup_{g \in \mathcal{G}} |\mathcal{R}_j(g) - \mathcal{R}_0(g)| < +\infty$.

Intuitively, for any $j \in [J]$, $\theta_j \geq 0$ models the bias incurred by having access to samples from P_j instead of the target distribution P_0 . More precisely, the risk excess associated with empirical risk minimization (ERM) based on samples from P_j is in the worst case at least θ_j . A poor sampling distribution P_j —which corresponds to a high θ_j in the previous expression—might be the consequence of low-quality sensors or degraded experimental conditions resulting in noisier data points.

The class of discrepancies appearing in **H2** has been considered in the domain adaptation literature [see, e.g., [Ben-David et al., 2010](#), [Kifer et al., 2004](#), [Konstantinov and Lampert, 2019](#)]. It provides a natural

framework to analyze the behavior of models trained on diverse distributions, and is practically appealing since it can be easily estimated in the context of classification [Ben-David et al., 2010].

We make the following assumptions on the quality indexes $(\theta_1, \dots, \theta_J)$, hereafter referred to as *types*.

H3. *There exists $(\underline{\theta}, \bar{\theta}) \in \mathbb{R}_+^2$ such that $\underline{\theta} \leq \theta_1 < \theta_2 < \dots < \theta_J \leq \bar{\theta}$.*

H3 The ordering assumption is just for ease of exposition but is neither used by the aggregator nor the agents. Our condition that types are strictly different is for convenience in simplifying the proofs.

Collaborative learning framework. We further suppose that agent $j \in [J]$ can fit a model $g \in \mathcal{G}$ based on i.i.d. samples $\{(X_1^j, Y_1^j), \dots, (X_{n_j}^j, Y_{n_j}^j)\}$ of size $n_j \geq 0$ from $P_j \in \mathcal{P}$ in one of two ways: they can compute on their own, or collaborate. This is captured by the two following options:

Option 1) Agent j performs ERM on their own samples:

$$\hat{g}_j = \operatorname{argmin}_{g \in \mathcal{G}} \hat{\mathcal{R}}_j(g), \quad \hat{\mathcal{R}}_j(g) = n_j^{-1} \sum_{i=1}^{n_j} \ell(g(X_i^j), Y_i^j). \quad (1)$$

This non-collaborative procedure is referred to as the *outside option*.

Option 2) Agent j can take part in a *coalition* orchestrated by a central data aggregator, encoded as $\mathbf{B} = (B_1, \dots, B_J) \in \{0, 1\}^J$, where $B_j = 1$ means that agent j is member of the coalition. We also write $\mathcal{B} = \{j \in [J] : B_j = 1\}$. In exchange for their samples, agent j gains access to the collaborative model trained over the concatenation of samples:

$$\hat{g}_{\mathbf{B}} = \operatorname{argmin}_{g \in \mathcal{G}} \hat{\mathcal{R}}_{\mathbf{B}}(g), \quad (2)$$

$$\text{where } \hat{\mathcal{R}}_{\mathbf{B}}(g) = N^{-1} \sum_{j \in [J]} B_j \sum_{i=1}^{n_j} \ell(g(X_i^j), Y_i^j) \quad \text{and} \quad N = \sum_{j \in [J]} B_j n_j.$$

Agent utilities. We assume that agents incur a unitary cost for sampling from their distribution and dislike statistical risk. Therefore, a baseline model for measuring the preferences associated with a model $g \in \mathcal{G}$ and a number of samples n is based on a linear map: $(g, n) \mapsto -a\mathcal{R}_0(g) - cn$ for $a, c > 0$. In practice, however, $\mathcal{R}_0(g)$ is typically unknown so agents instead can use a PAC bound of the form $\mathbb{P}(\mathcal{R}_0(\hat{g}_n) \leq \mathcal{R}_0^* + \varepsilon) \geq 1 - \delta$ to assess a model $\hat{g}_n \in \mathcal{G}$ trained over their samples, where $\varepsilon > 0$, $\delta \in (0, 1)$ and $\mathcal{R}_0^* = \inf_{g \in \mathcal{G}} \mathcal{R}_0(g)$. Our next result shows that **H 1** and **H 2** allow each agent to pin down such an $\varepsilon > 0$, under either **Option 1)** or **Option 2)**. Assuming **H1**, define the function ε for any $(\theta, n) \in \Theta \times \mathbb{R}_+$ by

$$\varepsilon(\theta, n) = 2[\alpha_\delta(1+n)^{-\gamma} + \beta + \theta]. \quad (3)$$

Lemma 1. *Assume **H1** and **H2**.*

(i) *Any agent $j \in [J]$ picking the outside **Option 1)** obtains a model $\hat{g}_j \in \mathcal{G}$ achieving*

$$\mathcal{R}_0(\hat{g}_j) \leq \mathcal{R}_0^* + \varepsilon(\theta_j, n_j) \quad \text{with probability } 1 - \delta.$$

(ii) *Any coalition $\mathbf{B} \in \{0, 1\}^J$ drawing $\mathbf{n} = (n_1, \dots, n_J) \in \mathbb{R}_+^J$, samples obtains a model $\hat{g}_{\mathbf{B}} \in \mathcal{G}$ achieving*

$$\mathcal{R}_0(\hat{g}_{\mathbf{B}}) \leq \mathcal{R}_0^* + \varepsilon(\vartheta, N) \quad \text{with probability } 1 - \delta,$$

where N is the total of samples and ϑ is the weighted average type within the coalition:

$$N(\mathbf{B}, \mathbf{n}) = \sum_{j \in [J]} B_j n_j, \quad \vartheta(\mathbf{B}, \mathbf{n}) = N^{-1} \sum_{j \in [J]} B_j n_j \theta_j. \quad (4)$$

When the context is clear, we write $\varepsilon(\mathbf{B}, \mathbf{n}) = \varepsilon(\vartheta(\mathbf{B}, \mathbf{n}), N(\mathbf{B}, \mathbf{n}))$ to lighten notation. Based on Lemma 1, we define the utility of agent $j \in [J]$ as

$$u_j : (\mathbf{B}, \mathbf{n}) \mapsto -a[\mathcal{R}_0^* + (1 - B_j)\varepsilon(\theta_j, n_j) + B_j\varepsilon(\mathbf{B}, \mathbf{n})] - cn_j. \quad (5)$$

Note that for any $j \in [J]$, we take $n_j \geq 0$ to be a real number for ease of presentation. In (5), $a/c > 0$ captures the extent to which individuals are willing to trade off model quality against sampling cost.

For any $j \in [J]$ and $\mathbf{B}_{-j} = (B_1, \dots, B_{j-1}, B_{j+1}, \dots, B_J) \in \{0, 1\}^{J-1}$, we denote by a slight abuse of notation $(B, \mathbf{B}_{-j}) = (B_1, \dots, B_{j-1}, B, B_{j+1}, \dots, B_J)$ for any $B \in \{0, 1\}$. Similarly for $\mathbf{n}_{-j} = (n_1, \dots, n_{j-1}, n_{j+1}, \dots, n_J) \in \mathbb{R}_+^{J-1}$, we write $(n, \mathbf{n}_{-j}) = (n_1, \dots, n_{j-1}, n, n_{j+1}, \dots, n_J)$ for any $n \geq 0$. We can then characterize the optimal behavior of any agent picking the outside option as follows.

Proposition 1. *Assume **H1** and **H2**. For any $j \in [J]$, $\mathbf{B}_{-j} \in \{0, 1\}^{J-1}$ and $\mathbf{n}_{-j} \in \mathbb{R}_+^{J-1}$, the optimal number of samples to draw under **Option 1** is $\arg\max_{n \geq 0} u_j((0, \mathbf{B}_{-j}), (n, \mathbf{n}_{-j}); \theta_j) = \underline{n}$ where*

$$\underline{n} = (2ac^{-1}\gamma\alpha_\delta)^{1/(\gamma+1)} - 1. \quad (6)$$

In what follows, we assume $2a/c > (\gamma\alpha_\delta)^{-1}$ to exclude the pathological case where no agent is willing to sample data points. From now on, we denote by

$$\underline{u}_j = u_j((0, \mathbf{B}_{-j}), (\underline{n}, \mathbf{n}_{-j})) \quad (7)$$

the best achievable utility under the outside option. Note that \underline{n} does not depend on $\theta_j \in \Theta$ (but \underline{u}_j does) so all agents outside of the coalition draw a same number of data points $\underline{n} > 0$. This result, which may be surprising at first glance, comes from the fact that all agents have the same accuracy-to-sampling-cost ratio a/c in their utility.

Aggregator. We finally assume that the aggregator acts benevolently to set up a Pareto-optimal collaboration, by maximizing the total welfare under individual rationality. In other words, they solve:

$$\begin{aligned} \text{maximize } W : (\mathbf{B}, \mathbf{n}) \in \{0, 1\}^J \times \mathbb{R}_+^J &\mapsto \sum_{j \in [J]} u_j(\mathbf{B}, \mathbf{n}) \\ \text{subject to } \min_{j \in [J]} u_j(\mathbf{B}, \mathbf{n}) - \underline{u}_j &\geq 0. \end{aligned} \quad (8)$$

In the Social Choice literature, W is referred to as the utilitarian social welfare function. The participation constraint ensures that no agent within the coalition finds it beneficial to switch to their outside option. Note that $\mathbf{n} \in \mathbb{R}_+^J$ is required to have non-negative entries, which prevents the aggregator from giving away data points to agents.

3 Full-Information Benchmark: First-Best Collaboration

In this section, we assume that the profile of types $(\theta_1, \dots, \theta_J) \in \Theta^J$ is public, and study how the aggregator can implement an optimal collaboration among agents under this most-favorable scenario.

Exact solution. We are looking for a solution to the aggregator's problem (8). For any $\mathbf{B} \in \{0, 1\}^J$ and $\mathbf{n} \in \mathbb{R}_+^J$, denote by

$$\begin{aligned} \bar{n}_j(\mathbf{B}, \mathbf{n}) &= \max\{n \geq 0 : u_j((1, \mathbf{B}_{-j}), (n, \mathbf{n}_{-j}); \theta_j) \geq \underline{u}_j\} \\ &= \underline{n} - (a/c)[\varepsilon(\mathbf{B}, \mathbf{n}) - \varepsilon(\theta_j, \underline{n})], \end{aligned} \quad (9)$$

the maximum number of samples that agent j can be asked to provide within the coalition under its participation constraint, where \underline{n} is defined as in (6), and \underline{u}_j as in (7). With this notation, problem (8) rewrites

$$\text{maximize } W(\mathbf{B}, \mathbf{n}) \quad \text{subject to } \min_{j \in [J]} \bar{n}_j(\mathbf{B}, \mathbf{n}) - n_j \geq 0. \quad (10)$$

Theorem 1. *Assume **H1**, **H2**, **H3**. Problem (8) admits a unique solution $(\mathbf{B}^{\text{opt}}, \mathbf{n}^{\text{opt}}(\boldsymbol{\theta})) \in \{0, 1\}^J \times \mathbb{R}_+^J$. Moreover,*

(i) $\mathbf{B}^{\text{opt}} = \mathbf{1} = (1, \dots, 1)$,

(ii) Denoting $\mathbf{n}^{\text{opt}}(\boldsymbol{\theta}) = (n_1^{\text{opt}}(\boldsymbol{\theta}), \dots, n_J^{\text{opt}}(\boldsymbol{\theta}))$, there exists $L^{\text{opt}} \in [J]$ such that for any $j \in [J]$,

$$n_j^{\text{opt}}(\boldsymbol{\theta}) \begin{cases} = \bar{n}_j(\mathbf{1}, \mathbf{n}^{\text{opt}}(\boldsymbol{\theta})) & \text{if } j < L^{\text{opt}}, \\ \in [0, \bar{n}_j(\mathbf{1}, \mathbf{n}^{\text{opt}}(\boldsymbol{\theta}))] & \text{if } j = L^{\text{opt}}, \\ = 0 & \text{otherwise.} \end{cases} \quad (11)$$

The couple $(\mathbf{B}^{\text{opt}}, \mathbf{n}^{\text{opt}})$ is referred to as the *optimal contribution scheme*. Although implicit, the condition (11) provides insights about the optimal scheme. The aggregator makes everyone enter the coalition, but only asks the $L^{\text{opt}} > 0$ first-best agents to contribute. This allows to obtain the best possible collaborative model while sparing any sampling cost to other agents. Moreover, the number of required samples $n_j^{\text{opt}}(\boldsymbol{\theta})$ slightly differs from \bar{n} according to the relative performance of the collaborative model with respect to agent j 's one: if the agent gets a better accuracy by collaborating, the aggregator can ask them for more data; if on the other hand the agent gets a worse model by collaborating (i.e., they are a contributor with very high quality data), the aggregator can only ask less data because of the participation constraint.

Relaxed solution. Working with the optimal scheme $(\mathbf{B}^{\text{opt}}, \mathbf{n}^{\text{opt}})$ is difficult because $\bar{n}_j(\mathbf{1}, \mathbf{n}^{\text{opt}}(\boldsymbol{\theta}))$ has no explicit expression. To make the analysis tractable, we slightly simplify the optimal contribution scheme in Theorem 1 and consider the simplified optimal contribution scheme $(\mathbf{B}^*, \mathbf{n}^*)$ where $\mathbf{B}^* = \mathbf{B}^{\text{opt}} = (1, \dots, 1)$ and for any $j \in [J]$,

$$n_j^*(\boldsymbol{\theta}) = \mathbb{1}_{\{j \leq L^*\}} \bar{n}_j(\mathbf{1}, \mathbf{n}^*(\boldsymbol{\theta})) \quad \text{and} \quad L^* = \min\{j \in [J] : \sum_{k \leq j} \bar{n}_k(\mathbf{1}, \mathbf{n}^*(\boldsymbol{\theta})) \geq \bar{N}\}, \quad (12)$$

with $\bar{N} = (\underline{n} + 1)J^{\frac{1}{1+\gamma}} - 1$.

Note that $(\mathbf{B}^*, \mathbf{n}^*)$ only differs from $(\mathbf{B}^{\text{opt}}, \mathbf{n}^{\text{opt}})$ in two ways. First, in $(\mathbf{B}^*, \mathbf{n}^*)$ all contributors' participation constraint bind, while in $(\mathbf{B}^{\text{opt}}, \mathbf{n}^{\text{opt}})$ the L^{opt} -the one could be slack. Second, the total number of data points required from the coalition is fixed and equal to $\bar{N} = \Theta(J^{\frac{1}{1+\gamma}})$. The quantity \bar{N} comes from a natural relaxation of the original problem Equation (10) where we leave aside an intricate term of the objective function. This relaxation, which is formally described in Appendix A, provides a good approximation of the exact solution in reasonable settings. Indeed, the following result establishes that applying $(\mathbf{B}^*, \mathbf{n}^*)$ instead of $(\mathbf{B}^{\text{opt}}, \mathbf{n}^{\text{opt}})$ comes at a negligible welfare cost when types are sufficiently evenly spaced.

Lemma 2. Assume **H1**, **H2** and $\theta_j - \theta_{j-1} = \mathcal{O}(1/J)$ for any $j \in \{2, \dots, J\}$. Then,

$$W(\mathbf{B}^{\text{opt}}, \mathbf{n}^{\text{opt}}(\boldsymbol{\theta})) = W(\mathbf{B}^*, \mathbf{n}^*(\boldsymbol{\theta})) + \mathcal{O}(J^{\frac{1}{1+\gamma}}).$$

Moreover, the following proposition shows that $\mathbf{n}^*(\boldsymbol{\theta})$ admits a workable expression.

Corollary 1. Assume **H1**, **H2** and **H3**. Then $L^* = \Theta(J^{\frac{1}{1+\gamma}})$ and for any $j \in [J]$,

$$n_j^*(\boldsymbol{\theta}) = \mathbb{1}_{\{j \leq L^*\}} \left[\frac{\bar{N}}{L^*} + \frac{2a}{c} \left(\theta_j - \frac{1}{L^*} \sum_{\ell=1}^{L^*} \theta_\ell \right) \right].$$

Since $(\mathbf{B}^*, \mathbf{n}^*)$ correctly approximates the optimal scheme while being more tractable, we work with it in the remainder to lighten proofs.

H4. The aggregator applies the simplified contribution scheme $(\mathbf{B}^*, \mathbf{n}^*)$.

4 Hidden information

The welfare-maximizing contribution scheme described in (12) depends explicitly on $\boldsymbol{\theta} \in \Theta^J$, so it is implementable only if types are public. This often unrealistic, either for legal or competitive reasons. We therefore turn to the problem of setting up a collaboration when types are private.

4.1 Naive aggregation and unravelling

A naive solution to coping with the private nature of $\theta \in \Theta^J$ is for the aggregator to ask agents to disclose their types, and apply the simplified optimal contribution scheme defined in (12). In this setting, however, agents may declare a type $\tilde{\theta}_j$ different from their true type θ_j .

This approach corresponds to a direct-revelation mechanism $\Gamma : (\mathbf{B}, \tilde{\boldsymbol{\theta}}) \mapsto \mathbf{n}^*(\tilde{\boldsymbol{\theta}})$ which unfolds as follows.

1. Any agent $j \in [J]$ declares a tuple $(B_j, \tilde{\theta}_j) \in \{0, 1\} \times \Theta \cup \{\dagger\}$. If $B_j = 1$, then agent j picks **Option 2**, and enters the coalition with type $\tilde{\theta}_j$. If $B_j = 0$, then agent j picks **Option 1**, their declared type $\tilde{\theta}_j$ is \dagger by convention.
2. Setting $\mathbf{B} = (B_1, \dots, B_J)$ and $\tilde{\boldsymbol{\theta}} \in (\Theta \cup \{\dagger\})^J$, then the aggregator applies the contribution scheme defined in (12), so the vector of number of contributions within the coalition is $\mathbf{n}^*(\tilde{\boldsymbol{\theta}})$.

Γ induces a direct revelation game $([J], \mathcal{S}^J, (v_j)_{j \in [J]})$ where the action space is $\mathcal{S} = \{(1, \tilde{\theta}) : \tilde{\theta} \in \Theta\} \cup \{(0, \dagger)\}$ and payoffs are for any $j \in [J]$ and $\mathbf{s} \in \mathcal{S}^J$,

$$v_j(s_j, \mathbf{s}_{-j}) = u_j(\mathbf{B}, \mathbf{n}^*(\tilde{\boldsymbol{\theta}})) = B_j \left[-a \left(\mathcal{R}_0^* + \varepsilon(\mathbf{B}, \mathbf{n}^*(\tilde{\boldsymbol{\theta}})) \right) - c n_j^*(\tilde{\boldsymbol{\theta}}) \right] + (1 - B_j) \underline{u}_j.$$

This mechanism is obviously vulnerable to strategic manipulation, since it disregards incentive compatibility. This has severe consequences for the coalition, as shown by the following proposition.

Theorem 2 (Unravelling). *Assume **H1**, **H2**, **H3**, and **H4**. Let $\mathcal{E} \subset \mathcal{S}^J$ be the set of pure-strategy Nash equilibria of the game induced by Γ . We have*

- (i) $\mathcal{E} \neq \emptyset$
- (ii) at any $s^* \in \mathcal{E}$, $\mathbf{B} = (0, \dots, 0)$ or $\mathbf{B} = (0, \dots, 0, 1)$.

Theorem 2 shows that under Γ , the coalition undergoes a full unravelling: it is either empty or comprised solely of the worst agent in any Nash equilibrium. Thus, collaborative learning is not immune to adverse selection, and may suffer from unravelling as any market characterized by information asymmetry.

Sketch of proof. The profile of actions $((0, \dagger), \dots, (0, \dagger))$ corresponding to $\mathbf{B} = (0, \dots, 0)$ is a pure Nash equilibrium, since forming a lone coalition cannot bring more utility than picking the outside option. Conversely, consider a pure-strategy Nash equilibrium $\mathbf{s} \in \mathcal{E}$ such that $\mathbf{B} \neq (0, \dots, 0)$. Denote by $\mathcal{C} = \{j \in [J] : B_j = 1 \text{ and } n_j^*(\tilde{\boldsymbol{\theta}}) > 0\}$ the set of contributors under this equilibrium. It can be shown that (i) for any $(j, k) \in \mathcal{C}^2$, $\theta_j - \tilde{\theta}_j = \theta_k - \tilde{\theta}_k$, and (ii) For any $j \in \mathcal{C}$, $(1, \tilde{\theta}_j)$ with $\tilde{\theta}_j > \underline{\theta}$ is strictly dominated by $(1, \underline{\theta})$ so $\tilde{\theta}_j = \underline{\theta}$ at the equilibrium. As a consequence, $\theta_j = \theta_k$ for any $(j, k) \in \mathcal{C}^2$, which implies by **H3** that $|\mathcal{C}| = 1$. From the definition of the contribution scheme (12), we can deduce that $\sum_{j \in [J]} B_j = |\mathcal{B}| = 1$, because $|\mathcal{B}| > 1$ would entail $|\mathcal{C}| > 1$. Finally, $B_J = 1$ because $v_J((1, \underline{\theta}), \mathbf{s}_{-J}) > v_J((0, \dagger), \mathbf{s}_{-J})$ and \mathbf{s} is a Nash equilibrium. This leads to $\mathbf{B} = (0, \dots, 0, 1)$. \square

4.2 Breaking unravelling

The previous results motivate the design of a more sophisticated aggregation scheme that addresses adverse selection. In this section, we discuss how to design such a procedure.

Is VCG available in our framework? Unravelling occurs under Γ because agents do not find it beneficial to declare their true type and eventually opt for their outside option. This could be avoided by modifying Γ to make it

- (i) *individually rational*, that is $v_j((1, \theta_j), \mathbf{s}_{-j}) \geq v_j((0, \dagger), \mathbf{s}_{-j})$ for any $j \in [J]$, $\mathbf{s}_{-j} \in \mathcal{S}^{J-1}$,
- (ii) and *incentive compatible*, that is $v_j((1, \theta_j), \mathbf{s}_{-j}) \geq v_j((1, \tilde{\theta}_j), \mathbf{s}_{-j})$ for any $\tilde{\theta}_j \in \Theta$.

Under these conditions, the truthful, optimal profile of actions $((1, \theta_1), \dots, (1, \theta_J))$ would emerge as a Nash equilibrium. Since the aggregator seeks to minimize the utilitarian function W , one option would be to rely on the VCG mechanism [Vickrey, 1961, Clarke, 1971, Groves, 1973], which is the direct-revelation mechanisms fulfilling these desiderata [Green and Laffont, 1977, Holmstrom, 1979]. Formally, the VCG mechanism writes $\Gamma^{\text{VCG}} : \tilde{\theta} \mapsto (\mathbf{n}^*(\tilde{\theta}), \mathbf{t}(\tilde{\theta}))$ where $\mathbf{t}(\tilde{\theta}) = (t_1(\tilde{\theta}), \dots, t_J(\tilde{\theta})) \in \mathbb{R}^J$ is a set of transfers satisfying for any $j \in [J]$:

$$t_j(\tilde{\theta}) = \sum_{k \neq j} u_k((0, \mathbf{1}_{-j}), \mathbf{n}^*(\tilde{\theta})) - \sum_{k \neq j} u_k(\mathbf{1}, \mathbf{n}^*(\tilde{\theta})) .$$

It re-establishes truthfulness as a dominant strategy by aligning individual payoffs $v_j^{\text{VCG}}(\tilde{\theta}) = u_j(\mathbf{1}, \mathbf{n}^*(\tilde{\theta})) - t_j(\tilde{\theta})$ with total social welfare. Unfortunately, the VCG approach is unavailable in our framework, because of the following observation.

Lemma 3. *There exists $j \in [J]$ such that $-t_j(\tilde{\theta}) > 0$.*

Lemma 3 shows that some agent would need to receive a strictly positive transfer. This is impossible without a monetary payment—utility can only be decreased by the aggregator, for instance through *accuracy shaping* [Karimireddy et al., 2022]—, which we exclude here.

A probabilistic verification-based mechanism. We now show how to design a mechanism that recovers the optimal collaboration as a Nash equilibrium in high probability without the need for transfers. Inspired by the probabilistic verification approach [Caragiannis et al., 2012, Ferraioli and Ventre, 2018, Ball and Kattwinkel, 2019], we assume that the aggregator can approximately estimate θ_j with few samples from P_j for any $j \in [J]$:

H5. *There exists a decreasing function $\eta_\delta : \mathbb{R}_+^* \rightarrow \mathbb{R}_+^*$, with $\delta \in (0, 1)$ defined in H1, such that for any $j \in [J]$ and i.i.d samples $(X_1^j, Y_1^j), \dots, (X_q^j, Y_q^j)$ of size $q > 0$ from P_j , there exists a $(X_1^j, Y_1^j), \dots, (X_q^j, Y_q^j)$ -measurable estimator $\hat{\theta}_j$ satisfying*

$$\mathbb{P}(|\hat{\theta}_j - \theta_j| \leq \eta_\delta(q)) \geq 1 - \delta .$$

In Section 4.2, we show how the aggregator can design such estimators. H5 allows us to consider a new mechanism $\hat{\Gamma} : \mathbf{B} \mapsto \mathbf{m}(\mathbf{B})$ as follows:

1. any agent $j \in [J]$ declares $B_j \in \{0, 1\}$. If $B_j = 1$, the principal asks for $q \leq \underline{n} - 2(a/c)(\bar{\theta} - \underline{\theta})$ i.i.d samples from P_j and estimates types as $\hat{\theta} = (\hat{\theta}_j)_{j \in \mathcal{B}}$ following H5.
2. Based on the estimated types $\hat{\theta} = (\hat{\theta}_j)_{j \in \mathcal{B}}$, the aggregator asks for $[n_j^*(\hat{\theta} + \boldsymbol{\eta}_j) - \underline{q}]_+$ additional samples from P_j , where $\mathbf{n}^*(\cdot)$ is defined as in (12) and

$$\boldsymbol{\eta}_j = \eta_{\delta/J}(q)\mathbf{1} - 2\delta_j\eta_{\delta/J}(q), \quad \text{with } \delta_j = (0, \dots, 0, 1, 0, \dots, 0)^T .$$

Thus, the number of draws required from agent j is $\max[\underline{q}, n_j^*(\hat{\theta} + \boldsymbol{\eta}_j)]$.

3. The aggregator keeps $m_j(\hat{\theta}) = \mathbb{1}\{n_j^*(\hat{\theta} + \boldsymbol{\eta}_j) > 0\} \max[\underline{q}, n_j^*(\hat{\theta} + \boldsymbol{\eta}_j)]$ samples from agent j , and trains a collaborative model with these pooled samples.

$\hat{\Gamma}$ induces a game $([J], \{0, 1\}^J, (\hat{v}_j)_{j \in [J]})$ where any agent $j \in [J]$ has a payoff function

$$\hat{v}_j : (B_j, \mathbf{B}_{-j}) \mapsto B_j \left[-a \left(\mathcal{R}_0^* + \varepsilon(\mathbf{B}, \mathbf{m}(\hat{\theta})) \right) - c \max[\underline{q}, n_j^*(\hat{\theta} + \boldsymbol{\eta}_j)] \right] + (1 - B_j) \underline{u}_j .$$

The rationale behind this mechanism is fairly intuitive: since $\hat{\theta}$ is a correct estimate of θ , $n_j^*(\hat{\theta} + \boldsymbol{\eta}_j) \approx n_j^*(\hat{\theta})$ correctly approximates the optimal contribution $n_j^*(\theta)$ for any contributor $j \in \mathcal{B}$. Note that type estimates are purposely biased by $\boldsymbol{\eta}_j$ when asking for samples. This is a safeguard against over-estimated types, which would lead to asking to many data points and could deter agents from participating in the coalition.

Critically, $\mathbf{m}(\hat{\theta})$ does not depend on declared type, so agents are no longer able to strategically manipulate the mechanism. Moreover, $\hat{\Gamma}$ does not require agents to know their own types, which would be an unrealistic assumption. Finally, observe that the number of data points asked to estimate types q is low enough to never deter agents from participating in the coalition.

Theorem 3. Assume **H1**, **H2**, **H3**, **H4** and **H5**. $B^* = (1, \dots, 1)^T$ is a Nash equilibrium under $\hat{\Gamma}$ with probability $1 - \delta$.

Theorem 3 shows that the optimal coalition is a sustainable equilibrium under $\hat{\Gamma}$, which effectively breaks unravelling: the set of (approximate) Nash equilibria is no more reduced to profiles of actions where the coalition is empty, or reduced to the worst agent.

Practical implementation. We now explain how to practically implement $\hat{\Gamma}$ in a collaborative learning setting. This requires defining a collection of estimators $(\hat{\theta}_j)_j$ satisfying **H5**. To this end, we assume that few samples from the target distribution are available.

H6. There are $q' > 0$ i.i.d samples $\{(X_1^0, Y_1^0), \dots, (X_{q'}^0, Y_{q'}^0)\}$ from P_0 available to the aggregator and agents.

Under **H6**, define $\hat{\mathcal{R}}_0(g) = q'^{-1} \sum_{i=1}^{q'} \ell(g(X_i^0), Y_i^0)$ for $g \in \mathcal{G}$. This allows us to devise suitable estimators $\hat{\theta}_j$ as follows.

Proposition 2. Assume **H1**, **H2** and **H6**. For any $j \in [J]$ the estimator

$$\hat{\theta}_{0,j}^{\text{ERM}} = \sup_{g \in \mathcal{G}} |\hat{\mathcal{R}}_j(g) - \hat{\mathcal{R}}_0(g)|,$$

satisfies **H5** with

$$\eta_\delta(q) = \alpha_{\delta/4} [(q+1)^{-\gamma} + (q'+1)^{-\gamma}] + 2\beta. \quad (13)$$

Proposition 2 shows that the empirical version of the \mathcal{G} -divergence defined in **H2** correctly estimate types. Note that the tighter the PAC bound in **H1**, the better the approximation term in (13). The type estimator $\hat{\theta}_{0,j}^{\text{ERM}}$ defined in **H5** can easily be computed in the classification case, as shown with the following example.

H7 (Classification setting). $\mathcal{Y} = \{-1, 1\}$, $\ell = \ell_{0,1} : (y, y') \in \mathcal{Y} \times \mathcal{Y} \mapsto \mathbb{1}_{\{yy' < 0\}}$, and \mathcal{G} is a symmetric ($g \in \mathcal{G}$ if and only if $-g \in \mathcal{G}$) class of classifiers.

Example 4. Assume **H1**, **H2**, **H6** and **H7**.

(i) Denoting $\hat{\mathcal{R}}_{j-}(g) = n_j^{-1} \sum_{i=1}^{n_j} \ell_{0,1}(g(X_i^j), -Y_i^j)$, we have

$$\hat{\theta}_{0,j}^{\text{ERM}} = 1 - \inf_{g \in \mathcal{G}} \left\{ \hat{\mathcal{R}}_0(g) + \hat{\mathcal{R}}_{j-}(g) \right\}.$$

(ii) In **H1**, assume $\alpha_\delta = \ln(1/\delta)^{1/2}$, $\beta = 2\text{RAD}(\mathcal{G})$ and $\gamma = 1$ [Bousquet et al., 2003]. With $\hat{\theta}_{0,j}^{\text{ERM}}$ defined in Proposition 2, we have

$$\eta_{\delta/J}(q) = \ln(4J/\delta)^{1/2} [(1+q)^{-\gamma} + (1+q')^{-\gamma}] + 2\text{Rad}(\mathcal{G}).$$

Example 4 shows that in the classification case, it is sufficient to flip the labels of the data received from each contributor, merge these samples with those from P_0 , and perform an empirical risk minimization to compute $\hat{\theta}_{0,j}^{\text{ERM}}$. The approximation error grows no more than logarithmically with the number of agents, while decreasing at rate γ with the number of samples used in the estimation.

5 Conclusion

In this work, we show that information asymmetry has deleterious consequences when strategic agents try to learn a collaborative model. More precisely, under a naive aggregation procedure, the ignorance of others' data quality leads the coalition of learners to be either empty or reduced to the lowest-quality agent. We introduce a transfer-free mechanism based on estimation of types. This effectively counteracts unravelling by letting the grand coalition ranks among the approximate Nash equilibria with high probability.

Several possible extensions can be considered. First, it would be interesting to relax the assumption that all agents have the same ratio a/c in their utility, and see how heterogeneity affects the results.

Second, the mechanism presented in Section 4 aims for individual rationality. A more desirable, yet difficult to achieve, property would be core stability, to ensure that no group of agents would benefit from a coordinated deviation, i.e., forming an alternative coalition. Finally, it would be interesting to check whether there exist mechanisms where the optimal collaboration not only emerges as a Nash equilibrium, but as a dominant equilibrium under imperfect information.

Acknowledgements

Funded by the European Union (ERC, Ocean, 101071601). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

References

- George A. Akerlof. The market for lemons: Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, 1970. URL <https://ideas.repec.org/a/oup/qjecon/v84y1970i3p488-500..html>.
- Nivasini Ananthkrishnan, Stephen Bates, Michael Jordan, and Nika Haghtalab. Delegating data collection in decentralized machine learning, 2023.
- Ian Ball and Deniz Kattwinkel. Probabilistic verification in mechanism design. September 2019.
- Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Vaughan. A theory of learning from different domains. *Machine Learning*, 79:151–175, 05 2010. doi: 10.1007/s10994-009-5152-4.
- Avrim Blum, Nika Haghtalab, Ariel Procaccia, and Mingda Qiao. Collaborative PAC learning. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*. Curran Associates, 2017.
- Avrim Blum, Nika Haghtalab, Richard Lanus Phillips, and Han Shao. One for one, or all for all: Equilibria and optimality of collaboration in federated learning. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 1005–1014. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/blum21a.html>.
- Olivier Bousquet, Stéphane Boucheron, Gábor Lugosi, Ulrike Luxburg, and Gunnar Rätsch. Introduction to statistical learning theory. *Advanced Lectures on Machine Learning*, 169-207 (2004), 01 2003. doi: 10.1007/978-3-540-28650-9_8.
- Ioannis Caragiannis, Edith Elkind, Mario Szegedy, and Lan Yu. Mechanism design: From partial to probabilistic verification. *Proceedings of the ACM Conference on Electronic Commerce*, 06 2012. doi: 10.1145/2229012.2229035.
- Edward Clarke. Multipart pricing of public goods. *Public Choice*, 11(1):17–33, September 1971. doi: 10.1007/BF01726210. URL <https://ideas.repec.org/a/kap/pubcho/v11y1971i1p17-33.html>.
- Kate Donahue and Jon Kleinberg. Optimality and stability in federated learning: A game-theoretic approach, 06 2021a.
- Kate Donahue and Jon Kleinberg. Model-sharing games: Analyzing federated learning under voluntary participation. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35:5303–5311, 05 2021b. doi: 10.1609/aaai.v35i6.16669.
- Florian E. Dorner, Nikola Konstantinov, Georgi Pashaliev, and Martin T. Vechev. Incentivizing honesty among competitors in collaborative learning and optimization. *ArXiv*, abs/2305.16272, 2023. URL <https://api.semanticscholar.org/CorpusID:258887502>.

- Liran Einav and Amy Finkelstein. Selection in insurance markets: Theory and empirics in pictures. *Journal of Economic Perspectives*, 25(1):115–38, March 2011. doi: 10.1257/jep.25.1.115. URL <https://www.aeaweb.org/articles?id=10.1257/jep.25.1.115>.
- Diodato Ferraioli and Carmine Ventre. Probabilistic verification for obviously strategyproof mechanisms, 2018.
- Lei Fu, Huanle Zhang, Ge Gao, Mi Zhang, and Xin Liu. Client selection in federated learning: Principles, challenges, and opportunities, 2023.
- Dashan Gao, Xin Yao, and Qiang Yang. A survey on heterogeneous federated learning. *arXiv preprint arXiv:2210.04505*, 2022.
- Jerry Green and Jean-Jacques Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45(2):427–38, 1977. URL <https://EconPapers.repec.org/RePEc:ecm:emetrp:v:45:y:1977:i:2:p:427-38>.
- Theodore Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973. URL <https://api.semanticscholar.org/CorpusID:264740987>.
- Nathaniel Hendren. Private Information and Insurance Rejections. *Econometrica*, 81(5):1713–1762, September 2013. doi: ECTA10931. URL <https://ideas.repec.org/a/ecm/emetrp/v81y2013i5p1713-1762.html>.
- Bengt Holmstrom. Groves’ scheme on restricted domains. *Econometrica*, 47(5):1137–44, 1979. URL <https://EconPapers.repec.org/RePEc:ecm:emetrp:v:47:y:1979:i:5:p:1137-44>.
- Baihe Huang, Sai Praneeth Karimireddy, and Michael I. Jordan. Evaluating and Incentivizing Diverse Data Contributions in Collaborative Learning. Papers 2306.05592, arXiv.org, June 2023. URL <https://ideas.repec.org/p/arx/papers/2306.05592.html>.
- Chao Huang, Shuqi Ke, Charles Kamhoua, Prasant Mohapatra, and Xin Liu. Incentivizing data contribution in cross-silo federated learning, 2022.
- Peter Kairouz, H. Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, Rafael G. L. D’Oliveira, Hubert Eichner, Salim El Rouayheb, David Evans, Josh Gardner, Zachary Garrett, Adrià Gascón, Badih Ghazi, Phillip B. Gibbons, Marco Gruteser, Zaid Harchaoui, Chaoyang He, Lie He, Zhouyuan Huo, Ben Hutchinson, Justin Hsu, Martin Jaggi, Tara Javidi, Gauri Joshi, Mikhail Khodak, Jakub Konevcný, Aleksandra Korolova, Farinaz Koushanfar, Sanmi Koyejo, Tancrede Lepoint, Yang Liu, Prateek Mittal, Mehryar Mohri, Richard Nock, Ayfer Özgür, Rasmus Pagh, Mariana Raykova, Hang Qi, Daniel Ramage, Ramesh Raskar, Dawn Song, Weikang Song, Sebastian U. Stich, Ziteng Sun, Ananda Theertha Suresh, Florian Tramèr, Praneeth Vepakomma, Jianyu Wang, Li Xiong, Zheng Xu, Qiang Yang, Felix X. Yu, Han Yu, and Sen Zhao. Advances and open problems in federated learning, 2021.
- Sai Praneeth Karimireddy, Wenshuo Guo, and Michael I. Jordan. Mechanisms that incentivize data sharing in federated learning, 2022.
- Daniel Kifer, Shai Ben-David, and Johannes Gehrke. Detecting change in data streams. pages 180–191, 04 2004.
- Nikola Konstantinov and Christoph Lampert. Robust learning from untrusted sources, 2019.
- Jean-Jacques Laffont and David Martimort. *The Theory of Incentives: The Principal-Agent Model*. Princeton University Press, Princeton, NJ, USA, 2001.
- Shutian Liu, Tao Li, and Quanyan Zhu. Game-theoretic distributed empirical risk minimization with strategic network design. *IEEE Transactions on Signal and Information Processing over Networks*, 9:542–556, 2023. doi: 10.1109/TSIPN.2023.3306106.
- Andreu Mas-Colell, Michael D. Whinston, and Jerry R. Green. *Microeconomic Theory*. Number 9780195102680 in OUP Catalogue. Oxford University Press, 1995. ISBN ARRAY(0x516030e0). URL <https://ideas.repec.org/b/oxp/obooks/9780195102680.html>.

Michael Rothschild and Joseph Stiglitz. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of Economics*, 90(4):629–649, 1976. URL <https://EconPapers.repec.org/RePEc:oup:qjecon:v:90:y:1976:i:4:p:629-649>.

Eden Saig, Inbal Talgam-Cohen, and Nir Rosenfeld. Delegated classification, 2023.

Vera Shalaeva, Alireza Fakhrizadeh Esfahani, Pascal Germain, and Mihály Petreczky. Improved pac-bayesian bounds for linear regression. In *AAAI Conference on Artificial Intelligence*, 2019. URL <https://api.semanticscholar.org/CorpusID:208857400>.

Nikita Tsoy and Nikola Konstantinov. Strategic data sharing between competitors. *Advances in Neural Information Processing Systems*, 36, 2024.

Xuezhen Tu, Kun Zhu, Nguyen Cong Luong, Dusit Niyato, Yang Zhang, and Juan Li. Incentive mechanisms for federated learning: From economic and game theoretic perspective. *IEEE Transactions on Cognitive Communications and Networking*, 8(3):1566–1593, 2022. doi: 10.1109/TCCN.2022.3177522.

William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16(1):8–37, 1961. URL <https://EconPapers.repec.org/RePEc:bla:jfinan:v:16:y:1961:i:1:p:8-37>.

Zhepei Wei, Chuanhao Li, Tianze Ren, Haifeng Xu, and Hongning Wang. Incentivized truthful communication for federated bandits, 2024.

Mariel Werner, Sai Praneeth Karimireddy, and Michael I. Jordan. Defection-free collaboration between competitors in a learning system, 2024. URL <https://arxiv.org/abs/2406.15898>.

Yizhou Yan, Xinyu Tang, Chao Huang, and Ming Tang. Price of stability in quality-aware federated learning. In *GLOBECOM 2023 - 2023 IEEE Global Communications Conference*, pages 734–739, 2023. doi: 10.1109/GLOBECOM54140.2023.10437743.

A Relaxation of the optimal aggregation problem.

Let $\mathbf{B} \in \{0, 1\}^J$ and the associated coalition $\mathcal{B} = \{j \in [J] : B_j = 1\}$ be fixed. We present in this appendix a natural relaxation of problem Equation (10) which motivates the choice $\sum_{j \in \mathcal{B}} n_j^*(\boldsymbol{\theta}) = \bar{N} = (\underline{n} + 1)|\mathcal{B}|^{\frac{1}{1+\gamma}} - 1$ in the simplified contribution scheme $(\mathbf{B}^*, \mathbf{n}^*)$. The optimal aggregation problem with \mathbf{B} fixed is

$$\text{maximize } \mathbf{n} \in \mathbb{R}^J \mapsto \sum_{j \in [J]} u_j(\mathbf{B}, \mathbf{n}) \quad \text{subject to } \begin{cases} \max_{j \in \mathcal{B}} n_j - \bar{n}(\mathbf{B}, \mathbf{n}) \leq 0 \\ \max_{j \in [J]} -n_j \leq 0, \end{cases}$$

Since agents outside of the coalition maximize their utility, the problem can equivalently be stated as

$$(\mathcal{T}_{\mathbf{B}}) : \quad \text{maximize } \mathbf{n} \in \mathbb{R}^{|\mathcal{B}|} \mapsto \sum_{j \in \mathcal{B}} u_j(\mathbf{B}, \mathbf{n}) + \zeta_{\mathbf{B}} \quad \text{subject to } \mathbf{n} \in \Xi_{\mathbf{B}},$$

where

$$\Xi_{\mathbf{B}} = \left\{ \mathbf{n} \in \mathbb{R}_+^{|\mathcal{B}|} : \begin{array}{l} \max_{j \in \mathcal{B}} \{n_j - \bar{n}_j(\mathbf{B}, \mathbf{n})\} \leq 0 \\ \max_{j \in \mathcal{B}} -n_j \leq 0 \end{array} \right\} \quad \text{and} \quad \zeta_{\mathbf{B}} = \sum_{j \notin \mathcal{B}} u_j. \quad (14)$$

We start by rewriting $(\mathcal{T}_{\mathbf{B}})$ in a more convenient way. Instead of working with $\mathbf{n} \in \mathbb{R}_+^{|\mathcal{B}|}$ directly, we make appear (i) the total number of samples and (ii) the sharing out of samples between agents. Formally, for any $\mathbf{n} \in \mathbb{R}_+^{|\mathcal{B}|}$ consider $N = \mathbf{n}^T \mathbf{B}$ and $\boldsymbol{\lambda} = N^{-1} \mathbf{n} \in \Delta_{|\mathcal{B}|}$, where $\Delta_{|\mathcal{B}|}$ is the simplex of dimension $|\mathcal{B}|$. Observe that the average type in the coalition reads $\vartheta(\mathbf{B}, \mathbf{n}) = \boldsymbol{\lambda}^T \boldsymbol{\theta}$. Moreover,

the welfare evaluated in (\mathbf{B}, \mathbf{n}) rewrites

$$\begin{aligned}
\sum_{j \in [J]} u_j(\mathbf{B}, \mathbf{n}) &= \sum_{j \in [J]} B_j [-a(\mathcal{R}_0^* + \varepsilon(\mathbf{B}, \boldsymbol{\lambda}N)) - cn_j] + \sum_{j \in [J]} (1 - B_j) \underline{u}_j \\
&= -a|\mathcal{B}|(\mathcal{R}_0^* + \varepsilon(\mathbf{B}, \boldsymbol{\lambda}N)) - cN + \zeta_{\mathbf{B}} , \\
&= -a|\mathcal{B}|[\mathcal{R}_0^* + 2(\alpha_\delta(1 + N)^{-\gamma} + \beta + \boldsymbol{\lambda}^T \boldsymbol{\theta})] - cN + \zeta_{\mathbf{B}} \\
&= \widetilde{W}_{\mathbf{B}}(\boldsymbol{\lambda}, N) .
\end{aligned}$$

We denote by $\widetilde{\Xi}_{\mathbf{B}} = \{(\boldsymbol{\lambda}, N) \in \Delta_{|\mathcal{B}|} \times \mathbb{R}_+ : \boldsymbol{\lambda}N \in \Xi_{\mathbf{B}}\}$ and $\widetilde{\mathcal{T}}_{\mathbf{B}}$ the problem:

$$(\widetilde{\mathcal{T}}_{\mathbf{B}}) : \text{minimize } (\boldsymbol{\lambda}, N) \in \Delta_{|\mathcal{B}|} \times \mathbb{R}_+ \mapsto -\widetilde{W}_{\mathbf{B}}(\boldsymbol{\lambda}, N) \text{ subject to } (\boldsymbol{\lambda}, N) \in \widetilde{\Xi}_{\mathbf{B}} . \quad (15)$$

By definition, if $(\boldsymbol{\lambda}, N)$ is a solution to $\widetilde{\mathcal{T}}_{\mathbf{B}}$, then $\mathbf{n} = \boldsymbol{\lambda}N$ is a solution to $\mathcal{T}_{\mathbf{B}}$. We can further decompose $\widetilde{\mathcal{T}}_{\mathbf{B}}$ by observing that

$$-\widetilde{W}_{\mathbf{B}}(\boldsymbol{\lambda}, N) = f(N) + g(\boldsymbol{\lambda}) ,$$

with

$$f(N) = a|\mathcal{B}|(\mathcal{R}_0^* + 2\alpha_\delta(1 + N)^{-\gamma} + 2\beta) + cN + \zeta_{\mathbf{B}} \quad \text{and} \quad g(\boldsymbol{\lambda}) = 2a|\mathcal{B}|\boldsymbol{\lambda}^T \boldsymbol{\theta} .$$

Finally, we denote a slice of $\widetilde{\Xi}_{\mathbf{B}}$ along $N \geq 0$ as $\widetilde{\Xi}_{\mathbf{B}}(N) = \{\boldsymbol{\lambda} \in \Delta_{|\mathcal{B}|} : (\boldsymbol{\lambda}, N) \in \widetilde{\Xi}_{\mathbf{B}}\}$ and $\mathcal{N} = \{N \geq 0 : \widetilde{\Xi}_{\mathbf{B}}(N) \neq \emptyset\}$. $\widetilde{\mathcal{T}}_{\mathbf{B}}$ comes down to

$$\min_{(\boldsymbol{\lambda}, N) \in \widetilde{\Xi}_{\mathbf{B}}} \{f(N) + g(\boldsymbol{\lambda})\} = \min_{N \in \mathcal{N}} \{f(N) + \min_{\boldsymbol{\lambda} \in \widetilde{\Xi}_{\mathbf{B}}(N)} g(\boldsymbol{\lambda})\} . \quad (16)$$

A strategy to solve Equation (16) is to (i) address the innermost problem $\min_{\boldsymbol{\lambda} \in \widetilde{\Xi}_{\mathbf{B}}(N)} g(\boldsymbol{\lambda})$ with $N \in \mathcal{N}$ fixed, and denoting $\boldsymbol{\lambda}^{(N)} \in \widetilde{\Xi}_{\mathbf{B}}(N)$ its solution, (ii) solve the outermost problem:

$$\min_{N \in \mathcal{N}} \{f(N) + g(\boldsymbol{\lambda}^{(N)})\} . \quad (17)$$

Point (i) is done in the proof of Theorem 1. However, the resulting problem in (ii) is hard to tackle because $\boldsymbol{\lambda}^{(N)}$ has no simple form. Therefore, we leave aside the term $g(\boldsymbol{\lambda}^{(N)})$ (which can be easily controlled, as shown in Lemma 2) to determine N and only consider the problem

$$\min_{N \geq 0} f(N) . \quad (18)$$

Since f is differentiable and strictly convex, its minimizer is uniquely defined by $f'(\bar{N}) = 0$, which gives

$$\bar{N} = (\underline{n} + 1)|\mathcal{B}|^{\frac{1}{1+\gamma}} - 1 , \quad (19)$$

where \underline{n} is defined in Proposition 1. As the solution of the relaxed problem (18), we take $\sum_{j \in [J]} n_j^*(\boldsymbol{\theta}) = \bar{N}$ is the simplified contribution scheme $(\mathbf{B}^*, \mathbf{n}^*)$. In many reasonable settings, this approximation is very satisfactory as shown by Lemma 2.

B Proofs

Lemma 1. Assume **H1** and **H2**.

(i) Any agent $j \in [J]$ picking the outside **Option 1** obtains a model $\widehat{g}_j \in \mathcal{G}$ achieving

$$\mathcal{R}_0(\widehat{g}_j) \leq \mathcal{R}_0^* + \varepsilon(\theta_j, n_j) \quad \text{with probability } 1 - \delta .$$

(ii) Any coalition $\mathbf{B} \in \{0, 1\}^J$ drawing $\mathbf{n} = (n_1, \dots, n_J) \in \mathbb{R}_+^J$, samples obtains a model $\widehat{g}_{\mathbf{B}} \in \mathcal{G}$ achieving

$$\mathcal{R}_0(\widehat{g}_{\mathbf{B}}) \leq \mathcal{R}_0^* + \varepsilon(\vartheta, N) \quad \text{with probability } 1 - \delta ,$$

where N is the total of samples and ϑ is the weighted average type within the coalition:

$$N(\mathbf{B}, \mathbf{n}) = \sum_{j \in [J]} B_j n_j , \quad \vartheta(\mathbf{B}, \mathbf{n}) = N^{-1} \sum_{j \in [J]} B_j n_j \theta_j . \quad (4)$$

Proof. (i) Let $\{(X_1^j, Y_1^j), \dots, (X_{n_j}^j, Y_{n_j}^j)\}$ be $n_j > 0$ i.i.d samples from $P_j \in \mathcal{P}$ and $\hat{g}_j = \inf_{g \in \mathcal{G}} \hat{\mathcal{R}}_j(g)$. First, observe that

$$\mathcal{R}_0(\hat{g}_j) \leq \mathcal{R}_j(\hat{g}_j) + \sup_{g \in \mathcal{G}} |\mathcal{R}_0(g) - \mathcal{R}_j(g)| = \mathcal{R}_j(\hat{g}_j) + \theta_j. \quad (20)$$

Let $v > 0$ and $g_j^v \in \mathcal{G}$ such that $\mathcal{R}_j(g_j^v) \leq \inf_{g \in \mathcal{G}} \mathcal{R}_j(g) + v$. Then with probability at least $1 - \delta$,

$$\begin{aligned} \mathcal{R}_j(\hat{g}_j) &= \mathcal{R}_j(\hat{g}_j) - \mathcal{R}_j(g_j^v) + \mathcal{R}_j(g_j^v) \\ &\leq \left(\hat{\mathcal{R}}_j(g_j^v) - \hat{\mathcal{R}}_j(\hat{g}_j) \right) + \mathcal{R}_j(\hat{g}_j) - \mathcal{R}_j(g_j^v) + \mathcal{R}_j(g_j^v) \\ &\leq 2 \sup_{g \in \mathcal{G}} \left| \mathcal{R}_j(g) - \hat{\mathcal{R}}_j(g) \right| + \inf_{g \in \mathcal{G}} \mathcal{R}_j(g) + v \\ &\leq 2 \left(\frac{\alpha \delta}{(n+1)^\gamma} + \beta \right) + \inf_{g \in \mathcal{G}} \mathcal{R}_j(g), \end{aligned} \quad (21)$$

where the last inequality holds taking the limit $v \rightarrow 0$ and using **H 1**. Finally, observing that $\inf_{g \in \mathcal{G}} \mathcal{R}_j(g) - \mathcal{R}_0^* = \inf_{g \in \mathcal{G}} \mathcal{R}_j(g) - \inf_{g \in \mathcal{G}} \mathcal{R}_0(g) \leq \sup_{g \in \mathcal{G}} |\mathcal{R}_j(g) - \mathcal{R}_0(g)| = \theta_j$, and combining (20) as well as (21) yields

$$\mathcal{R}_0(\hat{g}_j) \leq \mathcal{R}_0^* + 2\theta_j + 2 \left(\frac{\alpha \delta}{(n+1)^\gamma} + \beta \right) = \mathcal{R}_0^* + \varepsilon(\theta_j, n), \quad (22)$$

with probability at least $1 - \delta$.

(ii) Let $\mathbf{B} \in \{0, 1\}^J$ and $\mathbf{n} = (n_1, \dots, n_J) \in \mathbb{R}_+^J$. For any $j \in \mathbf{B}$, let $\{(X_1^j, Y_1^j), \dots, (X_{n_j}^j, Y_{n_j}^j)\}$ be a collection of $n_j > 0$ i.i.d samples from $P_j \in \mathcal{P}$. Denote by $N = \sum_{j \in [J]} B_j n_j$ and consider $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_J)$ such that $\lambda_j = B_j n_j / N$. Note that $\boldsymbol{\lambda}$ belongs to the J -dimensional simplex. For any $g \in \mathcal{G}$, the empirical risk over contributions is

$$\hat{\mathcal{R}}_{\mathbf{B}}(g) = \frac{1}{N} \sum_{j \in [J]} B_j \sum_{i=1}^{n_j} \ell(g(X_i^j), Y_i^j) = \sum_{j \in [J]} B_j \lambda_j \hat{\mathcal{R}}_j(g), \quad (23)$$

and the population risk is

$$\mathcal{R}_{\mathbf{B}}(g) = \mathbb{E} \left[\hat{\mathcal{R}}_{\mathbf{B}}(g) \right] = \sum_{j \in [J]} B_j \lambda_j \mathbb{E} \left[\hat{\mathcal{R}}_j(g) \right] = \sum_{j \in [J]} B_j \lambda_j \mathcal{R}_j(g). \quad (24)$$

On the one hand, the collaborative model $\hat{g}_{\mathbf{B}} \in \mathcal{G}$ satisfies:

$$\begin{aligned} \mathcal{R}_0(\hat{g}_{\mathbf{B}}) &\leq \mathcal{R}_{\mathbf{B}}(\hat{g}_{\mathbf{B}}) + \sup_{g \in \mathcal{G}} |\mathcal{R}_0(g) - \mathcal{R}_{\mathbf{B}}(g)| \leq \mathcal{R}_{\mathbf{B}}(\hat{g}_{\mathbf{B}}) + \sup_{g \in \mathcal{G}} \left\{ \sum_{j \in [J]} \lambda_j |\mathcal{R}_j(g) - \mathcal{R}_0(g)| \right\} \\ &\leq \mathcal{R}_{\mathbf{B}}(\hat{g}_{\mathbf{B}}) + \sum_{j \in [J]} \lambda_j \sup_{g \in \mathcal{G}} |\mathcal{R}_j(g) - \mathcal{R}_0(g)| \\ &= \mathcal{R}_{\mathbf{B}}(\hat{g}_{\mathbf{B}}) + \vartheta(\mathbf{B}, \mathbf{n}), \end{aligned} \quad (25)$$

with $\vartheta(\mathbf{B}, \mathbf{n}) = N^{-1} \sum_{j \in [J]} B_j n_j \theta_j = \sum_{j \in [J]} B_j \lambda_j \theta_j$. Now, let $v > 0$ and $g^v \in \mathcal{G}$ such that $\mathcal{R}_{\mathbf{B}}(g^v) \leq \inf_{g \in \mathcal{G}} \mathcal{R}_{\mathbf{B}}(g) + v$. We also have

$$\begin{aligned} \mathcal{R}_{\mathbf{B}}(\hat{g}_{\mathbf{B}}) &= \mathcal{R}_{\mathbf{B}}(\hat{g}_{\mathbf{B}}) - \mathcal{R}_{\mathbf{B}}(g^v) + \mathcal{R}_{\mathbf{B}}(g^v) \\ &\leq \left(\hat{\mathcal{R}}_{\mathbf{B}}(g^v) - \hat{\mathcal{R}}_{\mathbf{B}}(\hat{g}_{\mathbf{B}}) \right) + \mathcal{R}_{\mathbf{B}}(\hat{g}_{\mathbf{B}}) - \mathcal{R}_{\mathbf{B}}(g^v) + \mathcal{R}_{\mathbf{B}}(g^v) \\ &\leq 2 \sup_{g \in \mathcal{G}} \left| \mathcal{R}_{\mathbf{B}}(g) - \hat{\mathcal{R}}_{\mathbf{B}}(g) \right| + \inf_{g \in \mathcal{G}} \mathcal{R}_{\mathbf{B}}(g) + v \\ &\leq 2 \left(\frac{\alpha \delta}{(N+1)^\gamma} + \beta \right) + \inf_{g \in \mathcal{G}} \mathcal{R}_{\mathbf{B}}(g) \quad \text{with probability at least } 1 - \delta, \end{aligned} \quad (26)$$

where the first inequality comes from $\widehat{\mathcal{R}}_{\mathbf{B}}(\widehat{g}_{\mathbf{B}}) = \inf_{g \in \mathcal{G}} \widehat{\mathcal{R}}_{\mathbf{B}}(g) \leq \widehat{\mathcal{R}}_{\mathbf{B}}(g^v)$, and the last is obtained by taking $v \rightarrow 0$ and applying assumption **H 1**. Now, observe that $\inf_{g \in \mathcal{G}} \mathcal{R}_{\mathbf{B}}(g) - \mathcal{R}_0^* = \inf_{g \in \mathcal{G}} \mathcal{R}_{\mathbf{B}}(g) - \inf_{g \in \mathcal{G}} \mathcal{R}_0(g) \leq \sup_{g \in \mathcal{G}} |\mathcal{R}_{\mathbf{B}}(g) - \mathcal{R}_0(g)| \leq \sum_j B_j \lambda_j \sup_{g \in \mathcal{G}} |\mathcal{R}_j(g) - \mathcal{R}_0(g)| = \vartheta(\mathbf{B}, \mathbf{n})$. Combining this observation with Equation (25) and Equation (26) gives with probability $1 - \delta$

$$\mathcal{R}_0(\widehat{g}_{\mathbf{B}}) \leq \mathcal{R}_0^* + 2 \left(\frac{\alpha_\delta}{(N+1)^\gamma} + \beta \right) + 2\vartheta(\mathbf{B}, \mathbf{n}) = \mathcal{R}_0^* + \varepsilon(\vartheta(\mathbf{B}, \mathbf{n}), N). \quad (27)$$

□

Proposition 1. Assume **H1** and **H2**. For any $j \in [J]$, $\mathbf{B}_{-j} \in \{0, 1\}^{J-1}$ and $\mathbf{n}_{-j} \in \mathbb{R}_+^{J-1}$, the optimal number of samples to draw under **Option 1** is $\operatorname{argmax}_{n \geq 0} u_j((0, \mathbf{B}_{-j}), (n, \mathbf{n}_{-j}); \theta_j) = \underline{n}$ where

$$\underline{n} = (2ac^{-1}\gamma\alpha_\delta)^{1/(\gamma+1)} - 1. \quad (6)$$

Proof. For any $j \in [J]$ and $n \geq 0$, Lemma 1 gives that

$$-u_j((0, \mathbf{B}_{-j}), (n, \mathbf{n}_{-j}), \theta_j) = a[\mathcal{R}_0^* + 2(\alpha_\delta(n+1)^{-\gamma} + \beta + \theta_j)] + cn = f(n).$$

Since f is strictly convex and differentiable, it admits a unique maximizer $\underline{n} \geq 0$ determined by $f'(\underline{n}) = 0$. Simple algebra leads to $\underline{n} = (2a\gamma c^{-1}\alpha_\delta)^{1/(\gamma+1)} - 1$. □

Theorem 1. Assume **H1**, **H2**, **H3**. Problem (8) admits a unique solution $(\mathbf{B}^{\text{opt}}, \mathbf{n}^{\text{opt}}(\boldsymbol{\theta})) \in \{0, 1\}^J \times \mathbb{R}_+^J$. Moreover,

(i) $\mathbf{B}^{\text{opt}} = \mathbf{1} = (1, \dots, 1)$,

(ii) Denoting $\mathbf{n}^{\text{opt}}(\boldsymbol{\theta}) = (n_1^{\text{opt}}(\boldsymbol{\theta}), \dots, n_J^{\text{opt}}(\boldsymbol{\theta}))$, there exists $L^{\text{opt}} \in [J]$ such that for any $j \in [J]$,

$$n_j^{\text{opt}}(\boldsymbol{\theta}) \begin{cases} = \bar{n}_j(\mathbf{1}, \mathbf{n}^{\text{opt}}(\boldsymbol{\theta})) & \text{if } j < L^{\text{opt}}, \\ \in [0, \bar{n}_j(\mathbf{1}, \mathbf{n}^{\text{opt}}(\boldsymbol{\theta}))] & \text{if } j = L^{\text{opt}}, \\ = 0 & \text{otherwise.} \end{cases} \quad (11)$$

Proof. Recall that the optimal aggregation problem reads

$$\begin{aligned} & \text{maximize} && W(\mathbf{B}, \mathbf{n}) \in \{0, 1\}^J \times \mathbb{R}_+^J \mapsto \sum_{j \in [J]} u_j(\mathbf{B}, \mathbf{n}) \\ & \text{subject to} && \min_{j \in [J]} u_j(\mathbf{B}, \mathbf{n}) - \underline{u}_j \geq 0. \end{aligned} \quad (28)$$

Define for any $j \in [J]$, $\mathbf{B}_{-j} \in \{0, 1\}^{J-1}$, and $\mathbf{n}_{-j} \in \mathbb{R}_+^{J-1}$, the maximum number of samples agent $j \in [J]$ having $B_j = 1$ may be asked given their participation constraint:

$$\bar{n}_j(\mathbf{B}, \mathbf{n}) = \max\{n \geq 0 : u_j((1, \mathbf{B}_{-j}), (n, \mathbf{n}_{-j})) \geq \underline{u}_j\}.$$

Given Equation (5) and Proposition 1, we obtain

$$\bar{n}_j(\mathbf{B}, \mathbf{n}) = \underline{n} - ac^{-1}(\varepsilon(\mathbf{B}, \mathbf{n}) - \varepsilon(\theta_j, \underline{n})), \quad (29)$$

where $N = \sum_{j \in [J]} B_j n_j = \mathbf{n}^T \mathbf{B}$. Problem (28) rewrites in canonical form

$$\text{minimize} \quad -W(\mathbf{B}, \mathbf{n}) \quad \text{subject to} \quad \begin{cases} \max_{j \in \mathcal{B}} n_j - \bar{n}_j(\mathbf{B}, \mathbf{n}) \leq 0 \\ \max_{j \in [J]} -n_j \leq 0, \end{cases} \quad (30)$$

where $\mathcal{B} = \{j \in [J] : B_j = 1\}$. We first show that it admits a unique minimum.

Fix a $\mathbf{B} \in \{0, 1\}^J$ and consider problem (30) with respect to $\mathbf{n} \in \mathbb{R}_+^J$ only. We call this subproblem $\mathcal{T}_{\mathbf{B}}$. First, we show that it is enough to focus on $(n_j)_{j \in \mathcal{B}}$ rather than $(n_j)_{j \in [J]}$ in $\mathcal{T}_{\mathbf{B}}$. With $\mathbf{n}_{\mathcal{B}} = (n_j)_{j \in \mathcal{B}}$ and $\mathbf{n}_{\mathcal{B}^c} = (n_j)_{j \notin \mathcal{B}}$, note that $W(\mathbf{B}, \mathbf{n})$ is decomposable in $\mathbf{n}_{\mathcal{B}}$ and $\mathbf{n}_{\mathcal{B}^c}$:

$$W(\mathbf{B}, \mathbf{n}) = W(\mathbf{B}, \mathbf{n}_{\mathcal{B}}) + W(\mathbf{B}, \mathbf{n}_{\mathcal{B}^c}),$$

where, by a slight abuse of notation $W(\mathbf{B}, \mathbf{n}_{\mathcal{B}}) = \sum_{j \in \mathcal{B}} u_j(\mathbf{B}, \mathbf{n})$ and $W(\mathbf{B}, \mathbf{n}_{\mathcal{B}^c}) = \sum_{j \in \mathcal{B}^c} u_j(\mathbf{B}, \mathbf{n})$. With

$$\Xi_{\mathcal{B}} = \left\{ \mathbf{n}_{\mathcal{B}} \in \mathbb{R}_+^{\mathcal{B}} : \begin{array}{l} \max_{j \in \mathcal{B}} \{n_j - \bar{n}_j(\mathbf{B}, \mathbf{n})\} \leq 0 \\ \max_{j \in \mathcal{B}} -n_j \leq 0 \end{array} \right\}, \quad (31)$$

problem $\mathcal{T}_{\mathcal{B}}$ is equivalent to

$$\min_{\mathbf{n}_{\mathcal{B}} \in \Xi_{\mathcal{B}}} -W(\mathbf{B}, \mathbf{n}_{\mathcal{B}}) + \min_{\mathbf{n}_{\mathcal{B}^c} \in \mathbb{R}_+^{J-|\mathcal{B}|}} -W(\mathbf{B}, \mathbf{n}_{\mathcal{B}^c}) = \min_{\mathbf{n}_{\mathcal{B}} \in \Xi_{\mathcal{B}}} -W(\mathbf{B}, \mathbf{n}_{\mathcal{B}}) + \sum_{j \in [J]} -(1 - B_j) \underline{u}_j, \quad (32)$$

by Proposition 1. Since $\sum_{j \in [J]} -(1 - B_j) \underline{u}_j$ is constant, by Equation (32) it is enough to focus on the existence of $\min_{\mathbf{n}_{\mathcal{B}} \in \Xi_{\mathcal{B}}} -W(\mathbf{B}, \mathbf{n}_{\mathcal{B}})$. On the one hand, $\Xi_{\mathcal{B}}$ is bounded. Indeed for any $j \in \mathcal{B}$, by Equation (29) and Lemma 1

$$\begin{aligned} \bar{n}_j(\mathbf{B}, \mathbf{n}_{\mathcal{B}}) &= \underline{n} - \frac{a}{c} [\varepsilon(\mathbf{B}, \mathbf{n}_{\mathcal{B}}) - \varepsilon(\theta_j, \underline{n})] \\ &= \underline{n} - \frac{a}{c} \left[\mathcal{R}_0^* + 2 \left[\frac{\alpha_{\delta}}{(1 + \mathbf{n}_{\mathcal{B}}^{\top} \mathbf{B})^{\gamma}} + \beta + \vartheta(\mathbf{B}, \mathbf{n}) \right] - \mathcal{R}_0^* - 2 \left[\frac{\alpha_{\delta}}{(1 + \underline{n})^{\gamma}} - \beta - \theta_j \right] \right] \\ &= \underline{n} - \frac{2a}{c} \left[\alpha_{\delta} \left((1 + \mathbf{n}_{\mathcal{B}}^{\top} \mathbf{B})^{-\gamma} - (1 + \underline{n})^{-\gamma} \right) + (\vartheta(\mathbf{B}, \mathbf{n}) - \theta_j) \right] \\ &\leq \underline{n} - \frac{2a}{c} \left[\alpha_{\delta} (1 - (1 + \underline{n})^{-\gamma}) + \max_{j \in \mathcal{B}} \theta_j \right] \\ &= M \end{aligned} \quad (33)$$

Thus, $\Xi_{\mathcal{B}} \subset [0, M]^{|\mathcal{B}|}$. Moreover, $\Xi_{\mathcal{B}}$ is closed and convex. For any $j \in \mathcal{B}$, rewrite

$$\begin{aligned} \bar{n}_j(\mathbf{B}, \mathbf{n}_{\mathcal{B}}) &= \underline{n} - \frac{2a}{c} \left[\alpha_{\delta} \left((1 + \mathbf{n}_{\mathcal{B}}^{\top} \mathbf{B})^{-\gamma} - (1 + \underline{n})^{-\gamma} \right) + \left((\mathbf{n}_{\mathcal{B}}^{\top} \mathbf{B})^{-1} \mathbf{n}_{\mathcal{B}}^{\top} \boldsymbol{\theta} - \theta_j \right) \right] \\ &= g_j(\mathbf{n}_{\mathcal{B}}), \end{aligned}$$

and define $h_j : \mathbf{n}_{\mathcal{B}} \mapsto n_{\mathcal{B},j} - g_j(\mathbf{n}_{\mathcal{B}})$. Observe that $\Xi_{\mathcal{B}} = \{\mathbf{n}_{\mathcal{B}} \in \mathbb{R}_+^{|\mathcal{B}|} : \forall j \in \mathcal{B}, n_{\mathcal{B},j} \geq 0 \text{ and } h_j(\mathbf{n}_{\mathcal{B}}) \leq 0\}$, that is $\Xi_{\mathcal{B}} = A \cap B$ with $A = \mathbb{R}_+^{|\mathcal{B}|}$ and $B = h_1^{-1}((-\infty, 0]) \cap \dots \cap h_{|\mathcal{B}|}^{-1}((-\infty, 0])$. For any $j \in \mathcal{B}$, $h_j^{-1}((-\infty, 0])$ is convex, because so is h_j . Additionally, $h_j^{-1}((-\infty, 0])$ is closed as the inverse image of a closed set by a continuous function. It follows that $\Xi_{\mathcal{B}}$ is convex and closed as the intersection of convex and closed sets.

As a consequence $\Xi_{\mathcal{B}}$ is compact and convex, and $-W(\mathbf{B}, \cdot)$ is strictly convex so $\mathcal{T}_{\mathcal{B}}$ admits a unique minimizer $\mathbf{n}_{\mathcal{B}}^{\text{opt}} \in \Xi_{\mathcal{B}}$.

Since there are finitely many $\mathbf{B} \in \{0, 1\}^J$, $\min\{-W(\mathbf{B}, \mathbf{n}_{\mathcal{B}}^{\text{opt}}) + \sum_{j \in \mathcal{B}^c} \underline{u}_j, \mathbf{B} \in \{0, 1\}^J\} > 0$ exists and coincide with the minimum of problem (30). We show later in the proof that the optimal $\mathbf{B}^{\text{opt}} \in \{0, 1\}^J$ is unique (see *part I*), so the minimizer of Equation (30) is unique. This establishes the point (i) of the result.

The remainder of the proof proceeds as follows: we first characterize the optimal $\mathbf{B}^{\text{opt}} \in \{0, 1\}^J$, and then prove that the solution $\mathbf{n}_{\mathcal{B}^{\text{opt}}}^{\text{opt}} \in \mathbb{R}_+^{|\mathcal{B}|}$ of $\mathcal{T}_{\mathbf{B}^{\text{opt}}}$ has the form presented in point (ii) of the result.

Part I: Optimal $\mathbf{B} \in \{0, 1\}^J$

We show that $\mathbf{B}^{\text{opt}} = (1, \dots, 1)^{\top}$. By contradiction, assume there exists $\mathbf{B}' \neq (1, \dots, 1)^{\top}$ and $\mathbf{n}' \in \mathbb{R}_+^J$ such that for any $(\mathbf{B}, \mathbf{n}) \in \{0, 1\}^J \times \mathbb{R}_+^J$

$$-W(\mathbf{B}', \mathbf{n}') \leq -W(\mathbf{B}, \mathbf{n}). \quad (34)$$

Denote by $j \in [J]$ an agent such that $B'_j = 0$, and first assume $\theta_j \geq \vartheta(\mathbf{B}', \mathbf{n}')$. Consider the alternative allocation $(\mathbf{B}'', \mathbf{n}'') = ((1, \mathbf{B}'_{-j}), (0, \mathbf{n}'_{-j}))$. Observe that $B'_j n'_j = B''_j n''_j = 0$, so

$$N' = \mathbf{B}'^{\top} \mathbf{n}' = \mathbf{B}''^{\top} \mathbf{n}'' = N'' \quad \text{and} \quad \vartheta(\mathbf{B}', \mathbf{n}') = \vartheta(\mathbf{B}'', \mathbf{n}''). \quad (35)$$

This in particular implies

$$\begin{aligned}\varepsilon(\mathbf{B}', \mathbf{n}') &= 2[\alpha_\delta(1 + N')^{-\gamma} + \beta + \vartheta(\mathbf{B}', N')] \\ &= 2[\alpha_\delta(1 + N'')^{-\gamma} + \beta + \vartheta(\mathbf{B}'', N'')] = \varepsilon(\mathbf{B}'', \mathbf{n}'').\end{aligned}\quad (36)$$

Thus by (36)

$$\begin{aligned}-u_j(\mathbf{B}', \mathbf{n}') &= \underline{u}_j = a(\mathcal{R}_0^* + \varepsilon(\theta_j, \underline{n})) + c\underline{n} > a(\mathcal{R}_0^* + \varepsilon(\theta_j, \underline{n})) \\ &\geq a(\mathcal{R}_0^* + \varepsilon(\mathbf{B}', \mathbf{n}')) = a(\mathcal{R}_0^* + \varepsilon(\mathbf{B}'', \mathbf{n}'')) = -u_j(\mathbf{B}'', \mathbf{n}''),\end{aligned}\quad (37)$$

where the second inequality results from $\theta_j \geq \vartheta(\mathbf{B}', \mathbf{n}')$ and $N' \geq \underline{n}$. Finally by definition of $(\mathbf{B}'', \mathbf{n}'')$, (35) and (36), for any $k \neq j$ such that $B_k'' = 1$:

$$\begin{aligned}-u_k(\mathbf{B}'', \mathbf{n}'') &= a[\mathcal{R}_0^* + 2(\alpha_\delta(1 + N'')^{-\gamma} + \beta + \vartheta(\mathbf{B}'', \mathbf{n}''))] + cn_k'' \\ &= a[\mathcal{R}_0^* + 2(\alpha_\delta(1 + N')^{-\gamma} + \beta + \vartheta(\mathbf{B}', \mathbf{n}'))] + cn_k' = -u_k(\mathbf{B}', \mathbf{n}').\end{aligned}\quad (38)$$

Hence, (38) together with (37) give

$$\begin{aligned}-W(\mathbf{B}'', \mathbf{n}'') &= -\sum_{k \neq j} u_k(\mathbf{B}'', \mathbf{n}'') - u_j(\mathbf{B}'', \mathbf{n}'') = -\sum_{k \neq j} u_k(\mathbf{B}', \mathbf{n}') - u_j(\mathbf{B}'', \mathbf{n}'') \\ &< -\sum_{k \neq j} u_k(\mathbf{B}', \mathbf{n}') - u_j(\mathbf{B}', \mathbf{n}') = -W(\mathbf{B}', \mathbf{n}').\end{aligned}\quad (39)$$

This contradicts (34). Now assume $\theta_j < \vartheta(\mathbf{B}', \mathbf{n}')$, and let $R \in [J]$ be such that $\theta_R = \max\{\theta_k, k \in \mathcal{B}\}$. Consider $(\mathbf{B}'', \mathbf{n}'') \in \{0, 1\}^J \times \mathbb{R}_+^J$ where

$$\begin{aligned}\mathbf{B}'' &= (B_1', \dots, B_{j-1}', 1, B_{j+1}', \dots, B_J')^T \\ \text{and } \mathbf{n}'' &= (n_1', \dots, n_{j-1}', \bar{n}_j, n_{j+1}', \dots, n_{R-1}', \max(0, n_R' - \bar{n}_j), n_{R+1}', \dots, n_J')^T.\end{aligned}$$

which is feasible. Observe on the one hand that

$$N'' = N' \quad \text{and} \quad \vartheta(\mathbf{B}'', \mathbf{n}'') \leq \vartheta(\mathbf{B}', \mathbf{n}') + \frac{\bar{n}_j}{N'}(\theta_j - \theta_R) < \vartheta(\mathbf{B}', \mathbf{n}'),$$

because $\theta_j < \vartheta(\mathbf{B}', \mathbf{n}') \leq \theta_R$. In particular

$$\begin{aligned}\varepsilon(\mathbf{B}'', \mathbf{n}'') &= \mathcal{R}_0^* + 2[\alpha_\delta(N'' + 1)^{-\gamma} + \beta + \vartheta(\mathbf{B}'', \mathbf{n}'')] \\ &< \mathcal{R}_0^* + 2[\alpha_\delta(N' + 1)^{-\gamma} + \beta + \vartheta(\mathbf{B}', \mathbf{n}')] = \varepsilon(\mathbf{B}', \mathbf{n}').\end{aligned}\quad (40)$$

(40) in turn implies that for any $k \neq j$ such that $B_k'' = 1$:

$$-u_k(\mathbf{B}'', \mathbf{n}'') = a[\mathcal{R}_0^* + \varepsilon(\mathbf{B}'', \mathbf{n}'')] + cn_k'' < a[\mathcal{R}_0^* + \varepsilon(\mathbf{B}', \mathbf{n}')] + cn_k' = -u_k'(\mathbf{B}', \mathbf{n}').\quad (41)$$

because $n_k' = n_k''$ for $k \in \mathcal{B} \setminus \{j, R\}$ and $n_R'' < n_R'$. On the other hand, we have

$$u_j(\mathbf{B}'', \mathbf{n}'') = \underline{u}_j = u_j(\mathbf{B}', \mathbf{n}'),\quad (42)$$

by definition (29) of $\bar{n}_j(\mathbf{B}'', \mathbf{n}'')$. (41) and (42) together yield

$$\begin{aligned}-W(\mathbf{B}'', \mathbf{n}'') &= -\sum_{k \neq j} u_k(\mathbf{B}'', \mathbf{n}'') - u_j(\mathbf{B}'', \mathbf{n}'') = -\sum_{k \neq j} u_k(\mathbf{B}'', \mathbf{n}'') - u_j(\mathbf{B}', \mathbf{n}') \\ &< -\sum_{k \neq j} u_k(\mathbf{B}', \mathbf{n}') - u_j(\mathbf{B}', \mathbf{n}') = -W(\mathbf{B}', \mathbf{n}'),\end{aligned}$$

which once again violates (34). Thus, we obtain $\mathbf{B}^{\text{opt}} = (1, \dots, 1)^T$, that is $\mathcal{B}^{\text{opt}} = [J]$.

Part 2: Characterization of $\mathbf{n}_{[J]}^{\text{opt}} \in \mathbb{R}_+^J$

We now focus on the problem $\mathcal{T}_{\mathcal{B}^{\text{opt}}}$. To lighten notations, we write $\mathbf{n}^{\text{opt}} \in \mathbb{R}_+^J$ instead of $\mathbf{n}_{\mathcal{B}^{\text{opt}}}^{\text{opt}}$, and \mathcal{B} instead of $\mathcal{B}^{\text{opt}} = (1, \dots, 1)$. We first reformulate $\mathcal{T}_{\mathcal{B}}$ in a more convenient way. Recall that $\mathcal{T}_{\mathcal{B}}$ reads

$$(\mathcal{T}_{\mathcal{B}}) : \quad \text{minimize} \quad -W(\mathcal{B}, \mathbf{n}) \quad \text{subject to} \quad \mathbf{n} \in \Xi_{\mathcal{B}},\quad (43)$$

where Ξ_B is defined in (31). As explained in Appendix A, this problem can equivalently be stated as

$$(\tilde{\mathcal{T}}_B) : \quad \text{minimize} \quad -\tilde{W}_B(\boldsymbol{\lambda}, N) \quad \text{subject to} \quad (\boldsymbol{\lambda}, N) \in \tilde{\Xi}_B, \quad (44)$$

where

$$-\tilde{W}_B(\boldsymbol{\lambda}, N) = f(N) + g(\boldsymbol{\lambda}),$$

with

$$f(N) = aJ(\mathcal{R}_0^* + 2\alpha_\delta(1+N)^{-\gamma}) + cN \quad \text{and} \quad g(\boldsymbol{\lambda}) = 2aJ\boldsymbol{\lambda}^T\boldsymbol{\theta},$$

and

$$\tilde{\Xi}_B = \{(\boldsymbol{\lambda}, N) \in \Delta_{|B|} \times \mathbb{R}_+ : \boldsymbol{\lambda}N \in \Xi_B\}.$$

Writing $\tilde{\Xi}_B(N) = \{\boldsymbol{\lambda} \in \Delta_J : (\boldsymbol{\lambda}, N) \in \tilde{\Xi}_B\}$ for $N \geq 0$ and $\mathcal{N} = \{N \geq 0 : \tilde{\Xi}_B(N) \neq \emptyset\}$, $\tilde{\mathcal{T}}_B$ comes down to

$$\min_{(\boldsymbol{\lambda}, N) \in \tilde{\Xi}_B} \{f(N) + g(\boldsymbol{\lambda})\} = \min_{N \in \mathcal{N}} \{f(N) + \min_{\boldsymbol{\lambda} \in \tilde{\Xi}_B(N)} g(\boldsymbol{\lambda})\}.$$

In this proof, we address the innermost problem, which is enough to show that \mathbf{n}^{opt} satisfies the point (ii) of the result. Let $N \geq 0$ such that $\tilde{\Xi}_B(N) \neq \emptyset$, which exists because $\Xi_B \neq \emptyset$ (for instance, $((1, 0, \dots, 0), (\underline{n}, 0, \dots, 0)) \in \Xi_B$). By definition of $\tilde{\Xi}_B$, it reads

$$(\tilde{\mathcal{T}}^{(N)}) : \quad \text{minimize} \quad \boldsymbol{\lambda} \in \Delta_J \mapsto 2aJ\boldsymbol{\lambda}^T\boldsymbol{\theta} \quad \text{subject to} \quad \begin{cases} \forall j \in \mathcal{B} : \lambda_j \geq 0 \\ \forall j \in \mathcal{B} : \lambda_j N \leq \bar{n}_j(\mathbf{B}, \boldsymbol{\lambda}N) \\ \sum_{j \in \mathcal{B}} \lambda_j = 1, \end{cases} \quad (45)$$

The following lemma provides a necessary condition for any solution to problem (45).

Lemma 4. *Let $\boldsymbol{\lambda}^{(N)} = (\lambda_1^{(N)}, \dots, \lambda_J^{(N)}) \in \Delta_J$ be a solution to (45). If there exists $r \in [J]$ such that $\lambda_r^{(N)} > 0$, then $\lambda_k^{(N)} = N^{-1}\bar{n}_k(\mathbf{B}, N\boldsymbol{\lambda}^{(N)})$ for any $k < r$.*

Proof. We denote by $\bar{\mathbf{n}}(\mathbf{B}, \mathbf{n}) = (\bar{n}_1(\mathbf{B}, \mathbf{n}), \dots, \bar{n}_J(\mathbf{B}, \mathbf{n}))^T$ for any $\mathbf{n} \in \mathbb{R}_+^J$. The Lagrangian associated to problem (45), with $\boldsymbol{\mu}, \boldsymbol{\rho}$ and ν the associated dual variables, reads:

$$\begin{aligned} \mathcal{L}(\boldsymbol{\lambda}, \boldsymbol{\mu}, \boldsymbol{\rho}, \nu) &= 2aJ\boldsymbol{\lambda}^T\boldsymbol{\theta} - \boldsymbol{\mu}^T\boldsymbol{\lambda} + \nu(1 - \boldsymbol{\lambda}^T\mathbf{1}) + \boldsymbol{\rho}^T[N\boldsymbol{\lambda} - \bar{\mathbf{n}}(\mathbf{B}, N\boldsymbol{\lambda})] \\ &= 2aJ\boldsymbol{\lambda}^T\boldsymbol{\theta} + cN - \boldsymbol{\mu}^T\boldsymbol{\lambda} + \nu(1 - \boldsymbol{\lambda}^T\mathbf{1}) \\ &\quad + \boldsymbol{\rho}^T(N\boldsymbol{\lambda} - \underline{n}\mathbf{1} + \frac{2a}{c}[\alpha_\delta((N+1)^{-\gamma} - (\underline{n}+1)^{-\gamma})\mathbf{1} + ((\boldsymbol{\lambda}^T\boldsymbol{\theta})\mathbf{1} - \boldsymbol{\theta})]). \end{aligned}$$

Since the objective and the constraints are convex, \mathcal{L} admits a saddle point $(\boldsymbol{\lambda}^{(N)}, \boldsymbol{\mu}, \boldsymbol{\rho}, \nu) \in \Delta \times \mathbb{R}^J \times \mathbb{R}^J \times \mathbb{R}$ which is solution to problem (45) and verifies the following KKT conditions:

$$\begin{cases} \forall k \leq J : 2aJ\theta_k - \mu_k + \rho_k(N + 2ac^{-1}\theta_k) = \nu & (46) \end{cases}$$

$$\begin{cases} \forall k \leq J : \mu_k \lambda_k^{(N)} = 0, \quad \mu_k \geq 0 & (47) \end{cases}$$

$$\begin{cases} \forall k \leq J : \rho_k \left(\lambda_k^{(N)} N - \bar{n}_k(\mathbf{B}, \boldsymbol{\lambda}^{(N)}N) \right) = 0, \quad \rho_k \geq 0 & (48) \end{cases}$$

$$\begin{cases} \sum_{k \in \mathcal{B}} \lambda_k^{(N)} = 1. & (49) \end{cases}$$

Assume there exists $r \in \mathcal{B}$ such that $\lambda_r^{(N)} > 0$. By complementary slackness (47), $\mu_r = 0$, and (46) gives for any $k < r$:

$$2aJ\theta_k - \mu_k + \rho_k(N + 2ac^{-1}\theta_k) = 2aJ\theta_r + \rho_r(N + 2ac^{-1}\theta_r), \quad (50)$$

That is

$$\rho_k = \frac{2aJ(\theta_r - \theta_k) + \mu_k}{N + 2ac^{-1}\theta_k} + \rho_r \frac{N + 2ac^{-1}\theta_r}{N + 2ac^{-1}\theta_k}.$$

By assumption $\theta_r > \theta_k$, and since $\mu_k \geq 0$ as well as $\rho_r \geq 0$, we have $\rho_k > 0$. It follows from (48) for agent k that $\lambda_k^{\text{opt}} N = \bar{n}_k(\mathbf{B}, N\boldsymbol{\lambda}^{(N)})$.

□

Now, define

$$M_j = \sum_{k \leq j} \bar{n}_k(\mathbf{B}, \boldsymbol{\lambda}^{(N)} N) \quad \text{and} \quad L = \min\{j \in \mathcal{B} : M_j \geq N\}. \quad (51)$$

We show that

$$\lambda_k^{(N)} = \bar{n}_k(\mathbf{B}, N\boldsymbol{\lambda}^{(N)})N^{-1} \quad \text{for any } k < L \quad \text{and} \quad \lambda_k^{(N)} = 0 \quad \text{for any } k > L. \quad (52)$$

To prove the first point, assume by contradiction that there exists $\ell < L$ such that $\lambda_\ell^{(N)} < \bar{n}_\ell(\mathbf{B}, \boldsymbol{\lambda}^{(N)} N)N^{-1}$. By the contrapositive of Lemma 4, $\lambda_{\ell+1}^{\text{opt}} = \dots = \lambda_J^{\text{opt}} = 0$. Thus

$$\sum_{k \leq J} \lambda_k^{(N)} = \sum_{k \leq \ell} \lambda_k^{(N)} \leq \sum_{k \leq \ell} \bar{n}_k(\mathbf{B}, \boldsymbol{\lambda}^{(N)} N)N^{-1} = M_\ell N^{-1} < 1,$$

by definition of L Equation (51) and $\ell < L$. This violates the third constraint of problem (45). For the second point, assume there exists $\ell > L$ such that $\lambda_\ell^{(N)} > 0$. By Lemma 4, $\lambda_k^{(N)} = \bar{n}_k(\mathbf{B}, \boldsymbol{\lambda}^{(N)} N)N^{-1}$ for any $k \leq L < \ell$, so

$$\begin{aligned} \sum_{k \leq J} \lambda_k^{(N)} &\geq \sum_{k \leq L} \lambda_k^{(N)} + \lambda_\ell^{(N)} = \sum_{k \leq L} \bar{n}_k(\mathbf{B}, \boldsymbol{\lambda}^{(N)} N)N^{-1} + \lambda_\ell^{(N)} \\ &= M_L N^{-1} + \lambda_\ell^{(N)} \\ &\geq 1 + \lambda_\ell^{(N)} > 1, \end{aligned}$$

which once again violates the constraints of the problem. Finally by (52),

$$\lambda_{i_L}^{(N)} = 1 - \sum_{k < L} \lambda_k^{(N)} - \sum_{k > L} \lambda_k^{(N)} = 1 - M_{L-1} N^{-1}, \quad (53)$$

by (52). By (52) and (53) we obtain for any $j \in [J]$:

$$\lambda_j^{(N)} = N^{-1} \bar{n}_j(\mathbf{B}, N\boldsymbol{\lambda}^{(N)}) \mathbb{1}_{\{j < L\}} + (1 - N^{-1} M_{j-1}) \mathbb{1}_{\{j=L\}}, \quad (54)$$

Now, consider the solution $\mathbf{n}^{\text{opt}}(\boldsymbol{\theta}) \in \mathbb{R}^J$ to the initial problem (45), and define

$$\begin{cases} N^{\text{opt}} &= \sum_{j \in [J]} n_j^{\text{opt}}(\boldsymbol{\theta}) \\ \boldsymbol{\lambda}^{\text{opt}}(\boldsymbol{\theta}) &= \frac{1}{N^{\text{opt}}} \mathbf{n}^{\text{opt}}(\boldsymbol{\theta}) \\ L^{\text{opt}} &= \min\{j \in [J] : \sum_{k \leq j} \bar{n}_k(\mathbf{B}, \mathbf{n}^{\text{opt}}(\boldsymbol{\theta})) \geq N^{\text{opt}}\}. \end{cases}$$

$\boldsymbol{\lambda}^{\text{opt}}(\boldsymbol{\theta})$ satisfies (54), and multiplying by N^{opt} yields for any $j \in [J]$:

$$n_j^{\text{opt}}(\boldsymbol{\theta}) = \bar{n}_j(\mathbf{B}, \mathbf{n}^{\text{opt}}(\boldsymbol{\theta})) \mathbb{1}_{\{j < L^{\text{opt}}\}} + (N^{\text{opt}} - \sum_{k \leq j} n_k^{\text{opt}}(\boldsymbol{\theta})) \mathbb{1}_{\{j=L^{\text{opt}}\}}, \quad (55)$$

which establishes point (ii), and concludes the proof. \square

Corollary 1. Assume **H1**, **H2** and **H3**. Then $L^* = \Theta(J^{\frac{1}{1+\gamma}})$ and for any $j \in [J]$,

$$n_j^*(\boldsymbol{\theta}) = \mathbb{1}_{\{j \leq L^*\}} \left[\frac{\bar{N}}{L^*} + \frac{2a}{c} \left(\theta_j - \frac{1}{L^*} \sum_{\ell=1}^{L^*} \theta_\ell \right) \right].$$

Proof. By definition of the simplified scheme (12), $n_j^*(\boldsymbol{\theta}) = \mathbb{1}_{\{j \leq L^*\}} \bar{n}_j(\mathbf{1}, \mathbf{n}^*(\boldsymbol{\theta})) = \mathbb{1}_{\{j \leq L^*\}} [\underline{n} - (2a/c)(\varepsilon(\mathbf{1}, \mathbf{n}^*(\boldsymbol{\theta})) - \varepsilon(\theta_j, \underline{n}))]$ for any $j \leq L^*$. Since $\sum_{j \in [J]} n_j^*(\boldsymbol{\theta}) = \bar{N}$, we get

$$\begin{aligned} \bar{N} &= L^* \underline{n} - \frac{a}{c} \left[L^* \varepsilon(\mathbf{1}, \mathbf{n}^*(\boldsymbol{\theta})) - \sum_{j \leq L^*} \varepsilon(\theta_j, \underline{n}) \right] \\ &= L^* \left[\underline{n} - \frac{2a}{c} [\alpha_\delta((1 + \bar{N})^{-\gamma} - (1 + \underline{n})^{-\gamma}) + (\vartheta(\mathbf{1}, \mathbf{n}^*(\boldsymbol{\theta})) - \bar{\theta}_{L^*})] \right], \end{aligned}$$

By Lemma 1 and denoting $\bar{\theta}_{L^*} = L^{*-1} \sum_{j \leq L^*} \theta_j$. This yields

$$\vartheta(\mathbf{1}, \mathbf{n}^*(\boldsymbol{\theta})) = \bar{\theta}_{L^*} - \alpha_\delta((1 + \bar{N})^{-\gamma} - (1 + \underline{n})^{-\gamma}) + (c/2a)(\underline{n} - \bar{N}/L^*).$$

Plugging back this value in $\bar{n}_j^*(\mathbf{1}, \mathbf{n}^*(\boldsymbol{\theta})) = \underline{n} - (2a/c)[\alpha_\delta((1 + \bar{N})^{-\gamma} - (1 + \underline{n})^{-\gamma}) + (\vartheta(\mathbf{1}, \mathbf{n}^*(\boldsymbol{\theta})) - \theta_j)]$ gives for any $j \in [J]$:

$$n_j^*(\boldsymbol{\theta}) = \mathbb{1}_{\{j \leq L^*\}} \left[\frac{\bar{N}}{L^*} - \frac{2a}{c}(\bar{\theta}_{L^*} - \theta_j) \right]. \quad (56)$$

We now determine the order of magnitude of $L^* \in \{1, \dots, J\}$. Recall on the one hand that by (33), there exists $M > 0$ such that $0 \leq n_j^*(\boldsymbol{\theta}) \leq M$ for any $j \in [J]$, and on the other hand that $\bar{\theta}_{L^*} - \theta_j \leq \bar{\theta} - \underline{\theta} = \text{diam}(\Theta) < \infty$. Therefore by (56),

$$\frac{2a}{c} \text{diam}(\Theta) \leq \frac{\bar{N}}{L^*} \leq M + \frac{2a}{c} \text{diam}(\Theta) \quad \text{so} \quad \frac{\bar{N}}{L^*} = \Theta(1). \quad (57)$$

Since $\bar{N} = (\underline{n} + 1)|\mathcal{B}^*|^{\frac{1}{1+\gamma}} - 1$ and $\mathcal{B}^* = [J]$, (57) results in $L = \Theta(J^{\frac{1}{1+\gamma}})$. □

Lemma 2. Assume **H1**, **H2** and $\theta_j - \theta_{j-1} = \mathcal{O}(1/J)$ for any $j \in \{2, \dots, J\}$. Then,

$$W(\mathbf{B}^{\text{opt}}, \mathbf{n}^{\text{opt}}(\boldsymbol{\theta})) = W(\mathbf{B}^*, \mathbf{n}^*(\boldsymbol{\theta})) + \mathcal{O}(J^{\frac{1}{1+\gamma}}).$$

Proof. Recall that for any admissible $\mathbf{n} \in \Xi_{\mathbf{B}}$, we can define $N = \mathbf{1}^T \mathbf{n}$ and $\boldsymbol{\lambda} = N^{-1} \mathbf{n}$ so that the social cost rewrites

$$-W(\mathbf{B}, \mathbf{n}) = -\widetilde{W}_{\mathbf{B}}(\boldsymbol{\lambda}, N) = f(N) + g(\boldsymbol{\lambda}),$$

with

$$f(N) = aJ(\mathcal{R}_0^* + 2\alpha_\delta(1 + N)^{-\gamma}) + cN \quad \text{and} \quad g(\boldsymbol{\lambda}) = 2aJ\boldsymbol{\lambda}^T \boldsymbol{\theta}.$$

By definition, $\sum_{j \in [J]} n_j^*(\boldsymbol{\theta}) = \bar{N}$ where $\bar{N} = (\underline{n} + 1)J^{\frac{1}{1+\gamma}} - 1 = \text{argmin}_{N \geq 0} f(N)$. Hence,

$$\begin{aligned} W(\mathbf{B}^{\text{opt}}, \mathbf{n}^{\text{opt}}(\boldsymbol{\theta})) - W(\mathbf{B}^*, \mathbf{n}^*(\boldsymbol{\theta})) &= \widetilde{W}_{\mathbf{B}}(\boldsymbol{\lambda}^{\text{opt}}, N^{\text{opt}}) - \widetilde{W}_{\mathbf{B}}(\boldsymbol{\lambda}^*, \bar{N}) \\ &= f(\bar{N}) + g(\boldsymbol{\lambda}^*) - f(N^{\text{opt}}) - g(\boldsymbol{\lambda}^{\text{opt}}) \\ &\leq g(\boldsymbol{\lambda}^*) - g(\boldsymbol{\lambda}^{\text{opt}}) = 2aJ(\boldsymbol{\lambda}^* - \boldsymbol{\lambda}^{\text{opt}})^T \boldsymbol{\theta}. \end{aligned}$$

Since $\theta_j \leq \theta_{L^*}$ for any $j \in [J]$ such that $\lambda_j^* > 0$, and $\theta_j \geq \theta_1$ for any $j \in [J]$ such that $\lambda_j^{\text{opt}} > 0$:

$$W(\mathbf{B}^{\text{opt}}, \mathbf{n}^{\text{opt}}(\boldsymbol{\theta})) - W(\mathbf{B}^*, \mathbf{n}^*(\boldsymbol{\theta})) \leq 2aJ(\theta_{L^*} - \theta_1),$$

By assumption, for any $j \in \{2, \dots, J\}$ $\theta_j - \theta_{j-1} = \mathcal{O}(1/J)$ so there exists $k_j > 0$ and $R_j \geq 0$ such that $\theta_j - \theta_{j-1} \leq k_j J^{-1}$ for $J \geq R_j$. Denoting $R = \max_{j \in [J]} R_j$ and $k = \max_{j \in [J]} k_j$, for any $J \geq R$ we have $\theta_{L^*} - \theta_1 \leq kL^*$, so:

$$W(\mathbf{B}^{\text{opt}}, \mathbf{n}^{\text{opt}}(\boldsymbol{\theta})) - W(\mathbf{B}^*, \mathbf{n}^*(\boldsymbol{\theta})) \leq 2aJ \frac{kL^*}{J} = 2akL^*,$$

so $W(\mathbf{B}^{\text{opt}}, \mathbf{n}^{\text{opt}}(\boldsymbol{\theta})) - W(\mathbf{B}^*, \mathbf{n}^*(\boldsymbol{\theta})) = \mathcal{O}(L^*) = \mathcal{O}(J^{\frac{1}{1+\gamma}})$ by Corollary 1. □

Theorem 2 (Unravelling). Assume **H1**, **H2**, **H3**, and **H4**. Let $\mathcal{E} \subset \mathcal{S}^J$ be the set of pure-strategy Nash equilibria of the game induced by Γ . We have

- (i) $\mathcal{E} \neq \emptyset$
- (ii) at any $s^* \in \mathcal{E}$, $\mathbf{B} = (0, \dots, 0)$ or $\mathbf{B} = (0, \dots, 0, 1)$.

Proof. First, we show that the situation where the coalition is empty is a Nash equilibrium. Consider $\mathbf{s} = ((0, \dagger), \dots, (0, \dagger)) \in \mathcal{S}^N$. For any $j \in [J]$ and deviation $(1, \tilde{\theta}_j)$, $\mathcal{B} = \{j\}$ so $\hat{g}_{\mathcal{B}} = \hat{g}_j$ and $v_j((0, \dagger), \mathbf{s}_{-j}^*) = \underline{u}_j \geq v_j((1, \tilde{\theta}_j), \mathbf{s}_{-j})$ by Proposition 1. Thus, $\mathbf{s}^* \in \mathcal{E}$.

Now, consider a pure-strategy Nash equilibrium $\mathbf{s}^* \in \mathcal{E}$ such that $\mathcal{B} = \{j \in [J] : B_j = 1\} \neq \emptyset$. Denote by $\mathcal{C} = \{j \in \mathcal{B} : n_j^*(\tilde{\theta}) > 0\}$ the set of contributors. We start with two technical lemmas:

Lemma 5. *There exists $\Delta \in \mathbb{R}$ such that for any $(j, k) \in \mathcal{C}^2$*

$$\theta_j - \tilde{\theta}_j = \theta_k - \tilde{\theta}_k = \Delta .$$

Proof. To lighten notation, we write in this proof

$$\vartheta(\tilde{\theta}) = N^{-1} \sum_{j \in \mathcal{C}} n_j^*(\tilde{\theta}) \theta_j \quad \text{and} \quad \tilde{\vartheta}(\tilde{\theta}) = N^{-1} \sum_{j \in \mathcal{C}} n_j^*(\tilde{\theta}) \tilde{\theta}_j, \quad \text{where} \quad N = \sum_{k \in \mathcal{C}} n_k^*(\tilde{\theta}) .$$

By Equation (12) and H4, for any $j \in \mathcal{C}$, $n_j^*(\tilde{\theta}) = \bar{n}_j(\mathcal{B}, \mathbf{n}^*(\tilde{\theta})) = \underline{n} - \frac{a}{c}(\varepsilon(\tilde{\vartheta}(\tilde{\theta}), N) - \varepsilon(\tilde{\theta}_j, \underline{n}))$, so their payoff reads

$$\begin{aligned} v_j((1, \tilde{\theta}_j), \mathbf{s}_{-j}) &= -a(\mathcal{R}_0^* + \varepsilon(\vartheta(\tilde{\theta}), N)) - c \left[\underline{n} - \frac{a}{c}(\varepsilon(\tilde{\vartheta}(\tilde{\theta}), N) - \varepsilon(\tilde{\theta}_j, \underline{n})) \right] \\ &= -a(\mathcal{R}_0^* + \varepsilon(\theta_j, \underline{n})) - c\underline{n} \\ &\quad + a \left[(\varepsilon(\theta_j, \underline{n}) - \varepsilon(\tilde{\theta}_j, \underline{n})) - (\varepsilon(\vartheta(\tilde{\theta}), N) - \varepsilon(\tilde{\vartheta}(\tilde{\theta}), N)) \right] \\ &= \underline{u}_j + 2a \left[(\theta_j - \tilde{\theta}_j) - (\vartheta(\tilde{\theta}) - \tilde{\vartheta}(\tilde{\theta})) \right] . \end{aligned} \tag{58}$$

Since \mathbf{s} is a Nash equilibrium and $\underline{u}_j = v_j((0, \dagger), \mathbf{s}_{-j})$, we have in particular that

$$2a \left[(\theta_j - \tilde{\theta}_j) - (\vartheta(\tilde{\theta}) - \tilde{\vartheta}(\tilde{\theta})) \right] \geq 0 ,$$

that is $\theta_j - \tilde{\theta}_j = \Delta_j \geq \Delta = \vartheta(\tilde{\theta}) - \tilde{\vartheta}(\tilde{\theta})$ for any $j \in \mathcal{C}$. We now show that this holds with strict equality for any $j \in \mathcal{C}$. By contradiction, assume there exists $r \in \mathcal{C}$ such that $\Delta_r = \Delta + \chi$ with $\chi > 0$. Then

$$\begin{aligned} \Delta &= \vartheta(\tilde{\theta}) - \tilde{\vartheta}(\tilde{\theta}) = \sum_{j \in \mathcal{C}} \lambda_j^*(\tilde{\theta}) \Delta_j = \sum_{j \in \mathcal{C} \setminus \{r\}} \lambda_j^*(\tilde{\theta}) \Delta_j + \lambda_r^*(\tilde{\theta}) (\Delta + \chi) \\ &\geq \Delta + \lambda_r^*(\tilde{\theta}) \chi > \Delta , \end{aligned}$$

because $\lambda_r^*(\tilde{\theta}) > 0$ as $r \in \mathcal{C}$. This is a contradiction, and establishes the result. \square

Lemma 6. *For any $j \in \mathcal{C}$, $\tilde{\theta}_j = \underline{\theta}$.*

Proof. Let $j \in \mathcal{C}$, by (58):

$$\begin{aligned} v_j((1, \tilde{\theta}_j), \mathbf{s}_{-j}) &= \underline{u}_j + 2a \left[(\theta_j - \tilde{\theta}_j) - (\vartheta(\tilde{\theta}) - \tilde{\vartheta}(\tilde{\theta})) \right] \\ &= \underline{u}_j + 2a \left[(1 - \lambda_j^*(\tilde{\theta})) (\theta_j - \tilde{\theta}_j) - \sum_{k \in \mathcal{C} \setminus \{j\}} \lambda_k^*(\tilde{\theta}) (\theta_k - \tilde{\theta}_k) \right] \\ &= \underline{u}_j + 2a \left[\left(\sum_{k \in \mathcal{C} \setminus \{j\}} \lambda_k^*(\tilde{\theta}) \right) (\theta_j - \tilde{\theta}_j) - \sum_{k \in \mathcal{C} \setminus \{j\}} \lambda_k^*(\tilde{\theta}) (\theta_k - \tilde{\theta}_k) \right] \\ &= \underline{u}_j + 2a q_j(\tilde{\theta}_j, \tilde{\theta}_{-j}) , \end{aligned}$$

where

$$q_j(\tilde{\theta}_j, \tilde{\theta}_{-j}) = \sum_{k \in \mathcal{C} \setminus \{j\}} \lambda_k^*(\tilde{\theta}) [(\theta_j - \tilde{\theta}_j) - (\theta_k - \tilde{\theta}_k)] . \tag{59}$$

We prove that $v_j((1, \cdot), \mathbf{s}_{-j})$ is strictly decreasing in $\tilde{\theta}_j$ for $j \in \mathcal{C}$, by showing that $\partial v_j((1, \tilde{\theta}_j), \mathbf{s}_{-i_j}) / \partial \tilde{\theta}_j < 0$ For any $j \in \mathcal{C}$. For v_j to be differentiable, we need q_j to be differentiable, that is $\lambda_j^*(\tilde{\theta}) = (\sum_{k \in \mathcal{C}} n_k^*(\tilde{\theta}))^{-1} n_j^*(\tilde{\theta})$ to be differentiable for any $j \in \mathcal{C}$. We re-index \mathcal{C} as $\{i_1, \dots, i_{|\mathcal{C}|}\}$ so that $\tilde{\theta}_{i_1} < \dots < \tilde{\theta}_{i_{|\mathcal{C}|}}$. Observe that for $0 < h < \delta$ for $\delta > 0$ sufficiently small, $\tilde{\theta}_{i_{j-1}} < \tilde{\theta}_j + h < \tilde{\theta}_{i_{j+1}}$ because $\tilde{\theta}_{i_{j-1}} < \tilde{\theta}_j < \tilde{\theta}_{i_{j+1}}$ by Lemma 5 and H3. Hence, the bid ordering does not change for any infinitesimal variation $d\tilde{\theta}_j > 0$, nor does the indicator $\mathbb{1}\{i_j \leq i_{|\mathcal{C}|}\}$. Consequently by Corollary 1, $n_j^*(\tilde{\theta})$ is differentiable in $\tilde{\theta}_j$ and so is $\lambda^*(\tilde{\theta})$. We have for any $j \in \mathcal{C}$:

$$\begin{aligned} \frac{\partial v_j((1, \tilde{\theta}_j), \mathbf{s}_k)}{\partial \tilde{\theta}_j} &= 2a \frac{\partial q_j(\tilde{\theta}_j, \tilde{\theta}_{-j})}{\partial \tilde{\theta}_j} \\ &= 2a \sum_{k \in \mathcal{C} \setminus \{j\}} \left[\frac{\partial \lambda_k^*(\tilde{\theta})}{\partial \tilde{\theta}_j} \left((\theta_j - \tilde{\theta}_j) - (\theta_k - \tilde{\theta}_k) \right) - \lambda_k^*(\tilde{\theta}) \right] \end{aligned}$$

$$= 2a \sum_{k \in \mathcal{C} \setminus \{j\}} \frac{\partial \lambda_k^*(\tilde{\theta})}{\partial \tilde{\theta}_j} (\Delta - \Delta) - 2a \sum_{k \in \mathcal{C} \setminus \{j\}} \lambda_k^*(\tilde{\theta}) \quad (60)$$

$$< 0. \quad (61)$$

where we have used Lemma 5 at the third line. This implies

$$\tilde{\theta}_j = \underline{\theta} \quad \text{for any } j \in \mathcal{C}. \quad (62)$$

To see why, assume by contradiction that there exists $j \in \mathcal{C}$ such that $s_j = (1, \tilde{\theta}_j)$ with $\tilde{\theta}_j > \underline{\theta}$. Then for any $h \in (0, \tilde{\theta}_j - \underline{\theta}]$, by (61):

$$v_j((1, \tilde{\theta}_j - h), \mathbf{s}_{-j}) = v_j((1, \tilde{\theta}_j), \mathbf{s}_{-j}) - \int_{\tilde{\theta}_j - h}^{\tilde{\theta}_j} \frac{\partial v_j((1, t), \mathbf{s}_k)}{\partial t} dt > v_j((1, \tilde{\theta}_j), \mathbf{s}_k),$$

which contradicts \mathbf{s} being a Nash equilibrium. \square

Combining Lemma 5 and Lemma 6 gives for any $(j, k) \in \mathcal{C}^2$:

$$\theta_j = \theta_k. \quad (63)$$

Recall that by H3, $\theta_m = \theta_n$ if and only if $m = n$, so (63) implies $|\mathcal{C}| = 1$. This in turn implies $|\mathcal{B}| = 1$. Indeed, assume by contradiction $|\mathcal{C}| = 1$ and $|\mathcal{B}| \geq 2$. Denote $r \in \mathcal{B}$ the only contributing agent. They are asked $n_r^*(\tilde{\theta}) = \bar{n}_r(\tilde{\theta}) = \underline{n}$. Moreover by (19) $\bar{N} = |\mathcal{B}|^{1/1+\gamma} (\underline{n} + 1) - 1 > \underline{n}$, so by definition of the contribution scheme (12), there exists $k \in \mathcal{B} \setminus \{r\}$ such that $n_k^*(\tilde{\theta}) > 0$. This contradicts $|\mathcal{C}| = 1$.

We now show that $\mathcal{B} = \{J\}$. By contradiction, assume $\mathcal{B} = \{j\}$ with $j < J$. In particular, $s_J = (0, \dagger)$. Consider a deviation $s'_J = (1, \tilde{\theta}_J)$ with $\tilde{\theta}_J \in \Theta$, so $\mathcal{B} = \{j, J\}$ under (s'_J, \mathbf{s}_{-j}) . By Lemma 6, $\tilde{\theta}_j = \tilde{\theta}_J = \underline{\theta}$, so (58) rewrites:

$$v_k((1, \tilde{\theta}_k), \mathbf{s}_{-k}) = \underline{u}_k + 2a(\theta_k - \vartheta(\mathcal{B}, \mathbf{n}^*(\tilde{\theta}))), \quad (64)$$

for any $k \in \{j, J\}$. Since $\theta_j < \vartheta(\mathcal{B}, \mathbf{n}^*(\tilde{\theta})) < \theta_J$ by H3, we have

$$v_J((1, \tilde{\theta}_J), \mathbf{s}_{-J}) > \underline{u}_J = v_J((0, \dagger), \mathbf{s}_{-J}),$$

which contradicts \mathbf{s} being a Nash equilibrium. Hence, $\mathcal{B} = \{J\}$. This concludes the proof. \square

Lemma 3. *There exists $j \in [J]$ such that $-t_j(\tilde{\theta}) > 0$.*

Proof. Let $(\mathcal{B}, \tilde{\theta}) \in \{0, 1\}^J \times \Theta^N$ be an equilibrium of the game induced by Γ^{VCG} . Since the VCG mechanism is strategyproof in dominant strategy, $\tilde{\theta} = \theta$. For any $j \in [J]$, the VCG payment is

$$\begin{aligned} t_j(\tilde{\theta}) &= \sum_{k \neq j} u_k((0, \mathbf{1}_{-j}), \mathbf{n}^*(\theta)) - \sum_{k \neq j} u_k(\mathbf{1}, \mathbf{n}^*(\theta)) \\ &= \underbrace{W((0, \mathbf{1}_{-j}), \mathbf{n}^*(\theta)) - W(\mathbf{1}, \mathbf{n}^*(\theta))}_{\text{(I)}} - \underbrace{[u_j((0, \mathbf{1}_{-j}), \mathbf{n}^*(\theta)) - u_j(\mathbf{1}, \mathbf{n}^*(\theta))]}_{\text{(II)}}. \end{aligned} \quad (65)$$

Consider $j \in \{1, \dots, L^*\}$. By the contribution scheme Corollary 1,

$$u_j(\mathbf{1}, \mathbf{n}^*(\boldsymbol{\theta})) = -a(\mathcal{R}_0^* + \varepsilon(\mathbf{1}, \mathbf{n}^*(\boldsymbol{\theta}))) - cn_j^*(\boldsymbol{\theta}) = \underline{u}_j = u_j((0, \mathbf{1}_{-j}), \mathbf{n}^*(\boldsymbol{\theta})),$$

where we used in the third inequality that $n_j^*(\boldsymbol{\theta}) = \bar{n}_j(\mathbf{B}, \mathbf{n}^*(\boldsymbol{\theta}))$. Thus, $(\mathbf{II}) = 0$. Moreover by Theorem 1, $W((0, \mathbf{1}_{-j}), \mathbf{n}^*(\boldsymbol{\theta})) < W(\mathbf{1}, \mathbf{n}^*(\boldsymbol{\theta}))$ so $(\mathbf{I}) < 0$. Thus, by (65) we have $t_j(\tilde{\boldsymbol{\theta}}) < 0$. \square

Theorem 3. Assume **H1**, **H2**, **H3**, **H4** and **H5**. $\mathbf{B}^* = (1, \dots, 1)^T$ is a Nash equilibrium under $\hat{\Gamma}$ with probability $1 - \delta$.

Proof. Let $\mathbf{B}^* = (1, \dots, 1)^T$, so the aggregator has at her disposal $\hat{\theta}_1, \dots, \hat{\theta}_J$.

We first focus on the event under which $\hat{\theta}_j$ correctly approximates θ_j for any $j \leq J$. Define $\mathbf{E}_j(\delta) = \{\omega \in \Omega : |\hat{\theta}_j(\omega) - \theta_j| \leq \eta_\delta\}$ and $\mathbf{E}(\delta) = \bigcap_{j \leq J} \mathbf{E}_j(\delta)$. We have

$$\mathbb{P}(\mathbf{E}(\delta/J)) \geq 1 - \delta, \quad (66)$$

because $\mathbb{P}(\mathbf{E}(\delta)) = \mathbb{P}(\bigcap_{j \leq J} \mathbf{E}_j(\delta)) = 1 - \mathbb{P}(\bigcup_{j \leq J} \mathbf{E}_j^c(\delta)) \geq 1 - \sum_{j \leq J} \mathbb{P}(\mathbf{E}_j^c(\delta)) = 1 - J\delta$, and setting $\delta' = J\delta$ leads to Equation (66). In what follows, we assume that $\mathbf{E}(\delta/J)$ is true so $|\hat{\theta}_j - \theta_j| \leq \eta_{\delta/J}(\underline{q})$ for any $j \in [J]$. We denote by $\mathcal{C} = \{j \in [J] : m_j(\hat{\boldsymbol{\theta}}) = n_j^*(\hat{\boldsymbol{\theta}} + \boldsymbol{\eta}_j) > 0\}$ the set of contributors, and to lighten notations we write

$$\begin{aligned} \vartheta(\hat{\boldsymbol{\theta}}) &= \bar{N}^{-1} \mathbf{n}^*(\hat{\boldsymbol{\theta}} + \boldsymbol{\eta}_j)^T \boldsymbol{\theta} = \boldsymbol{\lambda}^*(\hat{\boldsymbol{\theta}} + \boldsymbol{\eta}_j)^T \boldsymbol{\theta} \\ \text{and } \hat{\vartheta}_j(\hat{\boldsymbol{\theta}}) &= \bar{N}^{-1} \mathbf{n}^*(\hat{\boldsymbol{\theta}} + \boldsymbol{\eta}_j)^T (\hat{\boldsymbol{\theta}} + \boldsymbol{\eta}_j) = \boldsymbol{\lambda}^*(\hat{\boldsymbol{\theta}} + \boldsymbol{\eta}_j)^T (\hat{\boldsymbol{\theta}} + \boldsymbol{\eta}_j), \\ \text{where } \boldsymbol{\eta}_j &= \eta_{\delta/J}(\underline{q}) \mathbf{1} - 2\delta_j \eta_{\delta/J}(\underline{q}), \quad \text{with } \boldsymbol{\delta}_j = (0, \dots, 0, 1, 0, \dots, 0), \end{aligned}$$

for any $j \in [J]$. Note that $\vartheta(\hat{\boldsymbol{\theta}})$ is the actual weighted type within the coalition, whereas $\hat{\vartheta}(\hat{\boldsymbol{\theta}})$ is the weighted type estimated by the aggregator. Consider first a contributor $j \in \mathcal{C}$ who is asked $n_j^*(\hat{\boldsymbol{\theta}} + \boldsymbol{\eta}_j) > 0$ samples. Their payoff under \mathbf{B}^* is

$$\begin{aligned} \hat{v}_j(1, \mathbf{1}_{-j}) &= -a(\mathcal{R}_0^* + \varepsilon(\vartheta(\hat{\boldsymbol{\theta}}), N)) - cn_j^*(\hat{\boldsymbol{\theta}} + \boldsymbol{\eta}_j) \\ &= -a(\mathcal{R}_0^* + \varepsilon(\vartheta(\hat{\boldsymbol{\theta}}), N)) - c \left[\underline{n} - \frac{a}{c} \left(\varepsilon(\hat{\vartheta}_j(\hat{\boldsymbol{\theta}}), N) - \varepsilon(\hat{\theta}_j - \eta_{\delta/J}(\underline{q}), \underline{n}) \right) \right] \\ &= -a(\mathcal{R}_0^* + \varepsilon(\theta_j, \underline{n})) - c\underline{n} \\ &+ a \left[\left(\varepsilon(\theta_j, \underline{n}) - \varepsilon(\hat{\theta}_j - \eta_{\delta/J}(\underline{q})) \right) - \left(\varepsilon(\vartheta(\hat{\boldsymbol{\theta}}), N) - \varepsilon(\hat{\vartheta}_j(\hat{\boldsymbol{\theta}}), N) \right) \right] \\ &= \underline{u}_j + 2a \left[\left(\theta_j - (\hat{\theta}_j - \eta_{\delta/J}(\underline{q})) \right) - \left(\vartheta(\hat{\boldsymbol{\theta}}) - \hat{\vartheta}_j(\hat{\boldsymbol{\theta}}) \right) \right] \\ &= \hat{v}_j(0, \mathbf{1}_{-j}) + 2a(1 - \lambda_j^*(\hat{\boldsymbol{\theta}} + \boldsymbol{\eta}_j))(\theta_j - (\hat{\theta}_j - \eta_{\delta/J}(\underline{q}))) \\ &- 2a \sum_{\ell \in \mathcal{C} \setminus \{j\}} \lambda_\ell^*(\hat{\boldsymbol{\theta}} + \boldsymbol{\eta}_j)(\theta_\ell - (\hat{\theta}_\ell + \eta_{\delta/J}(\underline{q}))), \end{aligned}$$

and since $\mathbf{E}(\delta/J)$ is true, $\theta_j \geq \hat{\theta}_j - \eta_{\delta/J}(\underline{q})$ and $\theta_\ell \leq \hat{\theta}_\ell + \eta_{\delta/J}(\underline{q})$ for any $\ell \in \mathcal{C} \setminus \{j\}$, so

$$\hat{v}_j(1, \mathbf{1}_{-j}) \geq \hat{v}_j(0, \mathbf{1}_{-j}). \quad (67)$$

Second, consider $j \notin \mathcal{C}$ who is asked $\underline{q} \leq \underline{n} - 2(a/c)(\bar{\theta} - \underline{\theta})$ samples, and denote by $r \in [J]$ the agent such that $\theta_r = \max_{k \in \mathcal{C}} \theta_k$. Observe that $\underline{q} \leq n_r^*(\hat{\boldsymbol{\theta}} + \boldsymbol{\eta}_r)$, so

$$\begin{aligned} \hat{v}_j(1, \mathbf{1}_{-j}) &= -a(\mathcal{R}_0^* + \varepsilon(\vartheta(\hat{\boldsymbol{\theta}}), N)) - c\underline{q} \\ &\geq -a(\mathcal{R}_0^* + \varepsilon(\vartheta(\hat{\boldsymbol{\theta}}), N)) - cn_r^*(\hat{\boldsymbol{\theta}} + \boldsymbol{\eta}_r) \\ &\geq \underline{u}_r, \end{aligned}$$

where the last inequality comes from $r \in \mathcal{C}$ and (67). Since $j \notin \mathcal{C}$, we have $\theta_j \geq \hat{\theta}_j - \eta_{\delta/J}(\underline{q}) \geq \hat{\theta}_r + \eta_{\delta/J}(\underline{q}) \geq \theta_r$ so $\underline{u}_r \geq \underline{u}_j$, and it follows that

$$\hat{v}_j(1, \mathbf{1}_{-j}) \geq \underline{u}_j = \hat{v}_j(0, \mathbf{1}_{-j}). \quad \square$$

Proposition 2. Assume **H1**, **H2** and **H6**. For any $j \in [J]$ the estimator

$$\hat{\theta}_{0,j}^{\text{ERM}} = \sup_{g \in \mathcal{G}} |\hat{\mathcal{R}}_j(g) - \hat{\mathcal{R}}_0(g)|,$$

satisfies **H5** with

$$\eta_\delta(q) = \alpha_{\delta/4} [(q+1)^{-\gamma} + (q'+1)^{-\gamma}] + 2\beta. \quad (13)$$

Proof. Let $j \in [J]$ and $g \in \mathcal{G}$, we have

$$\begin{aligned} \hat{\theta}_{0,j}^{\text{ERM}} &= \sup_{g \in \mathcal{G}} |\hat{\mathcal{R}}_j(g) - \hat{\mathcal{R}}_0(g)| \\ &\leq \sup_{g \in \mathcal{G}} |\hat{\mathcal{R}}_j(g) - \mathcal{R}_j(g)| + \sup_{g \in \mathcal{G}} |\mathcal{R}_j(g) - \mathcal{R}_0(g)| + \sup_{g \in \mathcal{G}} |\hat{\mathcal{R}}_0(g) - \mathcal{R}_0(g)| \\ &\leq \alpha_{\delta/4} [(q+1)^{-\gamma} + (q'+1)^{-\gamma}] + 2\beta + \theta_j, \end{aligned} \quad (68)$$

(69)

with probability $1 - \delta/2$ by **H2**, **H1**, **H6**, and a union bound. Similarly,

$$\begin{aligned} \theta_j &= \sup_{g \in \mathcal{G}} |\mathcal{R}_j(g) - \mathcal{R}_0(g)| \\ &\leq \sup_{g \in \mathcal{G}} |\mathcal{R}_j(g) - \hat{\mathcal{R}}_j(g)| + \sup_{g \in \mathcal{G}} |\hat{\mathcal{R}}_j(g) - \hat{\mathcal{R}}_0(g)| + \sup_{g \in \mathcal{G}} |\hat{\mathcal{R}}_0(g) - \mathcal{R}_0(g)| \\ &\leq \alpha_{\delta/4} [(q+1)^{-\gamma} + (q'+1)^{-\gamma}] + 2\beta + \hat{\theta}_{0,j}^{\text{ERM}}, \end{aligned} \quad (70)$$

with probability $1 - \delta/2$. Combining (68) and (70) along with an union bound yields

$$|\hat{\theta}_{0,j}^{\text{ERM}} - \theta_j| \leq \alpha_{\delta/4} [(q+1)^{-\gamma} + (q'+1)^{-\gamma}] + 2\beta,$$

with probability $1 - \delta$. \square

Example 4. Assume **H1**, **H2**, **H6** and **H7**.

(i) Denoting $\hat{\mathcal{R}}_{j-}(g) = n_j^{-1} \sum_{i=1}^{n_j} \ell_{0,1}(g(X_i^j), -Y_i^j)$, we have

$$\hat{\theta}_{0,j}^{\text{ERM}} = 1 - \inf_{g \in \mathcal{G}} \left\{ \hat{\mathcal{R}}_0(g) + \hat{\mathcal{R}}_{j-}(g) \right\}.$$

(ii) In **H1**, assume $\alpha_\delta = \ln(1/\delta)^{1/2}$, $\beta = 2\text{RAD}(\mathcal{G})$ and $\gamma = 1$ [[Bousquet et al., 2003](#)]. With $\hat{\theta}_{0,j}^{\text{ERM}}$ defined in Proposition 2, we have

$$\eta_{\delta/J}(q) = \ln(4J/\delta)^{1/2} [(1+q)^{-\gamma} + (1+q')^{-\gamma}] + 2\text{Rad}(\mathcal{G}).$$

Proof. (i) Let $j \in \mathcal{B}$. Observe that under **H7**, we have

$$\hat{\mathcal{R}}_j(-g) = n_j^{-1} \sum_{i=1}^{n_j} \mathbb{1}\{-g(X_i^j)Y_i^j < 0\} = n_j^{-1} \sum_{i=1}^{n_j} (1 - \mathbb{1}\{g(X_i^j)Y_i^j < 0\}) = 1 - \hat{\mathcal{R}}_j(g).$$

Since the hypothesis class \mathcal{G} is symmetric, we have

$$\begin{aligned} \hat{\theta}_j^{\text{ERM}} &= \sup_{g \in \mathcal{G}} |\hat{\mathcal{R}}_j(g) - \hat{\mathcal{R}}_0(g)| = \sup_{g \in \mathcal{G}} (\hat{\mathcal{R}}_j(g) - \hat{\mathcal{R}}_0(g)) \\ &= \sup_{g \in \mathcal{G}} \left(1 - (\hat{\mathcal{R}}_j(-g) + \hat{\mathcal{R}}_0(g)) \right) = 1 - \inf_{g \in \mathcal{G}} (\hat{\mathcal{R}}_{j-}(g) + \hat{\mathcal{R}}_0(g)). \end{aligned}$$

(ii) Let $j \in [J]$, we have

$$\begin{aligned} \eta_{\delta/J}(q) &= \alpha_{\delta/4J} [(q+1)^{-\gamma} + (1+q')^{-\gamma}] + 2\beta \\ &= \ln(4J/\delta)^{1/2} [(1+q)^{-\gamma} + (1+q')^{-\gamma}] + 2\text{Rad}(\mathcal{G}). \end{aligned}$$

\square

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 claimed contributions are supported by proven theorems.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

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

Answer: [Yes]

Justification: all the made assumptions are clearly highlighted and discussed.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

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: Assumptions are made clear and given for each theorem individually. Full proofs are given in the Appendix.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

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: [NA]

Justification: no experimental result in the paper

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

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: [NA]

Justification: no experimental result in the paper

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

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

Answer: [NA]

Justification: no experimental result in the paper

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

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: [NA]

Justification: no experimental result in the paper

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

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: [NA]

Justification: no experimental result in the paper

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

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: no data used for this work

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

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

Answer: [NA]

Justification: no societal impact of the work performed

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to

generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.

- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

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., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: no data

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

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: [NA]

Justification: no existing assets

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

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

Answer: [NA]

Justification: no new asset

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

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: [NA]

Justification: no crowdsourcing nor research with human subjects

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

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: [NA]

Justification: no crowdsourcing nor research with human subjects

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.