GRADIENT-VARIATION AND SMALL-LOSS REGRET FOR DECENTRALIZED ONLINE CONVEX OPTIMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

We study decentralized online convex optimization over a fixed network of N learners with spectral gap ρ, where learners communicate only with their neighbors. Existing algorithms achieve nearly optimal worst-case regret bounds, but do not adapt to benign loss sequences, such as those with small gradient variation or small comparator loss. To address this limitation, we propose Decentralized Optimistic Follow-the-Regularized-Leader with Gradient Tracking (DOFTRL-GT), which combines delay-compensated optimism with two-level gradient tracking via accelerated gossip. DOFTRL-GT achieves regret bounds of $\widetilde{\mathcal{O}}(N \min\{\rho^{-3/4}\sqrt{V_T}, \rho^{-1/4}\sqrt{F_T}\})$ for convex losses and $\widetilde{\mathcal{O}}(N\rho^{-1/2}\log(2+\min\{V_T, F_T\}))$ for strongly convex losses, up to additive terms independent of T, V_T, and F_T. Here, V_T denotes the gradient variation, F_T the cumulative loss of the best fixed comparator, and $\widetilde{\mathcal{O}}$ hides logarithmic factors in N and $1/\rho$. We establish lower bounds that match the corresponding upper bounds up to logarithmic factors in their respective problem-dependent regimes, showing that convex gradient-variation and small-loss guarantees necessarily incur distinct $\rho^{-3/4}$ and $\rho^{-1/4}$, respectively. Finally, applications to offline optimization and stochastic-adversarial learning further demonstrate its ability to exploit benign loss sequences.

1 INTRODUCTION

Decentralized online convex optimization (D-OCO) (Yan et al., 2012; Hosseini et al., 2013) provides a standard framework for sequential decision making over distributed networks, with applications including distributed online regression (Yuan et al., 2020), target tracking in sensor networks (Shahrampour & Jadbabaie, 2017), and online packet routing (Awerbuch & Kleinberg, 2008). Consider a network of N agents connected by a fixed undirected communication graph $\mathcal{G} = ([N], \mathcal{E})$ with edge set $\mathcal{E} \subseteq [N] \times [N]$. At each round $t \in [T]$, every agent $i \in [N]$ selects a decision $x_{t,i}$ from a convex decision set $\mathcal{X}$, observes its local convex loss function $f_{t,i}$, and communicates with neighboring agents before making its next decision. The objective is to minimize each agent's regret on the global loss $f_t(x) \triangleq \sum_{j=1}^{N} f_{t,j}(x)$ relative to the best fixed decision in hindsight.

Recent progress has established nearly optimal worst-case regret guarantees for D-OCO under both convex and strongly convex losses. In particular, by combining a blocking update mechanism with accelerated gossip (Liu & Morse, 2011), Wan et al. (2024; 2025) substantially improved the network dependence. Together with corresponding lower bounds, these results characterize the optimal worst-case dependence on the horizon and network parameters up to logarithmic factors. However, these worst-case guarantees scale with the horizon and do not adapt to benign loss sequences, such as those with small gradient variation or a small cumulative loss of the best fixed comparator.

In contrast, when $N = 1$, D-OCO reduces to classical OCO, where problem-dependent regret guarantees have been developed for both convex and strongly convex losses. In particular, gradient-variation bounds yield improved guarantees when the gradients change slowly over time (Chiang et al., 2012; Yan et al., 2023), while small-loss bounds improve when the best fixed comparator incurs a small cumulative loss (Srebro et al., 2010; Wang et al., 2020; Yan et al., 2023). These results show that centralized OCO can adapt to the realized difficulty of the loss sequence rather than being

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Method	Convex	Strongly convex
Yan et al. (2012)	$\mathcal{O}(N^{\frac{5}{4}}\rho^{-\frac{1}{2}}\sqrt{T})$	$\mathcal{O}(N^{\frac{3}{2}}\rho^{-1}\log T)$
Wan et al. (2024; 2025)	$\widetilde{\Theta}(N\rho^{-\frac{1}{4}}\sqrt{T})$	$\widetilde{\Theta}(N\rho^{-\frac{1}{2}}\log T)$
Our upper bounds	$\widetilde{\mathcal{O}}\big(N\min\{\rho^{-\frac{3}{4}}\sqrt{V_T}, \rho^{-\frac{1}{4}}\sqrt{F_T}\}\big)$ Theorem 1	$\widetilde{\mathcal{O}}\big(N\rho^{-\frac{1}{2}}\log\big(2+\min\{V_T, F_T\}\big)\big)$ Theorem 2
Our lower bounds	$\Omega\big(N\min\{\rho^{-\frac{1}{4}}\sqrt{T}, \rho^{-\frac{3}{4}}\sqrt{V_T}\}\big)$	$\Omega\big(N\rho^{-\frac{1}{2}}\log(2+V_T)\big)$
Gradient variation	Theorem 3	Theorem 4
	$\Omega\big(N\rho^{-\frac{1}{4}}\min\{\sqrt{T}, \sqrt{F_T}\}\big)$	$\Omega\big(N\rho^{-\frac{1}{2}}\log(2+\rho^{\frac{1}{2}}F_T)\big)$
Small loss	Theorem 5	Theorem 6

Table 1: Comparison of regret guarantees for D-OCO. We omit horizon-independent additive terms and treat problem-geometry and loss-scale parameters as constants. The notation $\widetilde{\mathcal{O}}$ and $\widetilde{\Theta}$ hide logarithmic factors in N and $1/\rho$.

governed solely by the horizon. The question is whether such sequence-dependent improvements extend to the decentralized setting. This leaves the following question open:

> *Can decentralized online convex optimization exploit benign loss sequences while retaining nearly optimal network dependence?*

Contributions. To address these difficulties, we develop a decentralized optimistic learning method with gradient-variation and small-loss regret guarantees for smooth convex and strongly convex losses. Here, V_T and F_T denote the gradient variation and the cumulative loss of the best fixed comparator, respectively, and can be much smaller than T on benign loss sequences. Our main contributions are summarized below:

- In Section 3, we propose Decentralized Optimistic Follow-the-Regularized-Leader with Gradient Tracking (DOFTRL-GT), which combines delay-compensated optimism with two-level gradient tracking via accelerated gossip. The two trackers estimate the block gradient and the cumulative gradient, respectively, while the optimistic update predicts the current block gradient and compensates for the block delay introduced by accelerated gossip. Together, these mechanisms control both prediction errors and tracking errors through the same block-gradient variation rather than gradient magnitudes. To the best of our knowledge, we establish the first unified gradient-variation and small-loss regret guarantees for the D-OCO problem, covering both convex and strongly convex losses. These guarantees achieve optimal dependence on V_T and F_T, respectively, together with nearly optimal dependence on N and ρ. See Table 1 for the resulting rates and comparison with prior work.
- In Section 4, we establish gradient-variation and small-loss lower bounds for both convex and strongly convex losses, matching the corresponding upper bounds up to logarithmic factors in their respective problem-dependent regimes. Notably, for convex losses, these bounds reveal an intrinsic separation in network dependence: the $\rho^{-3/4}$ factor multiplying $\sqrt{V_T}$ is unavoidable in the variation-dependent regime, whereas the $\sqrt{F_T}$ term retains the $\rho^{-1/4}$ factor of the optimal horizon-dependent rate (Wan et al., 2025).
- In Section 5, For decentralized offline optimization, DOFTRL-GT recovers the $\mathcal{O}(1/T)$ ergodic rate of offline gradient tracking (Qu & Li, 2017) without strong convexity or additional communication, improving over the $\mathcal{O}(T^{-1/2})$ rate implied by horizon-dependent decentralized regret bounds. The application to stochastic-adversarial learning further illustrates its ability to exploit the benign temporal structure.

1.1 Related Works

Decentralized online optimization. Early methods for decentralized online optimization combined online gradient or dual-averaging updates with consensus-based communication (Yan et al.,

2012; Hosseini et al., 2013). These methods achieve regret bounds of $\mathcal{O}(N^{5/4}\rho^{-1/2}\sqrt{T})$ for convex losses and $\mathcal{O}(N^{3/2}\rho^{-1}\log T)$ for strongly convex losses, where N is the number of agents and ρ the spectral gap of the communication matrix. To improve this network dependence, Wan et al. (2024; 2025) combined accelerated gossip (Liu & Morse, 2011) with a blocking update mechanism, achieving $\mathcal{O}(N\rho^{-1/4}\sqrt{T\log N})$ and $\mathcal{O}(N\rho^{-1/2}\log N\log T)$ regret for convex and strongly convex losses, respectively. The lower bounds of Wan et al. (2025) establish optimality in N, ρ, and T up to logarithmic factors. More recently, Wan (2025) developed black-box reductions from D-OCO to centralized OCO with delayed feedback, obtaining nearly optimal adaptive and dynamic regret bounds. However, the resulting bounds retain additional logarithmic dependence on T, and the reduction does not directly yield gradient-variation or small-loss guarantees. Existing worst-case static regret guarantees also remain horizon-dependent and do not adapt to benign loss sequences, such as those with small gradient variation or a small cumulative loss of the best fixed comparator.

Optimistic and problem-dependent online learning. In centralized OCO, problem-dependent regret bounds exploit benign loss sequences. For smooth convex losses, Chiang et al. (2012) obtained $\mathcal{O}(\sqrt{1+V_T})$ regret in terms of gradient variation V_T. Optimistic mirror descent and FTRL provide a general framework with regret controlled by gradient-prediction errors (Rakhlin & Sridharan, 2013). A complementary line establishes $\mathcal{O}(1+\sqrt{F_T})$ small-loss bounds for smooth nonnegative convex losses, where F_T is the cumulative loss of the best fixed comparator (Srebro et al., 2010). Subsequent universal algorithms extend small-loss guarantees (Wang et al., 2020) and gradient-variation guarantees (Yan et al., 2023; 2024) across curvature classes. Related extensions include dynamic regret (Zhao et al., 2020; 2024), the stochastically extended adversarial model (Chen et al., 2024; Wang et al., 2026), and relaxed smoothness conditions (Xie et al., 2024; Zhao et al., 2026). Extending these guarantees to decentralized learning requires handling both delayed information and consensus errors. Flaspohler et al. (2021) connected delayed feedback with optimistic learning, showing how predictions compensate for missing gradients, but did not address decentralized consensus errors.

2 PRELIMINARIES

To begin with, we introduce some notation used throughout the paper. For any positive integer K, let $[K] \triangleq \{1,\dots,K\}$. Unless otherwise specified, we use $\|\cdot\|_2$ for the Euclidean norm of vectors and $\|\cdot\|_{\mathrm{F}}$ for the Frobenius norm of matrices. Let $(a)_+ \triangleq \max\{a,0\}$ and let $\mathbf{1}\in\mathbb{R}^N$ denote the all-ones vector. The notation $\widetilde{\mathcal{O}}$ hides polylogarithmic factors in N and $1/\rho$.

Consider N agents that collaboratively solve an online optimization problem over a decision set $\mathcal{X}\subseteq\mathbb{R}^d$. Their communication is governed by a fixed undirected graph $\mathcal{G}=(\mathcal{N},\mathcal{E})$, with node set $\mathcal{N}=[N]$ and edge set $\mathcal{E}\subseteq\mathcal{N}\times\mathcal{N}$. At each round $t\in[T]$, every agent $i\in[N]$ selects a decision $x_{t,i}\in\mathcal{X}$. After the decisions are chosen, the environment reveals a local loss function $f_{t,i}:\mathbb{R}^d\to\mathbb{R}$ to agent i. The global loss at round t is defined as $f_t(x)\triangleq\sum_{i=1}^N f_{t,i}(x)$. Agents communicate with their neighbors according to the gossip protocol described below and update their next-round decisions using local feedback and information received from the network. The performance of agent i over T rounds is measured by its regret $\mathrm{Reg}_T(i)\triangleq\sum_{t=1}^T f_t(x_{t,i})-\min_{u\in\mathcal{X}}\sum_{t=1}^T f_t(u)$. Let $u^\star\in\arg\min_{u\in\mathcal{X}}\sum_{t=1}^T f_t(u)$ be an optimal comparator and $\mathrm{Reg}_T\triangleq\max_{i\in[N]}\mathrm{Reg}_T(i)$.

To characterize benign loss sequences, we consider the following two problem-dependent quantities:

$$F_T \triangleq \min_{u\in\mathcal{X}}\sum_{t=1}^T\frac{1}{N}\sum_{i=1}^N f_{t,i}(u) \;\text{ and }\; V_T\triangleq\sum_{t=2}^T\frac{1}{N}\sum_{i=1}^N\sup_{x\in\mathcal{X}}\|\nabla f_{t,i}(x)-\nabla f_{t-1,i}(x)\|_2^2. \tag{1}$$

Here, the *small loss* F_T measures the cumulative loss of the best fixed comparator across all agents, while the *gradient variation* V_T measures the temporal variation of local gradients across the network. F_T and V_T can be viewed as decentralized extensions of the small-loss and gradient-variation measures studied in OCO, respectively (Srebro et al., 2010; Chiang et al., 2012). Small F_T or V_T characterize benign loss sequences for which problem-dependent regret bounds can improve over worst-case guarantees.

In this paper, following standard settings in D-OCO (Yan et al., 2012; Tu et al., 2022; Wan et al., 2024), we make the following assumptions.

Assumption 1. *The set $\mathcal{X} \subset \mathbb{R}^d$ is nonempty, closed, and convex, contains the origin, and satisfies $\|x\|_2 \leq R$ for every $x \in \mathcal{X}$.*

Assumption 2. *For every $t \in [T]$ and $i \in [N]$, the loss function $f_{t,i}$ is differentiable, α-strongly convex, and has bounded gradients. Namely, for all $x, y \in \mathcal{X}$, $f_{t,i}(y) \geq f_{t,i}(x) + \langle \nabla f_{t,i}(x), y - x \rangle + \frac{\alpha}{2}\|y - x\|_2^2$, $\|\nabla f_{t,i}(x)\|_2 \leq G$.*

Note that Assumption 2 with $\alpha = 0$ reduces to the case with general convex loss functions. To establish gradient-variation regret bounds for first-order algorithms, we further impose the following smoothness assumption (Chiang et al., 2012; Yan et al., 2023; 2024).

Assumption 3 (Smoothness)**.** *For every $t \in [T]$ and $i \in [N]$, the loss function $f_{t,i}$ is L-smooth, i.e., $\|\nabla f_{t,i}(x) - \nabla f_{t,i}(y)\|_2 \leq L\|x - y\|_2$, $\forall x, y \in \mathbb{R}^d$.*

Gossip Protocol. We assume that agents communicate only with their neighbors through a gossip protocol (Yan et al., 2012; Hosseini et al., 2013; Wan et al., 2024). For each agent $i \in [N]$, let $\mathcal{N}_i \triangleq \{i\} \cup \{j \in [N] : (i, j) \in \mathcal{E}\}$ denote its inclusive neighborhood. Communication over $\mathcal{G}$ is characterized by a nonnegative mixing matrix $W \in \mathbb{R}^{N \times N}$ supported on $\mathcal{G}$, i.e., $W_{ij} > 0$ only if $i = j$ or $(i, j) \in \mathcal{E}$. We make the following standard assumption on the communication network.

Assumption 4 (Communication network)**.** *The graph $\mathcal{G}$ is connected, and the mixing matrix W satisfies $W = W^\top$, $W\mathbf{1} = \mathbf{1}$, $W \succeq 0$, and $\sigma_2(W) < 1$. Here, $\sigma_2(W)$ denotes the second-largest singular value of W, and $\rho \triangleq 1 - \sigma_2(W)$ denotes its spectral gap.*

Given local states $\{v_i\}_{i=1}^N$ in a common real vector space, the goal of gossip is to approximate their network average $\overline{v} \triangleq \frac{1}{N}\sum_{i=1}^N v_i$ using only neighbor-to-neighbor communication. To this end, we initialize $v_i^{-1} = v_i^0 = v_i$ and perform K rounds of accelerated gossip. At each communication round $k = 0, \ldots, K-1$, agent i updates its local state as

$$v_i^{k+1} = (1 + \theta)\sum_{j=1}^N W_{ij} v_j^k - \theta v_i^{k-1}, \tag{2}$$

where $\theta \geq 0$ is the mixing coefficient. When $\theta = 0$, the above iteration reduces to standard gossip (Xiao & Boyd, 2004); an appropriate choice of $\theta > 0$ yields accelerated gossip (Liu & Morse, 2011), which achieves faster average consensus. After K communication rounds, agent i uses v_i^K as an approximation of $\overline{v}$.

3 Decentralized Optimistic FTRL with Gradient Tracking

In this section, we present a unified decentralized optimistic FTRL algorithm for both convex and strongly convex losses (see Algorithm 1). To retain the nearly optimal network dependence of Wan et al. (2024; 2025), we build on their blocking update mechanism with accelerated gossip. Incorporating gradient-variation adaptivity into this framework introduces two difficulties. First, blockwise communication delays the network-wide gradient information, so the standard previous-gradient prediction used in optimistic FTRL (Rakhlin & Sridharan, 2013; Chiang et al., 2012) is not directly available. Second, the algorithmic design and analysis of Wan et al. (2024; 2025) control the approximation error in cumulative gradient estimates through gradient magnitudes, but such magnitude-based control does not by itself provide the variation-dependent error bound required to derive gradient-variation regret bounds.

Algorithm 1 addresses these two issues through the following coupled mechanisms.

- **Delay-compensated optimism.** The decision update in Line 4 combines the latest completed block-gradient and cumulative-gradient estimates to compensate for the unavailable previous-block gradient. This delay-compensated update yields an optimistic prediction whose error is controlled by consecutive block-gradient differences and vanishes when the block gradients remain unchanged.
- **Two-level gradient tracking through block-gradient differences.** We use two-level gradient tracking: the first tracker incorporates local block-gradient differences to estimate

Algorithm 1: Decentralized Optimistic FTRL with Gradient Tracking (DOFTRL-GT)

1 **Input:** Agent index i, gossip matrix W, block length B, mixing coefficient θ, curvature parameter $\alpha \geq 0$, and regularization parameter $h > 0$
2 **Initialization:** set $a_i(z) = y_i(z) = s_i(z) = 0$ for all $z \leq 0$
3 **for** $z = 1, \ldots, T/B$ **do**
4 $\quad x_i(z) = \arg\min_{x \in \mathcal{X}} \left\{ \langle s_i(z-2) + 2y_i(z-2), x \rangle + \frac{z\alpha + 2h}{2} \|x\|_2^2 \right\}$
5 $\quad$ Set $u_i^{-1} = u_i^0 = y_i(z-2) + a_i(z-1) - a_i(z-2)$ and $v_i^{-1} = v_i^0 = s_i(z-2)$
6 $\quad$ **for** $t = (z-1)B + 1, \ldots, zB$ **do**
7 $\quad\quad$ Play $x_{t,i} = x_i(z)$ and query $\nabla f_{t,i}(x_i(z))$
8 $\quad\quad$ Set $k = t - (z-1)B - 1$
9 $\quad\quad$ Update $u_i^{k+1} = (1+\theta) \sum_{j \in \mathcal{N}_i} W_{ij} u_j^k - \theta u_i^{k-1}$ and $v_i^{k+1} = (1+\theta) \sum_{j \in \mathcal{N}_i} W_{ij} v_j^k - \theta v_i^{k-1}$
10 $\quad$ Set $a_i(z) = \frac{1}{B} \sum_{t \in \mathcal{I}_z} (\nabla f_{t,i}(x_i(z)) - \alpha x_i(z))$
11 $\quad$ Set $y_i(z-1) = u_i^B$ and $s_i(z-1) = v_i^B + y_i(z-1)$

the network-average block gradient, while the second accumulates these estimates to track the cumulative gradient. This allows the tracking errors to be controlled by block-gradient variation rather than gradient magnitudes, which can subsequently be related to the original gradient variation V_T.

More specifically, assume that B divides T, and let $Z \triangleq T/B$ denote the total number of blocks. For each $z \in [Z]$, let $\mathcal{I}_z \triangleq \{(z-1)B + 1, \ldots, zB\}$ denote the rounds in block z. Define

$$a_i(z) \triangleq \frac{1}{B} \sum_{t \in \mathcal{I}_z} \left(\nabla f_{t,i}(x_i(z)) - \alpha x_i(z) \right), \quad g(z) \triangleq \frac{1}{N} \sum_{j=1}^{N} a_j(z), \quad G_z \triangleq \sum_{\tau=1}^{z} g(\tau), \tag{3}$$

where $a_i(z)$ is the local block gradient, including the curvature correction $-\alpha x_i(z)$ when $\alpha > 0$. The quantities $g(z)$ and G_z denote the network-average block gradient and its cumulative sum, respectively.

The algorithm operates iteratively over these Z blocks as follows.

- First, at the beginning of block z, the newest completed tracking states available to agent i are $y_i(z-2)$ and $s_i(z-2)$, which track the network-average block gradient $g(z-2)$ and the cumulative gradient G_{z-2}, respectively. Agent i uses these delayed states to compute the delay-compensated optimistic decision $x_i(z)$ in Line 4, which is then played throughout the block. In this update, the two copies of $y_i(z-2)$ jointly compensate for the unavailable previous-block gradient $g(z-1)$ and provide the optimistic prediction.
- Second, while playing $x_i(z)$ for all $t \in \mathcal{I}_z$, agent i collects the current-block gradients to form $a_i(z)$ and simultaneously performs two-level gradient tracking using the auxiliary communication states u_i^k and v_i^k, where k indexes gossip rounds within the block. As initialized in Line 5, $u_i^0 = y_i(z-2) + a_i(z-1) - a_i(z-2)$ incorporates the newly available local block-gradient difference into the previous block-gradient estimate, so that the u-states track $g(z-1)$. In parallel, the v-states are initialized with $v_i^0 = s_i(z-2)$ to mix the previous cumulative-gradient estimates and reduce their disagreement around G_{z-2}; their output will be combined with the updated block-gradient estimate to track G_{z-1}. At each online round, agent i exchanges the pair (u_i^k, v_i^k) with its neighbors and updates both states by accelerated gossip in Line 9, completing B communication rounds over the block.
- Finally, at the end of block z, the gossip outputs are used to form $y_i(z-1)$ and $s_i(z-1)$, which track $g(z-1)$ and G_{z-1}, respectively; see Line 11. The gradients collected during block z are also used to form $a_i(z)$ in Line 10.

3.1 Regret Guarantees

We now present the regret guarantees of Algorithm 1 for convex and strongly convex losses. Both regimes share the same delay-compensated optimism and two-level gradient-tracking architecture,

differing only in the curvature correction and regularization scale. We set the common mixing coefficient and block length as

$$\theta \triangleq \frac{1}{1+\sqrt{1-\sigma_2^2(W)}} \quad \text{and} \quad B \triangleq \left\lceil \frac{\sqrt{2}\log(14\sqrt{N})}{(\sqrt{2}-1)\sqrt{\rho}} \right\rceil. \tag{4}$$

For convex losses, we set $\alpha = 0$ and $h = (2\eta)^{-1}$. DOFTRL-GT admits both gradient-variation and small-loss regret bounds.

Theorem 1. *Suppose that Assumptions 1 to 4 hold with $\alpha = 0$. Assume further that $f_{t,i} \geq 0$ for all $t \in [T]$ and $i \in [N]$. Each agent $i \in [N]$ runs Algorithm 1 with $h = (2\eta)^{-1}$, θ and B as defined in Equation* (4)*, and $\eta = \min\left\{\frac{1}{800L}, \max\left\{\frac{R}{\sqrt{108(G^2+BV_T)}}, R\sqrt{\frac{B}{800LF_T}}\right\}\right\}$. Then, the regret satisfies*

$$\mathrm{Reg}_T = \mathcal{O}\left(NBR(LR+G) + NR\min\left\{B^{3/2}\sqrt{V_T}, \sqrt{BLF_T}\right\}\right). \tag{5}$$

Substituting B from Equation (4) gives $\widetilde{\mathcal{O}}\big(N\min\left\{\rho^{-3/4}\sqrt{V_T}, \rho^{-1/4}\sqrt{F_T}\right\}\big)$, up to additive terms independent of T, V_T, and F_T. Thus, Theorem 1 replaces the $\sqrt{T}$ dependence in the worst-case bound of Wan et al. (2025) with $\sqrt{V_T}$ or $\sqrt{F_T}$. Moreover, the lower bounds in Section 4 show that the $\rho^{-3/4}\sqrt{V_T}$ and $\rho^{-1/4}\sqrt{F_T}$ dependences are optimal up to logarithmic factors in their respective problem-dependent regimes. The nonnegativity assumption is required only for the F_T guarantee.

We remark that while Theorem 1 requires tuning the stepsize based on the unknown quantities V_T and F_T, a doubling trick based on the observed block-gradient variation $\widehat{V}_Z$ removes this requirement (Rakhlin & Sridharan, 2013). Block gossip allows agents to agree on restarts using one additional scalar per message and no extra communication rounds. Appendix F shows the same run retains both regret bounds and their network dependence, without an extra logarithmic factor in T.

For strongly convex losses, we set $h = \alpha/2$. DOFTRL-GT then admits gradient-variation and small-loss regret bounds with logarithmic dependence on V_T and F_T.

Theorem 2. *Suppose that Assumptions 1 to 4 hold with $\alpha > 0$. Assume further that $f_{t,i} \geq 0$ for all $t \in [T]$ and $i \in [N]$. Each agent $i \in [N]$ runs Algorithm 1 with $h = \alpha/2$ and θ and B as defined in Equation* (4)*. Then, the regret satisfies*

$$\mathrm{Reg}_T = \mathcal{O}\left(\frac{NB(G+\alpha R)^2}{\alpha}\left[\frac{L}{\alpha} + \log\left(1 + \frac{1}{(G+\alpha R)^2}\min\left\{BV_T, \frac{LF_T}{B}\right\}\right)\right]\right).$$

Substituting B from Equation (4) gives $\widetilde{\mathcal{O}}\big(N\rho^{-1/2}\log\left(2+\min\left\{V_T, F_T\right\}\right)\big)$, up to additive terms independent of T, V_T, and F_T. Thus, Theorem 2 replaces the $\log T$ dependence in the worst-case bound of Wan et al. (2025) with logarithmic dependence on V_T or F_T. Moreover, the lower bounds in Section 4 show that the $\rho^{-1/2}$ network dependence and logarithmic dependence on V_T and F_T are optimal up to logarithmic factors in their respective problem-dependent regimes. The nonnegativity assumption is required only for the F_T guarantee.

3.2 PROOF SKETCH

In this subsection, we sketch the proof of Theorem 1 and Theorem 2, focusing on how delay-compensated optimism and two-level gradient tracking jointly yield the gradient-variation guarantee. The full proof is deferred to the appendix.

Convex losses. By Lemma 2 with $\alpha = 0$ and $u = u^\star$, we first decompose the regret into a linearized regret term and a network disagreement term:

$$\frac{\mathrm{Reg}_T(i)}{NB} \leq \sum_{z=1}^{Z}\langle g(z), x_i(z) - u^\star\rangle + \frac{L}{2N}\sum_{z=1}^{Z}\sum_{j=1}^{N}\|x_i(z) - x_j(z)\|_2^2. \tag{6}$$

We first control the linearized regret term in Equation (6). The update in Line 4 of Algorithm 1 implements delay-compensated optimism and can be equivalently written as an optimistic FTRL

update (Rakhlin & Sridharan, 2013; Farina et al., 2019) with hint $\widetilde{m}_i(z) = s_i(z-2) + 2y_i(z-2) - \sum_{\tau=1}^{z-1} g(\tau)$. Applying Lemma 4, the linearized regret term can be bounded as

$$\sum_{z=1}^{Z} \langle g(z), x_i(z) - u^\star \rangle \le \frac{R^2}{2\eta} + \eta \sum_{z=1}^{Z} \|g(z) - \widetilde{m}_i(z)\|_2^2 - \frac{M_{Z,i}}{4\eta}, \tag{7}$$

where $M_{Z,i} = \sum_{z=2}^{Z} \|x_i(z) - x_i(z-1)\|_2^2$ measures agent i's cumulative decision movement. Optimism thus replaces gradient magnitudes with prediction errors. While negative stability terms are standard in optimistic analyses with gradient-variation guarantees (Zhao et al., 2020; Yan et al., 2023), retaining the term $-M_{Z,i}/(4\eta)$ here is crucial because it will cancel the decision-movement term that arises when relating $\widehat{V}_Z$ to V_T.

To control the prediction error together with the network disagreement, define the block-gradient variation $\widehat{V}_Z \triangleq N^{-1} \sum_{z=1}^{Z} \sum_{j=1}^{N} \|a_j(z) - a_j(z-1)\|_2^2$, where $a_j(z)$ denotes agent j's local block gradient. We next show that both the prediction error and network disagreement are controlled by $\widehat{V}_Z$ through the two-level gradient tracking mechanism. Recall that $y_i(z)$ tracks the network-average block gradient $g(z)$ using local block-gradient differences, while $s_i(z)$ tracks its cumulative sum $G_z = \sum_{\tau=1}^{z} g(\tau)$ by mixing its previous state and adding the updated $y_i(z)$. The horizon-independent block length $B = \Theta(\rho^{-1/2} \log N)$ in Equation (4) ensures a per-block gossip contraction factor $q < 1$. Applying this contraction to the two recursions yields the following bounds on the tracking errors (Lemma 5):

$$\sum_{z=1}^{Z-1} \sum_{j=1}^{N} \|y_j(z) - g(z)\|_2^2 \le \frac{Nq^2}{(1-q)^2} \widehat{V}_Z \quad \text{and} \quad \sum_{z=1}^{Z-1} \sum_{j=1}^{N} \|s_j(z) - G_z\|_2^2 \le \frac{Nq^2}{(1-q)^4} \widehat{V}_Z. \tag{8}$$

The analysis of Wan et al. (2025) bounds the error in estimating the cumulative network-average gradient in terms of gradient magnitudes rather than their temporal differences. Consequently, such magnitude-based control does not by itself yield a bound that adapts to small temporal gradient variation. In contrast, our difference-based two-level gradient tracking mechanism controls the tracking errors through the block-gradient variation $\widehat{V}_Z$, which depends on consecutive block-gradient differences rather than gradient magnitudes.

We next relate these tracking errors to the prediction error and network disagreement. The prediction error $g(z) - \widetilde{m}_i(z)$ decomposes into consecutive gradient differences and the errors $y_i(z-2) - g(z-2)$ and $s_i(z-2) - G_{z-2}$. For network disagreement, the agents share the same regularizer, while the differences between their linear coefficients $s_i(z-2) + 2y_i(z-2)$ are determined solely by the tracking errors. Combining Equation (8) with the η-Lipschitz property of the decision update in Line 4 of Algorithm 1, Lemma 6 gives

$$\sum_{z=1}^{Z} \|g(z) - \widetilde{m}_i(z)\|_2^2 \le 24 \widehat{V}_Z \ \text{ and } \ \frac{1}{N} \sum_{z=1}^{Z} \sum_{j=1}^{N} \|x_i(z) - x_j(z)\|_2^2 \le 8\eta^2 \widehat{V}_Z. \tag{9}$$

Thus, both the prediction error and network disagreement are controlled by the same block-gradient variation. Substituting Equation (9) into Equations (6) and (7) yields

$$\frac{\mathrm{Reg}_T(i)}{NB} \le \frac{R^2}{2\eta} + (24\eta + 4L\eta^2) \widehat{V}_Z - \frac{M_{Z,i}}{4\eta}. \tag{10}$$

It remains to relate the auxiliary block-gradient variation $\widehat{V}_Z$ to the problem-dependent quantities V_T and F_T. Lemma 7 shows that $\widehat{V}_Z$ is controlled by the original gradient variation V_T together with the decision movement $M_{Z,i}$, since consecutive block gradients are evaluated at different decisions. For the small-loss branch, Lemma 8 uses the self-bounding property (Srebro et al., 2010, Lemma 3.1) to relate $\widehat{V}_Z$ to F_T. Under the prescribed stepsize, these two relations are

$$\widehat{V}_Z \le \frac{32}{29}\Big(G^2 + 2BV_T + 6L^2 M_{Z,i}\Big) \ \text{ and } \ \widehat{V}_Z \le 33L\Big(\frac{F_T}{B} + \frac{R^2}{2\eta}\Big). \tag{11}$$

Substituting the first bound into Equation (10) yields

$$\frac{\mathrm{Reg}_T(i)}{NB} \le \frac{R^2}{2\eta} + 54\eta(G^2 + BV_T) + \underbrace{\Big(162\eta L^2 - \frac{1}{4\eta}\Big)}_{\le 0} M_{Z,i}. \tag{12}$$

This bound makes explicit how the negative stability induced by optimism absorbs the effect of decision movement on the regret, leaving a guarantee that depends on the original gradient variation V_T rather than the horizon T.

For the small-loss branch, substituting the second bound in Equation (11) into Equation (10) yields $\mathrm{Reg}_T(i)/(NB) \le \mathcal{O}(R^2/\eta + \eta L F_T/B)$. Combining this with Equation (12) and substituting the prescribed stepsize and block length gives the claimed gradient-variation and small-loss guarantees.

Strongly convex losses. The analysis shares the delay-compensated optimism and two-level gradient tracking framework of the convex case.The block-index-dependent curvature of the FTRL objective yields a logarithmic dependence on $\widehat{V}_Z$ and a horizon-independent network disagreement bound (Lemma 9). It also bounds the cumulative decision movement independently of the horizon, allowing Lemma 10 to relate $\widehat{V}_Z$ to V_T without the negative-stability cancellation used in the convex case. For the small-loss branch, Lemma 11 relates $\widehat{V}_Z$ to F_T and $\mathrm{Reg}_T(i)$ via self-bounding. Substituting this relation into the regret bound logarithmic in $\widehat{V}_Z$ completes the proof of Theorem 2.

4 LOWER BOUNDS

In this section, we establish lower bounds in terms of gradient variation and small loss for convex and strongly convex losses. Throughout this section, let $N = 8a$ and $T = mN/4$, where $a \ge 2$ and $m \ge 1$ are integers. The constructions and full proofs are given in Appendix D.

4.1 GRADIENT-VARIATION LOWER BOUNDS

We first establish gradient-variation lower bounds for convex and strongly convex losses, characterizing the fundamental network dependence of gradient-variation adaptivity.

Theorem 3. *Let $R, G, L > 0$ and $v \ge 0$. There exist a decision set $\mathcal{X} \subset \mathbb{R}^{m+1}$ and a network $(\mathcal{G}, W)$ satisfying Assumptions 1 and 4, with spectral gap $\rho = \sin^2(\pi/N)$, such that, for any possibly randomized D-OCO algorithm and any fixed agent $i \in [N]$, there exists a sequence of linear losses satisfying Assumptions 2 and 3 with $\alpha = 0$ and $V_T \le v$ for which $\mathbb{E}[\mathrm{Reg}_T(i)] \ge \Omega(NR \min\{G\rho^{-1/4}\sqrt{T},\ G\rho^{-1/2} + \rho^{-3/4}\sqrt{v}\})$.*

Theorem 4. *Let $\alpha, R > 0$ and $M \in [m]$. There exist a decision set $\mathcal{X} \subset \mathbb{R}^2$ and a network $(\mathcal{G}, W)$ satisfying Assumptions 1 and 4, with spectral gap $\rho = \sin^2(\pi/N)$, such that, for any possibly randomized D-OCO algorithm and any fixed agent $i \in [N]$, there exists a sequence of nonnegative, α-strongly convex losses satisfying Assumptions 2 and 3 with $L = \alpha$ and $G = \alpha R$ for which $\frac{M}{2} \le 1 + \frac{V_T}{\alpha^2 R^2} \le \frac{25M}{16}$, and $\mathbb{E}[\mathrm{Reg}_T(i)] \ge \Omega\big(N\alpha R^2 \rho^{-1/2} \log\big(2 + \frac{V_T}{\alpha^2 R^2}\big)\big)$.*

Theorem 3 shows that, although the optimal worst-case convex regret has a $\rho^{-1/4}$ network factor (Wan et al., 2025), achieving gradient-variation adaptivity necessarily incurs the stronger $\rho^{-3/4}$ dependence in the variation-dependent regime. Thus, the spectral-gap dependence of our convex gradient-variation bound is optimal up to logarithmic factors. For strongly convex losses, Theorem 4 shows that the $\rho^{-1/2}$ network dependence remains necessary, matching our upper bound up to logarithmic factors.

4.2 SMALL-LOSS LOWER BOUNDS

We next establish small-loss lower bounds for convex and strongly convex losses, characterizing the fundamental network dependence of small-loss adaptivity.

Theorem 5. *Let $R, G, L > 0$, $f \ge 0$, and $\lambda = \min\{L, G/(2R)\}$. There exist a decision set $\mathcal{X} \subset \mathbb{R}^{m+1}$ and a network $(\mathcal{G}, W)$ satisfying Assumptions 1 and 4, with spectral gap $\rho = \sin^2(\pi/N)$, such that, for any possibly randomized D-OCO algorithm and any fixed agent $i \in [N]$, there exists a sequence of losses, nonnegative on $\mathbb{R}^{m+1}$, satisfying Assumptions 2 and 3 with $\alpha = 0$ and $F_T \le f$ for which $\mathbb{E}[\mathrm{Reg}_T(i)] \ge \Omega(NR[\lambda R\rho^{-1/2} + \rho^{-1/4}\min\{G\sqrt{T}, \sqrt{Lf}\}])$.*

Theorem 6. *Let $\alpha, R > 0$ and $M \in [m]$. There exist a decision set $\mathcal{X} \subset \mathbb{R}^2$ and a network $(\mathcal{G}, W)$ satisfying Assumptions 1 and 4, with spectral gap $\rho = \sin^2(\pi/N)$, such that, for any possibly*

randomized D-OCO algorithm and any fixed agent $i \in [N]$, there exists a sequence of nonnegative, α-strongly convex losses satisfying Assumptions 2 and 3 with $L = \alpha$ and $G = \alpha R$ for which $\frac{M}{32} \le \frac{\sqrt{\rho} F_T}{\alpha R^2} \le \frac{3\pi M}{128}$, and $\mathbb{E}[\mathrm{Reg}_T(i)] \ge \Omega(N\alpha R^2 \rho^{-1/2} \log(2 + \frac{\sqrt{\rho} F_T}{\alpha R^2}))$.

Theorems 5 and 6 establish small-loss regret lower bounds of $\Omega(N\rho^{-1/4} \min\{\sqrt{T}, \sqrt{F_T}\})$ and $\Omega(N\rho^{-1/2} \log(2 + \rho^{1/2} F_T))$ for convex and strongly convex losses, respectively. Together with the gradient-variation lower bounds, these results show that the dependence on the sequence difficulty and the spectral gap in our upper bounds is optimal up to logarithmic factors in the corresponding problem-dependent regimes. In particular, the convex gradient-variation and small-loss regimes necessarily exhibit the different network factors $\rho^{-3/4}$ and $\rho^{-1/4}$, respectively, while the strongly convex regimes retain the $\rho^{-1/2}$ dependence.

5 APPLICATIONS

We present two applications of Algorithm 1 and establish the corresponding theoretical guarantees. Formal statements and proofs appear in Appendix E.

Decentralized Offline Convex Optimization. We first examine the performance of Algorithm 1 in decentralized offline optimization, viewed as a special case of the online setting with fixed local losses $f_{t,i} = f_i$. Let $\overline{f}(x) \triangleq N^{-1} \sum_{j=1}^{N} f_j(x)$ denote the network-averaged objective. Since $V_T = 0$, the following corollary gives a convergence guarantee for each agent's ergodic average $\widehat{x}_{T,i} \triangleq T^{-1} \sum_{t=1}^{T} x_{t,i} = Z^{-1} \sum_{z=1}^{Z} x_i(z)$.

Corollary 7 (Informal)**.** *Under the conditions of Theorem 1 specialized to the setting above, Algorithm 1 satisfies, for every $i \in [N]$, $\overline{f}(\widehat{x}_{T,i}) - \min_{x \in \mathcal{X}} \overline{f}(x) = \widetilde{\mathcal{O}}\left(\rho^{-1/2} T^{-1}\right)$.*

On a fixed network, existing worst-case D-OCO regret bounds imply only $\mathcal{O}(T^{-1/2})$ ergodic convergence for stationary losses (Yan et al., 2012; Wan et al., 2024; 2025), leaving a gap to the $\mathcal{O}(1/T)$ rate of smooth decentralized offline optimization with gradient tracking (Qu & Li, 2017). Corollary 7 closes this gap within the online framework, without strong convexity or additional communication. More generally, our gradient-variation bound yields $\mathcal{O}(1/T)$ average regret whenever $V_T = \mathcal{O}(1)$.

Stochastically Extended Adversarial Model (SEA). The SEA model captures both stochastic uncertainty and adversarial variation (Sachs et al., 2022; Yan et al., 2023). At round t, the environment chooses a history-dependent joint distribution D_t of the local losses, with expected losses $F_{t,i}(x) \triangleq \mathbb{E}_{D_t}[f_{t,i}(x)]$. We quantify these effects by the cumulative stochastic variance $\sigma_{1:T}^2 \triangleq \mathbb{E}[\sum_{t=1}^{T} \frac{1}{N} \sum_{i=1}^{N} \sup_{x \in \mathcal{X}} \mathbb{E}_{D_t} \|\nabla f_{t,i}(x) - \nabla F_{t,i}(x)\|_2^2]$ and cumulative adversarial variation $\Sigma_{1:T}^2 \triangleq \mathbb{E}[\sum_{t=2}^{T} \frac{1}{N} \sum_{i=1}^{N} \sup_{x \in \mathcal{X}} \|\nabla F_{t,i}(x) - \nabla F_{t-1,i}(x)\|_2^2]$, where the outer expectations account for the random history. Corollary 8 bounds the expected regret in terms of these two quantities.

Corollary 8 (Informal)**.** *Suppose Assumptions 1 to 4 hold almost surely. With the parameter choices in Corollary 10, Algorithm 1 guarantees expected regret $\widetilde{\mathcal{O}}(N\sqrt{\sigma_{1:T}^2 + \rho^{-3/2}\Sigma_{1:T}^2})$ for convex losses and $\widetilde{\mathcal{O}}(N\rho^{-1/2} \log(2 + \sigma_{1:T}^2 + \Sigma_{1:T}^2))$ for strongly convex losses.*

The convex regret bound is stated up to an additive term independent of T, $\sigma_{1:T}^2$, and $\Sigma_{1:T}^2$. These guarantees capture both extremes: when $\Sigma_{1:T}^2 = 0$, their dependence on the SEA statistics reduces to stochastic variance alone, while $\sigma_{1:T}^2 = 0$ yields dependence on adversarial variation alone.

6 CONCLUSION

We proposed DOFTRL-GT by combining delay-compensated optimism with two-level gradient tracking and accelerated gossip. For convex and strongly convex losses, we established gradient-variation and small-loss regret bounds with nearly matching lower bounds in the stated regimes, revealing an intrinsic separation in network dependence for convex losses. Applications to offline optimization and stochastic–adversarial learning demonstrate the benefits of benign sequences. Future work includes extending these guarantees to dynamic regret and adaptive regret.

AI USE STATEMENT

We used generative AI tools to assist with polishing the manuscript and exploring potential proof strategies, with the AI-suggested proof strategies requiring substantial revision and further development by the authors. The writing assistance focused on refining grammar and improving clarity, readability, and the presentation of technical arguments, while proof-related assistance involved discussing possible approaches and checking the algebraic correctness of mathematical derivations. The authors led the formulation of the research questions, the design of the methods and algorithms, and the development of the main theoretical proofs, with AI tools serving a supporting role. The authors reviewed and revised all AI-assisted material, independently checked the arguments and proofs, and take full responsibility for the accuracy and integrity of the paper.

ETHICS STATEMENT

This paper focuses on algorithm design and regret analysis for decentralized online convex optimization. No human-subject research or personal-data processing is involved. We are not aware of any specific ethical concerns arising directly from the theoretical results presented in this work.

REPRODUCIBILITY STATEMENT

The problem setup and assumptions are stated in Section 2, and DOFTRL-GT and its parameter choices are specified in Section 3. The appendix provides auxiliary results and complete proofs of the regret bounds and application guarantees, including explicit constructions for the lower bounds. Together, these materials enable readers to independently check the derivations and reproduce the theoretical results.

REFERENCES

Baruch Awerbuch and Robert Kleinberg. Online linear optimization and adaptive routing. *Journal of Computer and System Sciences*, 74(1):97–114, 2008.

Sijia Chen, Yu-Jie Zhang, Wei-Wei Tu, Peng Zhao, and Lijun Zhang. Optimistic online mirror descent for bridging stochastic and adversarial online convex optimization. *Journal of Machine Learning Research*, 25(178):1–62, 2024.

Chao-Kai Chiang, Tianbao Yang, Chia-Jung Lee, Mehrdad Mahdavi, Chi-Jen Lu, Rong Jin, and Shenghuo Zhu. Online optimization with gradual variations. In *Proceedings of the 25th Annual Conference on Learning Theory*, volume 23, pp. 6.1–6.20. JMLR Workshop and Conference Proceedings, 2012.

John C Duchi, Alekh Agarwal, and Martin J Wainwright. Dual averaging for distributed optimization: Convergence analysis and network scaling. *IEEE Transactions on Automatic control*, 57(3): 592–606, 2011.

Gabriele Farina, Christian Kroer, and Tuomas Sandholm. Optimistic regret minimization for extensive-form games via dilated distance-generating functions. *Advances in neural information processing systems*, 32, 2019.

Genevieve E Flaspohler, Francesco Orabona, Judah Cohen, Soukayna Mouatadid, Miruna Oprescu, Paulo Orenstein, and Lester Mackey. Online learning with optimism and delay. In *International Conference on Machine Learning*, pp. 3363–3373. PMLR, 2021.

Saghar Hosseini, Airlie Chapman, and Mehran Mesbahi. Online distributed optimization via dual averaging. In *52nd IEEE Conference on Decision and Control*, pp. 1484–1489. IEEE, 2013.

Anastasiia Koloskova, Tao Lin, and Sebastian U Stich. An improved analysis of gradient tracking for decentralized machine learning. *Advances in Neural Information Processing Systems*, 34: 11422–11435, 2021.

Huan Li and Zhouchen Lin. Revisiting extra for smooth distributed optimization. *SIAM Journal on Optimization*, 30(3):1795–1821, 2020.

Ji Liu and A Stephen Morse. Accelerated linear iterations for distributed averaging. *Annual Reviews in Control*, 35(2):160–165, 2011.

Angelia Nedic and Asuman Ozdaglar. Distributed subgradient methods for multi-agent optimization. *IEEE Transactions on automatic control*, 54(1):48–61, 2009.

Guannan Qu and Na Li. Harnessing smoothness to accelerate distributed optimization. *IEEE Transactions on Control of Network Systems*, 5(3):1245–1260, 2017.

Alexander Rakhlin and Karthik Sridharan. Online learning with predictable sequences. In *Conference on Learning Theory*, pp. 993–1019. PMLR, 2013.

Sarah Sachs, Hédi Hadiji, Tim van Erven, and Cristóbal Guzmán. Between stochastic and adversarial online convex optimization: Improved regret bounds via smoothness. *Advances in Neural Information Processing Systems*, 35:691–702, 2022.

Shahin Shahrampour and Ali Jadbabaie. Distributed online optimization in dynamic environments using mirror descent. *IEEE Transactions on Automatic Control*, 63(3):714–725, 2017.

Nathan Srebro, Karthik Sridharan, and Ambuj Tewari. Smoothness, low noise and fast rates. *Advances in neural information processing systems*, 23, 2010.

Zhipeng Tu, Xi Wang, Yiguang Hong, Lei Wang, Deming Yuan, and Guodong Shi. Distributed online convex optimization with compressed communication. *Advances in Neural Information Processing Systems*, 35:34492–34504, 2022.

Yuanyu Wan. Black-box reductions for decentralized online convex optimization in changing environments. In *COLT*, pp. 5605–5631, 2025.

Yuanyu Wan, Tong Wei, Mingli Song, and Lijun Zhang. Nearly optimal regret for decentralized online convex optimization. In *The Thirty Seventh Annual Conference on Learning Theory*, pp. 4862–4888. PMLR, 2024.

Yuanyu Wan, Tong Wei, Bo Xue, Mingli Song, and Lijun Zhang. Optimal and efficient algorithms for decentralized online convex optimization. *Journal of Machine Learning Research*, 26(135): 1–43, 2025.

Guanghui Wang, Shiyin Lu, Yao Hu, and Lijun Zhang. Adapting to smoothness: A more universal algorithm for online convex optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 6162–6169, 2020.

Shuche Wang, Adarsh Barik, Peng Zhao, and Vincent Tan. Parameter-free algorithms for the stochastically extended adversarial model. *Advances in Neural Information Processing Systems*, 38:64068–64102, 2026.

Lin Xiao and Stephen Boyd. Fast linear iterations for distributed averaging. *Systems & Control Letters*, 53(1):65–78, 2004.

Yan-Feng Xie, Peng Zhao, and Zhi-Hua Zhou. Gradient-variation online learning under generalized smoothness. *Advances in Neural Information Processing Systems*, 37:37865–37899, 2024.

Feng Yan, Shreyas Sundaram, SVN Vishwanathan, and Yuan Qi. Distributed autonomous online learning: Regrets and intrinsic privacy-preserving properties. *IEEE Transactions on Knowledge and Data Engineering*, 25(11):2483–2493, 2012.

Yu-Hu Yan, Peng Zhao, and Zhi-Hua Zhou. Universal online learning with gradient variations: A multi-layer online ensemble approach. *Advances in Neural Information Processing Systems*, 36: 37682–37715, 2023.

Yu-Hu Yan, Peng Zhao, and Zhi-Hua Zhou. A simple and optimal approach for universal online learning with gradient variations. *Advances in Neural Information Processing Systems*, 37:11132–11163, 2024.

Tianbao Yang, Mehrdad Mahdavi, Rong Jin, and Shenghuo Zhu. Regret bounded by gradual variation for online convex optimization. *Machine learning*, 95(2):183–223, 2014.

Haishan Ye, Luo Luo, Ziang Zhou, and Tong Zhang. Multi-consensus decentralized accelerated gradient descent. *Journal of machine learning research*, 24(306):1–50, 2023.

Deming Yuan, Alexandre Proutiere, and Guodong Shi. Distributed online linear regressions. *IEEE Transactions on Information Theory*, 67(1):616–639, 2020.

Kun Yuan, Qing Ling, and Wotao Yin. On the convergence of decentralized gradient descent. *SIAM Journal on Optimization*, 26(3):1835–1854, 2016.

Peng Zhao, Yu-Jie Zhang, Lijun Zhang, and Zhi-Hua Zhou. Dynamic regret of convex and smooth functions. *Advances in Neural Information Processing Systems*, 33:12510–12520, 2020.

Peng Zhao, Yu-Jie Zhang, Lijun Zhang, and Zhi-Hua Zhou. Adaptivity and non-stationarity: Problem-dependent dynamic regret for online convex optimization. *Journal of Machine Learning Research*, 25(98):1–52, 2024.

Yuheng Zhao, Yu-Hu Yan, Kfir Y Levy, and Peng Zhao. Gradient-variation online adaptivity for accelerated optimization with hölder smoothness. *Advances in Neural Information Processing Systems*, 38:3749–3779, 2026.

A Auxiliary Results

A.1 Accelerated gossip

Given vectors $v_1, \dots, v_N \in \mathbb{R}^d$, initialize $v_i^{-1} = v_i^0 = v_i$ for every $i \in [N]$ and let $\overline{v} \triangleq N^{-1} \sum_{i=1}^N v_i$. Stack the local iterates and their common average row-wise as

$$V^k \triangleq \begin{bmatrix} (v_1^k)^\top \\ \vdots \\ (v_N^k)^\top \end{bmatrix} \in \mathbb{R}^{N \times d}, \quad \overline{V} \triangleq \begin{bmatrix} \overline{v}^\top \\ \vdots \\ \overline{v}^\top \end{bmatrix} \in \mathbb{R}^{N \times d}.$$

Then, the accelerated gossip iteration in Equation (2) can be written as

$$V^{k+1} = (1+\theta) W V^k - \theta V^{k-1}, \quad k \geq 0. \tag{13}$$

We first recall the contraction property of accelerated gossip, which will be used repeatedly in the tracking analysis. The following lemma restates Proposition 1 of Ye et al. (2023) in our notation.

Lemma 1. *Under Assumption 4, let θ be defined in Equation* (4)*. Then, for every integer $K \geq 1$, the iterates generated by Equation* (13) *satisfy*

$$\frac{1}{N} \mathbf{1}^\top V^K = \overline{v}^\top, \tag{14}$$

$$\|V^K - \overline{V}\|_{\mathrm{F}} \leq \sqrt{14} \left(1 - \left(1 - \frac{1}{\sqrt{2}}\right) \sqrt{\rho}\right)^K \|V^0 - \overline{V}\|_{\mathrm{F}}. \tag{15}$$

A.2 Regret Decomposition

The following lemma bounds the blockwise loss difference by a linear term, a curvature correction, and a network disagreement term. This decomposition will be used in the regret analysis for both convex and strongly convex losses.

Lemma 2. *Under Assumptions 1 to 3, for every block z, agent $i \in [N]$, and comparator $u \in \mathcal{X}$, the following inequality holds:*

$$\frac{1}{NB} \sum_{t \in \mathcal{I}_z} \sum_{j=1}^N \big(f_{t,j}(x_i(z)) - f_{t,j}(u)\big) \leq \langle g(z), x_i(z) - u \rangle + \frac{\alpha}{2} \big(\|x_i(z)\|_2^2 - \|u\|_2^2\big) + \frac{(L-\alpha)_+}{2N} \sum_{j=1}^N \|x_i(z) - x_j(z)\|_2^2.$$

Proof. We follow an analysis similar to that of Wan et al. (2025, Equations (41)–(43)), but use smoothness to obtain a quadratic disagreement term instead of the linear term arising from their Lipschitz argument. Fix a block z, an agent $i \in [N]$, and a comparator $u \in \mathcal{X}$. For every $j \in [N]$, it follows from Assumptions 2 and 3 that

$$\begin{aligned}
&\frac{1}{B} \sum_{t \in \mathcal{I}_z} \big(f_{t,j}(x_i(z)) - f_{t,j}(u)\big) \\
&\leq \langle a_j(z) + \alpha x_j(z), x_i(z) - u \rangle + \frac{L}{2} \|x_i(z) - x_j(z)\|_2^2 - \frac{\alpha}{2} \|x_j(z) - u\|_2^2 \qquad \text{(by Equation (3))} \\
&= \langle a_j(z), x_i(z) - u \rangle + \frac{\alpha}{2} \big(\|x_i(z)\|_2^2 - \|u\|_2^2\big) + \frac{L - \alpha}{2} \|x_i(z) - x_j(z)\|_2^2,
\end{aligned}$$

where the equality follows from the identity $\langle x_j(z), x_i(z) - u \rangle - \frac{1}{2}\|x_j(z) - u\|_2^2 = \frac{1}{2}\big(\|x_i(z)\|_2^2 - \|u\|_2^2 - \|x_i(z) - x_j(z)\|_2^2\big)$. Averaging over $j \in [N]$, using $g(z) = N^{-1} \sum_{j=1}^N a_j(z)$, and replacing $L - \alpha$ by $(L-\alpha)_+$ yields the desired bound. □

A.3 Optimistic FTRL Analysis

The following two lemmas collect standard optimistic FTRL regret bounds in the particular forms needed in our analysis. Lemma 3 follows by specializing the general optimistic FTRL regret decomposition of Chen et al. (2024, Lemma 15). Lemma 4 is a direct consequence of the RVU bound for optimistic FTRL in Farina et al. (2019, Theorem 2), specialized to the Euclidean distance-generating function $d(x) = \frac{1}{2}\|x\|^2$. We provide the details below to make the correspondence between the cited results and our notation explicit.

Lemma 3. *Suppose that Assumption 1 holds. Let $Z \geq 1$ be an integer, let $\alpha \geq 0$ and $h > 0$, and let $g(1), \dots, g(Z), m_1, \dots, m_Z \in \mathbb{R}^d$ be arbitrary vectors. Set $G_0 \triangleq 0$ and define $\ell_z(x) \triangleq \langle g(z), x\rangle + \frac{\alpha}{2}\|x\|_2^2$, $G_z \triangleq \sum_{\tau=1}^{z} g(\tau)$, and $x_z \triangleq \arg\min_{x\in\mathcal{X}}\{\langle G_{z-1} + m_z, x\rangle + (\frac{z\alpha}{2} + h)\|x\|_2^2\}$ for every $z \in [Z]$. Then, the following inequality holds:*

$$\sum_{z=1}^{Z}\big(\ell_z(x_z) - \ell_z(u)\big) \leq h\|u\|_2^2 + \sum_{z=1}^{Z}\frac{\|g(z) - m_z\|_2^2}{2(z\alpha + 2h)}, \quad \forall u \in \mathcal{X}. \tag{16}$$

Proof. This is a direct specialization of the standard optimistic FTRL bound; see, e.g., Chen et al. (2024, Lemma 15). Specifically, take $q_z(x) = \langle g(z), x\rangle$, $\widetilde{q}_z(x) = \langle m_z, x\rangle$, and $\Psi_z(x) = (\frac{z\alpha}{2} + h)\|x\|_2^2$ for $z \in [Z]$, which recover the update in the lemma. For the terminal terms in the cited bound, set $\Psi_{Z+1} = \Psi_Z$ and $\widetilde{q}_{Z+1} = 0$; these choices do not affect the decisions $x_1, \dots, x_Z$. Since Ψ_z is $(z\alpha + 2h)$-strongly convex, the cited bound and Young's inequality give

$$\begin{aligned}\sum_{z=1}^{Z}\big(q_z(x_z) - q_z(u)\big) &\leq \Psi_Z(u) - \Psi_1(x_1) - \frac{\alpha}{2}\sum_{z=2}^{Z}\|x_z\|_2^2 + \sum_{z=1}^{Z}\frac{\|g(z) - m_z\|_2^2}{2(z\alpha + 2h)} \\ &= \left(\frac{Z\alpha}{2} + h\right)\|u\|_2^2 - \frac{\alpha}{2}\sum_{z=1}^{Z}\|x_z\|_2^2 - h\|x_1\|_2^2 + \sum_{z=1}^{Z}\frac{\|g(z) - m_z\|_2^2}{2(z\alpha + 2h)}.\end{aligned}$$

Using $\ell_z(x) = q_z(x) + \frac{\alpha}{2}\|x\|_2^2$ and dropping the nonpositive term $-h\|x_1\|_2^2$ yields Equation (16). □

Lemma 4. *Suppose that Assumption 1 holds. Let $Z \geq 1$ be an integer, let $\eta > 0$, and let $g(1), \dots, g(Z), m_1, \dots, m_Z \in \mathbb{R}^d$ be arbitrary vectors. Set $G_0 \triangleq 0$ and define $G_z \triangleq \sum_{\tau=1}^{z} g(\tau)$ and $x_z \triangleq \arg\min_{x\in\mathcal{X}}\left\{\langle G_{z-1} + m_z, x\rangle + \frac{1}{2\eta}\|x\|_2^2\right\}$ for every $z \in [Z]$. Then, the following inequality holds:*

$$\sum_{z=1}^{Z}\langle g(z), x_z - u\rangle \leq \frac{R^2}{2\eta} + \eta\sum_{z=1}^{Z}\|g(z) - m_z\|_2^2 - \frac{1}{4\eta}\sum_{z=2}^{Z}\|x_z - x_{z-1}\|_2^2, \quad \forall u \in \mathcal{X}. \tag{17}$$

Proof. The result follows directly from the RVU bound for optimistic FTRL in Farina et al. (2019, Theorem 2). Take the distance-generating function $d(x) = \frac{1}{2}\|x\|_2^2$. Then the optimistic FTRL update in their theorem becomes

$$x_z = \arg\min_{x\in\mathcal{X}}\left\{\langle G_{z-1} + m_z, x\rangle + \frac{1}{2\eta}\|x\|_2^2\right\},$$

which is exactly our update. By Assumption 1, $0 \in \mathcal{X}$ and $\|x\|_2 \leq R$, and hence $\Delta_d \triangleq \max_{x,y\in\mathcal{X}}\{d(x) - d(y)\} \leq \frac{R^2}{2}$. Therefore, Farina et al. (2019, Theorem 2) gives

$$\sum_{z=1}^{Z}\langle g(z), x_z - u\rangle \leq \frac{R^2}{2\eta} + \eta\sum_{z=1}^{Z}\|g(z) - m_z\|_2^2 - \frac{1}{4\eta}\sum_{z=2}^{Z}\|x_z - x_{z-1}\|_2^2.$$

□

B Omitted Proof Details for Theorem 1

In this section, we provide the proof of Theorem 1. We first bound the tracking errors, prediction errors, and network disagreement in terms of the block-gradient variation $\widehat{V}_Z$, and then relate this quantity to V_T and F_T.

Recall that the two-level gradient tracking mechanism in Algorithm 1 uses $y_i(z)$ and $s_i(z)$ to track the network-average block gradient $g(z)$ and its cumulative sum $G_z = \sum_{\tau=1}^{z} g(\tau)$, respectively. For each agent $i \in [N]$, define the corresponding tracking errors as $e_i(z) \triangleq y_i(z) - g(z)$ and $r_i(z) \triangleq s_i(z) - G_z$. Lemma 5 bounds the cumulative squared norms of these errors in terms of the block-gradient variation, which we recall for convenience:

$$\widehat{V}_Z = \frac{1}{N}\sum_{z=1}^{Z}\sum_{i=1}^{N}\|a_i(z) - a_i(z-1)\|_2^2,$$

where $a_i(0) = 0$. We use the convention $g(z) = G_z = e_i(z) = r_i(z) = 0$ for $z \le 0$.

Lemma 5. *Under Assumptions 1 and 4, let $\alpha \ge 0$ and $h > 0$, let θ and B be defined in Equation (4), and set $q \triangleq \sqrt{14}(1 - (1 - \frac{1}{\sqrt{2}})\sqrt{\rho})^B$ and $\eta \triangleq (2h)^{-1}$. Then, the states generated by Algorithm 1 satisfy*

$$\sum_{z=1}^{Z-1}\sum_{i=1}^{N}\|y_i(z) - g(z)\|_2^2 \le \frac{Nq^2}{(1-q)^2}\widehat{V}_Z, \quad \sum_{z=1}^{Z-1}\sum_{i=1}^{N}\|s_i(z) - G_z\|_2^2 \le \frac{Nq^2}{(1-q)^4}\widehat{V}_Z, \tag{18}$$

Proof. For every $z \in \{1, \dots, Z-1\}$, Algorithm 1 forms $y_i(z)$ and $s_i(z)$ during block $z+1$. In this block, the communication states are initialized as $u_i^{-1} = u_i^0 = y_i(z-1) + a_i(z) - a_i(z-1)$ and $v_i^{-1} = v_i^0 = s_i(z-1)$, and the outputs satisfy $y_i(z) = u_i^B$ and $s_i(z) = v_i^B + y_i(z)$. By the average-preservation property in Lemma 1 and the zero initialization, induction gives

$$\begin{aligned}\frac{1}{N}\sum_{i=1}^{N} y_i(z) &= \frac{1}{N}\sum_{i=1}^{N}\big(y_i(z-1) + a_i(z) - a_i(z-1)\big) = g(z),\\ \frac{1}{N}\sum_{i=1}^{N} s_i(z) &= \frac{1}{N}\sum_{i=1}^{N}\big(s_i(z-1) + y_i(z)\big) = G_z.\end{aligned} \tag{19}$$

We first bound the block-gradient tracking errors $y_i(z) - g(z)$. Applying Lemma 1, we have

$$\begin{aligned}\Big(\sum_{i=1}^{N}\|y_i(z) - g(z)\|_2^2\Big)^{\frac{1}{2}} &= \Big(\sum_{i=1}^{N}\|e_i(z)\|_2^2\Big)^{\frac{1}{2}}\\ &\le q\Big(\sum_{i=1}^{N}\|e_i(z-1) + a_i(z) - a_i(z-1) - (g(z) - g(z-1))\|_2^2\Big)^{\frac{1}{2}}.\end{aligned}$$

Applying Minkowski's inequality to the preceding bound, we further obtain

$$\begin{aligned}\Big(\sum_{i=1}^{N}\|e_i(z)\|_2^2\Big)^{\frac{1}{2}} &\le q\Big(\sum_{i=1}^{N}\|e_i(z-1)\|_2^2\Big)^{\frac{1}{2}} + q\Big(\sum_{i=1}^{N}\|a_i(z) - a_i(z-1) - (g(z) - g(z-1))\|_2^2\Big)^{\frac{1}{2}}\\ &= q\Big(\sum_{i=1}^{N}\|e_i(z-1)\|_2^2\Big)^{\frac{1}{2}} + q\Big(\sum_{i=1}^{N}\|a_i(z) - a_i(z-1)\|_2^2 - N\|g(z) - g(z-1)\|_2^2\Big)^{\frac{1}{2}}\\ &\le q\Big(\sum_{i=1}^{N}\|e_i(z-1)\|_2^2\Big)^{\frac{1}{2}} + q\Big(\sum_{i=1}^{N}\|a_i(z) - a_i(z-1)\|_2^2\Big)^{\frac{1}{2}},\end{aligned} \tag{20}$$

where the equality follows from $g(z) - g(z-1) = N^{-1}\sum_{i=1}^{N}(a_i(z) - a_i(z-1))$. We next bound the sum of squared errors over $z \in \{1, \dots, Z-1\}$. Combining Equation (20) with $e_i(0) = 0$ and the definition of $\widehat{V}_Z$ yields

$$\Big(\sum_{z=1}^{Z-1}\sum_{i=1}^{N}\|e_i(z)\|_2^2\Big)^{\frac{1}{2}} \le \Big[\sum_{z=1}^{Z-1}\Big(q\Big(\sum_{i=1}^{N}\|e_i(z-1)\|_2^2\Big)^{\frac{1}{2}} + q\Big(\sum_{i=1}^{N}\|a_i(z) - a_i(z-1)\|_2^2\Big)^{\frac{1}{2}}\Big)^2\Big]^{\frac{1}{2}}$$

$$\leq q\Big(\sum_{z=1}^{Z-1}\sum_{i=1}^{N}\|e_i(z-1)\|_2^2\Big)^{\frac{1}{2}} + q\Big(\sum_{z=1}^{Z-1}\sum_{i=1}^{N}\|a_i(z)-a_i(z-1)\|_2^2\Big)^{\frac{1}{2}}$$

(by Minkowski's inequality)

$$\leq q\Big(\sum_{z=1}^{Z-1}\sum_{i=1}^{N}\|e_i(z)\|_2^2\Big)^{\frac{1}{2}} + q\sqrt{N\widehat{V}_Z}.$$

Since the choice of B ensures $q<1$, rearranging the preceding inequality gives

$$\Big(\sum_{z=1}^{Z-1}\sum_{i=1}^{N}\|e_i(z)\|_2^2\Big)^{\frac{1}{2}} \leq \frac{q}{1-q}\sqrt{N\widehat{V}_Z}. \tag{21}$$

Squaring both sides of Equation (21) and substituting $e_i(z)=y_i(z)-g(z)$ proves the first inequality in Equation (18).

Next, we next bound the cumulative-gradient tracking errors $r_i(z)$. By the identities $s_i(z)=v_i^B+y_i(z)$ and $G_z=G_{z-1}+g(z)$, we have

$$\begin{aligned}\Big(\sum_{i=1}^{N}\|r_i(z)\|_2^2\Big)^{\frac{1}{2}} &= \Big(\sum_{i=1}^{N}\|v_i^B-G_{z-1}+e_i(z)\|_2^2\Big)^{\frac{1}{2}}\\ &\leq \Big(\sum_{i=1}^{N}\|v_i^B-G_{z-1}\|_2^2\Big)^{\frac{1}{2}} + \Big(\sum_{i=1}^{N}\|e_i(z)\|_2^2\Big)^{\frac{1}{2}}\\ &\leq q\Big(\sum_{i=1}^{N}\|r_i(z-1)\|_2^2\Big)^{\frac{1}{2}} + \Big(\sum_{i=1}^{N}\|e_i(z)\|_2^2\Big)^{\frac{1}{2}}.\end{aligned} \tag{22}$$

The last inequality follows from Lemma 1, since $v_i^0=s_i(z-1)$ and the average of these initial states is G_{z-1} by Equation (19). Similarly, combining Equation (22) with $r_i(0)=0$ and Equation (21), we obtain

$$\Big(\sum_{z=1}^{Z-1}\sum_{i=1}^{N}\|r_i(z)\|_2^2\Big)^{\frac{1}{2}} \leq q\Big(\sum_{z=1}^{Z-1}\sum_{i=1}^{N}\|r_i(z-1)\|_2^2\Big)^{\frac{1}{2}} + \Big(\sum_{z=1}^{Z-1}\sum_{i=1}^{N}\|e_i(z)\|_2^2\Big)^{\frac{1}{2}}$$

(by Minkowski's inequality)

$$\leq q\Big(\sum_{z=1}^{Z-1}\sum_{i=1}^{N}\|r_i(z)\|_2^2\Big)^{\frac{1}{2}} + \frac{q}{1-q}\sqrt{N\widehat{V}_Z}.$$

By rearranging the preceding inequality, squaring both sides, and substituting $r_i(z)=s_i(z)-G_z$, we obtain the second inequality in Equation (18), completing the proof. □

To identify the corresponding optimistic FTRL prediction, we express the linear coefficient of the decision update in Line 4 of Algorithm 1 as

$$s_i(z-2)+2y_i(z-2) = G_{z-2}+2g(z-2)+r_i(z-2)+2e_i(z-2) = \sum_{\tau=1}^{z-1}g(\tau)+\widetilde{m}_i(z),$$

where $\widetilde{m}_i(z)\triangleq 2g(z-2)-g(z-1)+r_i(z-2)+2e_i(z-2)$ serves as the optimistic prediction of $g(z)$ in the analysis. With this representation, we next analyze the optimistic prediction error and the network disagreement. The following lemma shows that both are bounded in terms of the same block-gradient variation $\widehat{V}_Z$.

Lemma 6. *Under the conditions of Lemma 5, the states generated by Algorithm 1 satisfy*

$$\sum_{z=1}^{Z}\|g(z)-\widetilde{m}_i(z)\|_2^2 \leq 24\widehat{V}_Z, \tag{23}$$

$$\frac{1}{N}\sum_{z=1}^{Z}\sum_{j=1}^{N}\Big\|x_i(z)-x_j(z)\Big\|_2^2 \leq 8\eta^2\widehat{V}_Z, \quad \forall i\in[N]. \tag{24}$$

Proof. We first decompose the prediction error $g(z) - \widetilde{m}_i(z)$ into consecutive block-gradient differences and tracking errors. Using the definitions of $\widetilde{m}_i(z)$, $r_i(z)$, and $e_i(z)$, we obtain

$$\sum_{z=1}^{Z} \|g(z)-\widetilde{m}_i(z)\|_2^2 = \sum_{z=1}^{Z} \Big\| (g(z)-g(z-1)) + 2(g(z-1)-g(z-2)) - (r_i(z-2)+2e_i(z-2)) \Big\|_2^2$$
$$\leq 4\sum_{z=1}^{Z} \|g(z)-g(z-1)\|_2^2 + 16\sum_{z=1}^{Z} \|g(z-1)-g(z-2)\|_2^2 + 2\sum_{z=1}^{Z} \|r_i(z-2)+2e_i(z-2)\|_2^2. \tag{25}$$

Here, the last inequality follows by applying $\|a+b\|_2^2 \leq 2\|a\|_2^2 + 2\|b\|_2^2$ successively.

The first two terms in Equation (25) are controlled by the block-gradient variation $\widehat{V}_Z$. Since $g(-1) = g(0) = 0$, we have

$$\sum_{z=1}^{Z} \|g(z-1) - g(z-2)\|_2^2 \leq \sum_{z=1}^{Z} \|g(z) - g(z-1)\|_2^2$$
$$= \sum_{z=1}^{Z} \Big\| \frac{1}{N} \sum_{j=1}^{N} (a_j(z) - a_j(z-1)) \Big\|_2^2 \qquad \text{(by Equation (3))}$$
$$\leq \frac{1}{N} \sum_{z=1}^{Z} \sum_{j=1}^{N} \|a_j(z) - a_j(z-1)\|_2^2 = \widehat{V}_Z. \tag{26}$$

To bound the contribution of the tracking errors in Equation (25), note that the choice of B in Equation (4) ensures that the per-block gossip contraction factor q satisfies

$$q \leq \sqrt{14} \exp\Big(-\Big(1 - \frac{1}{\sqrt{2}}\Big) B\sqrt{\rho} \Big) \qquad \text{(by } 1 - x \leq e^{-x}\text{)}$$
$$\leq \sqrt{14} \exp\Big(-\log(14\sqrt{N}) \Big) = \frac{1}{\sqrt{14N}} < \frac{1}{3}, \tag{27}$$

where the last strict inequality follows from $N \geq 2$. By Lemma 5, we further have

$$\sum_{\tau=1}^{Z-1} \sum_{j=1}^{N} \|r_j(\tau) + 2e_j(\tau)\|_2^2 \leq \Big[\Big(\sum_{\tau=1}^{Z-1} \sum_{j=1}^{N} \|r_j(\tau)\|_2^2 \Big)^{\frac{1}{2}} + 2\Big(\sum_{\tau=1}^{Z-1} \sum_{j=1}^{N} \|e_j(\tau)\|_2^2 \Big)^{\frac{1}{2}} \Big]^2$$
(by Minkowski's inequality)
$$\leq Nq^2 \Big(\frac{1}{(1-q)^2} + \frac{2}{1-q} \Big)^2 \widehat{V}_Z \qquad \text{(by Equation (18))}$$
$$\leq \frac{1}{14} \Big(\frac{9}{4} + 3 \Big)^2 \widehat{V}_Z = \frac{63}{32} \widehat{V}_Z \leq 2\widehat{V}_Z, \tag{28}$$

where the last inequality chain uses $Nq^2 \leq 1/14$ and $q < 1/3$ from Equation (27). The zero initialization then gives

$$\sum_{z=1}^{Z} \|r_i(z-2) + 2e_i(z-2)\|_2^2 \leq \sum_{\tau=1}^{Z-1} \sum_{j=1}^{N} \|r_j(\tau) + 2e_j(\tau)\|_2^2 \leq 2\widehat{V}_Z. \tag{29}$$

Substituting Equation (29) and Equation (26) into Equation (25) proves Equation (23).

We next bound the network disagreement in Equation (24). Applying the Lipschitz bound for regularized projections in Duchi et al. (2011, Lemma 2) to the decision update in Line 4 of Algorithm 1, with $\psi(x) = \|x\|_2^2/2$ and parameter $(z\alpha + 2h)^{-1}$, gives, for every $z \in [Z]$ and $i, j \in [N]$,

$$\|x_i(z) - x_j(z)\|_2 \leq \frac{\|s_i(z-2) + 2y_i(z-2) - s_j(z-2) - 2y_j(z-2)\|_2}{z\alpha + 2h}$$
$$\leq \eta \|r_i(z-2) + 2e_i(z-2) - r_j(z-2) - 2e_j(z-2)\|_2, \tag{30}$$

where the last inequality uses $(z\alpha + 2h)^{-1} \leq \eta$ and the definitions of the tracking errors.

Squaring Equation (30), summing over blocks $z \in [Z]$, and averaging over agents $j \in [N]$ yields

$$\begin{aligned}
&\frac{1}{N}\sum_{z=1}^{Z}\sum_{j=1}^{N}\|x_i(z)-x_j(z)\|_2^2 \\
&\le \frac{\eta^2}{N}\sum_{z=1}^{Z}\sum_{j=1}^{N}\|r_i(z-2)+2e_i(z-2)-r_j(z-2)-2e_j(z-2)\|_2^2 \\
&\le 2\eta^2\sum_{z=1}^{Z}\|r_i(z-2)+2e_i(z-2)\|_2^2+\frac{2\eta^2}{N}\sum_{z=1}^{Z}\sum_{j=1}^{N}\|r_j(z-2)+2e_j(z-2)\|_2^2 \\
&\qquad\qquad\qquad

Since $\alpha = 0$, adding and subtracting the previous block's averaged gradient at $x_j(z)$ yields

$$\begin{aligned}
&\|a_j(z) - a_j(z-1)\|^2 \\
=&\Big\|a_j(z) - \frac{1}{B}\sum_{t\in\mathcal{I}_{z-1}} \nabla f_{t,j}(x_j(z)) + \frac{1}{B}\sum_{t\in\mathcal{I}_{z-1}} \nabla f_{t,j}(x_j(z)) - a_j(z-1)\Big\|^2 \\
\le& 2\Big\|a_j(z) - \frac{1}{B}\sum_{t\in\mathcal{I}_{z-1}} \nabla f_{t,j}(x_j(z))\Big\|^2 + 2\Big\|\frac{1}{B}\sum_{t\in\mathcal{I}_{z-1}} \nabla f_{t,j}(x_j(z)) - a_j(z-1)\Big\|^2 \\
\le& 2\sup_{x\in\mathcal{X}}\Big\|\frac{1}{B}\sum_{t\in\mathcal{I}_z} \nabla f_{t,j}(x) - \frac{1}{B}\sum_{t\in\mathcal{I}_{z-1}} \nabla f_{t,j}(x)\Big\|^2 + 2L^2\|x_j(z) - x_j(z-1)\|^2, \qquad (34)
\end{aligned}$$

where the last inequality follows from the definition of $a_j(z)$ and Assumption 3.

Moreover, Assumption 2 implies $\|a_j(1)\| \le B^{-1}\sum_{t\in\mathcal{I}_1}\|\nabla f_{t,j}(x_j(1))\| \le G$. Using this bound and $a_j(0) = 0$, one can derive

$$\begin{aligned}
\widehat{V}_Z &= \frac{1}{N}\sum_{j=1}^N \|a_j(1)\|^2 + \frac{1}{N}\sum_{z=2}^Z\sum_{j=1}^N \|a_j(z) - a_j(z-1)\|^2 \\
&\le G^2 + \frac{2}{N}\sum_{z=2}^Z\sum_{j=1}^N \sup_{x\in\mathcal{X}}\Big\|\frac{1}{B}\sum_{t\in\mathcal{I}_z} \nabla f_{t,j}(x) - \frac{1}{B}\sum_{t\in\mathcal{I}_{z-1}} \nabla f_{t,j}(x)\Big\|^2 \\
&\quad + \frac{2L^2}{N}\sum_{z=2}^Z\sum_{j=1}^N \|x_j(z) - x_j(z-1)\|^2 \qquad \text{(by Equation (34))} \\
&\le G^2 + 2BV_T + \frac{2L^2}{N}\sum_{z=2}^Z\sum_{j=1}^N \|x_j(z) - x_j(z-1)\|^2, \qquad (35)
\end{aligned}$$

where the last inequality uses Equation (33).

We next bound the last term on the right-hand side of Equation (35) as follows:

$$\begin{aligned}
&\frac{1}{N}\sum_{z=2}^Z\sum_{j=1}^N \|x_j(z) - x_j(z-1)\|^2 \\
=&\frac{1}{N}\sum_{z=2}^Z\sum_{j=1}^N \Big\|(x_j(z) - x_i(z)) + (x_i(z) - x_i(z-1)) + (x_i(z-1) - x_j(z-1))\Big\|^2 \\
\le& 3\sum_{z=2}^Z \|x_i(z) - x_i(z-1)\|^2 + \frac{3}{N}\sum_{z=2}^Z\sum_{j=1}^N \|x_j(z) - x_i(z)\|^2 \\
&+ \frac{3}{N}\sum_{z=2}^Z\sum_{j=1}^N \|x_i(z-1) - x_j(z-1)\|^2 \qquad \text{(by } \|\textstyle\sum_{k=1}^3 v_k\|^2 \le 3\sum_{k=1}^3\|v_k\|^2\text{)} \\
\le& 3\sum_{z=2}^Z \|x_i(z) - x_i(z-1)\|^2 + \frac{6}{N}\sum_{z=1}^Z\sum_{j=1}^N \|x_j(z) - x_i(z)\|^2 \\
\le& 3\sum_{z=2}^Z \|x_i(z) - x_i(z-1)\|^2 + 48\eta^2\widehat{V}_Z. \qquad \text{(by Equation (24))}
\end{aligned}$$

Substituting this inequality into Equation (35) gives

$$\widehat{V}_Z \le G^2 + 2BV_T + 6L^2\sum_{z=2}^Z \|x_i(z) - x_i(z-1)\|^2 + 96\eta^2L^2\widehat{V}_Z.$$

Since $\eta L \le 1/32$, we have $96\eta^2 L^2 \le 96/32^2 = 3/32$. Rearranging therefore yields

$$\frac{29}{32}\widehat{V}_Z \le (1 - 96\eta^2 L^2)\widehat{V}_Z \le G^2 + 2BV_T + 6L^2 \sum_{z=2}^{Z} \|x_i(z) - x_i(z-1)\|^2,$$

which completes the proof. □

For the small-loss guarantee, the following lemma relates the block-gradient variation $\widehat{V}_Z$ to the small loss F_T.

Lemma 8. *Suppose that Assumptions 1 to 4 hold with $\alpha = 0$. Assume further that $f_{t,i} \ge 0$ for all $t \in [T]$ and $i \in [N]$. Each agent $i \in [N]$ runs Algorithm 1 with $h = (2\eta)^{-1}$, θ and B as defined in Equation* (4)*, and $0 < \eta \le 1/(800L)$. Then, it holds that*

$$\widehat{V}_Z \le 33L\left(\frac{F_T}{B} + \frac{R^2}{2\eta}\right).$$

Proof. By Jensen's inequality and the self-bounding property of nonnegative smooth losses (Srebro et al., 2010, Lemma 3.1), Assumption 3 implies

$$\left\|\frac{1}{B}\sum_{t\in\mathcal{I}_z} \nabla f_{t,j}(x)\right\|_2^2 \le \frac{1}{B}\sum_{t\in\mathcal{I}_z} \|\nabla f_{t,j}(x)\|_2^2 \le \frac{4L}{B}\sum_{t\in\mathcal{I}_z} f_{t,j}(x), \quad x \in \mathcal{X}. \tag{36}$$

Fix an agent $i \in [N]$. Since $\alpha = 0$, the weighted squared triangle inequality and Assumption 3 give

$$\begin{aligned}
\|a_j(z)\|^2 &= \left\|\frac{1}{B}\sum_{t\in\mathcal{I}_z} \nabla f_{t,j}(x_i(z)) + \frac{1}{B}\sum_{t\in\mathcal{I}_z} \big(\nabla f_{t,j}(x_j(z)) - \nabla f_{t,j}(x_i(z))\big)\right\|^2 \\
&\le \frac{33}{32}\left\|\frac{1}{B}\sum_{t\in\mathcal{I}_z} \nabla f_{t,j}(x_i(z))\right\|^2 + 33\left\|\frac{1}{B}\sum_{t\in\mathcal{I}_z} \big(\nabla f_{t,j}(x_j(z)) - \nabla f_{t,j}(x_i(z))\big)\right\|^2 \\
&\le \frac{33L}{8B}\sum_{t\in\mathcal{I}_z} f_{t,j}(x_i(z)) + 33L^2\|x_j(z) - x_i(z)\|^2.
\end{aligned} \tag{37}$$

Combining Lemmas 2 and 4 with $u = u^\star$, and dropping the nonpositive stability term, we obtain

$$\begin{aligned}
&\frac{1}{NB}\sum_{z=1}^{Z}\sum_{j=1}^{N}\sum_{t\in\mathcal{I}_z} \big(f_{t,j}(x_i(z)) - f_{t,j}(u^\star)\big) \\
&\le \frac{R^2}{2\eta} + \eta\sum_{z=1}^{Z} \|g(z) - \widetilde{m}_i(z)\|_2^2 + \frac{L}{2N}\sum_{z=1}^{Z}\sum_{j=1}^{N} \|x_i(z) - x_j(z)\|_2^2 \\
&\le \frac{R^2}{2\eta} + (24\eta + 4L\eta^2)\widehat{V}_Z,
\end{aligned} \tag{38}$$

where the last inequality follows from Lemma 6. Using $a_j(0) = 0$ and the network disagreement bound in Lemma 6, we obtain

$$\begin{aligned}
\widehat{V}_Z &= \frac{1}{N}\sum_{z=1}^{Z}\sum_{j=1}^{N} \|a_j(z) - a_j(z-1)\|^2 \\
&\le \frac{4}{N}\sum_{z=1}^{Z}\sum_{j=1}^{N} \|a_j(z)\|^2 \\
&\le \frac{33L}{2NB}\sum_{z=1}^{Z}\sum_{j=1}^{N}\sum_{t\in\mathcal{I}_z} f_{t,j}(x_i(z)) + \frac{132L^2}{N}\sum_{z=1}^{Z}\sum_{j=1}^{N} \|x_j(z) - x_i(z)\|^2 \qquad \text{(by Equation (37))}
\end{aligned}$$

$$\leq \frac{33L}{2NB}\sum_{z=1}^{Z}\sum_{j=1}^{N}\sum_{t\in\mathcal{I}_z} f_{t,j}(x_i(z)) + 1056\eta^2L^2\widehat{V}_Z \qquad \text{(by Equation (24))}$$

$$= \frac{33L}{2}\left[\frac{F_T}{B} + \frac{1}{NB}\sum_{z=1}^{Z}\sum_{j=1}^{N}\sum_{t\in\mathcal{I}_z}\left(f_{t,j}(x_i(z)) - f_{t,j}(u^\star)\right)\right] + 1056\eta^2L^2\widehat{V}_Z$$

$$\leq \frac{33L}{2}\left(\frac{F_T}{B} + \frac{R^2}{2\eta}\right) + \left(396\eta L + 1122\eta^2L^2\right)\widehat{V}_Z,$$

where the last equality follows from $\frac{1}{NB}\sum_{z=1}^{Z}\sum_{j=1}^{N}\sum_{t\in\mathcal{I}_z} f_{t,j}(u^\star) = F_T/B$, and the last inequality follows from Equation (38). Since $\eta L \leq 1/800$, we can further obtain

$$396\eta L + 1122\eta^2L^2 \leq \frac{396}{800} + \frac{1122}{800^2} < \frac{1}{2}.$$

Moving this term to the left completes the proof of the lemma. □

With the preceding ingredients in place, we are ready to prove the regret guarantee for convex losses. For completeness, we restate Theorem 1 and combine the gradient-variation and small-loss arguments developed above.

Theorem 1. *Suppose that Assumptions 1 to 4 hold with $\alpha = 0$. Assume further that $f_{t,i} \geq 0$ for all $t \in [T]$ and $i \in [N]$. Each agent $i \in [N]$ runs Algorithm 1 with $h = (2\eta)^{-1}$, θ and B as defined in Equation* (4)*, and $\eta = \min\left\{\frac{1}{800L}, \max\left\{\frac{R}{\sqrt{108(G^2+BV_T)}}, R\sqrt{\frac{B}{800LF_T}}\right\}\right\}$. Then, the regret satisfies*

$$\mathrm{Reg}_T = \mathcal{O}\left(NBR(LR+G) + NR\min\left\{B^{3/2}\sqrt{V_T}, \sqrt{BLF_T}\right\}\right). \tag{5}$$

Proof. We first derive a regret bound in terms of $\widehat{V}_Z$ that holds for every $\eta > 0$. Since $x_{t,i} = x_i(z)$ for $t \in \mathcal{I}_z$, we have, for every $i \in [N]$,

$$\frac{\mathrm{Reg}_T(i)}{NB} = \frac{1}{NB}\sum_{z=1}^{Z}\sum_{j=1}^{N}\sum_{t\in\mathcal{I}_z}\left(f_{t,j}(x_i(z)) - f_{t,j}(u^\star)\right)$$

$$\leq \sum_{z=1}^{Z}\langle g(z), x_i(z) - u^\star\rangle + \frac{L}{2N}\sum_{z=1}^{Z}\sum_{j=1}^{N}\|x_i(z) - x_j(z)\|_2^2, \tag{39}$$

where the inequality follows from Lemma 2 with $\alpha = 0$. We bound the two terms on the right-hand side separately.

For the linearized regret term, the decision update in Line 4 of Algorithm 1 can be written as the optimistic FTRL update in Lemma 4, with $x_z = x_i(z)$ and $m_z = \widetilde{m}_i(z)$. Applying this lemma with $u = u^\star$, we obtain

$$\sum_{z=1}^{Z}\langle g(z), x_i(z) - u^\star\rangle \leq \frac{R^2}{2\eta} + \eta\sum_{z=1}^{Z}\|g(z) - \widetilde{m}_i(z)\|_2^2 - \frac{1}{4\eta}\sum_{z=2}^{Z}\|x_i(z) - x_i(z-1)\|_2^2$$

$$\leq \frac{R^2}{2\eta} + 24\eta\widehat{V}_Z - \frac{1}{4\eta}\sum_{z=2}^{Z}\|x_i(z) - x_i(z-1)\|_2^2, \tag{40}$$

where the last inequality follows from Equation (23) in Lemma 6.

The second term in Equation (39) measures the disagreement between agents' decisions. The bounds on $y_j(z) - g(z)$ and $s_j(z) - G_z$ in Lemma 5 control this disagreement through Equation (24) in Lemma 6. Consequently,

$$\frac{L}{2N}\sum_{z=1}^{Z}\sum_{j=1}^{N}\|x_i(z) - x_j(z)\|_2^2 \leq 4L\eta^2\widehat{V}_Z. \tag{41}$$

Substituting Equations (40) and (41) into Equation (39) yields

$$\frac{\mathrm{Reg}_T(i)}{NB} \le \frac{R^2}{2\eta} + (24\eta + 4L\eta^2)\widehat{V}_Z - \frac{1}{4\eta}\sum_{z=2}^{Z}\|x_i(z) - x_i(z-1)\|_2^2. \tag{42}$$

We first derive the gradient-variation bound. Since $\eta \le 1/(800L) \le 1/(32L)$, substituting Lemma 7 into Equation (42) yields

$$\begin{aligned}\frac{\mathrm{Reg}_T(i)}{NB} &\le \frac{R^2}{2\eta} + \frac{32}{29}(24\eta + 4L\eta^2)(G^2 + 2BV_T) \\ &\quad + \left[\frac{192L^2}{29}(24\eta + 4L\eta^2) - \frac{1}{4\eta}\right]\sum_{z=2}^{Z}\|x_i(z) - x_i(z-1)\|_2^2 \\ &\le \frac{R^2}{2\eta} + 27\eta(G^2 + 2BV_T) + \left(162\eta L^2 - \frac{1}{4\eta}\right)\sum_{z=2}^{Z}\|x_i(z) - x_i(z-1)\|_2^2 \\ &\le \frac{R^2}{2\eta} + 54\eta(G^2 + BV_T). \end{aligned} \tag{43}$$

The second inequality uses $\frac{32}{29}(24+4\eta L) \le \frac{772}{29} < 27$. The last inequality follows from $162\eta^2L^2 \le 162/32^2 < 1/4$ and $G^2 + 2BV_T \le 2(G^2 + BV_T)$.

For the small-loss bound, applying Lemma 8 to Equation (42) and dropping the nonpositive stability term, we can obtain

$$\begin{aligned}\frac{\mathrm{Reg}_T(i)}{NB} &\le \frac{R^2}{2\eta} + 33L(24\eta + 4L\eta^2)\left(\frac{F_T}{B} + \frac{R^2}{2\eta}\right) \\ &\le \frac{R^2}{2\eta} + \frac{800\eta L F_T}{B} + 400LR^2 \\ &\le \frac{R^2}{\eta} + \frac{800\eta L F_T}{B}. \end{aligned} \tag{44}$$

Here, the second inequality uses $33(24 + 4\eta L) \le 33(24 + 1/200) < 800$, and the last follows from $400LR^2 \le R^2/(2\eta)$.

Both Equation (43) and Equation (44) hold for the stepsize specified in the theorem. Consequently,

$$\begin{aligned}\frac{\mathrm{Reg}_T(i)}{NB} &\le \min\left\{\frac{R^2}{2\eta} + 54\eta(G^2 + BV_T), \frac{R^2}{\eta} + \frac{800\eta L F_T}{B}\right\} \\ &\le \frac{R^2}{\eta} + \eta\min\left\{54(G^2 + BV_T), \frac{800LF_T}{B}\right\}. \end{aligned} \tag{45}$$

The prescribed stepsize satisfies $\frac{1}{\eta} = \max\{800L, R^{-1}\min\{\sqrt{108(G^2 + BV_T)}, \sqrt{800LF_T/B}\}\}$, where inverse-square-root terms with zero denominators in the original stepsize are interpreted as $+\infty$. Therefore, we have

$$\frac{R^2}{\eta} \le 800LR^2 + R\min\left\{\sqrt{108(G^2 + BV_T)}, \sqrt{\frac{800LF_T}{B}}\right\}. \tag{46}$$

Moreover,

$$\begin{aligned}\eta\min\left\{54(G^2 + BV_T), \frac{800LF_T}{B}\right\} &\le \eta\min\left\{108(G^2 + BV_T), \frac{800LF_T}{B}\right\} \\ &\le R\min\left\{\sqrt{108(G^2 + BV_T)}, \sqrt{\frac{800LF_T}{B}}\right\}. \end{aligned} \tag{47}$$

Substituting Equation (46) and Equation (47) into Equation (45) yields

$$\mathrm{Reg}_T(i) \le 800NBLR^2 + 2NBR\min\left\{\sqrt{108(G^2 + BV_T)}, \sqrt{\frac{800LF_T}{B}}\right\}$$

$$\leq 800NBLR^2 + 2\sqrt{108}\,NBRG + 2NR\min\left\{\sqrt{108}\,B^{3/2}\sqrt{V_T}, \sqrt{800BLF_T}\right\}$$
$$= \mathcal{O}\Big(NBR(LR+G) + NR\min\left\{B^{3/2}\sqrt{V_T}, \sqrt{BLF_T}\right\}\Big).$$

The second inequality uses $\sqrt{G^2 + BV_T} \leq G + \sqrt{BV_T}$ and $\min\{a+b, c\} \leq a + \min\{b, c\}$ for nonnegative a, b, c. Taking the maximum over $i \in [N]$ proves Theorem 1. □

C Omitted Proof Details for Theorem 2

We now turn to the strongly convex case, where the curvature of the FTRL objective increases with the block index. The following lemma uses this growing curvature to bound the weighted sum of squared prediction errors logarithmically in $\widehat{V}_Z$ and the network disagreement independently of the horizon.

Lemma 9. *Under Assumptions 1, 2 and 4 with $\alpha > 0$, each agent $i \in [N]$ runs Algorithm 1 with $h = \alpha/2$ and θ and B as defined in Equation* (4)*. Then, for every $i \in [N]$, it holds that*

$$\sum_{z=1}^{Z} \frac{\|g(z) - \widetilde{m}_i(z)\|_2^2}{2\alpha(z+1)} \leq \frac{50(G+\alpha R)^2}{\alpha} \log\left(1 + \frac{6\widehat{V}_Z}{25(G+\alpha R)^2}\right), \tag{48}$$

$$\frac{1}{N}\sum_{z=1}^{Z}\sum_{j=1}^{N} \|x_i(z) - x_j(z)\|_2^2 \leq \frac{64(G+\alpha R)^2}{\alpha^2}. \tag{49}$$

Proof. We first derive uniform bounds on $y_i(z) - g(z)$ and $s_i(z) - G_z$, which will be combined with the increasing curvature $\alpha(z+1)$ of the FTRL objective. By the definition of $a_i(z)$,

$$\|a_i(z)\|_2 = \Big\|\frac{1}{B}\sum_{t\in\mathcal{I}_z} \nabla f_{t,i}(x_i(z)) - \alpha x_i(z)\Big\|_2 \leq G + \alpha R,$$

where the last inequality follows from Assumptions 1 and 2. Moreover, $\|g(z)\|_2 \leq N^{-1}\sum_{j=1}^{N}\|a_j(z)\|_2 \leq G + \alpha R$. For every $z \in \{1, \ldots, Z-1\}$, iterating Equation (20) from the zero initialization, one can derive

$$\Big(\sum_{j=1}^{N}\|y_j(z) - g(z)\|_2^2\Big)^{\frac{1}{2}} \leq 2\sqrt{N}(G+\alpha R)\sum_{\tau=1}^{z} q^{z-\tau+1}$$
$$\leq \frac{2q\sqrt{N}}{1-q}(G+\alpha R) \leq G + \alpha R.$$

The last inequality follows from Equation (27), since $q\sqrt{N} \leq 1/\sqrt{14}$ and $q < 1/3$ imply $2q\sqrt{N}/(1-q) \leq 3/\sqrt{14} < 1$. Similarly, Equation (22) implies

$$\Big(\sum_{j=1}^{N}\|s_j(z) - G_z\|_2^2\Big)^{\frac{1}{2}} \leq \sum_{\tau=1}^{z} q^{z-\tau}\Big(\sum_{j=1}^{N}\|y_j(\tau) - g(\tau)\|_2^2\Big)^{\frac{1}{2}}$$
$$\leq \frac{G+\alpha R}{1-q} \leq 2(G+\alpha R).$$

Consequently, for every $i \in [N]$ and $z \in \{0, \ldots, Z-1\}$, it follows that

$$\|r_i(z) + 2e_i(z)\|_2 = \|s_i(z) - G_z + 2(y_i(z) - g(z))\|_2 \leq 4(G+\alpha R). \tag{50}$$

To establish Equation (48), the decomposition used in Equation (25) and the bound in Equation (50) imply

$$\|g(z) - \widetilde{m}_i(z)\|_2 \leq \|g(z) - g(z-1)\|_2 + 2\|g(z-1) - g(z-2)\|_2 + \|r_i(z-2) + 2e_i(z-2)\|_2$$
$$\leq 10(G+\alpha R), \quad z \in [Z]. \tag{51}$$

We next combine the uniform bound in Equation (51) with the cumulative bound in Equation (23) to obtain logarithmic dependence on $\widehat{V}_Z$ under the weights $1/(z+1)$. For a fixed agent $i \in [N]$, set $b_z \triangleq \|g(z) - \widetilde{m}_i(z)\|_2^2/[100(G+\alpha R)^2]$. By Equation (51), $0 \le b_z \le 1$, and hence $\sum_{\tau=1}^{z} b_\tau \le z$. Therefore,

$$\sum_{z=1}^{Z} \frac{b_z}{z+1} \le \sum_{z=1}^{Z} \frac{b_z}{1+\sum_{\tau=1}^{z} b_\tau} \le \sum_{z=1}^{Z} \log\left(\frac{1+\sum_{\tau=1}^{z} b_\tau}{1+\sum_{\tau=1}^{z-1} b_\tau}\right)$$
$$= \log\left(1+\sum_{z=1}^{Z} b_z\right) \le \log\left(1+\frac{6\widehat{V}_Z}{25(G+\alpha R)^2}\right),$$

where the second inequality uses $x/(1+x) \le \log(1+x)$ with $x = b_z/(1+\sum_{\tau=1}^{z-1} b_\tau)$, and the last follows from Equation (23). Multiplying the preceding bound by $50(G+\alpha R)^2/\alpha$ proves Equation (48).

We next establish Equation (49). Unlike the convex case, the curvature of the FTRL objective increases with the block index, reducing the effect of differences between the agents' linear terms. Using $z\alpha + 2h = \alpha(z+1)$ in the first inequality of Equation (30) and the definitions of $r_i(z)$ and $e_i(z)$, we obtain

$$\|x_i(z) - x_j(z)\|_2 \le \frac{\|r_i(z-2) + 2e_i(z-2) - r_j(z-2) - 2e_j(z-2)\|_2}{\alpha(z+1)}$$
$$\le \frac{8(G+\alpha R)}{\alpha(z+1)}, \quad \forall z \in [Z], \quad i,j \in [N], \tag{52}$$

where the last inequality follows from Equation (50). The resulting $1/(z+1)$ decay makes the squared decision differences summable over blocks, leading to a bound independent of both Z and $\widehat{V}_Z$. Squaring both sides of Equation (52), summing over blocks $z \in [Z]$, and averaging over agents $j \in [N]$, we obtain

$$\frac{1}{N}\sum_{z=1}^{Z}\sum_{j=1}^{N} \|x_i(z) - x_j(z)\|_2^2 \le \frac{64(G+\alpha R)^2}{\alpha^2} \sum_{z=1}^{Z} \frac{1}{(z+1)^2}$$
$$\le \frac{64(G+\alpha R)^2}{\alpha^2} \sum_{z=1}^{Z} \left(\frac{1}{z} - \frac{1}{z+1}\right).$$

Since $\sum_{z=1}^{Z}(1/z - 1/(z+1)) = 1 - 1/(Z+1) \le 1$, the preceding bound establishes Equation (49) and completes the proof. □

We next relate $\widehat{V}_Z$ to the gradient variation V_T in the strongly convex case. The increasing curvature of the FTRL objective allows the contribution of decision movement to be bounded independently of the horizon.

Lemma 10. *Under the conditions of Lemma 9, assume further that Assumption 3 holds. Then, the following inequality holds:*

$$\widehat{V}_Z \le G^2 + 2BV_T + \frac{2(L+\alpha)^2(13G+14\alpha R)^2}{\alpha^2}.$$

Proof. As in the proof of Lemma 7, we first decompose $\widehat{V}_Z$ using $a_j(0) = 0$:

$$\widehat{V}_Z = \frac{1}{N}\sum_{j=1}^{N} \|a_j(1)\|_2^2 + \frac{1}{N}\sum_{z=2}^{Z}\sum_{j=1}^{N} \|a_j(z) - a_j(z-1)\|_2^2. \tag{53}$$

Since $x_j(1) = 0$ by $s_j(-1) = y_j(-1) = 0$ and $0 \in \mathcal{X}$, the correction $-\alpha x_j(1)$ vanishes. Assumption 2 therefore implies $\|a_j(1)\|_2 \le G$ and hence $N^{-1}\sum_{j=1}^{N} \|a_j(1)\|_2^2 \le G^2$.

We next bound the second term in Equation (53). To account for the curvature correction $-\alpha x_j(z)$, we adapt the derivation of Equation (34) in Lemma 7. For every $z \in \{2, \dots, Z\}$ and $j \in [N]$, we obtain

$$\begin{aligned}
&\|a_j(z) - a_j(z-1)\|_2^2 \\
&\le 2\Big\|\frac{1}{B}\sum_{t\in\mathcal{I}_z}\nabla f_{t,j}(x_j(z)) - \frac{1}{B}\sum_{t\in\mathcal{I}_{z-1}}\nabla f_{t,j}(x_j(z))\Big\|_2^2 \\
&\quad + 2\Big\|\frac{1}{B}\sum_{t\in\mathcal{I}_{z-1}}\big(\nabla f_{t,j}(x_j(z)) - \nabla f_{t,j}(x_j(z-1))\big) - \alpha\big(x_j(z) - x_j(z-1)\big)\Big\|_2^2 \\
&\le 2\sup_{x\in\mathcal{X}}\Big\|\frac{1}{B}\sum_{t\in\mathcal{I}_z}\nabla f_{t,j}(x) - \frac{1}{B}\sum_{t\in\mathcal{I}_{z-1}}\nabla f_{t,j}(x)\Big\|_2^2 + 2(L+\alpha)^2\|x_j(z) - x_j(z-1)\|_2^2, \qquad (54)
\end{aligned}$$

where the last inequality follows from Assumption 3 and the triangle inequality. Thus, the curvature correction replaces the factor L^2 in the decision-movement term of Equation (34) by $(L+\alpha)^2$.

On the other hand, the correction $-\alpha x$ cancels when the two block-averaged gradients are evaluated at the same point x. Therefore, the bound in Equation (33) remains unchanged. Substituting Equation (54) into Equation (53) and applying Equation (33), we obtain

$$\widehat{V}_Z \le G^2 + 2BV_T + \frac{2(L+\alpha)^2}{N}\sum_{z=2}^{Z}\sum_{j=1}^{N}\|x_j(z) - x_j(z-1)\|_2^2. \qquad (55)$$

We next bound the cumulative decision movement for each agent. For every $i \in [N]$ and $z \in \{2, \dots, Z\}$, the change in the linear coefficient of the decision update satisfies

$$\begin{aligned}
&\|s_i(z-2) + 2y_i(z-2) - s_i(z-3) - 2y_i(z-3)\|_2 \\
&= \|3g(z-2) - 2g(z-3) + r_i(z-2) + 2e_i(z-2) - r_i(z-3) - 2e_i(z-3)\|_2 \\
&\le 3\|g(z-2)\|_2 + 2\|g(z-3)\|_2 + \|r_i(z-2) + 2e_i(z-2)\|_2 + \|r_i(z-3) + 2e_i(z-3)\|_2 \\
&\le 13(G + \alpha R), \qquad (56)
\end{aligned}$$

where the equality follows from the definitions of $r_i(z)$ and $e_i(z)$, and the last inequality follows from Equation (50) and $\|g(z)\|_2 \le G + \alpha R$, with all quantities at nonpositive indices equal to zero.

The decision objectives at blocks $z-1$ and z have curvatures αz and $\alpha(z+1)$, respectively. Combining the first-order optimality conditions for these two updates, we obtain

$$\begin{aligned}
&\alpha z\|x_i(z) - x_i(z-1)\|_2^2 \\
&\le \langle s_i(z-3) + 2y_i(z-3) - s_i(z-2) - 2y_i(z-2) - \alpha x_i(z), x_i(z) - x_i(z-1)\rangle \\
&\le \big(\|s_i(z-2) + 2y_i(z-2) - s_i(z-3) - 2y_i(z-3)\|_2 + \alpha\|x_i(z)\|_2\big)\,\|x_i(z) - x_i(z-1)\|_2 \\
&\le (13G + 14\alpha R)\|x_i(z) - x_i(z-1)\|_2, \qquad \text{(by Equation (56) and Assumption 1)}
\end{aligned}$$

Dividing the above inequality by $\alpha z\|x_i(z) - x_i(z-1)\|_2$ when the decision difference is nonzero yields $\|x_i(z) - x_i(z-1)\|_2 \le (13G + 14\alpha R)/(\alpha z)$, and the bound holds trivially otherwise. Thus, the increasing curvature makes the squared decision movements summable, and for every $i \in [N]$,

$$\begin{aligned}
\sum_{z=2}^{Z}\|x_i(z) - x_i(z-1)\|_2^2 &\le \frac{(13G + 14\alpha R)^2}{\alpha^2}\sum_{z=2}^{Z}\frac{1}{z^2} \\
&\le \frac{(13G + 14\alpha R)^2}{\alpha^2}\sum_{z=2}^{Z}\left(\frac{1}{z-1} - \frac{1}{z}\right) \\
&\le \frac{(13G + 14\alpha R)^2}{\alpha^2}, \qquad (57)
\end{aligned}$$

where $\sum_{z=2}^{Z}(1/(z-1) - 1/z) = 1 - 1/Z \le 1$. Averaging Equation (57) over agents and substituting into Equation (55) completes the proof. □

In the strongly convex case, the following lemma relates $\widehat{V}_Z$ to F_T and $\mathrm{Reg}_T(i)$.

Lemma 11. *Under the conditions of Lemma 10, assume further that $f_{t,i} \geq 0$ for all $t \in [T]$ and $i \in [N]$. Then, for every $i \in [N]$, the following inequality holds:*

$$\widehat{V}_Z \leq 64L\left(\frac{F_T}{B} + \frac{\mathrm{Reg}_T(i)}{NB}\right) + \frac{3072L^2(G+\alpha R)^2}{\alpha^2} + 2(13G+14\alpha R)^2.$$

Proof. To apply the self-bounding estimate Equation (36) used in the proof of Lemma 8, we first separate the curvature correction $-\alpha(x_j(z) - x_j(z-1))$ from $a_j(z) - a_j(z-1)$. For every $z \in \{2, \ldots, Z\}$ and $j \in [N]$, the definition of $a_j(z)$ implies

$$\begin{aligned}
&\|a_j(z) - a_j(z-1)\|_2^2 \\
&\leq 2\Big\|\frac{1}{B}\sum_{t\in\mathcal{I}_z}\nabla f_{t,j}(x_j(z)) - \frac{1}{B}\sum_{t\in\mathcal{I}_{z-1}}\nabla f_{t,j}(x_j(z-1))\Big\|_2^2 + 2\alpha^2\|x_j(z) - x_j(z-1)\|_2^2 \\
&\leq 4\Big\|\frac{1}{B}\sum_{t\in\mathcal{I}_z}\nabla f_{t,j}(x_j(z))\Big\|_2^2 + 4\Big\|\frac{1}{B}\sum_{t\in\mathcal{I}_{z-1}}\nabla f_{t,j}(x_j(z-1))\Big\|_2^2 + 2\alpha^2\|x_j(z) - x_j(z-1)\|_2^2,
\end{aligned} \tag{58}$$

where both inequalities follow from $\|u+v\|_2^2 \leq 2\|u\|_2^2 + 2\|v\|_2^2$. Using $x_j(1) = 0$ and the decomposition in Equation (53), we sum Equation (58) over blocks and agents to obtain

$$\begin{aligned}
\widehat{V}_Z &\leq \frac{8}{N}\sum_{z=1}^{Z}\sum_{j=1}^{N}\Big\|\frac{1}{B}\sum_{t\in\mathcal{I}_z}\nabla f_{t,j}(x_j(z))\Big\|_2^2 + \frac{2\alpha^2}{N}\sum_{z=2}^{Z}\sum_{j=1}^{N}\|x_j(z) - x_j(z-1)\|_2^2 \\
&\leq \frac{32L}{NB}\sum_{z=1}^{Z}\sum_{j=1}^{N}\sum_{t\in\mathcal{I}_z} f_{t,j}(x_j(z)) + \frac{2\alpha^2}{N}\sum_{z=2}^{Z}\sum_{j=1}^{N}\|x_j(z) - x_j(z-1)\|_2^2,
\end{aligned} \tag{59}$$

where the last inequality follows by applying the self-bounding estimate for block-averaged gradients in Equation (36).

We next compare $f_{t,j}(x_j(z))$ with $f_{t,j}(x_i(z))$ to express the accumulated losses in terms of $\mathrm{Reg}_T(i)$. For every $i, j \in [N]$, $z \in [Z]$, and $t \in \mathcal{I}_z$, Assumption 3 implies

$$\begin{aligned}
f_{t,j}(x_j(z)) &\leq f_{t,j}(x_i(z)) + \langle\nabla f_{t,j}(x_i(z)), x_j(z) - x_i(z)\rangle + \frac{L}{2}\|x_j(z) - x_i(z)\|_2^2 \\
&\leq f_{t,j}(x_i(z)) + \frac{1}{4L}\|\nabla f_{t,j}(x_i(z))\|_2^2 + \frac{3L}{2}\|x_j(z) - x_i(z)\|_2^2 \quad \text{(by Young's inequality)} \\
&\leq 2f_{t,j}(x_i(z)) + \frac{3L}{2}\|x_j(z) - x_i(z)\|_2^2,
\end{aligned} \tag{60}$$

where the last inequality follows from the self-bounding property (Srebro et al., 2010, Lemma 3.1).

Substituting Equation (60) into Equation (59), we obtain, for every $i \in [N]$,

$$\begin{aligned}
\widehat{V}_Z &\leq \frac{64L}{NB}\sum_{z=1}^{Z}\sum_{j=1}^{N}\sum_{t\in\mathcal{I}_z} f_{t,j}(x_i(z)) + \frac{48L^2}{N}\sum_{z=1}^{Z}\sum_{j=1}^{N}\|x_j(z) - x_i(z)\|_2^2 \\
&\quad + \frac{2\alpha^2}{N}\sum_{z=2}^{Z}\sum_{j=1}^{N}\|x_j(z) - x_j(z-1)\|_2^2 \\
&= 64L\Big(\frac{F_T}{B} + \frac{\mathrm{Reg}_T(i)}{NB}\Big) + \frac{48L^2}{N}\sum_{z=1}^{Z}\sum_{j=1}^{N}\|x_j(z) - x_i(z)\|_2^2 + \frac{2\alpha^2}{N}\sum_{z=2}^{Z}\sum_{j=1}^{N}\|x_j(z) - x_j(z-1)\|_2^2 \\
&\leq 64L\Big(\frac{F_T}{B} + \frac{\mathrm{Reg}_T(i)}{NB}\Big) + \frac{3072L^2(G+\alpha R)^2}{\alpha^2} + 2(13G+14\alpha R)^2,
\end{aligned}$$

where the equality follows from the definitions of F_T and $\mathrm{Reg}_T(i)$, and the last inequality follows from Equations (49) and (57). This completes the proof. □

We now restate Theorem 2 and provide its proof.

Theorem 2. *Suppose that Assumptions 1 to 4 hold with $\alpha > 0$. Assume further that $f_{t,i} \geq 0$ for all $t \in [T]$ and $i \in [N]$. Each agent $i \in [N]$ runs Algorithm 1 with $h = \alpha/2$ and θ and B as defined in Equation* (4)*. Then, the regret satisfies*

$$\mathrm{Reg}_T = \mathcal{O}\left(\frac{NB(G+\alpha R)^2}{\alpha}\left[\frac{L}{\alpha} + \log\left(1 + \frac{1}{(G+\alpha R)^2}\min\left\{BV_T, \frac{LF_T}{B}\right\}\right)\right]\right).$$

Proof. We first derive a regret bound in terms of $\widehat{V}_Z$, following the argument for Theorem 1. For $\alpha > 0$, we retain the quadratic term in Lemma 2 and apply Lemma 3 with $u = u^\star$ and $h = \alpha/2$ to obtain, for every $i \in [N]$,

$$\frac{\mathrm{Reg}_T(i)}{NB} \leq \frac{\alpha R^2}{2} + \sum_{z=1}^{Z}\frac{\|g(z) - \widetilde{m}_i(z)\|_2^2}{2\alpha(z+1)} + \frac{(L-\alpha)_+}{2N}\sum_{z=1}^{Z}\sum_{j=1}^{N}\|x_i(z) - x_j(z)\|_2^2, \tag{61}$$

where $\|u^\star\|_2 \leq R$ by Assumption 1. Substituting Equations (48) and (49) from Lemma 9 into Equation (61), we can further obtain

$$\frac{\mathrm{Reg}_T(i)}{NB} \leq \frac{\alpha R^2}{2} + \frac{32(L-\alpha)_+(G+\alpha R)^2}{\alpha^2} + \frac{50(G+\alpha R)^2}{\alpha}\log\Big(1 + \frac{6\widehat{V}_Z}{25(G+\alpha R)^2}\Big). \tag{62}$$

We next derive the gradient-variation bound. Substituting Lemma 10 into Equation (62), we obtain

$$\begin{aligned}\frac{\mathrm{Reg}_T(i)}{NB} &\leq \frac{33L(G+\alpha R)^2}{\alpha^2} + \frac{50(G+\alpha R)^2}{\alpha}\log\Big[\frac{512L^2}{\alpha^2}\Big(1 + \frac{BV_T}{(G+\alpha R)^2}\Big)\Big] \\ &= \mathcal{O}\Big(\frac{(G+\alpha R)^2}{\alpha}\Big[\frac{L}{\alpha} + \log\Big(1 + \frac{BV_T}{(G+\alpha R)^2}\Big)\Big]\Big),\end{aligned} \tag{63}$$

where we use $13G + 14\alpha R \leq 14(G + \alpha R)$ and $\alpha \leq L$. The latter follows from Assumptions 2 and 3 unless $\mathcal{X}$ is a singleton, in which case the regret is zero.

For the small-loss bound, it suffices to consider $\mathrm{Reg}_T(i) > 0$. Substituting Lemma 11 into Equation (62) and separating the terms inside the logarithm yields

$$\begin{aligned}\frac{\mathrm{Reg}_T(i)}{NB} \leq{}& \frac{33L(G+\alpha R)^2}{\alpha^2} + \frac{50(G+\alpha R)^2}{\alpha}\log\Big(\frac{1024L^2}{\alpha^2}\Big) \\ &+ \frac{50(G+\alpha R)^2}{\alpha}\log\Big(1 + \frac{LF_T}{B(G+\alpha R)^2}\Big) \\ &+ \frac{50(G+\alpha R)^2}{\alpha}\log\Big(1 + \frac{384L}{25(G+\alpha R)^2}\frac{\mathrm{Reg}_T(i)}{NB}\Big).\end{aligned} \tag{64}$$

To control the last term in Equation (64), we further bound

$$\log\Big(1 + \frac{384L}{25(G+\alpha R)^2}\frac{\mathrm{Reg}_T(i)}{NB}\Big) \leq \frac{\alpha}{100(G+\alpha R)^2}\frac{\mathrm{Reg}_T(i)}{NB} + \log\Big(1 + \frac{1536L}{\alpha}\Big),$$

where the inequality follows from $\log(1 + ab) \leq a + \log(1 + b)$ for $a, b \geq 0$. Substituting into Equation (64) and rearranging, we obtain

$$\begin{aligned}\frac{\mathrm{Reg}_T(i)}{NB} \leq{}& \frac{66L(G+\alpha R)^2}{\alpha^2} + \frac{100(G+\alpha R)^2}{\alpha}\log\Big[\frac{1024L^2}{\alpha^2}\Big(1 + \frac{1536L}{\alpha}\Big)\Big] \\ &+ \frac{100(G+\alpha R)^2}{\alpha}\log\Big(1 + \frac{LF_T}{B(G+\alpha R)^2}\Big) \\ ={}& \mathcal{O}\Big(\frac{(G+\alpha R)^2}{\alpha}\Big[\frac{L}{\alpha} + \log\Big(1 + \frac{LF_T}{B(G+\alpha R)^2}\Big)\Big]\Big).\end{aligned} \tag{65}$$

The last step uses $L/\alpha \geq 1$, which implies $\log\Big[\frac{1024L^2}{\alpha^2}\Big(1 + \frac{1536L}{\alpha}\Big)\Big] \leq \log(1024 \cdot 1537) + 3\log\Big(\frac{L}{\alpha}\Big) = \mathcal{O}\Big(\frac{L}{\alpha}\Big)$. Taking the smaller of the bounds in Equations (63) and (65), using the monotonicity of the logarithm, and maximizing over $i \in [N]$, we complete the proof. □

D Omitted Proof Details for Lower Bounds

In this section, we prove the four lower bounds stated in Section 4. Before presenting the proofs, we specify the common network and decision sets and describe the communication-delay argument used below.

Let $N = 8a$ and $T = mN/4$, where $a \geq 2$ and $m \geq 1$ are integers, and set $K = N/4 = 2a$. Let C_N denote the cycle graph with vertices $1, \ldots, N$ in cyclic order. Choose $\mathcal{G} = C_N$ and

$$W = I - \frac{1}{4} L_{C_N}, \quad \rho = 1 - \sigma_2(W) = \sin^2\Big(\frac{\pi}{N}\Big), \tag{66}$$

where L_{C_N} is the graph Laplacian of C_N. The matrix W is nonnegative, symmetric, and doubly stochastic, with eigenvalues $\cos^2(\pi k/N)$ for $k = 0, \ldots, N-1$, and thus satisfies Assumption 4. Moreover, it follows from $2/N \leq \sin(\pi/N) \leq \pi/N$ that

$$\frac{1}{2}\rho^{-1/2} \leq K \leq \frac{\pi}{4}\rho^{-1/2}. \tag{67}$$

Following the communication-delay argument of Wan et al. (2024; 2025), fix a target agent $i \in [N]$ and choose S as a cyclic shift of $\{1, \ldots, a\}$ whose shortest graph distance from i exceeds K. Then $|S| = N/8$. Partition the horizon into blocks of K rounds:

$$\mathcal{J}_z = \{(z-1)K + 1, \ldots, zK\}, \quad z \in [m]. \tag{68}$$

Under the protocol in Section 2, information travels across at most one edge per online round. Therefore, loss information first revealed at S during a block cannot reach agent i within that block, even when the algorithm chooses a different decision at every round.

For the convex bounds, use the decision set

$$\mathcal{X} = \left[-\frac{R}{\sqrt{2}}, \frac{R}{\sqrt{2}}\right] \times \left[-\frac{R}{\sqrt{2m}}, \frac{R}{\sqrt{2m}}\right]^m \subset \mathbb{R}^{m+1}, \tag{69}$$

with coordinates indexed by $0, 1, \ldots, m$. For the strongly convex bounds, use $\mathcal{X} = [0, R/\sqrt{2}]^2$. Both sets satisfy Assumption 1.

In the randomized constructions below, $\mathbb{E}$ initially averages over both the sampled losses and the algorithm's independent randomization. We then fix a loss realization, leaving expectation only over the algorithm's randomization. Throughout this appendix, parenthesized superscripts denote vector coordinates: $x_{t,i}^{(k)}$ and $(u^\star)^{(k)}$ denote coordinate k of $x_{t,i}$ and $u^\star$, respectively.

D.1 Omitted Proof of Theorem 3

We now turn to the proof of the lower bound in Theorem 3. For completeness, we restate the theorem and give its proof.

Theorem 3. *Let $R, G, L > 0$ and $v \geq 0$. There exist a decision set $\mathcal{X} \subset \mathbb{R}^{m+1}$ and a network $(\mathcal{G}, W)$ satisfying Assumptions 1 and 4, with spectral gap $\rho = \sin^2(\pi/N)$, such that, for any possibly randomized D-OCO algorithm and any fixed agent $i \in [N]$, there exists a sequence of linear losses satisfying Assumptions 2 and 3 with $\alpha = 0$ and $V_T \leq v$ for which $\mathbb{E}[\mathrm{Reg}_T(i)] \geq \Omega(NR\min\{G\rho^{-1/4}\sqrt{T},\ G\rho^{-1/2} + \rho^{-3/4}\sqrt{v}\})$.*

Proof. Fix an arbitrary possibly randomized algorithm and a target agent i, and choose S as above. We give two linear-loss constructions on the decision set in Equation (69). Both have gradient norms at most G and satisfy Assumptions 2 and 3 with $\alpha = 0$ and any prescribed $L > 0$.

A sequence with zero gradient variation. Draw a uniform sign $\varepsilon_0 \in \{-1, 1\}$ independently of the algorithm and let $f_{t,j}(x) = G\varepsilon_0 x^{(0)}$ for $j \in S$ and $f_{t,j}(x) = 0$ otherwise, for all $t \in [T]$. Then $V_T = 0$, and a comparator with $(u^\star)^{(0)} = -R\varepsilon_0/\sqrt{2}$ minimizes every loss. All instantaneous excess losses are nonnegative. The communication-delay observation ensures that the first K decisions of agent i are independent of ε_0. Therefore,

$$\mathbb{E}[\mathrm{Reg}_T(i)] = |S| G \sum_{t=1}^{T} \mathbb{E}\left[\varepsilon_0 x_{t,i}^{(0)} + \frac{R}{\sqrt{2}}\right]$$

$$\geq |S|G\sum_{t=1}^{K}\mathbb{E}\left[\varepsilon_0 x_{t,i}^{(0)}+\frac{R}{\sqrt{2}}\right]$$
$$=\frac{|S|GRK}{\sqrt{2}}=\frac{NGRK}{8\sqrt{2}}. \tag{70}$$

Here, the inequality follows from the nonnegativity of each instantaneous excess loss, and the last line uses $\mathbb{E}[\varepsilon_0 x_{t,i}^{(0)}]=0$ for $t\leq K$ and $|S|=N/8$.

A sequence with controlled variation. Draw independent uniform signs $\varepsilon_1,\ldots,\varepsilon_m$, independently of the algorithm. For an amplitude $A\in[0,G]$, define

$$c_r=\frac{2A}{K}\min\{r-1,K-r\},\quad f_{(z-1)K+r,j}(x)=\begin{cases}c_r\varepsilon_z x^{(0)}, & j\in S,\\ 0, & j\notin S,\end{cases}\quad z\in[m],\ r\in[K]. \tag{71}$$

Since K is even, the coefficients satisfy the following two identities:

$$\sum_{r=1}^{K}c_r=\frac{A(K-2)}{2},\quad \sum_{r=2}^{K}(c_r-c_{r-1})^2=\frac{4A^2(K-2)}{K^2}. \tag{72}$$

The endpoint coefficients c_1 and c_K both vanish, so changes of sign between blocks contribute no gradient variation. Applying the definition in Equation (1) and the second identity in Equation (72), we obtain the following expression for the gradient variation:

$$V_T=\frac{|S|}{N}m\sum_{r=2}^{K}(c_r-c_{r-1})^2=\frac{mA^2(K-2)}{2K^2}. \tag{73}$$

Suppose $A>0$. Consider a centralized online linear problem on $[-R/\sqrt{2},R/\sqrt{2}]$ with losses $h_z(u)=\varepsilon_z u$. At round z, define the induced decision as the weighted average $y_z=\frac{2}{A(K-2)}\sum_{r=1}^{K}c_r x_{(z-1)K+r,i}^{(0)}$. By the first identity in Equation (72), the weights are nonnegative and sum to one. Hence, y_z belongs to the decision interval $[-R/\sqrt{2},R/\sqrt{2}]$. The communication-delay observation ensures that y_z depends only on previous signs and the algorithm's independent randomization, making it a valid online decision against h_z. Since the local losses are linear and depend only on coordinate 0, the decentralized regret and the regret of this centralized learner satisfy the following exact relation:

$$\mathrm{Reg}_T(i)=\frac{|S|A(K-2)}{2}\left(\sum_{z=1}^{m}h_z(y_z)-\min_{|u|\leq R/\sqrt{2}}\sum_{z=1}^{m}h_z(u)\right). \tag{74}$$

We use the standard lower bound for online linear optimization, also invoked in the tightness discussions of Yang et al. (2014, Remark 5) and Chiang et al. (2012, Section 4.1). For the random-sign losses considered here, independence gives $\mathbb{E}[\varepsilon_z y_z]=0$, and the Khintchine inequality yields

$$\mathbb{E}\left[\sum_{z=1}^{m}h_z(y_z)-\min_{|u|\leq R/\sqrt{2}}\sum_{z=1}^{m}h_z(u)\right]=\frac{R}{\sqrt{2}}\mathbb{E}\left[\left|\sum_{z=1}^{m}\varepsilon_z\right|\right]\geq\frac{R\sqrt{m}}{2}. \tag{75}$$

Combining Equations (74) and (75) and using $|S|=N/8$ and $K\geq 4$, we obtain

$$\begin{aligned}\mathbb{E}[\mathrm{Reg}_T(i)]&\geq\frac{NRA(K-2)\sqrt{m}}{32}\geq\frac{NRAK\sqrt{m}}{64}\\&=\frac{NRK}{64}\sqrt{\frac{2K^2V_T}{K-2}}\geq\Omega\Big(NR\rho^{-3/4}\sqrt{V_T}\Big).\end{aligned} \tag{76}$$

Here, the equality in the second line follows from Equation (73), and the last inequality uses $K-2\leq K$ and Equation (67). If $A=0$, both the regret and the variation are zero. Choose $A=\min\{G,\sqrt{2K^2v/(m(K-2))}\}$, which ensures $V_T\leq v$. Substituting this choice and using $T=mK$, at least one of the two constructions satisfies

$$\mathbb{E}[\mathrm{Reg}_T(i)]\geq\Omega\Big(NR\max\Big\{GK,\min\Big\{G\sqrt{KT},K^{3/2}\sqrt{v}\Big\}\Big\}\Big)\quad\text{(by Equations (70) and (76))}$$

$$\geq \Omega\Big(NR\min\Big\{G\sqrt{KT}, GK + K^{3/2}\sqrt{v}\Big\}\Big). \tag{77}$$

The last inequality uses $\max\{b, \min\{a, c\}\} \geq \frac{1}{2}\min\{a, b+c\}$ for nonnegative a, b, c.

In each construction, V_T is deterministic and independent of the signs. Revealing its actual value before the first round therefore leaves the independence arguments unchanged. Averaging over the signs yields a fixed loss sequence satisfying Equation (77), with expectation only over the algorithm's randomization. Finally, Equation (67) converts this bound into Theorem 3. □

D.2 Omitted Proof of Theorem 5

We begin by recalling the small-loss lower bound for convex losses.

Theorem 5. *Let $R, G, L > 0$, $f \geq 0$, and $\lambda = \min\{L, G/(2R)\}$. There exist a decision set $\mathcal{X} \subset \mathbb{R}^{m+1}$ and a network $(\mathcal{G}, W)$ satisfying Assumptions 1 and 4, with spectral gap $\rho = \sin^2(\pi/N)$, such that, for any possibly randomized D-OCO algorithm and any fixed agent $i \in [N]$, there exists a sequence of losses, nonnegative on $\mathbb{R}^{m+1}$, satisfying Assumptions 2 and 3 with $\alpha = 0$ and $F_T \leq f$ for which $\mathbb{E}[\mathrm{Reg}_T(i)] \geq \Omega(NR[\lambda R\rho^{-1/2} + \rho^{-1/4}\min\{G\sqrt{T}, \sqrt{Lf}\}])$.*

Proof. Fix an arbitrary possibly randomized algorithm and a target agent i, and choose S as above. Use the decision set in Equation (69) and $\lambda = \min\{L, G/(2R)\}$ from the theorem. We again give two constructions, now with losses that are convex, nonnegative, and L-smooth on $\mathbb{R}^{m+1}$.

A sequence with zero comparator loss. Draw a uniform sign $\varepsilon \in \{-1, 1\}$ independently of the algorithm and set $f_{t,j}(x) = \frac{\lambda}{2}(x^{(0)} - \varepsilon R/\sqrt{2})^2$ for $t \leq K$ and $j \in S$, and $f_{t,j}(x) = 0$ otherwise. Since $\lambda \leq L$ and $\sqrt{2}\lambda R \leq G$, these losses satisfy the required smoothness and gradient bounds. A comparator with $(u^\star)^{(0)} = \varepsilon R/\sqrt{2}$ incurs zero cumulative loss, so $F_T = 0$. The first K decisions of agent i are independent of ε by the communication-delay observation. Consequently,

$$\mathbb{E}[\mathrm{Reg}_T(i)] = \frac{\lambda}{2}|S|\sum_{t=1}^{K}\mathbb{E}\left[\left(x_{t,i}^{(0)} - \frac{\varepsilon R}{\sqrt{2}}\right)^2\right] \geq \frac{N\lambda R^2 K}{32}. \tag{78}$$

A sequence with controlled comparator loss. Set $g = \min\{G, 4\sqrt{Lf/T}\}$ and $r = R/\sqrt{2m}$. When $g > 0$, define the nonnegative convex function

$$h(s) = \begin{cases} 0, & s \leq 0, \\ \dfrac{L}{2}s^2, & 0 < s \leq g/L, \\ gs - \dfrac{g^2}{2L}, & s > g/L. \end{cases}$$

Its derivative $h'(s) = \min\{g, \max\{0, Ls\}\}$ is L-Lipschitz and bounded by g. Draw independent uniform signs $\varepsilon_1, \ldots, \varepsilon_m$, independently of the algorithm, and, for $t \in \mathcal{J}_z$, let $f_{t,j}(x) = h(r + g/L + \varepsilon_z x^{(z)})$ if $j \in S$ and $f_{t,j}(x) = 0$ otherwise. These losses are convex, nonnegative, and L-smooth on the whole space, with gradient norms at most $g \leq G$. Since $|x^{(z)}| \leq r$ on $\mathcal{X}$, the nonzero losses have the affine restrictions

$$f_{t,j}(x) = gr + \frac{g^2}{2L} + g\varepsilon_z x^{(z)}, \quad j \in S,\ t \in \mathcal{J}_z. \tag{79}$$

The comparator with $(u^\star)^{(0)} = 0$ and $(u^\star)^{(z)} = -\varepsilon_z r$ for $z \in [m]$ minimizes every local loss. Therefore, $F_T = \frac{|S|}{N}T\frac{g^2}{2L} = \frac{Tg^2}{16L} \leq f$. For $t \in \mathcal{J}_z$, the decision $x_{t,i}$ depends only on previous signs and the algorithm's independent randomization. In particular, $\mathbb{E}[\varepsilon_z x_{t,i}^{(z)}] = 0$. Using Equation (79), we can conclude that

$$\begin{aligned}\mathbb{E}[\mathrm{Reg}_T(i)] &= |S|g\sum_{z=1}^{m}\sum_{t\in\mathcal{J}_z}\mathbb{E}\Big[\varepsilon_z x_{t,i}^{(z)} + r\Big] \\ &= \frac{NgR\sqrt{KT}}{8\sqrt{2}} \geq \frac{NR\sqrt{K}}{8\sqrt{2}}\min\{G\sqrt{T}, \sqrt{Lf}\}.\end{aligned} \tag{80}$$

If $g = 0$, use zero losses, for which the same lower bound holds.

In each construction, F_T is deterministic and independent of the signs, so revealing its actual value before play does not change the argument. Averaging over the signs gives a fixed loss sequence attaining the corresponding lower bound, with expectation only over the algorithm's randomization. Choosing the better of the two constructions yields at least half the sum of the lower bounds in Equations (78) and (80). Applying Equation (67) completes the proof of Theorem 5. □

D.3 Omitted Proofs of Theorems 4 and 6

We recall the gradient-variation and small-loss lower bounds for strongly convex losses and prove them using a common quadratic construction.

Theorem 4. *Let $\alpha, R > 0$ and $M \in [m]$. There exist a decision set $\mathcal{X} \subset \mathbb{R}^2$ and a network $(\mathcal{G}, W)$ satisfying Assumptions 1 and 4, with spectral gap $\rho = \sin^2(\pi/N)$, such that, for any possibly randomized D-OCO algorithm and any fixed agent $i \in [N]$, there exists a sequence of nonnegative, α-strongly convex losses satisfying Assumptions 2 and 3 with $L = \alpha$ and $G = \alpha R$ for which $\frac{M}{2} \le 1 + \frac{V_T}{\alpha^2 R^2} \le \frac{25M}{16}$, and $\mathbb{E}[\mathrm{Reg}_T(i)] \ge \Omega\big(N\alpha R^2 \rho^{-1/2} \log\big(2 + \frac{V_T}{\alpha^2 R^2}\big)\big)$.*

Theorem 6. *Let $\alpha, R > 0$ and $M \in [m]$. There exist a decision set $\mathcal{X} \subset \mathbb{R}^2$ and a network $(\mathcal{G}, W)$ satisfying Assumptions 1 and 4, with spectral gap $\rho = \sin^2(\pi/N)$, such that, for any possibly randomized D-OCO algorithm and any fixed agent $i \in [N]$, there exists a sequence of nonnegative, α-strongly convex losses satisfying Assumptions 2 and 3 with $L = \alpha$ and $G = \alpha R$ for which $\frac{M}{32} \le \frac{\sqrt{\rho} F_T}{\alpha R^2} \le \frac{3\pi M}{128}$, and $\mathbb{E}[\mathrm{Reg}_T(i)] \ge \Omega(N\alpha R^2 \rho^{-1/2} \log(2 + \frac{\sqrt{\rho} F_T}{\alpha R^2}))$.*

Proof. The two results follow from the same quadratic construction. Fix an arbitrary possibly randomized algorithm, a target agent i, and the prescribed $M \in [m]$. Choose S as above and use $\mathcal{X} = [0, R/\sqrt{2}]^2$.

Construction. Let $q = |S|/N = 1/8$ and $r = R/\sqrt{2}$. For $z \in [M]$ and $j \in [N]$, define $b_{z,j} = \mathbf{1}_{\{j \le N/2\}}$ when z is odd and $b_{z,j} = 1 - \mathbf{1}_{\{j \le N/2\}}$ when z is even, where $\mathbf{1}$ denotes the indicator. Thus, in every block, half the agents have $b_{z,j} = 1$, and every agent's value flips between successive blocks. Draw p uniformly on $[0, 1]$ and, conditionally on p, draw independent $w_1, \dots, w_M \sim \mathrm{Bernoulli}(p)$. Draw all these variables before play, independently of the algorithm's random seed ξ. Define the comparator $u(p) = (qrp, r/2) \in \mathcal{X}$ and the losses $f_{t,j}(x) = \frac{\alpha}{2}\|x - c_{t,j}\|_2^2$, where

$$c_{t,j} = \begin{cases} r\big(w_z \mathbf{1}_{\{j \in S\}}, b_{z,j}\big), & t \in \mathcal{J}_z,\ z \in [M], \\ u(p), & t > MK. \end{cases}$$

The final $T - MK$ rounds form a stationary tail; this tail is absent when $M = m$. Every center lies in $\mathcal{X} = [0, r]^2$, whose diameter is R. Hence the losses are globally nonnegative, α-strongly convex, and α-smooth, and their gradient norms on $\mathcal{X}$ are at most αR. Thus they satisfy the required assumptions with $L = \alpha$ and $G = \alpha R$.

Expected regret. During an active block $\mathcal{J}_z$, the current observation w_z affects only losses at S. The communication-delay observation implies that the target agent's decisions within that block are functions only of $w_1, \dots, w_{z-1}$, ξ, and deterministic problem data. In particular, the agent has no access to p except through past observations before the tail begins. Define its block average $\overline{x}_i(z) = K^{-1} \sum_{t \in \mathcal{J}_z} x_{t,i} \in \mathcal{X}$. For $t \in \mathcal{J}_z$, the network-average loss is

$$\frac{1}{N} f_t(x) = \frac{\alpha}{2}\left[(x^{(1)} - qrw_z)^2 + q(1-q)r^2 w_z + (x^{(2)} - r/2)^2 + \frac{r^2}{4}\right], \tag{81}$$

where $x^{(1)}$ and $x^{(2)}$ are the two coordinates of x. Conditionally on $p, w_1, \dots, w_{z-1}, \xi$, the average $\overline{x}_i(z)$ is fixed and the conditional mean of w_z is p, so the conditional expected network-average center is $u(p)$. The best fixed comparator incurs no more loss than $u(p)$, and every loss in the stationary tail is minimized at $u(p)$. Since the losses are constant within each block, we obtain

$$\mathbb{E}[\mathrm{Reg}_T(i)] \ge \sum_{z=1}^{M} \mathbb{E}\left[\sum_{t \in \mathcal{J}_z} \big(f_t(x_{t,i}) - f_t(u(p))\big)\right]$$

$$\geq K \sum_{z=1}^{M} \mathbb{E}[f_{zK}(\overline{x}_i(z)) - f_{zK}(u(p))] \qquad \text{(by convexity)}$$

$$= \frac{N\alpha K}{2} \sum_{z=1}^{M} \mathbb{E}\left[\|\overline{x}_i(z) - u(p)\|_2^2\right], \tag{82}$$

where the equality follows by taking conditional expectations in Equation (81).

To lower bound the estimation error, fix $n \in \{0, \dots, M-1\}$ and let $s = \sum_{z=1}^{n} w_z$, with the empty sum equal to zero. The uniform prior gives the posterior $p \mid w_1, \dots, w_n \sim \mathrm{Beta}(s+1, n-s+1)$, whose variance is $\mathrm{Var}(p \mid w_1, \dots, w_n) = (s+1)(n-s+1)/((n+2)^2(n+3))$. Moreover, $\Pr(s=k) = \binom{n}{k} \int_0^1 p^k (1-p)^{n-k}\, dp = 1/(n+1)$ for $k \in \{0, \dots, n\}$. Averaging the posterior variance therefore yields

$$\mathbb{E}[\mathrm{Var}(p \mid w_1, \dots, w_n)] = \frac{\sum_{s=0}^{n}(s+1)(n-s+1)}{(n+1)(n+2)^2(n+3)} = \frac{1}{6(n+2)}.$$

Conditioning additionally on the independent random seed ξ leaves this posterior unchanged. Since $\overline{x}_i(z)$ is a function of the previous observations and ξ, the conditional bias–variance decomposition bounds its expected squared error from below by $q^2 r^2 \mathbb{E}[\mathrm{Var}(p \mid w_1, \dots, w_{z-1})]$. Combining this observation with Equation (82) and using $q = 1/8$ and $r^2 = R^2/2$, we obtain

$$\mathbb{E}[\mathrm{Reg}_T(i)] \geq \frac{N\alpha K q^2 r^2}{2} \sum_{z=1}^{M} \mathbb{E}[\mathrm{Var}(p \mid w_1, \dots, w_{z-1})]$$

$$= \frac{N\alpha R^2 K}{1536} \sum_{z=1}^{M} \frac{1}{z+1} \geq \frac{N\alpha R^2 K}{3072} \log(M+1), \tag{83}$$

where the last inequality uses $1/(z+1) \geq 1/(2z)$ and $\sum_{z=1}^{M} 1/z \geq \log(M+1)$.

Comparator loss and gradient variation. We next bound both statistics for every realization of $p, w_1, \dots, w_M$. By Equation (81), the second coordinate contributes at least $\alpha r^2/8 = \alpha R^2/16$ to the network-average loss in every active round, regardless of the comparator. This gives $F_T \geq \alpha R^2 MK/16$. For the upper bound, observe that

$$\frac{1}{N} f_t(u(p)) = \frac{\alpha r^2}{2} \left[q w_z - 2q^2 p w_z + q^2 p^2 + \frac{1}{4} \right], \quad t \in \mathcal{J}_z,\ z \leq M.$$

For $p \in [0,1]$ and $w_z \in \{0,1\}$, the quantity $q w_z - 2q^2 p w_z + q^2 p^2$ lies in $[0, q]$. Since $u(p)$ incurs zero loss in the tail, we have

$$\frac{\alpha R^2 MK}{16} \leq F_T \leq \frac{\alpha r^2 MK}{2} \left(q + \frac{1}{4} \right) = \frac{3\alpha R^2 MK}{32}. \tag{84}$$

For the gradient variation, $\nabla f_{t,j}(x) - \nabla f_{t-1,j}(x) = \alpha(c_{t-1,j} - c_{t,j})$, so the supremum over x in the definition of V_T disappears. At the transition from block $z-1$ to block z, the second coordinate of every center flips, while first coordinates change only at S. The contribution to V_T is therefore $\alpha^2 r^2 [1 + q(w_z - w_{z-1})^2]$, which lies between $\alpha^2 r^2$ and $\alpha^2 r^2 (1+q)$. If $M < m$, the transition to the tail contributes

$$\frac{\alpha^2}{N} \sum_{j=1}^{N} \|c_{MK,j} - u(p)\|_2^2 = \alpha^2 r^2 \left[q w_M - 2q^2 p w_M + q^2 p^2 + \frac{1}{4} \right] \leq \alpha^2 r^2 \left(q + \frac{1}{4} \right).$$

There are no other contributions, since the centers are constant within each block and within the tail, and V_T starts at round 2. It follows that

$$\frac{\alpha^2 R^2}{2}(M-1) \leq V_T \leq \frac{\alpha^2 R^2}{2} \left[\frac{9}{8}(M-1) + \frac{3}{8} \mathbf{1}_{\{M<m\}} \right] \leq \frac{9}{16} \alpha^2 R^2 M. \tag{85}$$

Since Equations (84) and (85) hold for every realization, an averaging argument yields a fixed loss sequence satisfying these bounds and the final lower bound in Equation (83).

For Theorem 4, Equation (85) and $M \geq 1$ give $M/2 \leq 1 + V_T/(\alpha^2 R^2) \leq 25M/16$ and $\log(2 + V_T/(\alpha^2 R^2)) \leq \log(M+2) \leq 2\log(M+1)$. Combining this comparison with Equations (67) and (83), we obtain

$$\mathbb{E}[\mathrm{Reg}_T(i)] \geq \frac{N\alpha R^2 K}{3072} \log(M+1) \geq \frac{N\alpha R^2}{12288\sqrt{\rho}} \log\left(2 + \frac{V_T}{\alpha^2 R^2}\right),$$

which completes the proof of Theorem 4.

For the small-loss bound in Theorem 6, Equations (67) and (84) imply that the same sequence satisfies $M/32 \leq \sqrt{\rho} F_T/(\alpha R^2) \leq 3\pi M/128 \leq M$. Hence, $\log(2 + \sqrt{\rho} F_T/(\alpha R^2)) \leq \log(M+2) \leq 2\log(M+1)$. Applying the same argument gives

$$\mathbb{E}[\mathrm{Reg}_T(i)] \geq \frac{N\alpha R^2 K}{3072} \log(M+1) \geq \frac{N\alpha R^2}{12288\sqrt{\rho}} \log\left(2 + \frac{\sqrt{\rho} F_T}{\alpha R^2}\right),$$

which completes the proof of Theorem 6. □

E Omitted Proof Details for Applications

In this section, we provide detailed proofs of the application guarantees for decentralized offline convex optimization and the stochastically extended adversarial model. We also compare the resulting rates with existing guarantees in the corresponding settings.

E.1 Decentralized Offline Convex Optimization

For a fixed network, a horizon-based $\mathcal{O}(\sqrt{T})$ regret bound yields only an $\mathcal{O}(T^{-1/2})$ ergodic objective-gap guarantee through the usual online-to-offline conversion (Wan et al., 2025), whereas offline gradient-tracking methods can achieve $\mathcal{O}(1/T)$ convergence for smooth convex objectives (Qu & Li, 2017). Our gradient-variation analysis bridges this gap in the offline regime: the decentralized online algorithm itself recovers an $\mathcal{O}(1/T)$ ergodic rate at every agent, without strong convexity or additional communication rounds between online decisions. When the local losses do not change over time, this follows from the gradient-variation argument for Theorem 1. Consider the network objective $\overline{f}(x) \triangleq N^{-1} \sum_{j=1}^{N} f_j(x)$ and the local time averages $\widehat{x}_{T,i} \triangleq T^{-1} \sum_{t=1}^{T} x_{t,i} = Z^{-1} \sum_{z=1}^{Z} x_i(z)$.

Corollary 9. *Under Assumptions 1 to 4 with $\alpha = 0$, suppose that $f_{t,i} = f_i$ for all $t \in [T]$ and $i \in [N]$. Each agent runs Algorithm 1 with $h = (2\eta)^{-1}$, θ and B as defined in Equation* (4)*, and $\eta = \min\{1/(800L), R/(\sqrt{108}\, G)\}$. Then, for every $i \in [N]$, it holds that*

$$\overline{f}(\widehat{x}_{T,i}) - \min_{x \in \mathcal{X}} \overline{f}(x) = \mathcal{O}\left(\frac{R(LR+G)\log N}{\sqrt{\rho}\, T}\right).$$

Proof. Since Assumption 2 with $\alpha = 0$ implies that $\overline{f}$ is convex, Jensen's inequality yields

$$\overline{f}(\widehat{x}_{T,i}) - \min_{x \in \mathcal{X}} \overline{f}(x) \leq \frac{1}{T} \sum_{t=1}^{T} \Big(\overline{f}(x_{t,i}) - \min_{x \in \mathcal{X}} \overline{f}(x)\Big) = \frac{\mathrm{Reg}_T(i)}{NT}.$$

Since $f_{t,i} = f_i$ for all $t \in [T]$ and $i \in [N]$, we have $V_T = 0$. Applying the gradient-variation bound Equation (43) in the proof of Theorem 1 with the prescribed stepsize, we obtain

$$\frac{\mathrm{Reg}_T(i)}{NT} \leq \frac{B}{T}\Big(\frac{R^2}{2\eta} + 54\eta G^2\Big) \leq \frac{B}{T}\Big(400LR^2 + 6\sqrt{3}\, RG\Big).$$

Combining these estimates and substituting $B = \mathcal{O}(\rho^{-1/2} \log N)$ completes the proof. □

For general convex losses, existing decentralized online methods achieve $\mathcal{O}(\sqrt{T})$ regret on a fixed network (Yan et al., 2012; Hosseini et al., 2013; Wan et al., 2024). This dependence on T is minimax optimal, as established by the $\Omega(\sqrt{T})$ lower bound of Wan et al. (2025). Consequently, these

Method	Ergodic convergence rate
Gradient tracking (Qu & Li, 2017)	$\mathcal{O}(\rho^{-2}T^{-1})$
Gradient tracking (Koloskova et al., 2021)	$\mathcal{O}(\rho^{-1}T^{-1}\log(2/\rho))$
EXTRA (Li & Lin, 2020, Lemma 3)	$\mathcal{O}(\rho^{-1/2}T^{-1})$
Corollary 9	$\mathcal{O}(\rho^{-1/2}T^{-1}\log N)$

Table 2: Comparison of ergodic convergence rates for decentralized offline smooth convex optimization. The bound of Koloskova et al. (2021) assumes a positive semidefinite mixing matrix.

worst-case guarantees alone yield only an $\mathcal{O}(T^{-1/2})$ ergodic convergence rate for decentralized offline convex optimization. In decentralized offline optimization, plain DGD can retain a nonzero error with a fixed stepsize when local objectives are heterogeneous (Yuan et al., 2016). Classical distributed subgradient analysis yields an $\mathcal{O}(T^{-1/2})$ ergodic guarantee on a fixed network with an appropriately chosen horizon-dependent stepsize (Nedic & Ozdaglar, 2009). Methods such as EXTRA and gradient tracking introduce correction mechanisms that permit exact convergence with fixed stepsizes and exploit smoothness to obtain $\mathcal{O}(T^{-1})$ guarantees in their respective convex settings (Li & Lin, 2020; Qu & Li, 2017). More specifically, classical gradient tracking achieves $\mathcal{O}(\rho^{-2}T^{-1})$ (Qu & Li, 2017), refined analysis improves this to $\mathcal{O}(\rho^{-1}T^{-1}\log(2/\rho))$ for positive semidefinite mixing matrices (Koloskova et al., 2021), and appropriately tuned EXTRA achieves $\mathcal{O}(\rho^{-1/2}T^{-1})$ (Li & Lin, 2020, Lemma 3). These offline results concern unconstrained optimization.

Corollary 9 guarantees $\mathcal{O}(\rho^{-1/2}T^{-1}\log N)$ for every agent's uniform time average over the convex constraint set $\mathcal{X}$, matching the T and ρ exponents of the above EXTRA guarantee up to $\log N$ and improving the spectral-gap dependence in the above gradient-tracking bounds. Our decentralized online algorithm recovers this offline guarantee directly from $V_T = 0$, without strong convexity, a separate offline algorithm, or additional communication rounds. More generally, its gradient-variation guarantee yields $\mathcal{O}(1/T)$ average regret on a fixed network whenever $V_T = \mathcal{O}(1)$, even when the local objectives vary over time.

E.2 Stochastically Extended Adversarial Model

We consider a decentralized extension of the stochastically extended adversarial (SEA) model (Sachs et al., 2022; Yan et al., 2023), which interpolates between stochastic and adversarial online convex optimization. At each round $t \in [T]$, every agent $i \in [N]$ selects a decision $x_{t,i} \in \mathcal{X}$, while the environment chooses a history-dependent joint distribution D_t over the local loss functions. The local losses are then sampled, and each agent observes its local loss $f_{t,i}$ and communicates with its neighbors according to the decentralized online protocol in Section 2.

Let $\mathcal{F}_{t-1}$ denote the full history before the losses at round t are sampled. The decisions $\{x_{t,i}\}_{i=1}^N$ and the distribution D_t are $\mathcal{F}_{t-1}$-measurable, and $(f_{t,1}, \dots, f_{t,N}) \mid \mathcal{F}_{t-1} \sim D_t$. No independence across agents is required. Define the expected local loss function of agent $i \in [N]$ as $F_{t,i}(x) \triangleq \mathbb{E}_{D_t}[f_{t,i}(x)] = \mathbb{E}[f_{t,i}(x) \mid \mathcal{F}_{t-1}]$ for all $x \in \mathcal{X}$. We assume that differentiation commutes with conditional expectation, so that $\nabla F_{t,i}(x) = \mathbb{E}_{D_t}[\nabla f_{t,i}(x)]$.

Existing SEA analyses bound regret in terms of cumulative stochastic variance and cumulative adversarial variation (Sachs et al., 2022; Yan et al., 2023). We define their decentralized counterparts by averaging the corresponding per-round quantities over agents:

$$\begin{aligned} \sigma_t^2 &\triangleq \frac{1}{N}\sum_{i=1}^N \sup_{x\in\mathcal{X}} \mathbb{E}_{D_t}\left[\|\nabla f_{t,i}(x) - \nabla F_{t,i}(x)\|_2^2\right], \quad \sigma_{1:T}^2 \triangleq \mathbb{E}\left[\sum_{t=1}^T \sigma_t^2\right], \\ \Sigma_t^2 &\triangleq \frac{1}{N}\sum_{i=1}^N \sup_{x\in\mathcal{X}} \|\nabla F_{t,i}(x) - \nabla F_{t-1,i}(x)\|_2^2, \qquad \Sigma_{1:T}^2 \triangleq \mathbb{E}\left[\sum_{t=2}^T \Sigma_t^2\right]. \end{aligned} \tag{86}$$

Here, Σ_t^2 is defined for $t \geq 2$, with $\Sigma_1^2 \triangleq 0$. The expectation $\mathbb{E}_{D_t}$ is taken over the current local losses conditional on the history, while the outer expectations account for the random

Work	Convex	Strongly convex
Centralized online learning		
Sachs et al. (2022)	$\mathcal{O}(\sqrt{\sigma_{1:T}^2 + \Sigma_{1:T}^2})$	$\mathcal{O}((\sigma_{\max}^2 + \Sigma_{\max}^2)\log(1+T))$
Chen et al. (2024)	$\mathcal{O}(\sqrt{\sigma_{1:T}^2 + \Sigma_{1:T}^2})$	$\mathcal{O}\Big((\sigma_{\max}^2 + \Sigma_{\max}^2)\log\Big(1 + \frac{\sigma_{1:T}^2+\Sigma_{1:T}^2}{\sigma_{\max}^2+\Sigma_{\max}^2}\Big)\Big)$
Yan et al. (2023)	$\widehat{\mathcal{O}}(\sqrt{\sigma_{1:T}^2 + \Sigma_{1:T}^2})$	$\mathcal{O}((\sigma_{\max}^2 + \Sigma_{\max}^2)\log(1+\sigma_{1:T}^2+\Sigma_{1:T}^2))$
Yan et al. (2024)	$\mathcal{O}(\sqrt{\sigma_{1:T}^2 + \Sigma_{1:T}^2})$	$\mathcal{O}\Big((\sigma_{\max}^2 + \Sigma_{\max}^2)\log\Big(1 + \frac{\sigma_{1:T}^2+\Sigma_{1:T}^2}{\sigma_{\max}^2+\Sigma_{\max}^2}\Big)\Big)$
Decentralized online learning		
Wan et al. (2025)	$\mathcal{O}(\sqrt{T})$	$\mathcal{O}(\log(1+T))$
Corollary 10	$\mathcal{O}(\sqrt{\sigma_{1:T}^2 + \Sigma_{1:T}^2})$	$\mathcal{O}\Big((\sigma_{\max}^2 + \Sigma_{\max}^2)\log\Big(1 + \frac{\sigma_{1:T}^2+\Sigma_{1:T}^2}{\sigma_{\max}^2+\Sigma_{\max}^2}\Big)\Big)$

Table 3: Comparison of expected-regret guarantees in the SEA model. We suppress dependence on the network parameters N and ρ and omit additive terms independent of T and the cumulative SEA statistics. The notation $\widehat{\mathcal{O}}$ hides logarithmic factors in the cumulative SEA statistics, and $\log\log T$ factors are omitted. Each bound uses the corresponding work's assumptions and statistical definitions.

history on which D_t depends. For the strongly convex guarantee, we additionally define the maximal stochastic variance and maximal adversarial variation as $\sigma_{\max}^2 \triangleq \max_{t\in[T]} \mathbb{E}[\sigma_t^2]$ and $\Sigma_{\max}^2 \triangleq \max_{t\in[T]} \mathbb{E}[\Sigma_t^2]$, respectively. The following corollary bounds the expected regret in terms of these SEA statistics.

Corollary 10. *Suppose that Assumptions 1 to 4 hold almost surely. Each agent runs Algorithm 1 with θ, B as defined in Equation (4). For $\alpha = 0$, set $h = (2\eta)^{-1}$ and $\eta = \min\{1/(800L), R/(10\sqrt{G^2 + 8\sigma_{1:T}^2/B^2 + 4B\Sigma_{1:T}^2})\}$. Then, it holds that*

$$\mathbb{E}\,\mathrm{Reg}_T = \mathcal{O}\left(\frac{NR(LR+G)\log N}{\sqrt{\rho}} + NR\sqrt{\sigma_{1:T}^2 + \frac{(\log N)^3}{\rho^{3/2}}\Sigma_{1:T}^2}\right). \tag{87}$$

For $\alpha > 0$, set $h = \alpha/2$. Define $\Gamma \triangleq \sigma_{\max}^2 + \min\{G^2\rho^{-1/2}\log N, \rho^{-3/2}(\log N)^3\Sigma_{\max}^2\}$. Then, the expected regret satisfies

$$\mathbb{E}\,\mathrm{Reg}_T = \mathcal{O}\left(\frac{N}{\alpha}\left[\frac{L(G+\alpha R)^2\log N}{\alpha\sqrt{\rho}} + \Gamma\log\left(1 + \frac{\sqrt{\rho}\,\sigma_{1:T}^2}{\Gamma\log N} + \frac{(\log N)^2\Sigma_{1:T}^2}{\rho\Gamma}\right)\right]\right). \tag{88}$$

Sachs et al. (2022) introduced the SEA model to interpolate between stochastic and adversarial online convex optimization, establishing expected-regret guarantees in terms of stochastic variance and adversarial variation. Chen et al. (2024) sharpened the strongly convex guarantee by replacing the horizon-dependent logarithm with one depending on the cumulative SEA statistics, normalized by the corresponding maximal statistics. Yan et al. (2023; 2024) obtained SEA guarantees through gradient-variation analyses, with the latter removing the additional logarithmic factor in the convex bound. Corollary 10 extends these guarantees to decentralized online learning: the leading terms involving the SEA statistics in Equations (87) and (88) can be simplified to

$$\mathcal{O}\left(\sqrt{\sigma_{1:T}^2 + \Sigma_{1:T}^2}\right) \quad \text{and} \quad \mathcal{O}\left((\sigma_{\max}^2 + \Sigma_{\max}^2)\log\left(1 + \frac{\sigma_{1:T}^2 + \Sigma_{1:T}^2}{\sigma_{\max}^2 + \Sigma_{\max}^2}\right)\right),$$

for convex and strongly convex losses, respectively. As summarized in Table 3, these bounds match the leading dependence on the SEA statistics in the centralized results of Chen et al. (2024); Yan et al. (2024) on the common class of smooth convex or strongly convex sampled losses.

By contrast, directly applying the decentralized methods of Wan et al. (2025) to the SEA model gives the horizon-based guarantees $\widetilde{\mathcal{O}}\big(N\rho^{-1/4}\sqrt{T}\big)$ and $\widetilde{\mathcal{O}}\big(N\rho^{-1/2}\log(1+T)\big)$ for convex and strongly convex losses, respectively. These guarantees do not improve with smaller SEA statistics, whereas Corollary 10 replaces the horizon dependence with dependence on cumulative stochastic variance and cumulative adversarial variation. Restoring the network dependence, Equation (87)

gives $\widetilde{\mathcal{O}}(N\sqrt{\sigma_{1:T}^2+\rho^{-3/2}\Sigma_{1:T}^2})$ for convex losses, as stated in Corollary 8. For strongly convex losses, the coefficient of the logarithmic term in Equation (96) is

$$\Gamma_B = \sigma_{\max}^2 + \min\{BG^2,\ B^3\Sigma_{\max}^2\}, \quad B = \Theta(\rho^{-1/2}\log N).$$

The variation-sensitive branch $B^3\Sigma_{\max}^2$ retains dependence on the maximal adversarial variation, with an $\widetilde{\mathcal{O}}(\rho^{-3/2})$ network factor. This factor does not determine the worst-case network dependence, however, because the other branch caps the adversarial contribution at BG^2. Indeed, $\sigma_{\max}^2 \le G^2$ implies $\Gamma_B \le (B+1)G^2$. Together with $\sigma_{1:T}^2 \le TG^2$ and $\Sigma_{1:T}^2 \le 4(T-1)G^2$, Equation (96) therefore yields $\mathbb{E}\,\mathrm{Reg}_T = \widetilde{\mathcal{O}}\big(N\rho^{-1/2}\log(1+T)\big)$ in the worst case. Thus, the minimum in Γ_B allows the strongly convex guarantee to exploit small maximal adversarial variation while retaining the nearly optimal worst-case dependence on N, ρ, and T established by Wan et al. (2025).

Proof of Corollary 10. In the SEA model, the local block gradient $a_i(z)$ in Equation (3) can be decomposed into the block-averaged expected gradients, the curvature correction, and a stochastic error:

$$a_i(z) = \frac{1}{B}\sum_{t\in\mathcal{I}_z}\nabla F_{t,i}(x_i(z)) - \alpha x_i(z) + \varepsilon_i(z), \quad \varepsilon_i(z) \triangleq \frac{1}{B}\sum_{t\in\mathcal{I}_z}\big[\nabla f_{t,i}(x_i(z)) - \nabla F_{t,i}(x_i(z))\big].$$

To control the stochastic contribution to $\widehat{V}_Z$, define $q_z \triangleq N^{-1}\sum_{i=1}^N \|\varepsilon_i(z)\|_2^2$. Indeed, the consecutive error differences satisfy $N^{-1}\sum_{i=1}^N \|\varepsilon_i(z)-\varepsilon_i(z-1)\|_2^2 \le 2(q_z+q_{z-1})$ for $z \ge 2$. Since $x_i(z)$ is fixed before block z, each gradient difference in $\varepsilon_i(z)$ has conditional mean zero given $\mathcal{F}_{t-1}$. Thus, the cross terms between different rounds vanish when expanding $\mathbb{E}\|\varepsilon_i(z)\|_2^2$, and

$$\mathbb{E}q_z \le \frac{1}{B^2}\sum_{t\in\mathcal{I}_z}\mathbb{E}\sigma_t^2 \le \frac{\sigma_{\max}^2}{B}, \quad \sum_{z=1}^Z \mathbb{E}q_z \le \frac{\sigma_{1:T}^2}{B^2}. \tag{89}$$

The expected losses $F_{t,i}$ inherit the gradient bound and smoothness in Assumptions 2 and 3. Applying the telescoping argument for Equation (33) to these losses bounds the variation of their block-averaged gradients at a common decision by $B\sum_{t=2}^T \Sigma_t^2$. Separating this variation from the stochastic gradient differences and decision movement, and using $x_i(1) = 0$, we obtain

$$\widehat{V}_Z \le G^2 + 8\sum_{z=1}^Z q_z + 4B\sum_{t=2}^T \Sigma_t^2 + \frac{4(L+\alpha)^2}{N}\sum_{z=2}^Z\sum_{i=1}^N \|x_i(z)-x_i(z-1)\|_2^2. \tag{90}$$

For convex losses, applying the movement comparison in the proof of Lemma 7 to Equation (90) and substituting the resulting bound into Equation (42) allows the movement term to be absorbed by the negative stability term under $\eta \le 1/(800L)$. Taking expectations and using Equation (89), we conclude that

$$\begin{aligned}\frac{\mathbb{E}\,\mathrm{Reg}_T}{NB} &\le \frac{R^2}{2\eta} + 50\eta\left(G^2 + \frac{8\sigma_{1:T}^2}{B^2} + 4B\Sigma_{1:T}^2\right)\\ &= \mathcal{O}\left(R(LR+G) + \frac{R}{B}\sqrt{\sigma_{1:T}^2 + B^3\Sigma_{1:T}^2}\right),\end{aligned}$$

where the last step uses the stepsize specified in Corollary 10. Multiplying by NB and substituting the block length in Equation (4) proves Equation (87).

For strongly convex losses, we start from Equation (61) and bound the weighted prediction errors in terms of the SEA statistics. Since Equation (23) in Lemma 6 holds for every prefix, summation by parts yields

$$\sum_{z=1}^Z \frac{\|g(z)-\widetilde{m}_i(z)\|_2^2}{z+1} \le \frac{24}{N}\sum_{z=1}^Z \frac{\sum_{j=1}^N \|a_j(z)-a_j(z-1)\|_2^2}{z+1}, \quad i\in[N]. \tag{91}$$

To control the local block-gradient differences on the right-hand side, for $z \ge 2$, define the variation of the expected block gradients and a bound on the decision-movement contribution as

$$d_z \triangleq \frac{1}{N}\sum_{i=1}^N \sup_{x\in\mathcal{X}}\left\|\frac{1}{B}\sum_{t\in\mathcal{I}_z}\big[\nabla F_{t,i}(x) - \nabla F_{t-B,i}(x)\big]\right\|_2^2,$$

$$v_z \triangleq \min\left\{4(G+\alpha R)^2, \frac{(L+\alpha)^2}{N}\sum_{i=1}^{N}\|x_i(z)-x_i(z-1)\|_2^2\right\},$$

and set $d_1 = v_1 = 0$. Adding and subtracting the expected gradients in Equation (3) and applying Assumptions 1 to 3, we obtain $\frac{1}{N}\sum_{i=1}^{N}\|a_i(z)-a_i(z-1)\|_2^2 \le 4(q_z+q_{z-1})+4d_z+4v_z$ for $z \ge 2$.

We next sum the preceding bound with weights $1/(z+1)$ to control the right-hand side of Equation (91). Since $a_i(0) = 0$ and $x_i(1) = 0$, Assumption 2 implies $\|a_i(1)\|_2 \le G$. Moreover, $\sum_{z=2}^{Z} q_{z-1}/(z+1) \le \sum_{z=1}^{Z} q_z/(z+1)$. Consequently, we have

$$\frac{1}{N}\sum_{z=1}^{Z}\frac{\sum_{i=1}^{N}\|a_i(z)-a_i(z-1)\|_2^2}{z+1} \le \frac{G^2}{2} + 8\sum_{z=1}^{Z}\frac{q_z}{z+1} + 4\sum_{z=1}^{Z}\frac{d_z+v_z}{z+1}.$$

Substituting this estimate into Equation (91), we obtain

$$\sum_{z=1}^{Z}\frac{\|g(z)-\widetilde{m}_i(z)\|_2^2}{z+1} \le 12G^2 + 192\sum_{z=1}^{Z}\frac{q_z+d_z+v_z}{z+1}, \quad i \in [N]. \tag{92}$$

It remains to bound the expectations of the three weighted sums on the right-hand side of Equation (92).

The telescoping argument for Equation (33) and the gradient bound in Assumption 2 imply

$$\sum_{z=1}^{Z}\mathbb{E}d_z \le B\Sigma_{1:T}^2, \quad \mathbb{E}d_z \le \min\{4G^2, B^2\Sigma_{\max}^2\}. \tag{93}$$

Similarly, it follows from the definition of v_z and the movement bound Equation (57) that $v_z \le 4(G+\alpha R)^2$ and $\sum_{z=1}^{Z} v_z \le (L+\alpha)^2(13G+14\alpha R)^2/\alpha^2$. Applying the logarithmic estimate from the proof of Lemma 9 to $\mathbb{E}v_z$, we obtain

$$\sum_{z=1}^{Z}\frac{\mathbb{E}v_z}{z+1} = \mathcal{O}\left((G+\alpha R)^2\log\left(1+\frac{L^2}{\alpha^2}\right)\right) = \mathcal{O}\left(\frac{L(G+\alpha R)^2}{\alpha}\right), \tag{94}$$

where the last step uses $\alpha \le L$. To combine the stochastic and adversarial contributions, write $\Gamma_B \triangleq \sigma_{\max}^2 + \min\{BG^2, B^3\Sigma_{\max}^2\}$. Then, Equations (89) and (93) imply $\mathbb{E}q_z + \mathbb{E}d_z \le \frac{4\Gamma_B}{B}$ and $\sum_{z=1}^{Z}(\mathbb{E}q_z + \mathbb{E}d_z) \le \sigma_{1:T}^2/B^2 + B\Sigma_{1:T}^2$. Applying the same logarithmic estimate to $\mathbb{E}q_z + \mathbb{E}d_z$, we conclude that

$$\sum_{z=1}^{Z}\frac{\mathbb{E}q_z+\mathbb{E}d_z}{z+1} = \mathcal{O}\left(\frac{\Gamma_B}{B}\log\left(1+\frac{\sigma_{1:T}^2}{B\Gamma_B}+\frac{B^2\Sigma_{1:T}^2}{\Gamma_B}\right)\right). \tag{95}$$

We combine Equations (61) and (92) with the network disagreement bound in Lemma 9. We then take expectations and apply Equations (94) and (95) to obtain

$$\mathbb{E}\,\mathrm{Reg}_T = \mathcal{O}\left(\frac{N}{\alpha}\left[\frac{BL(G+\alpha R)^2}{\alpha} + \Gamma_B\log\left(1+\frac{\sigma_{1:T}^2}{B\Gamma_B}+\frac{B^2\Sigma_{1:T}^2}{\Gamma_B}\right)\right]\right), \tag{96}$$

where $\alpha \le L$ absorbs the initialization terms, as in Equation (63). The statistical term is interpreted as zero when $\Gamma_B = 0$. Substituting $B = \Theta(\rho^{-1/2}\log N)$ from Equation (4) and $\Gamma_B = \Theta(\Gamma)$ proves Equation (88). In both cases, the bounds before taking expectations have a common right-hand side for all agents and therefore apply to $\mathrm{Reg}_T = \max_{i\in[N]}\mathrm{Reg}_T(i)$. □

F A Decentralized Doubling Trick for Convex Losses

The stepsize in Theorem 1 uses V_T and F_T, neither of which is available before play. Following the doubling trick of Rakhlin & Sridharan (2013, Appendix B), we instead monitor the observed block-gradient variation $\widehat{V}_Z$. Its local contributions are available to the agents, but a restart must be agreed upon across the network. We first show how block gossip permits this decision, and then apply the analysis of Theorem 1 within each epoch.

Let P denote one B-round application of accelerated gossip in Equation (2), initialized with $v_i^{-1} = v_i^0 = v_i$. By Lemma 1 and Equation (27), it preserves averages and contracts network disagreement by the factor $q \le (14N)^{-1/2}$ from Lemma 5. Each application reinitializes both starting states to its input, so P^2 applies this B-round map twice.

Lemma 12 (Common restart test)**.** *Under Assumption 4, let θ, B be as in Equation* (4)*. Given nonnegative local scalars c_i and a common budget $Q > 0$, write $c = (c_1, \dots, c_N)^\top$ and set $\widehat{c} = P^2 c$, $b_i = \mathbf{1}_{\{\widehat{c}_i > Q\}}$, and $p = P^2 b$. Then all agents agree on the test $p_i > 1/(2N)$. Writing $\overline{c} = N^{-1}\sum_i c_i$, if the test holds, then $\overline{c} > 14Q/15$; otherwise, $\overline{c} \le Q$.*

Proof. Average preservation and the contraction bound imply

$$\max_{i\in[N]} |(P^2 c)_i - \overline{c}| \le q^2 \|c - \overline{c}\mathbf{1}\|_2 \le \frac{\|c\|_2}{14N} \le \frac{\overline{c}}{14}, \tag{97}$$

where the last inequality uses $c_i \ge 0$. If $b = 0$, then $p = 0$. Otherwise, let $\overline{b} \triangleq N^{-1}\sum_{i=1}^N b_i$. Then $\overline{b} \ge 1/N$, and the same estimate gives $p_i \ge 13\overline{b}/14 > 1/(2N)$ for every $i \in [N]$. Thus, the test holds exactly when some $\widehat{c}_i > Q$, in which case $15\overline{c}/14 \ge \widehat{c}_i > Q$. If the test does not hold, then $\widehat{c}_i \le Q$ for all $i \in [N]$, and average preservation gives $\overline{c} = N^{-1}\sum_i \widehat{c}_i \le Q$. □

Doubling trick. Set $Q_0 = \max\{64G^2, (800LR)^2\}$ and divide the run into epochs separated by restarts. In epoch $\ell = 0, 1, \dots$, use the budget $Q_\ell = 2^\ell Q_0$ and run Algorithm 1 on the subsequent online rounds with $\alpha = 0$, $\eta_\ell = R/\sqrt{Q_\ell}$, and $h_\ell = (2\eta_\ell)^{-1}$. Block indices start at one in each epoch, with $\mathcal{I}_z$ denoting its z-th block. Agent i maintains its contribution to the block-gradient variation, $c_i(z) = \sum_{s=1}^z \|a_i(s) - a_i(s-1)\|_2^2$, where $a_i(z)$ is the local block gradient in Equation (3) and $a_i(0) = c_i(0) = 0$. Thus, $N^{-1}\sum_i c_i(z)$ is the epoch's block-gradient variation.

At the start of each group of four blocks, the agents retain a copy of their current local sums $c_i(z)$. They apply P twice to these copies during the first two blocks, form the indicators in Lemma 12 with $Q = Q_\ell$, and apply P twice to the indicators during the next two blocks. Throughout this test, Algorithm 1 and the local sums continue to update. At the group's end, if $p_i > 1/(2N)$, all agents start epoch $\ell + 1$ at the next round, reinitializing Algorithm 1 as in Line 2, resetting $c_i(0) = 0$, and restarting the block index at one. Otherwise, they start the next test with the updated local sums. The first test uses zero local sums and cannot trigger a restart; an unfinished final test is simply left incomplete. The scalar communication state is appended to the original messages, requiring one additional scalar per message and no additional communication rounds or gradient evaluations.

Proposition 1. *Suppose that Assumptions 1 to 4 hold with $\alpha = 0$ and $R, L > 0$. If the local losses are nonnegative, then the doubling trick above, with θ, B from Equation* (4)*, satisfies*

$$\mathrm{Reg}_T = \mathcal{O}\Big(NBR(LR + G) + NR\min\Big\{B^{3/2}\sqrt{V_T}, \sqrt{BLF_T}\Big\}\Big). \tag{98}$$

Without nonnegativity, the gradient-variation bound $\mathrm{Reg}_T = \mathcal{O}(NBR(LR+G) + NRB^{3/2}\sqrt{V_T})$ still holds. Neither bound requires prior knowledge of T, V_T, or F_T.

Proof. We first consider complete blocks. Let κ be the number of nonempty epochs and Z_ℓ the number of blocks in epoch ℓ. For brevity, write $E_\ell = N^{-1}\sum_i c_i(Z_\ell)$ for its block-gradient variation and $E = \sum_\ell E_\ell$. Since $\|a_i(z)\|_2 \le G$ by Assumption 2, each block increases the network average of the local sums by at most $4G^2$. By Lemma 12, a test that triggers a restart uses local sums with average exceeding $14Q_\ell/15$. The sums used by the most recent test that did not trigger a restart have average at most Q_ℓ and were stored at most eight blocks before the epoch's end. If no test has completed, the initial zero sums provide the same upper bound. This also covers an unfinished final group. Hence, we have

$$E_\ell \le Q_\ell + 32G^2 \le 2Q_\ell, \quad E_\ell > \frac{14}{15}Q_\ell \quad \text{for } \ell < \kappa - 1. \tag{99}$$

Summing the completed budgets and using their geometric growth, we obtain

$$Q_{\kappa-1} \le Q_0 + \frac{15E}{14}, \quad \kappa \le 1 + \frac{15E}{14Q_0}, \quad \sum_{\ell=0}^{\kappa-1}\sqrt{Q_\ell} \le \frac{\sqrt{Q_{\kappa-1}}}{1 - 2^{-1/2}}. \tag{100}$$

Fix $i \in [N]$ and use the optimal comparator $u^\star$ from Section 2 in every epoch. Let M_i be the sum of the epochwise cumulative decision movements $M_{Z_\ell,i}$ from the proof sketch of Theorem 1, excluding movements across restart boundaries. The derivation of Equation (42) applies to every comparator and every finite prefix. Applying it in each epoch, using $\eta_\ell \le \eta_0 \le 1/(800L)$, and retaining the negative stability term, the following formula holds:

$$\begin{aligned}\frac{\mathrm{Reg}_T(i)}{NB} &\le \sum_{\ell=0}^{\kappa-1}\left[\frac{R^2}{2\eta_\ell} + (24\eta_\ell + 4L\eta_\ell^2)E_\ell\right] - \frac{M_i}{4\eta_0} \\ &\le CR\sqrt{Q_0+E} - \frac{M_i}{4\eta_0}, \end{aligned} \tag{101}$$

where C is universal. The second line follows from $24\eta_\ell + 4L\eta_\ell^2 \le 25\eta_\ell$ and Equations (99) and (100). Since these inequalities hold for every prefix of the loss sequence, they also hold at the restart times.

For the gradient-variation bound, apply Lemma 7 within each epoch. The original gradient variations within epochs sum to at most V_T, so it holds that

$$E \le \frac{32}{29}\left(\kappa G^2 + 2BV_T + 6L^2M_i\right) \le \frac{32}{29}\left(G^2 + 2BV_T + 6L^2M_i\right) + \frac{32}{29}\frac{15G^2}{14Q_0}E.$$

Here, the second inequality uses Equation (100). Since $Q_0 \ge 64G^2$, the last coefficient is less than $1/2$, and therefore $E = \mathcal{O}(G^2 + BV_T + L^2M_i)$. Substituting this estimate into Equation (101) yields, for a universal constant C_1, we conclude that

$$\begin{aligned}\frac{\mathrm{Reg}_T(i)}{NB} &\le \mathcal{O}\left(R\sqrt{Q_0} + R\sqrt{BV_T}\right) + C_1RL\sqrt{M_i} - \frac{M_i}{4\eta_0} \\ &\le \mathcal{O}\left(R(LR+G) + R\sqrt{BV_T}\right).\end{aligned}$$

The last step uses $a\sqrt{m} - bm \le a^2/(4b)$ for $a \ge 0$, $b > 0$, together with $\sqrt{Q_0} = \mathcal{O}(G + LR)$ and $\eta_0 L \le 1/800$.

For the small-loss bound, the sum of the epochwise small losses, defined as in Equation (1), is at most F_T: each epoch's minimum is no larger than the network-average cumulative loss of $u^\star$ on that epoch. Applying Lemma 8 epochwise and then Equation (100), we obtain

$$E \le 33L\left(\frac{F_T}{B} + \frac{R^2}{2}\sum_{\ell=0}^{\kappa-1}\eta_\ell^{-1}\right) \le \frac{33LF_T}{B} + C_2LR\sqrt{Q_0+E},$$

where C_2 is universal. Solving this quadratic inequality in $\sqrt{Q_0+E}$ gives $\sqrt{Q_0+E} = \mathcal{O}(\sqrt{Q_0} + LR + \sqrt{LF_T/B})$. Substitution into Equation (101), now dropping its nonpositive stability term, yields $\mathrm{Reg}_T(i) = \mathcal{O}(NBR(LR+G) + NR\sqrt{BLF_T})$. Both bounds hold for the same run and uniformly over i, so taking their minimum and then the maximum over $i \in [N]$ proves Equation (98).

If T is not divisible by B, apply the argument to its complete blocks with the same comparator $u^\star$. Their gradient variation is at most V_T, and their network-average cumulative comparator loss is at most F_T under nonnegativity. The remaining fewer than B rounds contribute at most $2NBGR$ by Assumptions 1 and 2, which is absorbed by the first term in Equation (98). This also covers $T < B$. □

Thus, the doubling trick preserves Equation (5) up to universal constants, without an additional logarithmic factor in T. The dependence on V_T, F_T, N, and ρ is unchanged, so the comparisons with the lower bounds in Section 4 continue to hold without prior knowledge of V_T and F_T.