

Robust Stochastic Optimization via Gradient Quantile Clipping

Anonymous authors
Paper under double-blind review

Abstract

We introduce a clipping strategy for Stochastic Gradient Descent (SGD) which uses quantiles of the gradient norm as clipping thresholds. We prove that this new strategy provides a robust and efficient optimization algorithm for smooth objectives (convex or non-convex), that tolerates heavy-tailed samples (including infinite variance) and a fraction of outliers in the data stream akin to Huber contamination. Our mathematical analysis leverages the connection between constant step size SGD and Markov chains and handles the bias introduced by clipping in an original way. For strongly convex objectives, we prove that the iteration converges to a concentrated distribution and derive high probability bounds on the final estimation error. In the non-convex case, we prove that the limit distribution is localized on a neighborhood with low gradient. We propose an implementation of this algorithm using rolling quantiles which leads to a highly efficient optimization procedure with strong robustness properties, as confirmed by our numerical experiments.

Contents

1	Introduction	2
2	Preliminaries	4
3	Strongly Convex Objectives	6
4	Non-Convex Objectives	9
5	Implementation and Numerical Experiments	10
6	Conclusion	12
A	Additional experimental results	19
B	Experimental details	19
	B.1 Mean estimation	20
	B.2 Linear regression	20
	B.3 Logistic regression	21
	B.4 Single hidden layer neural network classifier	21
C	Geometric convergence speed and relation to step size	22

D Proofs	23
D.1 Preliminary lemmas	23
D.2 Preliminary properties for Markov chains	25
D.3 Proof of Theorem 1	26
D.4 Proof of Proposition 1	28
D.5 The case $q \in (1, 2)$	28
D.6 Proof of Proposition 2	31
D.7 Unimprovability of the sub-exponential property for $\eta > 0$	32
D.8 Proof of Corollary 1	33
D.9 Proof of Corollary 2	34
D.10 Proof of Theorem 2	34
D.11 Proof of Proposition 3	36

1 Introduction

Stochastic gradient descent (SGD) (Robbins & Monro, 1951) is the core optimization algorithm at the origin of most stochastic optimization procedures (Kingma & Ba, 2014; Defazio et al., 2014; Johnson & Zhang, 2013). SGD and its variants are ubiquitously employed in machine learning in order to train most models (Kushner & Yin, 2003; Benveniste et al., 2012; Lan, 2020; Shalev-Shwartz et al., 2007; Bottou et al., 2018; Ma et al., 2018). The convergence properties of SGD are therefore subjects of major interest.

Early studies of SGD convergence generally relied on strong assumptions such as bounded domain (Shalev-Shwartz et al., 2009) or uniformly bounded gradient variance (Rakhlin et al., 2011) and obtained error bounds in expectation. With the recent resurgence of interest for robust statistics (Hsu & Sabato, 2016; Diakonikolas et al., 2019; Lecué & Lerasle, 2017; Prasad et al., 2018), variants of SGD based on clipping are shown to be robust to heavy-tailed gradients (Gorbunov et al., 2020; Tsai et al., 2022), where the gradient samples are only required to have a finite variance. The latter requirement has been further weakened to the existence of a q -th moment for some $q > 1$ in (Sadiev et al., 2023; Nguyen et al., 2023). In this paper, we go further and show that another variant of clipped SGD with proper thresholds is robust both to heavy tails *and* outliers in the data stream.

Robust statistics appeared in the 60s with the pioneering works of Huber, Tukey and others (Tukey, 1960; Huber, 1992; 1972; Rousseeuw & Hubert, 2011; Hampel, 1971). More recently, the field found new momentum thanks to a series of works about robust scalar mean estimation (Catoni, 2012; Alon et al., 1996; Jerrum et al., 1986; Lugosi & Mendelson, 2021) and the more challenging multidimensional case (Hopkins, 2020; Catoni & Giulini, 2018; Lugosi & Mendelson, 2019; Minsker, 2015; Cherapanamjeri et al., 2019; Depersin & Lecué, 2022; Lei et al., 2020; Diakonikolas et al., 2020). These paved the way to the elaboration of a host of robust learning algorithms (Holland & Ikeda, 2019; Prasad et al., 2018; Lecué & Lerasle, 2017; Liu et al., 2020; Pensia et al., 2020) which have to date overwhelmingly focused on the batch learning setting. We consider the setting of streaming stochastic optimization (Bottou & Cun, 2003; Bottou & Lecun, 2005; McMahan et al., 2013), which raises an additional difficulty coming from the fact that algorithms can see each sample only once and must operate under an $\mathcal{O}(d)$ memory and complexity constraint for d -dimensional optimization problems. A limited number of papers (Tsai et al., 2022; Nazin et al., 2019; Diakonikolas et al., 2022) propose theoretical guarantees for robust algorithms learning from streaming data.

This work introduces such an algorithm that learns from data on the fly and is robust both to heavy tails and outliers, with minimal computational overhead and sound theoretical guarantees.

We consider the problem of minimizing a smooth objective

$$\min_{\theta \in \mathbb{R}^d} \mathcal{L}(\theta) := \mathbb{E}_{\zeta}[\ell(\theta, \zeta)] \quad (1)$$

using observations $G(\theta, \zeta_t)$ of the unknown gradient $\nabla \mathcal{L}(\theta)$, based on samples $(\zeta_t)_{t \geq 0}$ received sequentially that include corruptions with probability $\eta < 1/2$. Formulation (1) is common to numerous machine learning problems where ℓ is a loss function evaluating the fit of a model with parameters θ on a sample ζ , the expectation \mathbb{E} is w.r.t the unknown uncorrupted sample distribution.

We introduce quantile-clipped SGD (QC-SGD) which uses the iteration

$$\theta_{t+1} = \theta_t - \alpha_{\theta_t} \beta G(\theta_t, \zeta_t) \quad \text{with} \quad \alpha_{\theta_t} = \min \left(1, \frac{\tau_{\theta_t}}{\|G(\theta_t, \zeta_t)\|} \right), \quad (2)$$

where $\beta > 0$ is a constant step size and α_{θ_t} is the clipping factor with threshold chosen as the p -th quantile $\tau_{\theta_t} = Q_p(\|\tilde{G}(\theta_t, \zeta_t)\|)$ with $\tilde{G}(\theta_t, \zeta_t)$ an uncorrupted sample of $\nabla \mathcal{L}(\theta_t)$ and $p \in (0, 1)$ (details will follow). Quantiles are a natural choice of clipping threshold which allows to handle heavy tails (Rothenberg et al., 1964; Bloch, 1966) and corrupted data. For instance, the trimmed mean offers a robust and computationally efficient estimator of a scalar expectation (Lugosi & Mendelson, 2021). Since the quantile $Q_p(\|\tilde{G}(\theta_t, \zeta_t)\|)$ is non-observable, we introduce a method based on rolling quantiles in Section 5 which keeps the procedure $\mathcal{O}(d)$ both memory and complexity-wise. The main benefit of QC-SGD 2 is to grant robustness to the presence of a proportion $\eta < 1/2$ of corruptions in the stream of gradient samples. This could not be achieved by previous clipped SGD methods (Gorbunov et al., 2020; Tsai et al., 2022; Sadiev et al., 2023; Nguyen et al., 2023). We also show that iteration (2) is adaptive to heavy-tailed gradient variance and converges to a limit distribution with strong concentration properties.

Contributions. Our main contributions are as follows:

- For small enough η and well-chosen p , we show that, whenever the optimization objective is smooth and strongly convex, QC-SGD converges *geometrically* to a limit distribution such that the deviation around the optimum achieves the *optimal* dependence on η .
- In the non-corrupted case $\eta = 0$ and with a strongly convex objective, we prove that a coordinated choice of β and p ensures that the limit distribution is sub-Gaussian with constant of order $\mathcal{O}(\sqrt{\beta})$. In the corrupted case $\eta > 0$, the limit distribution is sub-exponential.
- For a smooth objective (non-convex) whose gradient satisfies an identifiability condition, we prove that the total variation distance between QC-SGD iterates and its limit distribution vanishes sub-linearly. In this case, the limit distribution is such that the deviation of the objective gradient is optimally controlled in terms of η .
- Finally, we provide experiments to demonstrate that QC-SGD can be easily and efficiently implemented by estimating $Q_p(\|\tilde{G}(\theta_t, \zeta_t)\|)$ with rolling quantiles. In particular, we show that the iteration is indeed robust to heavy tails and corruption on multiple stochastic optimization tasks.

Our theoretical results are derived thanks to a modelling through Markov chains and hold under an L_q assumption on the gradient distribution with $q > 1$.

Related works. Convergence in distribution of the Markov chain generated by constant step size SGD, relatively to the Wasserstein metric, was first established in (Dieuleveut et al., 2020). Another geometric convergence result was derived in (Yu et al., 2021) for non-convex, non-smooth, but quadratically growing objectives, where a convergence statement relatively to a weighted total variation distance is given and a CLT is established. These papers do not consider robustness to heavy tails or outliers. Early works proposed stochastic optimization and parameter estimation algorithms which are robust to a wide class of noise of distributions (Martin & Masreliez, 1975; Polyak & Tsytkin, 1979; 1981; Price & VandeLinde, 1979; Stanković & Kovačević, 1986; Chen et al., 1987; Chen & Gao, 1989; Nazin et al., 1992), where asymptotic convergence

guarantees are stated for large sample sizes. Initial evidence of the robustness of clipped SGD to heavy tails was given by (Zhang et al., 2020) who obtained results in expectation. Subsequent works derived high-confidence sub-Gaussian performance bounds under a finite variance assumption (Gorbunov et al., 2020; Tsai et al., 2022) and later under an L_q assumption (Sadiev et al., 2023; Nguyen et al., 2023) with $q > 1$. A similar SGD clipping scheme to (2) is presented in (Seetharaman et al., 2020), however, in contrast to our work, they do not consider the robust setting and focus on experimental study while we also provide theoretical guarantees.

Robust versions of Stochastic Mirror Descent (SMD) are introduced in (Nazin et al., 2019; Juditsky et al., 2023). For a proper choice of the mirror map, SMD is shown to handle infinite variance gradients without any explicit clipping (Nemirovskij & Yudin, 1983; Vural et al., 2022). Finally, (Diakonikolas et al., 2022) study heavy-tailed and outlier robust streaming estimation algorithms of the expectation and covariance. On this basis, robust algorithms for linear and logistic regression are derived. However, the involved filtering procedure is hard to implement in practice and no numerical evaluation of the considered approach is proposed.

Agenda. In Section 2 we set notations, state the assumptions required by our theoretical results and provide some necessary background on continuous state Markov chains. In Section 3, we state our results for strongly convex objectives including geometric ergodicity of QC-SGD (Theorem 1), characterizations of the limit distribution and deviation bounds on the final estimate. In Section 4, we remove the convexity assumption and obtain a weaker ergodicity result (Theorem 2) and characterize the limit distribution in terms of the deviations of the objective gradient. Finally, we present a rolling quantile procedure in Section 5 and demonstrate its performance through a few numerical experiments on synthetic and real data.

2 Preliminaries

The model parameter space is \mathbb{R}^d endowed with the Euclidean norm $\|\cdot\|$, $\mathcal{B}(\mathbb{R}^d)$ is the Borel σ -algebra of \mathbb{R}^d and we denote by $\mathcal{M}_1(\mathbb{R}^d)$ the set of probability measures over \mathbb{R}^d . We assume throughout the paper that the objective \mathcal{L} is smooth.

Assumption 1. *The objective \mathcal{L} is L -Lipschitz-smooth, namely*

$$\mathcal{L}(\theta') \leq \mathcal{L}(\theta) + \langle \nabla \mathcal{L}(\theta), \theta' - \theta \rangle + \frac{L}{2} \|\theta - \theta'\|^2$$

with $L < +\infty$ for all $\theta, \theta' \in \mathbb{R}^d$.

The results from Section 3 below use the following

Assumption 2. *The objective \mathcal{L} is μ -strongly convex, namely*

$$\mathcal{L}(\theta') \geq \mathcal{L}(\theta) + \langle \nabla \mathcal{L}(\theta), \theta' - \theta \rangle + \frac{\mu}{2} \|\theta - \theta'\|^2$$

with $\mu > 0$ for all $\theta, \theta' \in \mathbb{R}^d$.

An immediate consequence of Assumption 2 is the existence of a unique minimizer $\theta^* = \arg \min_{\theta \in \mathbb{R}^d} \mathcal{L}(\theta)$. The next assumption formalizes our corruption model.

Assumption 3 (η -corruption). *The gradients $(G(\theta_t, \zeta_t))_{t \geq 0}$ used in Iteration (2) are sampled as $G(\theta_t, \zeta_t) = U_t \tilde{G}(\theta_t) + (1 - U_t) \tilde{G}(\theta_t, \zeta_t)$ where U_t are i.i.d Bernoulli random variables with parameter $\eta < 1/2$, $\tilde{G}(\theta_t) \sim \mathcal{D}_{\mathcal{O}}(\theta_t)$ with $\mathcal{D}_{\mathcal{O}}(\theta_t)$ an arbitrary distribution and $\tilde{G}(\theta_t, \zeta_t) \sim \mathcal{D}_{\mathcal{I}}(\theta_t)$ follows the true gradient distribution and is independent from the past given θ_t .*

Assumption 3 is an online analog of the Huber contamination model (Huber, 1965; 1992) where corruptions occur with probability η and where the distribution of corrupted samples $\mathcal{D}_{\mathcal{O}}(\theta_t)$ (outliers) is not fixed and may depend on the current iterate θ_t . On the other hand, $\mathcal{D}_{\mathcal{I}}(\theta_t)$ denotes the distribution of inliers. This notation and dichotomy between inliers and outliers follows the example of (Lecué & Lerasle, 2017; Lecué &

Lerasle, 2019). Assumption 3 corresponds to *additive* contamination (Diakonikolas & Kane, 2023, Section 1.2.2) where corruptions are only added to the data. A more general TV-contamination model allowing for true samples to be adversely removed is used in (Diakonikolas et al., 2022). Note however, that the latter mainly focuses on mean estimation. Note also that additive contamination remains realistic since it accounts for invalid entries occurring in the data stream even if it doesn't support entries being targeted and censored. The next assumption requires the true gradient distribution to be unbiased and diffuse.

Assumption 4. For all θ , non-corrupted gradient samples $\tilde{G}(\theta, \zeta) \sim \mathcal{D}_{\mathcal{I}}(\theta)$ are such that

$$\tilde{G}(\theta, \zeta) = \nabla \mathcal{L}(\theta) + \varepsilon_{\theta}, \quad (3)$$

where ε_{θ} is a centered noise $\mathbb{E}[\varepsilon_{\theta} | \theta] = 0$ with distribution $\delta \nu_{\theta,1} + (1 - \delta) \nu_{\theta,2}$ where $\delta > 0$ and $\nu_{\theta,1}, \nu_{\theta,2}$ are distributions over \mathbb{R}^d such that $\nu_{\theta,1}$ admits a density h_{θ} w.r.t. the Lebesgue measure satisfying

$$\inf_{\|\omega\| \leq R} h_{\theta}(\omega) > \varkappa(R) > 0$$

for all $R > 0$, where $\varkappa(\cdot)$ is independent of θ .

In addition to the unbiased property, Assumption 4 imposes that the noise distribution be expressible as the combination of two components, one of which must be diffuse with density satisfying a minorization inequality. Note that this is a weak constraint since it is satisfied, for example, as soon as the noise ε_{θ} admits a density w.r.t. Lebesgue's measure which is positive everywhere. This condition is similar to (Yu et al., 2021, Assumption 2.3) since both find their origin in Markov chain minorization conditions (Meyn & Tweedie, 1993, Section 5.2). These ensure that a chain properly explores its state space and are a common way to prove Markov chain convergence (Rosenthal, 1995b;a; Douc et al., 2004; Meyn & Tweedie, 1994; Baxendale, 2005). Our last assumption formalizes the requirement of a finite moment for the gradient error.

Assumption 5. There is $q > 1$ such that for $\tilde{G}(\theta, \zeta) \sim \mathcal{D}_{\mathcal{I}}(\theta)$, we have

$$\mathbb{E}[\|\varepsilon_{\theta}\|^q | \theta]^{1/q} = \mathbb{E}[\|\tilde{G}(\theta, \zeta) - \nabla \mathcal{L}(\theta)\|^q | \theta]^{1/q} \leq A_q \|\theta - \theta^*\| + B_q \quad (4)$$

for all $\theta \in \mathbb{R}^d$, where $A_q, B_q > 0$. When \mathcal{L} is not strongly convex, we further assume that $A_q = 0$.

The bound (4) captures the case of arbitrarily high noise magnitude through the dependence on $\|\theta - \theta^*\|$. This is consistent with convex optimization problems with L -Lipschitz-smooth objectives (Assumption 1) where the norm of the gradient $\|\nabla \mathcal{L}(\theta)\|$ is bounded by $L \cdot \|\theta - \theta^*\|$. Assumption 5 improves upon the conditions used in (Gorbunov et al., 2020; Tsai et al., 2022; Gorbunov et al., 2023; Nguyen et al., 2023) since these either required a uniformly constant upperbound (independent of θ) or only considered the case $q = 2$ (finite variance). For non-strongly convex \mathcal{L} , we require $A_q = 0$ since θ^* may not exist.

Definition 1. If X is a real random variable, we say that X is K -sub-Gaussian for $K > 0$ if

$$\mathbb{E} \exp(\lambda^2 X^2) \leq e^{\lambda^2 K^2} \quad \text{for } |\lambda| \leq 1/K. \quad (5)$$

We say that X is K -sub-exponential for $K > 0$ if

$$\mathbb{E} \exp(\lambda |X|) \leq \exp(\lambda K) \quad \text{for all } 0 \leq \lambda \leq 1/K. \quad (6)$$

The convergence results presented in this paper use the following formalism of continuous state Markov chains. Given a step size $\beta > 0$ and a quantile $p \in (0, 1)$, we denote by $P_{\beta,p}$ the Markov transition kernel governing the Markov chain $(\theta_t)_{t \geq 0}$ generated by QC-SGD, so that

$$\mathbb{P}(\theta_{t+1} \in A | \theta_t) = P_{\beta,p}(\theta_t, A)$$

for $t \geq 0$ and $A \in \mathcal{B}(\mathbb{R}^d)$. The transition kernel $P_{\beta,p}$ acts on probability distributions $\nu \in \mathcal{M}_1(\mathbb{R}^d)$ through the mapping $\nu \rightarrow \nu P_{\beta,p}$ which is defined, for all $A \in \mathcal{B}(\mathbb{R}^d)$, by $\nu P_{\beta,p}(A) = \int_A P_{\beta,p}(\theta, A) d\nu(\theta) = \mathbb{P}(\theta_{t+1} \in A | \theta_t \sim \nu)$. For $n \geq 1$, we similarly define the multi-step transition kernel $P_{\beta,p}^n$ which is such that $P_{\beta,p}^n(\theta_t, A) =$

$\mathbb{P}(\theta_{t+n} \in A \mid \theta_t)$ and acts on probability distributions $\nu \in \mathcal{M}_1(\mathbb{R}^d)$ through $\nu P_{\beta,p}^n = (\nu P_{\beta,p}) P_{\beta,p}^{n-1}$. Finally, we define the total variation (TV) norm of a signed measure ν as

$$2\|\nu\|_{\text{TV}} = \sup_{f:|f|\leq 1} \int f(\theta)\nu(d\theta) = \sup_{A \in \mathcal{B}(\mathbb{R}^d)} \nu(A) - \inf_{A \in \mathcal{B}(\mathbb{R}^d)} \nu(A). \quad (7)$$

In particular, we recover the TV *distance* between $\nu_1, \nu_2 \in \mathcal{M}_1(\mathbb{R}^d)$ as

$$d_{\text{TV}}(\nu_1, \nu_2) = \|\nu_1 - \nu_2\|_{\text{TV}} = \sup_{A \in \mathcal{B}(\mathbb{R}^d)} |\nu_1(A) - \nu_2(A)|.$$

The second equality reflects the fact that the TV distance between two probability measures corresponds to the largest absolute difference between the probabilities they assign to the same event. The TV distance is a broadly used metric to quantify the convergence of Markov chains (Levin & Peres, 2017; Baxendale, 2005; Meyn & Tweedie, 1993; Rosenthal, 1995a) besides the Wasserstein distance (Dieuleveut et al., 2020).

In the next section, we will prove that the Markov chain defined by iteration (2) converges to unique invariant distribution in TV distance. This convergence mode will allow us to extrapolate the properties of the limit distribution on the iterates θ_t and thus derive non-asymptotic concentration bounds for them, see Corollaries 2 and 1 below.

3 Strongly Convex Objectives

We are ready to state our convergence result for the stochastic optimization of a strongly convex objective using QC-SGD with η -corrupted samples.

Theorem 1 (Geometric ergodicity). *Let Assumptions 1-5 hold and assume there is a quantile $p \in [\eta, 1 - \eta]$ such that*

$$\kappa := (1 - \eta)p\mu - \eta L - (1 - p)^{-\frac{1}{q}} A_q (1 - p(1 - \eta)) > 0. \quad (8)$$

Then, for a step size β satisfying

$$\beta < \frac{1}{4} \frac{\kappa}{\mu^2 + 24\eta L^2 + 28A_q^2} \wedge \frac{2}{\mu + L}, \quad (9)$$

the Markov chain $(\theta_t)_{t \geq 0}$ generated by QC-SGD with parameters β and p converges geometrically to a unique invariant measure $\pi_{\beta,p}$: for any initial $\theta_0 \in \mathbb{R}^d$, there is $\rho < 1$ and $M < \infty$ such that after T iterations

$$\|\delta_{\theta_0} P_{\beta,p}^T - \pi_{\beta,p}\|_{\text{TV}} \leq M \rho^T (1 + \|\theta_0 - \theta^*\|^2),$$

where δ_{θ_0} is the Dirac measure located at θ_0 .

The proof of Theorem 1 is given in Appendix D.3 and relies on the geometric ergodicity result of (Meyn & Tweedie, 1993, Chapter 15) for Markov chains with a geometric drift property. A similar result for quadratically growing objectives was established by (Yu et al., 2021) and convergence w.r.t. Wasserstein's metric was shown in (Dieuleveut et al., 2020) assuming gradient co-coercivity. However, robustness was not considered in these works. Theorem 1 establishes the iteration's convergence to a unique invariant measure $\pi_{\beta,p}$. The properties of this limit distribution will be explored in the sequel. The restriction $p \in [\eta, 1 - \eta]$ comes from the consideration that other quantiles are not estimable in the event of η -corruption. Condition (8) is best interpreted for the choice $p = 1 - \eta$ in which case it translates into $\eta^{1-1/q} \leq \mathcal{O}(\mu/(L + A_q))$ implying that it is verified for η small enough within a limit fixed by the problem conditioning. A similar condition with $q = 2$ appears in (Diakonikolas et al., 2022, Theorem E.9) which uses a finite variance assumption.

When (8) is satisfied, one clearly has that $\kappa = \mathcal{O}(\mu)$. Considering $q = 2$ for simplicity and taking the maximum allowed corruption rate in this case $\eta = \mathcal{O}(\mu^2/(L + A_q)^2)$ leads to an upperbound on the step-size β of order $\mathcal{O}(\mu/(\mu^2 + A_q^2) \wedge 1/L)$. While the condition $\beta = \mathcal{O}(1/L)$ is standard in smooth optimization, the additional condition in terms of A_q ensures that the noise introduced to the iteration by the gradient samples does not cause it to diverge.

The constants M and ρ controlling the geometric convergence speed in Theorem 1 depend on the parameters β, p and the initial θ_0 . Among choices fulfilling the convergence conditions, it is straightforward that greater step size β and θ_0 closer to θ^* lead to faster convergence. However, the dependence in p is more intricate and should be evaluated through the resulting value of κ . We provide a more detailed discussion about the value of ρ in Appendix C.

The choice $p = 1 - \eta$ appears to be ideal since it leads to optimal deviation of the invariant distribution around the optimum θ^* which is the essence of our next statement.

Proposition 1. *Assume the same as in Theorem 1 and condition (8) with the choice $p = 1 - \eta$. For step size β satisfying (9), $q \geq 2$, and additionally:*

$$\beta \leq \eta^{2-2/q}/\kappa, \quad (10)$$

for $\theta \sim \pi_{\beta, 1-\eta}$, we have the following upper bound:

$$\mathbb{E}\|\theta - \theta^*\|^2 \leq \left(\frac{6\eta^{1-1/q}B_q}{\kappa} \right)^2.$$

Proposition 1 is proven in Appendix D.4. An analogous result holds for $q \in (1, 2)$ but requires a different proof and can be found in Appendix D.5. Proposition 1 may be compared to (Yu et al., 2021, Theorem 3.1) which shows that the asymptotic estimation error can be reduced arbitrarily using a small step size. However, this is impossible in our case since we consider corrupted gradients. The performance of Proposition 1 is best discussed in the specific context of linear regression where gradients are given as $G(\theta, (X, Y)) = X(X^\top\theta - Y)$ for samples $X, Y \in \mathbb{R}^d \times \mathbb{R}$ such that $Y = X^\top\theta^* + \epsilon$ with ϵ a centered noise. In this case, a finite moment of order k for the data implies order $k/2$ for the gradient corresponding to an $\eta^{1-2/k}$ rate in Proposition 1. Since Assumption 5 does not include independence of the noise ϵ from X , this corresponds to the negatively correlated moments assumption of (Bakshi & Prasad, 2021) being unsatisfied. Consequently, Proposition 1 is information-theoretically optimal in η based on (Bakshi & Prasad, 2021, Corollary 4.2). Nonetheless, the dimension dependence through B_q remains poor since we have $B_q \sim \sqrt{d}$ in general because the Euclidean norm is used in Assumption 5. This dimension dependence may be improvable by using the quantiles $\sup_{\|v\|=1} Q_p(|\langle \tilde{G}(\theta_t, \zeta_t), v \rangle|)$ as clipping thresholds and adapting ideas from (Catoni & Giulini, 2018) in the analysis. However, exploring this method is beyond our scope as the involved estimations for all $\|v\| = 1$ would be excessively sample hungry and computationally heavy for stochastic optimization. If the gradient is sub-Gaussian with constant K , we would have $B_q \lesssim K\sqrt{q}$ for $q \geq 1$ (see (Vershynin, 2018) for a reference), in which case, the choice $q = \log(1/\eta)$ recovers the optimal rate in $\eta\sqrt{\log(1/\eta)}$ for the Gaussian case.

We now turn to showing strong concentration properties for the invariant distribution $\pi_{\beta, p}$. For this purpose, we restrict the optimization to a bounded and convex set $\Theta \subset \mathbb{R}^d$ and replace Iteration (2) by the projected iteration

$$\theta_{t+1} = \Pi_\Theta(\theta_t - \alpha_{\theta_t}\beta G(\theta_t, \zeta_t)), \quad (11)$$

where Π_Θ is the projection onto Θ . Assuming that the latter contains the optimum $\theta^* \in \Theta$, one can check that the previous results continue to hold thanks to the inequality

$$\|\Pi_\Theta(\theta) - \theta^*\| = \|\Pi_\Theta(\theta) - \Pi_\Theta(\theta^*)\| \leq \|\theta - \theta^*\|,$$

which results from the convexity of Θ . The restriction of the optimization to a bounded set allows us to uniformly bound the clipping threshold τ_θ , which is indispensable for the following result.

Proposition 2. *In the setting of Theorem 1, consider projected QC-SGD (11) and let $\bar{\tau} = \sup_{\theta \in \Theta} \tau_\theta$, $D = \text{diam}(\Theta)$ the diameter of Θ and $\bar{B}_q = A_q D + B_q$.*

- *Consider the non-corrupted case $\eta = 0$ and set the quantile p such that $p \geq 1 - (\beta\mu)^{\frac{q}{2(q-1)}}$. Then, for $\theta \sim \pi_{\beta, p}$, the variable $\|\theta - \theta^*\|$ is sub-Gaussian in the sense of Definition 1 with constant*

$$K = 4\sqrt{\frac{2\beta(\bar{B}_q^2 + \bar{\tau}^2)}{p\mu}}.$$

- Consider the corrupted case $\eta > 0$, and set the quantile $p \in [\eta, 1 - \eta]$ such that Inequality (8) holds. Then, for $\theta \sim \pi_{\beta,p}$, the variable $\|\theta - \theta^*\|$ is sub-exponential in the sense of Definition 1 with constant

$$K = \frac{7\bar{\tau} + (1-p)^{1-1/q}\bar{B}_q}{p\mu}.$$

The proof can be found in Appendix D.6. The strong concentration properties given by Proposition 2 for the invariant distribution appear to be new. Still, the previous result remains asymptotic in nature. High confidence deviation bounds for an iterate θ_t can be derived by leveraging the convergence in Total Variation distance given by Theorem 1 leading to the following result.

Corollary 1. *In the setting of Proposition 2, in the absence of corruption $\eta = 0$, after T iterations, for $\delta > 0$, we have*

$$\mathbb{P}\left(\|\theta_T - \theta^*\| > 4\sqrt{\bar{B}_q^2 + \bar{\tau}^2} \sqrt{\frac{2\beta \log(e/\delta)}{p\mu}}\right) \leq \delta + \rho^T M(1 + \|\theta_0 - \theta^*\|^2).$$

Choosing a smaller step size β in Corollary 1 allows to improve the deviation bound. However, this comes at the cost of weaker confidence because of slower convergence due to a greater ρ . See Appendix C for a discussion including a possible compromise. Corollary 1 may be compared to the results of (Gorbunov et al., 2020; Tsai et al., 2022; Sadiev et al., 2023; Nguyen et al., 2023) which correspond to $\beta \approx 1/T$ and have a similar dependence on the dimension through the gradient variance. Although their approach is also based on gradient clipping, they use different thresholds and proof methods. In the presence of corruption, the invariant distribution is not sub-Gaussian. This can be seen by considering the following toy Markov chain:

$$X_{t+1} = \begin{cases} \alpha X_t + \xi & \text{w.p. } 1 - \eta \\ X_t + \tau & \text{w.p. } \eta \end{cases}$$

where $\alpha < 1, \tau > 0$ are constants and ξ is a positive random noise. Using similar methods to the proof of Theorem 1, one can show that $(X_t)_{t \geq 0}$ converges (for any initial X_0) to an invariant distribution whose moments can be shown to grow linearly, indicating a sub-exponential distribution and excluding a sub-Gaussian one. We provide additional details for the underlying argument in Appendix D.7. For the corrupted case, the sub-exponential property stated in Proposition 2 holds with a constant K of order $\bar{\tau}/\mu$, which is not satisfactory and leaves little room for improvement due to the inevitable bias introduced by corruption. Therefore, we propose the following procedure in order to obtain a high confidence estimate, similarly to Corollary 1.

Algorithm 1: Aggregation of cycling iterates

Input: Step size $\beta > 0$, quantile index $p \in (0, 1)$, initial parameter $\theta_0 \in \Theta$, horizon T and number of concurrent iterates $N \geq 1$.

- 1 Optimize multiple parameters $\theta_t^{(1)}, \dots, \theta_t^{(N)}$ starting from a common $\theta_0 = \theta_0^{(n)}$ for $n \in \llbracket N \rrbracket =: \{1, \dots, N\}$ and T steps $t = 0, \dots, T$ using the following cycling iteration:

$$\theta_{t+1}^{(n)} = \begin{cases} \theta_t^{(n)} - \alpha_{\theta_t^{(n)}} \beta G(\theta_t^{(n)}, \zeta_t) & \text{if } t \equiv n - 1 \pmod{N}, \\ \theta_t^{(n)} & \text{otherwise.} \end{cases} \quad (12)$$

- 2 Compute the pairwise distances $r_{i,j} = \|\theta_T^{(i)} - \theta_T^{(j)}\|$ for $i, j \in \llbracket N \rrbracket$.
 - 3 For $i \in \llbracket N \rrbracket$, let $r^{(i)} \in \mathbb{R}_+^N$ be the vector $r_{i,:} := [r_{i,1}, \dots, r_{i,N}]$ sorted in non decreasing order.
 - 4 Compute the aggregated estimator as $\hat{\theta} = \theta_T^{(\hat{i})}$ with $\hat{i} = \arg \min_{i \in \llbracket N \rrbracket} r_{\lfloor N/2 \rfloor}^{(i)}$.
 - 5 **return** $\hat{\theta}$
-

Algorithm 1 uses ideas from (Hsu & Sabato, 2016) (see also (Minsker, 2015; Juditsky et al., 2023)) and combines the collection of *weak* estimators $(\theta_T^{(i)})_{i \in \llbracket N \rrbracket}$ (only satisfying L_2 bounds) into a strong one with

sub-exponential deviation. This is done by picking $\theta_T^{(i)}$ which is such that the median of its distances to other estimators $r_{\lfloor N/2 \rfloor}^{(i)}$ is minimal. The aggregated estimator $\hat{\theta}$ satisfies the high probability bound given in the next result.

Corollary 2. *Assume the same as in Theorem 1 and Proposition 1. Consider $\hat{\theta}$ given by Algorithm 1, with the assumption that the gradient sample sets used for each $(\theta_T^{(n)})_{n \in \llbracket N \rrbracket}$ in Equation (12) are independent. For $\delta > 0$, if $N \geq 16 \log(1/\delta)$ and T satisfies*

$$T \geq N \log(15M(1 + \|\theta_0 - \theta^*\|^2)) / \log(1/\rho),$$

then, with probability at least $1 - \delta$, we have

$$\|\hat{\theta} - \theta^*\| \leq \frac{27\eta^{1-\frac{1}{q}} \overline{B}_q}{\kappa}. \quad (13)$$

We obtain a high confidence version of the bound in expectation previously stated in Proposition 1. As argued before, the above bound depends optimally on η . Similar bounds to (13) are obtained for $q = 2$ in (Diakonikolas et al., 2022) for streaming mean estimation, linear and logistic regression. Their results enjoy better dimension dependence but are less general than ours since we handle the case $q \in (1, 2)$ and consider strongly convex objectives more broadly. In addition, our results further extend to non-convex objectives as detailed in the next section. Finally, the implementation of the algorithm in (Diakonikolas et al., 2022) is not straightforward whereas our method is quite easy to use (see Section 5).

4 Non-Convex Objectives

In this section, we drop Assumption 2 and consider the optimization of possibly non-convex objectives. Consequently, the existence of a unique optimum θ^* and the quadratic growth of the objective are no longer guaranteed. This motivates us to use a uniform version of Assumption 5 with $A_q = 0$ since the gradient is no longer assumed coercive and its deviation moments can be taken as bounded. In this context, we obtain the following weaker (compared to Theorem 1) ergodicity result for QC-SGD.

Theorem 2 (Ergodicity). *Let Assumptions 1, 3, 4 and 5 hold with $A_q = 0$ (uniformly bounded moments) and let \mathcal{L} be an objective such that $\inf_{\theta} \mathcal{L}(\theta) > -\infty$ is finite. Let $(\theta_t)_{t \geq 0}$ be the Markov chain generated by QC-SGD with step size β and quantile $p \in [\eta, 1 - \eta]$. Assume that p and β are such that $3p(1 - \eta)/4 > L\beta + \eta$ and that the subset of \mathbb{R}^d given by*

$$\left\{ \theta : \frac{1}{2} \|\nabla \mathcal{L}(\theta)\|^2 \leq \frac{B_q^2 ((1-p)^{-\frac{2}{q}} (L\beta + 2\eta^2) + 2\eta^{2-\frac{2}{q}})}{p(1-\eta)(3p(1-\eta)/4 - L\beta - \eta)} \right\} \quad (14)$$

is bounded. Then, for any initial $\theta_0 \in \mathbb{R}^d$, there exists $M < +\infty$ such that after T iterations

$$\|\delta_{\theta_0} P_{\beta,p}^T - \pi_{\beta,p}\|_{\text{TV}} \leq \frac{M}{T}, \quad (15)$$

where $\pi_{\beta,p}$ is a unique invariant measure and where δ_{θ_0} is the Dirac measure located at θ_0 .

The proof is given in Appendix D.10 and uses ergodicity results from (Meyn & Tweedie, 1993, Chapter 13). Theorem 2 provides convergence conditions for an SGD Markov chain on a smooth objective in a robust setting. We are unaware of anterior results of this kind in the literature. Condition (14) requires that the set where the true gradient norm is smaller than the estimation error is bounded. This aims to exclude the possibility that the iteration gets trapped within this set and keep using unreliable gradient estimates causing it to diverge. The result is stronger when the upperbound in (14) is smaller. Note that setting p close to $1 - \eta$ increases the clipping threshold and the estimation error as a consequence, making this condition harder to satisfy. On the other hand, using $\beta = \mathcal{O}(1/L)$ and a more conservative value of p makes the upperbound of order $\mathcal{O}(B_q^2)$ and condition 14 easier to satisfy. Observe that, for no corruption ($\eta = 0$), the condition is always fulfilled for some β and p . Note also that without strong convexity (Assumption 2),

convergence occurs at a slower sublinear rate which is consistent with the optimization rate expected for a smooth objective (see (Bubeck, 2015, Theorem 3.3)).

As previously, we complement Theorem 1 with a characterization of the invariant distribution.

Proposition 3. *Under the conditions of Theorem 2, assume that the choice $p = 1 - \eta$ is such that the set (14) is bounded. For step size $\beta \leq \eta^2/L$, the stationary measure $\theta \sim \pi_{\beta, 1-\eta}$ satisfies*

$$\mathbb{E}\|\nabla\mathcal{L}(\theta)\|^2 \leq \frac{5\eta^{2-\frac{2}{q}}B_q^2}{p(1-\eta)(3p(1-\eta)/4 - L\beta - \eta)}. \quad (16)$$

The statement of Proposition 3 is clearly less informative than Propositions 1 and 2 since it only pertains to the gradient rather than, for example, the excess risk. This is due to the weaker assumptions that do not allow to relate these quantities. Still, the purpose remains to find a critical point and is achieved up to $\mathcal{O}(\eta^{1-1/q})$ precision according to this result. Due to corruption, the estimation error on the gradient cannot be reduced beyond $\Omega(\eta^{1-1/q})$ (Prasad et al., 2020; Hopkins & Li, 2018; Diakonikolas & Kane, 2019). Therefore, one may draw a parallel with a corrupted mean estimation task, in which case, the previous rate is, in fact, information-theoretically optimal.

5 Implementation and Numerical Experiments

The use of the generally unknown quantile $Q_p(\|\tilde{G}(\theta_t, \zeta_t)\|)$ in QC-SGD constitutes the main obstacle to its implementation. For strongly convex objectives, one may use a proxy such as $a\|\theta_t - \theta_{\text{ref}}\| + b$ with positive a, b and $\theta_{\text{ref}} \in \mathbb{R}^d$ an approximation of θ^* serving as reference point. This choice is consistent with Assumptions 1 and 5, see Lemma 2 in Appendix D. For instances of Problem (1) defined with an

Algorithm 2: Rolling QC-SGD

Input: Step size $\beta > 0$, quantile index $p \in (0, 1)$, initial parameter $\theta_0 \in \mathbb{R}^d$, $\tau_{\text{unif}} > 0$, buffer B of size S and horizon T .

- 1 Fill B with $S - 1$ values equal to τ_{unif} .
- 2 **for** $t = 0 \dots T - 1$ **do**
- 3 Draw a sample $G(\theta_t, \zeta_t)$ and add $\|G(\theta_t, \zeta_t)\|$ to B .
- 4 $\hat{Q}_p \leftarrow \lfloor pS \rfloor$ rank element of B .
- 5 $\theta_{t+1} \leftarrow \theta_t - \beta \text{clip}(G(\theta_t, \zeta_t), \hat{Q}_p)$
- 6 Delete the oldest value in B .
- 7 **return** θ_T

asymptotically linear function ℓ such as the logistic, hinge or Huber’s loss, a constant threshold can be used since the gradient is a priori uniformly bounded, implying the same for the quantiles of its deviations. In practice, we propose a simpler and more direct approach: we use a rolling quantile procedure, described in Algorithm 2. The latter stores the values $(\|G(\theta_{t-j}, \zeta_{t-j})\|)_{1 \leq j \leq S}$ in a buffer of size $S \in \mathbb{N}^*$ and replaces $Q_p(\|\tilde{G}(\theta_t, \zeta_t)\|)$ in QC-SGD by an estimate \hat{Q}_p which is the $\lfloor pS \rfloor$ -th order statistic in the buffer. Note that only the norms of previous gradients are stored in the buffer, limiting the memory overhead to $\mathcal{O}(S)$. The computational cost of \hat{Q}_p can also be kept to $\mathcal{O}(S)$ per iteration thanks to a bookkeeping procedure (see Appendix B).

Note that, since Algorithm 2 uses the corrupted samples $G(\theta_t, \zeta_t)$ rather than the true ones $\tilde{G}(\theta_t, \zeta_t)$ to estimate the quantiles, a more conservative upperbound of roughly $p \leq 1 - 2\eta$ should be respected when an estimate of η is available. Otherwise, one may default to $p = 1/2$ as an initial guess and adapt based on performance. In practice, our experiments show that relatively low values within $p \in [0.1, 0.2]$ are best for strongly convex objectives while higher values are affordable in other cases. See Appendix B for more details.

We implement this procedure for a few tasks and compare its performance with relevant baselines. We do not include a comparison with (Diakonikolas et al., 2022) whose procedure has no implementation we are aware of and is difficult to use in practice. Indeed, the algorithm in question heavily depends on several problem parameters and involves a filtering procedure which requires multiple passes on large data mini-batches making it impractical for the streaming setting. Moreover, a number of special methods are required to mitigate the costs of the matrix operations needed in the original procedure making the algorithm’s implementation even more involved.

Our experiments on synthetic data consider an infinite horizon, dimension $d = 128$, and a constant step size for all methods.

Linear regression. We consider least-squares linear regression and compare RQC-SGD with Huber’s estimator (Huber, 1973) and clipped SGD (designated as CClip(λ)) with three clipping levels $\lambda\sigma_{\max}\sqrt{d}$ for $\lambda \in \{0.8, 1.0, 1.2\}$ where σ_{\max} is a fixed data scaling factor. These thresholds provide a rough estimate of the gradient norm near the optimum θ^* . We generate covariates X and labels Y both heavy-tailed and corrupted. Corruption in the data stream is generated according to Assumption 3 with outliers represented either by aberrant values or *fake* samples $Y = X^\top\theta_{\text{fake}} + \epsilon$ using a false parameter θ_{fake} , see Appendix B for further details on data generation and fine tuning of the Huber parameter. All methods are run with constant step size and averaged results over 100 runs are displayed on Figure 1 (top row).

As anticipated, Huber’s loss function is not robust to corrupted covariates. In contrast, using gradient clipping allows convergence to meaningful estimates. Although this holds true for a constant threshold, Figure 1 shows it may considerably slow the convergence if started away from the optimum. In addition, the clipping level also affects the final estimation precision and requires tuning. Both of the previous issues are well addressed by RQC-SGD whose adaptive clipping level allows fast progress of the optimization and accurate convergence towards a small neighborhood of the optimum.

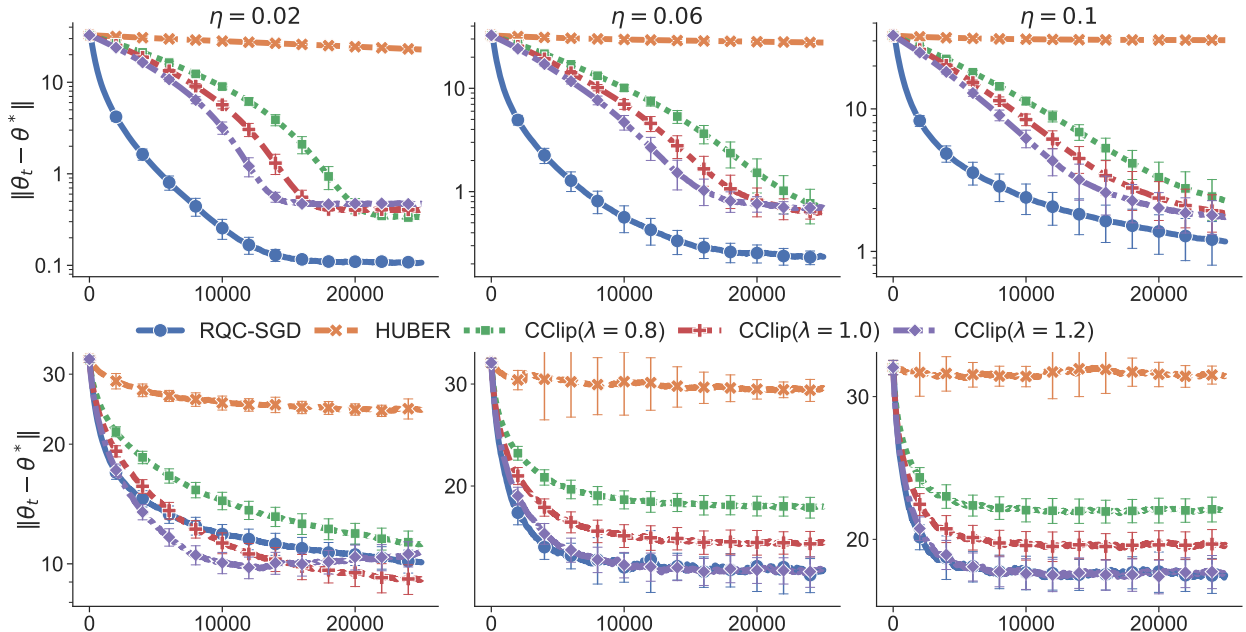


Figure 1: Evolution of $\|\theta_t - \theta^*\|$ on the tasks of linear regression (top row) and logistic regression (bottom row) averaged over 100 runs at increasing corruption levels (error bars represent half the standard deviation). Estimators based on Huber’s loss are strongly affected by data corruption. SGD with constant clipping thresholds is robust but slow to converge for linear regression and requires tuning for better final precision. RQC-SGD combines fast convergence with good final precision thanks to its adaptive clipping strategy.




Figure 2: Evolution of the test loss (y -axis) against iteration t (x -axis) for the training of a single hidden layer network on different real world classification datasets (average over 20 runs). We observe more consistent and stable objective decrease for RQC-SGD whereas constant clipping baselines are slower and may fail to converge.

Logistic regression. We test the same methods on logistic regression. Huber’s baseline is represented by the modified Huber loss (also known as quadratic SVM (Zhang, 2004)). We generate data similarly to the previous task except for the labels which follow $Y \sim \text{Bernoulli}(\sigma(X^\top \theta^*))$ with σ the sigmoid function. Corrupted labels are either uninformative, flipped or obtained with a fake θ_{fake} (see details in Appendix B). Results are displayed on the bottom row of Figure 1.

As previously, Huber’s estimator performs poorly with corruption. However, constant clipping appears to be better suited when the gradient is bounded, so that the optimization is less affected by its underestimation. We observe, nonetheless, that a higher clipping level may lead to poor convergence properties, even at a low corruption rate. Note also that the constant levels we use are based on prior knowledge about the data distribution and would have to be fine tuned in practice. Meanwhile, the latter issue is well addressed by quantile clipping. Finally, notice that no algorithm truly approaches the true solution for this task. This reflects the difficulty of improving upon Proposition 3 which only states convergence to a neighborhood where the objective gradient is comparable to the estimation error in magnitude.

Classification with shallow networks. Finally, we evaluate the performance on the task of training a single hidden layer neural network classifier on some real datasets which corresponds to a non-convex optimization problem. To handle multiclass data, we use the cross entropy loss and replace Huber’s baseline with plain SGD for simplicity. We define constant clipping baselines using thresholds given by the quantiles of order $p = 0.25, 0.5,$ and 0.75 of the norms of a batch of gradients at the beginning of the optimisation. Due to the greater sensitivity to corruption observed in this case, we set $\eta = 0.02$ and use $p = 0.9$ for RQC-SGD. We train all methods with one sample per iteration using equal step sizes and evaluate them through the test loss. We provide further results and experimental details in Appendix B. Results are displayed on Figure 2.

Unsurprisingly, standard SGD is not robust to corrupted samples and, while using a constant clipping level helps keep the optimisation on track, the experiments show that careful tuning may sometimes be necessary to prevent divergence. On the other hand, the adaptive clipping levels used by RQC-SGD allow to make the iteration faster *and* more resilient to corruption. This leads to an optimization path with a more consistent decrease of the objective. Moreover, we also observe that RQC-SGD allows for a better control of the asymptotic variance of the optimized parameter compared to constant clipping.

6 Conclusion

We introduced a new clipping strategy for SGD and proved that it defines a stochastic optimization procedure which is robust to both heavy tails and outliers in the data stream. We also provided an efficient rolling

quantile procedure to implement it and demonstrated its performance through numerical experiments on synthetic and real data. Future research directions include improving the dimension dependence in our bounds, possibly by using sample rejection rules or by considering stochastic mirror descent (Nemirovskij & Yudin, 1983; Beck & Teboulle, 2003) clipped with respect to a non Euclidean norm. This may also procure robustness to higher corruption rates. Another interesting research track is the precise quantification of the geometric convergence speed of the Markov chain generated by constant step size SGD on a strongly convex objective.

Broader Impact Statement

In this optional section, TMLR encourages authors to discuss possible repercussions of their work, notably any potential negative impact that a user of this research should be aware of. Authors should consult the TMLR Ethics Guidelines available on the TMLR website for guidance on how to approach this subject.

Author Contributions

If you'd like to, you may include a section for author contributions as is done in many journals. This is optional and at the discretion of the authors. Only add this information once your submission is accepted and deanonymized.

Acknowledgments

Use unnumbered third level headings for the acknowledgments. All acknowledgments, including those to funding agencies, go at the end of the paper. Only add this information once your submission is accepted and deanonymized.

References

- Noga Alon, Yossi Matias, and Mario Szegedy. The space complexity of approximating the frequency moments. In *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, pp. 20–29, 1996.
- Ainesh Bakshi and Adarsh Prasad. Robust linear regression: Optimal rates in polynomial time. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pp. 102–115, 2021.
- Martyna Bator. Dataset for Sensorless Drive Diagnosis. UCI Machine Learning Repository, 2015. DOI: <https://doi.org/10.24432/C5VP5F>.
- Peter H Baxendale. Renewal theory and computable convergence rates for geometrically ergodic Markov chains. *The Annals of Applied Probability*, 15(1B):700–738, 2005.
- Amir Beck and Marc Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Research Letters*, 31(3):167–175, 2003.
- Witold Bednorz. The Kendall's Theorem and its Application to the Geometric Ergodicity of Markov Chains. *arXiv preprint arXiv:1301.1481*, 2013.
- Albert Benveniste, Michel Métivier, and Pierre Priouret. *Adaptive Algorithms and Stochastic Approximations*, volume 22. Springer Science & Business Media, 2012.
- Kenneth S Berenhaut and Robert Lund. Geometric renewal convergence rates from hazard rates. *Journal of Applied Probability*, 38(1):180–194, 2001.
- Jock Blackard. Covertypes. UCI Machine Learning Repository, 1998. DOI: <https://doi.org/10.24432/C50K5N>.
- Joseph K Blitzstein and Jessica Hwang. *Introduction to probability*. Crc Press, 2019.
- D. A. Bloch. A note on the estimation of the location parameter of the cauchy distribution. *Journal of the American Statistical Association*, 61:852–855, 1966.

- Léon Bottou and Yann Cun. Large scale online learning. In S. Thrun, L. Saul, and B. Schölkopf (eds.), *Advances in Neural Information Processing Systems*, volume 16. MIT Press, 2003. URL https://proceedings.neurips.cc/paper_files/paper/2003/file/9fb7b048c96d44a0337f049e0a61ff06-Paper.pdf.
- Léon Bottou, Frank E Curtis, and Jorge Nocedal. Optimization methods for large-scale machine learning. *SIAM review*, 60(2):223–311, 2018.
- Léon Bottou and Yann Lecun. On-line learning for very large data sets. *Applied Stochastic Models in Business and Industry*, 21:137 – 151, 03 2005. doi: 10.1002/asmb.538.
- Sébastien Bubeck. Convex optimization: Algorithms and complexity. *Foundations and Trends® in Machine Learning*, 8(3-4):231–357, 2015.
- Hervé Cardot, Peggy Cénac, and Pierre-André Zitt. Efficient and fast estimation of the geometric median in Hilbert spaces with an averaged stochastic gradient algorithm. *Bernoulli*, 19(1):18 – 43, 2013. doi: 10.3150/11-BEJ390. URL <https://doi.org/10.3150/11-BEJ390>.
- Hervé Cardot, Peggy Cénac, and Antoine Godichon-Baggioni. Online estimation of the geometric median in Hilbert spaces: Nonasymptotic confidence balls. *The Annals of Statistics*, 45(2):591 – 614, 2017. doi: 10.1214/16-AOS1460. URL <https://doi.org/10.1214/16-AOS1460>.
- Olivier Catoni. Challenging the empirical mean and empirical variance: A deviation study. *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*, 48(4):1148 – 1185, 2012. doi: 10.1214/11-AIHP454. URL <https://doi.org/10.1214/11-AIHP454>.
- Olivier Catoni and Iaria Giulini. Dimension-free pac-bayesian bounds for the estimation of the mean of a random vector. *arXiv preprint arXiv:1802.04308*, 2018.
- Han-fu Chen and AI-JUN Gao. Robustness analysis for stochastic approximation algorithms. *Stochastics and Stochastic Reports*, 26(1):3–20, 1989.
- Han-Fu Chen, Lei Guo, and Ai-Jun Gao. Convergence and robustness of the robbins-monro algorithm truncated at randomly varying bounds. *Stochastic Processes and their Applications*, 27:217–231, 1987.
- Yeshwanth Cherapanamjeri, Nicolas Flammarion, and Peter L. Bartlett. Fast mean estimation with sub-gaussian rates. In *Conference on Learning Theory*, pp. 786–806. PMLR, 2019.
- Aaron Defazio, Francis Bach, and Simon Lacoste-Julien. Saga: A fast incremental gradient method with support for non-strongly convex composite objectives. *Advances in Neural Information Processing Systems*, 27, 2014.
- Jules Depersin and Guillaume Lecué. Robust sub-Gaussian estimation of a mean vector in nearly linear time. *The Annals of Statistics*, 50(1):511–536, 2022. doi: 10.1214/21-AOS2118. URL <https://doi.org/10.1214/21-AOS2118>.
- Ilias Diakonikolas and Daniel M Kane. Recent advances in algorithmic high-dimensional robust statistics. *arXiv preprint arXiv:1911.05911*, 2019.
- Ilias Diakonikolas and Daniel M Kane. *Algorithmic high-dimensional robust statistics*. Cambridge university press, 2023.
- Ilias Diakonikolas, Gautam Kamath, Daniel Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Robust estimators in high-dimensions without the computational intractability. *SIAM Journal on Computing*, 48(2):742–864, 2019.
- Ilias Diakonikolas, Daniel M Kane, and Ankit Pensia. Outlier robust mean estimation with subgaussian rates via stability. *Advances in Neural Information Processing Systems*, 33:1830–1840, 2020.

- Ilias Diakonikolas, Daniel M Kane, Ankit Pensia, and Thanasis Pittas. Streaming algorithms for high-dimensional robust statistics. In *International Conference on Machine Learning*, pp. 5061–5117. PMLR, 2022.
- Aymeric Dieuleveut, Alain Durmus, and Francis Bach. Bridging the gap between constant step size stochastic gradient descent and Markov chains. *The Annals of Statistics*, 48(3):1348 – 1382, 2020. doi: 10.1214/19-AOS1850. URL <https://doi.org/10.1214/19-AOS1850>.
- R. Douc, E. Moulines, and Jeffrey S. Rosenthal. Quantitative bounds on convergence of time-inhomogeneous Markov chains. *The Annals of Applied Probability*, 14(4):1643 – 1665, 2004. doi: 10.1214/105051604000000620. URL <https://doi.org/10.1214/105051604000000620>.
- Eduard Gorbunov, Marina Danilova, and Alexander Gasnikov. Stochastic optimization with heavy-tailed noise via accelerated gradient clipping. *Advances in Neural Information Processing Systems*, 33:15042–15053, 2020.
- Eduard Gorbunov, Abdurakhmon Sadiev, Marina Danilova, Samuel Horváth, Gauthier Gidel, Pavel Dvurechensky, Alexander Gasnikov, and Peter Richtárik. High-probability convergence for composite and distributed stochastic minimization and variational inequalities with heavy-tailed noise. *arXiv preprint arXiv:2310.01860*, 2023.
- Frank R Hampel. A general qualitative definition of robustness. *The Annals of Mathematical Statistics*, 42(6):1887–1896, 1971.
- Abdelhakim Hannousse and Salima Yahiouche. Web page phishing detection. *Mendeley Data*, 2, 2020.
- Matthew Holland and Kazushi Ikeda. Better generalization with less data using robust gradient descent. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 2761–2770. PMLR, 09–15 Jun 2019. URL <https://proceedings.mlr.press/v97/holland19a.html>.
- Samuel B. Hopkins. Mean estimation with sub-Gaussian rates in polynomial time. *The Annals of Statistics*, 48(2):1193–1213, 2020. doi: 10.1214/19-AOS1843. URL <https://doi.org/10.1214/19-AOS1843>.
- Samuel B Hopkins and Jerry Li. Mixture models, robustness, and sum of squares proofs. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 1021–1034, 2018.
- Daniel Hsu and Sivan Sabato. Loss minimization and parameter estimation with heavy tails. *The Journal of Machine Learning Research*, 17(1):543–582, 2016.
- Peter J Huber. A robust version of the probability ratio test. *The Annals of Mathematical Statistics*, pp. 1753–1758, 1965.
- Peter J Huber. The 1972 wald lecture robust statistics: A review. *The Annals of Mathematical Statistics*, 43(4):1041–1067, 1972.
- Peter J. Huber. Robust Regression: Asymptotics, Conjectures and Monte Carlo. *The Annals of Statistics*, 1(5):799 – 821, 1973. doi: 10.1214/aos/1176342503. URL <https://doi.org/10.1214/aos/1176342503>.
- Peter J Huber. Robust estimation of a location parameter. *Breakthroughs in statistics: Methodology and distribution*, pp. 492–518, 1992.
- Daniel C Jerison. Quantitative convergence rates for reversible Markov chains via strong random times. *arXiv preprint arXiv:1908.06459*, 2019.
- Mark R Jerrum, Leslie G Valiant, and Vijay V Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical Computer Science*, 43:169–188, 1986.
- Rie Johnson and Tong Zhang. Accelerating stochastic gradient descent using predictive variance reduction. *Advances in Neural Information Processing Systems*, 26, 2013.

- Anatoli Juditsky, Andrei Kulunchakov, and Hlib Tsyntseus. Sparse recovery by reduced variance stochastic approximation. *Information and Inference: A Journal of the IMA*, 12(2):851–896, 2023.
- Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *International Conference on Learning Representations*, 12 2014.
- Harold J. Kushner and G. George Yin. *Stochastic Approximation and Recursive Algorithms and Applications*, volume 35. Springer New York, NY, 2003.
- Guanghui Lan. *First-Order and Stochastic Optimization Methods for Machine Learning*, volume 1. Springer, 2020.
- Guillaume Lecu e and Matthieu Lerasle. Learning from mom’s principles: Le cam’s approach. *Stochastic Processes and their applications*, 129(11):4385–4410, 2019.
- Guillaume Lecu e and Matthieu Lerasle. Robust machine learning by median-of-means: Theory and practice. *The Annals of Statistics*, 48, 11 2017. doi: 10.1214/19-AOS1828.
- Zhixian Lei, Kyle Luh, Prayaag Venkat, and Fred Zhang. A fast spectral algorithm for mean estimation with sub-Gaussian rates. In *Conference on Learning Theory*, pp. 2598–2612. PMLR, 2020.
- David A Levin and Yuval Peres. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.
- Liu Liu, Yanyao Shen, Tianyang Li, and Constantine Caramanis. High dimensional robust sparse regression. In *International Conference on Artificial Intelligence and Statistics*, pp. 411–421. PMLR, 2020.
- G abor Lugosi and Shahar Mendelson. Sub-gaussian estimators of the mean of a random vector. *The Annals of Statistics*, 47(2):783–794, 2019.
- G abor Lugosi and Shahar Mendelson. Robust multivariate mean estimation: The optimality of trimmed mean. *The Annals of Statistics*, 49(1):393 – 410, 2021. doi: 10.1214/20-AOS1961. URL <https://doi.org/10.1214/20-AOS1961>.
- Robert B Lund, Sean P Meyn, and Richard L Tweedie. Computable exponential convergence rates for stochastically ordered Markov processes. *The Annals of Applied Probability*, 6(1):218–237, 1996.
- Siyuan Ma, Raef Bassily, and Mikhail Belkin. The power of interpolation: Understanding the effectiveness of SGD in modern over-parametrized learning. In *International Conference on Machine Learning*, pp. 3325–3334. PMLR, 2018.
- Rainer Martin and Carl Masreli e. Robust estimation via stochastic approximation. *IEEE Transactions on Information Theory*, 21(3):263–271, 1975. doi: 10.1109/TIT.1975.1055386.
- H Brendan McMahan, Gary Holt, David Sculley, Michael Young, Dietmar Ebner, Julian Grady, Lan Nie, Todd Phillips, Eugene Davydov, Daniel Golovin, et al. Ad click prediction: a view from the trenches. In *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1222–1230, 2013.
- Sean P Meyn and Richard L Tweedie. *Markov Chains and Stochastic Stability*. Springer London, 1993.
- Sean P Meyn and Robert L Tweedie. Computable bounds for geometric convergence rates of Markov chains. *The Annals of Applied Probability*, pp. 981–1011, 1994.
- Stanislav Minsker. Geometric median and robust estimation in Banach spaces. *Bernoulli*, 21(4):2308 – 2335, 2015. doi: 10.3150/14-BEJ645. URL <https://doi.org/10.3150/14-BEJ645>.
- Alexander V Nazin, Boris T Polyak, and Alexandre B Tsybakov. Optimal and robust kernel algorithms for passive stochastic approximation. *IEEE Transactions on Information Theory*, 38(5):1577–1583, 1992.

- Alexander V Nazin, Arkadi S Nemirovsky, Alexandre B Tsybakov, and Anatoli B Juditsky. Algorithms of robust stochastic optimization based on mirror descent method. *Automation and Remote Control*, 80: 1607–1627, 2019.
- Arkadij Semenovič Nemirovskij and David Borisovich Yudin. *Problem Complexity and Method Efficiency in Optimization*. A Wiley-Interscience publication. Wiley, 1983. ISBN 9780471103455. URL <https://books.google.fr/books?id=6ULvAAAAAAAJ>.
- Yurii Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*. Springer Publishing Company, Incorporated, 1 edition, 2014. ISBN 1461346916.
- Ta Duy Nguyen, Thien Hang Nguyen, Alina Ene, and Huy Le Nguyen. High probability convergence of clipped-SGD under heavy-tailed noise. *arXiv preprint arXiv:2302.05437*, 2023.
- Ankit Pensia, Varun Jog, and Po-Ling Loh. Robust regression with covariate filtering: Heavy tails and adversarial contamination. *arXiv preprint arXiv:2009.12976*, 2020.
- Boris Teodorovich Polyak and Yakov Zalmanovich Tsypkin. Adaptive estimation algorithms: convergence, optimality, stability. *Automation and Remote Control*, 40(3):378–389, 1979.
- Boris Teodorovich Polyak and Yakov Zalmanovich Tsypkin. Robust pseudogradient adaptation algorithms. *Automation and Remote Control*, 41(10):1404–1409, 1981.
- Adarsh Prasad, Arun Sai Suggala, Sivaraman Balakrishnan, and Pradeep Ravikumar. Robust estimation via robust gradient estimation. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 82, 2018.
- Adarsh Prasad, Sivaraman Balakrishnan, and Pradeep Ravikumar. A robust univariate mean estimator is all you need. In *International Conference on Artificial Intelligence and Statistics*, pp. 4034–4044. PMLR, 2020.
- E Price and V VandeLinde. Robust estimation using the robbins-monro stochastic approximation algorithm. *IEEE Transactions on Information Theory*, 25(6):698–704, 1979.
- Alexander Rakhlin, Ohad Shamir, and Karthik Sridharan. Making gradient descent optimal for strongly convex stochastic optimization. *arXiv preprint arXiv:1109.5647*, 2011.
- Herbert Robbins and Sutton Monro. A stochastic approximation method. *The Annals of Mathematical Statistics*, pp. 400–407, 1951.
- Gareth O Roberts and Richard L Tweedie. Rates of convergence of stochastically monotone and continuous time Markov models. *Journal of Applied Probability*, 37(2):359–373, 2000.
- Byron Roe. MiniBooNE particle identification. UCI Machine Learning Repository, 2010. DOI: <https://doi.org/10.24432/C5QC87>.
- Jeffrey S Rosenthal. Convergence rates for Markov chains. *Siam Review*, 37(3):387–405, 1995a.
- Jeffrey S Rosenthal. Minorization conditions and convergence rates for markov chain monte carlo. *Journal of the American Statistical Association*, 90(430):558–566, 1995b.
- Thomas J. Rothenberg, Franklin M. Fisher, and C. B. Tilanus. A note on estimation from a cauchy sample. *Journal of the American Statistical Association*, 59(306):460–463, 1964. doi: 10.1080/01621459.1964.10482170. URL <https://www.tandfonline.com/doi/abs/10.1080/01621459.1964.10482170>.
- Peter J Rousseeuw and Mia Hubert. Robust statistics for outlier detection. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, 1(1):73–79, 2011.

- Abdurakhmon Sadiev, Marina Danilova, Eduard Gorbunov, Samuel Horváth, Gauthier Gidel, Pavel Dvurechensky, Alexander Gasnikov, and Peter Richtárik. High-probability bounds for stochastic optimization and variational inequalities: the case of unbounded variance. In *International Conference on Machine Learning*, pp. 29563–29648. PMLR, 2023.
- Prem Seetharaman, Gordon Wichern, Bryan Pardo, and Jonathan Le Roux. Autoclip: Adaptive gradient clipping for source separation networks. In *2020 IEEE 30th International Workshop on Machine Learning for Signal Processing (MLSP)*, pp. 1–6. IEEE, 2020.
- Shai Shalev-Shwartz, Yoram Singer, and Nathan Srebro. Pegasos: Primal estimated sub-gradient solver for svm. In *Proceedings of the 24th International Conference on Machine Learning*, pp. 807–814, 2007.
- Shai Shalev-Shwartz, Ohad Shamir, Nathan Srebro, and Karthik Sridharan. Stochastic convex optimization. In *Conference on Learning Theory*, volume 2, pp. 5, 2009.
- Srdjan S Stanković and Branko D Kovačević. Analysis of robust stochastic approximation algorithms for process identification. *Automatica*, 22(4):483–488, 1986.
- Che-Ping Tsai, Adarsh Prasad, Sivaraman Balakrishnan, and Pradeep Ravikumar. Heavy-tailed streaming statistical estimation. In *International Conference on Artificial Intelligence and Statistics*, pp. 1251–1282. PMLR, 2022.
- John Wilder Tukey. A survey of sampling from contaminated distributions. *Contributions to Probability and Statistics*, pp. 448–485, 1960.
- Andrew V Uzilov, Joshua M Keegan, and David H Mathews. Detection of non-coding rnas on the basis of predicted secondary structure formation free energy change. *BMC bioinformatics*, 7(1):1–30, 2006.
- Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*, volume 47. Cambridge university press, 2018.
- Nuri Mert Vural, Lu Yu, Krishna Balasubramanian, Stanislav Volgushev, and Murat A Erdogdu. Mirror descent strikes again: Optimal stochastic convex optimization under infinite noise variance. In *Conference on Learning Theory*, pp. 65–102. PMLR, 2022.
- Lu Yu, Krishnakumar Balasubramanian, Stanislav Volgushev, and Murat A Erdogdu. An analysis of constant step size SGD in the non-convex regime: Asymptotic normality and bias. *Advances in Neural Information Processing Systems*, 34:4234–4248, 2021.
- Jingzhao Zhang, Sai Praneeth Karimireddy, Andreas Veit, Seungyeon Kim, Sashank Reddi, Sanjiv Kumar, and Suvrit Sra. Why are adaptive methods good for attention models? *Advances in Neural Information Processing Systems*, 33:15383–15393, 2020.
- Tong Zhang. Solving large scale linear prediction problems using stochastic gradient descent algorithms. In *Proceedings of the Twenty-First International Conference on Machine Learning*, pp. 116, 2004.

Supplementary Material

Robust Stochastic Optimization via Gradient Quantile Clipping

A Additional experimental results

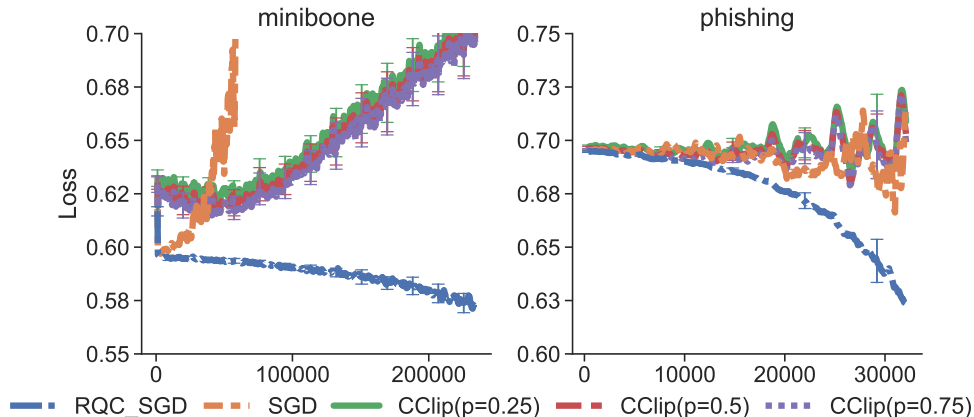


Figure 3: Evolution of the test loss (y -axis) against iteration t (x -axis) for the training of a single hidden layer network on additional real world classification datasets (average over 20 runs).

Classification with shallow networks. We performed the same experiment using two additional datasets. The results are displayed on Figure 3 and corroborate our statements in the main paper.

Expectation estimation. We estimate the expectation of a random vector X by minimizing the objective $\mathcal{L}(\theta) = \frac{1}{2}\|\theta - \theta^*\|^2$ with $\theta^* = \mathbb{E}[X]$ using a stream of both corrupted and heavy-tailed samples, see Appendix B for details. We run RQC-SGD (Algorithm 2) and compare it to an online version of geometric and coordinate-wise Median-Of-Means (GMOM and CMOM (Cardot et al., 2017; 2013)) which use block sample means to minimize an L_1 objective (see Appendix B). Although these estimators are a priori not robust to η -corruption, we ensure that their estimates are meaningful by limiting η to 4% and using blocks of 10 samples. Thus, blocks are corrupted with probability $< 1/2$ so that the majority contains only true samples. Figure 4 displays the evolution of $\|\theta_t - \theta^*\|$ for each method averaged over 100 runs for increasing η and constant step size. We also display a single run for $\eta = 0.04$. We observe that RQC-SGD is only weakly affected by the increasing corruption whereas the performance of GMOM and CMOM quickly degrades with η , leading to unstable estimates.

B Experimental details

As previously mentioned, the dimension is set to $d = 128$ in our experiments with synthetic data. We also set $\sigma_{\min} = 1$ and $\sigma_{\max} = 5$ as minimum and maximum scaling factors. For all tasks and algorithms, the optimization starts from $\theta_0 = 0$.

Bookkeeping in RQC-SGD The buffer in Algorithm 2 stores values in sorted order along with their “ages”. The most recent and oldest values have ages 0 and $S - 1$ respectively. At each iteration, a new gradient is received, all ages are incremented and the oldest value is replaced by the new one with age 0. The latter is then sorted using one iteration of insertion sort. The estimate \hat{Q}_p is retrieved at each iteration as the value at position $\lfloor pS \rfloor$.




Figure 4: Evolution of $\|\theta_t - \theta^*\|$ (y -axis) against iteration t (x -axis) for the expectation estimation task, averaged over 100 runs at different corruption levels η (bands widths correspond to the standard deviation of the 100 runs). For $\eta = 0.04$, the evolution on a single run is also displayed. We observe good performance for RQC-SGD for increasing η while CMOM and GMOM are more sensitive.

B.1 Mean estimation

Data generation We compute a matrix $\Sigma = (AA^\top + A^\top A)/2$ where $A \in \mathbb{R}^{d \times d}$ is a random matrix with i.i.d centered Gaussian entries with variance $1/d$ sampled once and for all. We generate true samples as $X = \mathbf{1} + \Sigma V$ where V is a vector of i.i.d symmetrized Pareto random variables with parameter 2 and $\mathbf{1} \in \mathbb{R}^d$ denotes the vector with all entries equal to 1.

We draw corrupted samples as $\check{X} = 10\check{V} - 100 \times \mathbf{1}$ where \check{V} is a vector of i.i.d symmetrized Pareto variables with parameter 1.5. We use step size $\beta = 10^{-3}$.

GMOM and CMOM The geometric and coordinatewise Median-Of-Means estimators (GMOM and CMOM) optimize the following objectives respectively:

$$\mathbb{E}\|\theta - \bar{X}_{N_b}\|_2 \quad \text{and} \quad \mathbb{E}\|\theta - \bar{X}_{N_b}\|_1,$$

where \bar{X}_{N_b} is the average of N_b independent copies of X . The block size is set to $N_b = 10$ in the whole experiment. The above objectives are optimized by computing samples of \bar{X}_{N_b} in a streaming fashion so that one step is made for each N_b samples. In order to compensate for this inefficiency we multiply the step size by N_b for both GMOM and CMOM. For GMOM, we additionally multiply the step size by \sqrt{d} in order to compensate the normalization included in the gradient formula.

RQC-SGD For mean estimation, we implement RQC-SGD (Algorithm 2) with buffer size $S = 100$, $p = 0.2$ and $\tau_{\text{unif}} = 10$.

B.2 Linear regression

Data generation We choose the true parameter θ^* by independently sampling its coordinate uniformly in the interval $[-5, +5]$. The true covariates are sampled as $X = \Sigma V$ where Σ is a diagonal matrix with entries sampled uniformly in the interval $[\sigma_{\min}, \sigma_{\max}]$ (once and for all) and V is a vector of i.i.d symmetrized Pareto random variables with parameter 2. The labels are sampled as $Y = X^\top \theta^* + \epsilon$ where ϵ is a symmetrized Pareto random variable with parameter 2.

The corrupted samples are obtained according to one of the following possibilities with equal probability:

- $X = 1000(\max_i \Sigma_{ii})v + W$ where v is a fixed unit vector and W is a standard Gaussian vector and $Y \sim \text{Bernoulli}(1/2)$.
- $X = 1000(\max_i \Sigma_{ii})V$ with V a unit norm random vector with uniform distribution and $Y = 1000(Z + U)$ where Z is a random sign and U is uniform over $[-1/5, 1/5]$.

- $X = 10V$ with V a random vector with i.i.d entries following a symmetrized Log-normal distribution and $Y = X^\top \theta_{\text{fake}} + \epsilon$ with θ_{fake} a fake parameter drawn similarly to θ^* once and for all and ϵ a standard Gaussian variable.

We use step size $\beta = 10^{-3}$.

Huber parameter In order to tune the parameter δ of Huber’s loss function, we proceed as follows:

- For each corruption level η , we consider 10 candidate values $\delta_j = 10^{j/2-5}$ for $0 \leq j < 10$.
- For each candidate δ_j , we train 250 estimators $(\hat{\theta}_{\delta_j}^{(i)})_{i \in \llbracket 250 \rrbracket}$ using 1000 samples each.
- We choose \hat{j} for which the average $\frac{1}{250} \sum_{i \in \llbracket 250 \rrbracket} \|\hat{\theta}_{\delta_j}^{(i)} - \theta^*\|$ is minimal and use $\hat{\delta} = \delta_{\hat{j}}$ as parameter.

RQC-SGD For linear regression, we run RQC-SGD with buffer size $S = 100$ and $\tau_{\text{unif}} = 10$. The quantile value was set to $p = 0.1$ for $\eta \in \{0.02, 0.06\}$ and $p = 0.05$ for $\eta = 0.1$.

B.3 Logistic regression

Data generation The true parameter θ^* and covariates X are chosen similarly to linear regression. Given X , the label Y is set to $+1$ with probability $\sigma(X^\top \theta^*)$ where σ is the sigmoid function $\sigma(x) = (1 + e^{-x})^{-1}$ and to -1 otherwise.

The corrupted covariates are determined similarly to linear regression while the labels are set as follows in each respective case:

- Y is set to $+1$ or -1 with equal probability.
- $Y = -\text{sign}(X^\top \theta^*)$.
- $Y = \text{sign}(X^\top \theta_{\text{fake}})$ with θ_{fake} a fake parameter drawn similarly to θ^* once and for all.

We use step size $\beta = 6 \times 10^{-3}$.

Huber parameter The same procedure is used to tune the parameter of the modified Huber loss as for linear regression.

RQC-SGD For logistic regression, we run RQC-SGD with buffer size $S = 100$ and $\tau_{\text{unif}} = 10$. The quantile value was set to $p = 1 - \eta - 0.1$ for $\eta = 0.02$ and $p = 1 - \eta - 0.05$ otherwise.

B.4 Single hidden layer neural network classifier

We train a single hidden layer neural network classifier with 100 hidden neurones for all datasets. We use one sample per iteration and step size $\beta = 10^{-2}$ for all methods.

As previously, RQC-SGD is run with buffer size $S = 100$ and $\tau_{\text{unif}} = 10$. The quantile value was set to $p = 0.9$. We compute the gradient norms over a batch of samples of size S at the beginning of the optimization and use the quantiles of order $p = 0.25, 0.5$ and 0.75 as the clipping level for the constant clipping baselines.

Data We used publicly available datasets for our experiments. We provide details about their characteristics and sources in Table 1.

We use a 10% share of each dataset as a test set in order to compute the test loss plotted in Figures 2 and 3. We also ensure the test set contains at least 5000 elements. Optimization is run using the remaining train set which is corrupted as specified next. The results are averaged over 20 runs for each datasets.

Dataset	# Samples	# Features	# Classes	Source
Codrna (Uzilov et al., 2006)	488,565	8	2	OpenML
Sensorless (Bator, 2015)	58,509	48	11	UCI
Covtype (Blackard, 1998)	581,012	52	7	scikit-learn
Miniboone (Roe, 2010)	130,065	50	2	UCI
Phishing (Hannousse & Yahiouche, 2020)	11,430	87	2	Kaggle

Table 1: Main characteristics of the data sets used in experiments, including number of samples, number of features, number of classes and sources.

Data corruption We corrupt train data samples at each iteration uniformly at random with probability $\eta = 0.02$.

Although we run the optimization using one sample per iteration, the datasets we use are available offline so that we have a data matrix denoted $\mathbf{X} \in \mathbb{R}^{n \times (d+1)}$ whose last column represents the labels. This corresponds to n samples and d features.

For each feature $j \in \llbracket d \rrbracket$, we compute $\hat{\mu}_j$ and $\hat{\sigma}_j$ the empirical mean and standard deviation respectively. We also sample a random unit vector u of size d and introduce corruption as follows:

- For the label column, we introduce corruption by changing the value uniformly at random among the other possible modalities.
- For features, we introduce corruption by replacing the original values with one of the following possibilities with equal probability:
 - a vector ξ sampled coordinatewise according to $\xi_j = r_j + 1000 \times \hat{\sigma}_j \nu$ where r_j is a value randomly picked in the column $\mathbf{X}_{\bullet, j}$ and ν is a sample from the Student distribution with 2.1 degrees of freedom.
 - a vector ξ sampled coordinatewise according to $\xi_j = \hat{\mu}_j + 1000 \times \hat{\sigma}_j u_j + z$ where z is a standard Gaussian.
 - a vector ξ sampled according to $\xi = \hat{\mu} + 1000 \times \hat{\sigma} \otimes w$ where w is a uniformly sampled unit vector.

C Geometric convergence speed and relation to step size

The geometric Markov chain convergence stated in Theorem 1 occurs at a speed determined by the contraction factor ρ which mainly depends on the step size β and quantile p defining the iteration. Therefore, an explicit formulation of this dependency is necessary to precisely quantify the convergence speed. This question is lightly touched upon in (Yu et al., 2021) whose Proposition 2.1 is an analogous SGD ergodicity result. Like Theorem 1, the latter relies on the Markov chain theory presented in (Meyn & Tweedie, 1993). It is argued in (Yu et al., 2021) that a vanishing step size $\beta \rightarrow 0$ causes ρ to be close to one, leading to slow convergence but with smaller bias. However, these considerations remain asymptotic and do not quite address the convergence speed issue.

More generally, the precise estimation of the factor ρ goes back to the evaluation of the convergence speed of a Markov chain satisfying a geometric drift property. Near optimal results exist for chains with particular properties such as stochastic order (Lund et al., 1996; Roberts & Tweedie, 2000), reversibility (Jerison, 2019) or special assumptions on the renewal distribution (Berenhaut & Lund, 2001). Unfortunately, such properties do not hold for SGD. Let $(\theta_t)_{t \geq 0}$ be a Markov chain satisfying the drift property:

$$\Delta V(\theta) \leq \begin{cases} (1 - \lambda)V(\theta) & \text{for } \theta \notin \mathcal{C} \\ b & \text{for } \theta \in \mathcal{C} \end{cases}$$

with $\lambda \in (0, 1)$, $b < +\infty$, V a real function such that $V(\theta) \geq 1$ for all θ and \mathcal{C} a (bounded) small set (see (Meyn & Tweedie, 1993, Chapter 5)). Then, based on the available literature (Baxendale, 2005; Bednorz, 2013),

$(\theta_t)_{t \geq 0}$ converges as in Theorem 1 with $\rho \approx 1 - \lambda^3$, the latter estimation being unimprovable without further information on $(\theta_t)_{t \geq 0}$ (see the discussion following (Baxendale, 2005, Theorem 3.2)). For the specific setting of Theorem 1 (and more generally for SGD by setting $p = 1$), this only yields an excessively pessimistic estimate

$$\rho \approx 1 - (p\beta\mu)^3 \quad (17)$$

whereas it is reasonable to conjecture that $\rho \approx 1 - p\beta\mu$. The suboptimality of (17) is felt in the uncorrupted case in Proposition 2 and Corollary 1 where one is tempted to set β of order $1/T$, with T the horizon, reducing the bias to $\mathcal{O}(1/\sqrt{T})$. However, this results in an unacceptable sample cost of order T^3 before convergence occurs. On the other hand, assuming the estimate $\rho \approx 1 - p\beta\mu$ holds, using a step size of order $\log(T)/T$ allows to combine fast convergence and near optimal statistical performance. Finally, note that in the corrupted case, the optimal statistical rate is $\mathcal{O}(\eta^{1-1/q})$ so that striking such a compromise is unnecessary.

D Proofs

D.1 Preliminary lemmas

Lemma 1. *Grant Assumptions 1 and 2. For any $\theta, \theta' \in \mathbb{R}^d$ and $\beta \leq \frac{2}{\mu+L}$ we have :*

$$\|\theta - \beta\nabla\mathcal{L}(\theta) - (\theta' - \beta\nabla\mathcal{L}(\theta'))\|^2 \leq (1 - \beta\mu)^2 \|\theta - \theta'\|^2 \quad (18)$$

Proof. For $\beta \leq \frac{2}{\mu+L}$, we have:

$$\begin{aligned} & \|\theta - \beta\nabla\mathcal{L}(\theta) - (\theta' - \beta\nabla\mathcal{L}(\theta'))\|^2 \\ &= \|\theta - \theta'\|^2 - 2\beta\langle\theta - \theta', \nabla\mathcal{L}(\theta) - \nabla\mathcal{L}(\theta')\rangle + \beta^2\|\nabla\mathcal{L}(\theta) - \nabla\mathcal{L}(\theta')\|^2 \\ &\leq (1 - \beta^2\mu L)\|\theta - \theta'\|^2 - \beta(2 - \beta(\mu + L))\langle\nabla\mathcal{L}(\theta) - \nabla\mathcal{L}(\theta'), \theta - \theta'\rangle \\ &\leq (1 - \beta^2\mu L)\|\theta - \theta'\|^2 - \beta(2 - \beta(\mu + L))\mu\|\theta - \theta'\|^2 \\ &= (1 - \beta^2\mu L - 2\beta\mu + \beta^2\mu(\mu + L))\|\theta - \theta'\|^2 \\ &= (1 - \beta\mu)^2\|\theta - \theta'\|^2, \end{aligned}$$

where we used the inequalities :

$$\|\nabla\mathcal{L}(\theta) - \nabla\mathcal{L}(\theta')\|^2 \leq (\mu + L)\langle\nabla\mathcal{L}(\theta) - \nabla\mathcal{L}(\theta'), \theta - \theta'\rangle - \mu L\|\theta - \theta'\|^2 \quad (19)$$

$$\mu\|\theta - \theta'\|^2 \leq \langle\nabla\mathcal{L}(\theta) - \nabla\mathcal{L}(\theta'), \theta - \theta'\rangle, \quad (20)$$

valid for all θ, θ' . Inequality (19) is stated, for example, in (Nesterov, 2014, Theorem 2.1.12) (see also (Bubeck, 2015, Lemma 3.11) and (20) is just a characterization of strong convexity (see for instance (Nesterov, 2014, Theorem 2.1.9)). \square

Note that Lemma 1 is a simple generalization of the usual contraction property in strongly-convex optimization where θ' is usually taken as θ^* . In that case, one has $\nabla\mathcal{L}(\theta') = 0$. An example of such a result is given by (Bubeck, 2015, Theorem 3.10). A similar inequality to Lemma 1 was previously obtained in (Dieuleveut et al., 2020, Proposition 2) where convergence of SGD as Markov chain was studied.

In the sequel we will write gradient samples $G(\theta, \zeta)$ and $\tilde{G}(\theta, \zeta)$ simply as $G(\theta)$ and $\tilde{G}(\theta)$ respectively in order to lighten notation. The following lemma will be needed in the proof of Theorem 1.

Lemma 2. *Let Assumptions 4 and 5 hold. Let $\theta \in \mathbb{R}^d$ be fixed and let $\tilde{G}(\theta) \sim \mathcal{D}_{\mathcal{I}}(\theta)$ be a non corrupted gradient sample. Choosing the clipping threshold as $\tau_\theta = Q_p(\|\tilde{G}(\theta)\|)$ for some $p \in (0, 1)$ and denoting $\alpha_\theta = \min\left(1, \frac{\tau_\theta}{\|\tilde{G}(\theta)\|}\right)$ the clipping factor and its average $\bar{\alpha}_\theta = \mathbb{E}[\alpha_\theta | \theta, G(\theta) = \tilde{G}(\theta) \sim \mathcal{D}_{\mathcal{I}}(\theta)]$ we have:*

$$\|\mathbb{E}[\alpha_\theta \tilde{G}(\theta)] - \bar{\alpha}_\theta \nabla\mathcal{L}(\theta)\| \leq (1 - p)^{1-1/q} (A_q \|\theta - \theta^*\| + B_q), \quad (21)$$

$$\begin{aligned}\tau_\theta &\leq \|\nabla\mathcal{L}(\theta)\| + Q_p(\|\varepsilon_\theta\|) \\ &\leq \|\nabla\mathcal{L}(\theta)\| + (1-p)^{-1/q}(A_q\|\theta - \theta^*\| + B_q).\end{aligned}\tag{22}$$

If Assumption 5 holds with $q \geq 2$ then we also have

$$\mathbb{E}\|\alpha_\theta\tilde{G}(\theta) - \mathbb{E}[\alpha_\theta\tilde{G}(\theta)]\|^2 \leq (A_q\|\theta - \theta^*\| + B_q)^2 + 5(1-p)\tau_\theta^2.\tag{23}$$

Proof. We condition on the event that the sample $G(\theta)$ is not corrupted i.e. $G(\theta) = \tilde{G}(\theta) \sim \mathcal{D}_{\mathcal{I}}(\theta)$. Noticing that $\mathbf{1}_{\|\tilde{G}(\theta)\| \leq \tau_\theta} = 1 - \mathbf{1}_{\|\tilde{G}(\theta)\| > \tau_\theta}$ and using the equality $\mathbb{E}[\tilde{G}(\theta)] = \nabla\mathcal{L}(\theta)$, we find :

$$\begin{aligned}\mathbb{E}[\alpha_\theta\tilde{G}(\theta)] - \bar{\alpha}_\theta\nabla\mathcal{L}(\theta) &= \mathbb{E}[(\alpha_\theta - \bar{\alpha}_\theta)(\tilde{G}(\theta) - \nabla\mathcal{L}(\theta))] \\ &= \mathbb{E}[(1 - \bar{\alpha}_\theta)(\tilde{G}(\theta) - \nabla\mathcal{L}(\theta))\mathbf{1}_{\|\tilde{G}(\theta)\| \leq \tau_\theta}] + \mathbb{E}[(\alpha_\theta - \bar{\alpha}_\theta)(\tilde{G}(\theta) - \nabla\mathcal{L}(\theta))\mathbf{1}_{\|\tilde{G}(\theta)\| > \tau_\theta}] \\ &= (\bar{\alpha}_\theta - 1)\mathbb{E}[(\tilde{G}(\theta) - \nabla\mathcal{L}(\theta))\mathbf{1}_{\|\tilde{G}(\theta)\| > \tau_\theta}] - \bar{\alpha}_\theta\mathbb{E}[(\tilde{G}(\theta) - \nabla\mathcal{L}(\theta))\mathbf{1}_{\|\tilde{G}(\theta)\| > \tau_\theta}] \\ &\quad + \mathbb{E}[(\tau_\theta/\|\tilde{G}(\theta)\|)(\tilde{G}(\theta) - \nabla\mathcal{L}(\theta))\mathbf{1}_{\|\tilde{G}(\theta)\| > \tau_\theta}] \\ &= -\mathbb{E}[(1 - \tau_\theta/\|\tilde{G}(\theta)\|)(\tilde{G}(\theta) - \nabla\mathcal{L}(\theta))\mathbf{1}_{\|\tilde{G}(\theta)\| > \tau_\theta}].\end{aligned}$$

Using our choice of τ_θ and Hölder's inequality, we find :

$$\begin{aligned}\|\mathbb{E}[\alpha_\theta\tilde{G}(\theta)] - \bar{\alpha}_\theta\nabla\mathcal{L}(\theta)\| &\leq \mathbb{E}[\mathbf{1}_{\|\tilde{G}(\theta)\| > \tau_\theta} |1 - \tau_\theta/\|\tilde{G}(\theta)\||\|\tilde{G}(\theta) - \nabla\mathcal{L}(\theta)\|] \\ &\leq \mathbb{E}[\mathbf{1}_{\|\tilde{G}(\theta)\| > \tau_\theta} \|\tilde{G}(\theta) - \nabla\mathcal{L}(\theta)\|] \\ &\leq (1-p)^{1-1/q}\mathbb{E}[\|\tilde{G}(\theta) - \nabla\mathcal{L}(\theta)\|^q]^{1/q} \\ &\leq (1-p)^{1-1/q}(A_q\|\theta - \theta^*\| + B_q),\end{aligned}$$

where the second step corresponds to the inequality $|1 - \tau_\theta/\|\tilde{G}(\theta)\|| \leq 1$ under the event $\{\|\tilde{G}(\theta)\| > \tau_\theta\}$. Note also that we used Assumption 5 since θ is fixed. Inequality (21) is now proven. To show (22), we first write the inequality:

$$\tau_\theta = Q_p(\|\tilde{G}(\theta)\|) = Q_p(\|\nabla\mathcal{L}(\theta) + \varepsilon_\theta\|) \leq \|\nabla\mathcal{L}(\theta)\| + Q_p(\|\varepsilon_\theta\|),$$

which holds since $\|\tilde{G}(\theta)\|$ is a positive random variable. Further, using Assumption 5, we have:

$$1 - p = \mathbb{P}(\|\varepsilon_\theta\| > Q_p(\|\varepsilon_\theta\|)) \leq \frac{\mathbb{E}[\|\varepsilon_\theta\|^q]}{Q_p(\|\varepsilon_\theta\|)^q} \leq \left(\frac{A_q\|\theta - \theta^*\| + B_q}{Q_p(\|\varepsilon_\theta\|)}\right)^q.$$

It only remains to take the q -th root and plug the obtained bound on $Q_p(\|\varepsilon_\theta\|)$ back above to obtain (22).

To show (23), we define the event $\mathcal{E} = \{\|\tilde{G}(\theta)\| \leq \tau_\theta\}$ and denote $\bar{\mathcal{E}}$ its complement such that $\mathbb{P}(\mathcal{E}) = p = 1 - \mathbb{P}(\bar{\mathcal{E}})$. We write

$$\begin{aligned}\mathbb{E}\|\alpha_\theta\tilde{G}(\theta) - \mathbb{E}[\alpha_\theta\tilde{G}(\theta)]\|^2 &= p\mathbb{E}[\|\alpha_\theta\tilde{G}(\theta) - \mathbb{E}[\alpha_\theta\tilde{G}(\theta)]\|^2|\mathcal{E}] \\ &\quad + (1-p)\mathbb{E}[\|\alpha_\theta\tilde{G}(\theta) - \mathbb{E}[\alpha_\theta\tilde{G}(\theta)]\|^2|\bar{\mathcal{E}}] \\ &\leq p\mathbb{E}[\|\alpha_\theta\tilde{G}(\theta) - \mathbb{E}[\alpha_\theta\tilde{G}(\theta)]\|^2|\mathcal{E}] + 4(1-p)\tau_\theta^2 \\ &= p\mathbb{E}[\|\alpha_\theta\tilde{G}(\theta) - \mathbb{E}[\alpha_\theta\tilde{G}(\theta)|\mathcal{E}]\|^2|\mathcal{E}] \\ &\quad + p\|\mathbb{E}[\alpha_\theta\tilde{G}(\theta)|\mathcal{E}] - \mathbb{E}[\alpha_\theta\tilde{G}(\theta)]\|^2 + 4(1-p)\tau_\theta^2 \\ &= p\mathbb{E}[\|\alpha_\theta\tilde{G}(\theta) - \mathbb{E}[\alpha_\theta\tilde{G}(\theta)|\mathcal{E}]\|^2|\mathcal{E}] \\ &\quad + p\|(1-p)\mathbb{E}[\alpha_\theta\tilde{G}(\theta)|\mathcal{E}] - (1-p)\mathbb{E}[\alpha_\theta\tilde{G}(\theta)|\bar{\mathcal{E}}]\|^2 + 4(1-p)\tau_\theta^2 \\ &\leq p\mathbb{E}[\|\alpha_\theta\tilde{G}(\theta) - \mathbb{E}[\alpha_\theta\tilde{G}(\theta)|\mathcal{E}]\|^2|\mathcal{E}] + 4p(1-p)^2\tau_\theta^2 + 4(1-p)\tau_\theta^2 \\ &\leq p\mathbb{E}[\|\alpha_\theta\tilde{G}(\theta) - \mathbb{E}[\alpha_\theta\tilde{G}(\theta)|\mathcal{E}]\|^2|\mathcal{E}] + 5(1-p)\tau_\theta^2\end{aligned}$$

where we used the identity $\mathbb{E}[\alpha_\theta \tilde{G}(\theta)] = p\mathbb{E}[\alpha_\theta \tilde{G}(\theta)|\mathcal{E}] + (1-p)\mathbb{E}[\alpha_\theta \tilde{G}(\theta)|\bar{\mathcal{E}}]$ and the inequalities $\|\alpha_\theta \tilde{G}(\theta)\| \leq \tau_\theta$ and $p(1-p) \leq 1/4$ which hold in all cases. In addition, we have

$$\begin{aligned} p\mathbb{E}[\|\alpha_\theta \tilde{G}(\theta) - \mathbb{E}[\alpha_\theta \tilde{G}(\theta)|\mathcal{E}]\|^2|\mathcal{E}] &= p\mathbb{E}[\|\tilde{G}(\theta) - \mathbb{E}[\tilde{G}(\theta)|\mathcal{E}]\|^2|\mathcal{E}] \leq \mathbb{E}[\|\tilde{G}(\theta) - \mathbb{E}[\tilde{G}(\theta)]\|^2] \\ &\leq \mathbb{E}[\|\tilde{G}(\theta) - \nabla\mathcal{L}(\theta)\|^q]^{2/q} \leq (A_q\|\theta - \theta^*\| + B_q)^2. \end{aligned}$$

The first inequality is obtained by applying Lemma 3 to each coordinate of $\tilde{G}(\theta)$ while conditioning on θ . The second one uses Jensen's inequality and the third results from Assumption 5. \square

The statement of Lemma 2 appears to be novel and is enabled by the specific choice of the gradient norm quantile as clipping threshold.

Lemma 3. *Let Y be a real random variable and \mathcal{E} an event, then we have the inequality*

$$\mathbb{P}(\mathcal{E})\mathbb{E}[(Y - \mathbb{E}[Y|\mathcal{E}])^2|\mathcal{E}] \leq \mathbb{E}(Y - \mathbb{E}[Y])^2.$$

Proof. Define the conditional variance of a real random variable Y w.r.t. another variable X as $\text{Var}(Y|X) = \mathbb{E}[(Y - \mathbb{E}[Y|X])^2|X]$.

By Eve's law (see for instance (Blitzstein & Hwang, 2019, Theorem 9.5.4)) we have the identity

$$\text{Var}(Y) = \mathbb{E}[\text{Var}(Y|X)] + \text{Var}(\mathbb{E}[Y|X]),$$

which entails the inequality $\text{Var}(Y) \geq \mathbb{E}[\text{Var}(Y|X)]$. Applying the latter with $X = \mathbf{1}_\mathcal{E}$ yields

$$\mathbb{E}(Y - \mathbb{E}[Y])^2 = \text{Var}(Y) \geq \mathbb{P}(\mathcal{E})\mathbb{E}[(Y - \mathbb{E}[Y|\mathcal{E}])^2|\mathcal{E}] + (1 - \mathbb{P}(\mathcal{E}))\mathbb{E}[(Y - \mathbb{E}[Y|\bar{\mathcal{E}}])^2|\bar{\mathcal{E}}],$$

which implies the result. \square

We show the geometric ergodicity of the SGD Markov chain $(\theta_t)_{t \geq 0}$ by relying on (Meyn & Tweedie, 1993, Theorem 15.0.1). We will show that the following function :

$$V(\theta) := 1 + \|\theta - \theta^*\|^2,$$

satisfies a *geometric drift* property. We define the action of the transition kernel $P_{\beta,p}$ on real integrable functions f over \mathbb{R}^d through:

$$P_{\beta,p}f(\theta) = \int f(\theta')P_{\beta,p}(\theta, d\theta') = \mathbb{E}[f(\theta - \alpha_\theta\beta G(\theta))].$$

We also define the variation operator:

$$\Delta f(\theta) := P_{\beta,p}f(\theta) - f(\theta).$$

In many of our proofs, we will make use of the following adjustable bound.

Fact 1. *For any real numbers a, b and positive ϵ , we have the inequality*

$$2ab \leq a^2\epsilon + b^2/\epsilon.$$

D.2 Preliminary properties for Markov chains

In order to prove Theorem 1, we set up the formalism we need from (Meyn & Tweedie, 1993) through the following definitions.

Definition 2 (Small set). *A set $\mathcal{C} \in \mathcal{B}(\mathbb{R}^d)$ is called a small set if there exists an $m > 0$, and a non-trivial measure ν on $\mathcal{B}(\mathbb{R}^d)$ such that for all $\theta \in \mathcal{C}$ and $B \in \mathcal{B}(\mathbb{R}^d)$,*

$$P_{\beta,p}^m(\theta, B) \geq \nu(B).$$

When this is the case, we say that \mathcal{C} is (m, ν) -small.

Minorization properties such as the one above are useful for proving the convergence of Markov chains. We define an analogous one via the notion of sampled chain.

Definition 3 (Sampled chain). *Let a be a probability measure on \mathbb{Z}_+ i.e. such that $a(n) \geq 0$ for $n \geq 0$ and $\sum_{n=0}^{\infty} a(n) = 1$. Then the transition kernel for the sampled Markov chain w.r.t. the distribution a is*

$$K_a(\theta, A) = \sum_{n=0}^{\infty} P_{\beta,p}^n(\theta, A) a(n) \quad \text{for } x \in \mathbb{R}^d, A \in \mathcal{B}(\mathbb{R}^d).$$

We define the notion of a *petite set* for a sampled Markov chain.

Definition 4 (Petite set). *Let a be a probability measure on \mathbb{Z}_+ defining a sampled chain. A set $C \in \mathcal{B}(\mathbb{R}^d)$ is called a (a, ν) -petite if*

$$K_a(\theta, B) \geq \nu(B),$$

for all $\theta \in C$ and $B \in \mathcal{B}(\mathbb{R}^d)$ where ν is a non-trivial measure on $\mathcal{B}(\mathbb{R}^d)$.

Note that an (m, ν) -small set is also (a, ν) -petite with $a = \delta_m$. Finally, we define the norm of a measure w.r.t. a potential function.

Definition 5 (f -norm). *Let $f : \mathbb{R}^d \rightarrow \mathbb{R}$ be a function such that $f \geq 1$ and let ν be a signed measure on \mathbb{R}^d . We define the f -norm of ν as*

$$\|\nu\|_f = \sup_{g:|g|\leq f} |\nu(g)| = \sup_{g:|g|\leq f} \left| \int g(\theta) \nu(d\theta) \right|.$$

In particular, we have $\|\nu\|_f \geq \|\nu\|_{\text{TV}}$.

D.3 Proof of Theorem 1

We will assume in this proof that $q \geq 2$. The case $q \in (1, 2)$ will be treated in Appendix D.5.

First, we define $\underline{\tau} = \inf_{\theta} \tau_{\theta}$. Note that Assumption 4 excludes the existence of θ such that $\tilde{G}(\theta) = 0$ almost surely, therefore, we have $\underline{\tau} > 0$.

Thanks to Assumption 4 and conditioning on $\theta_t = \theta \in \mathbb{R}^d$ for $t \geq 0$, the distribution of θ_{t+1} has a strictly positive density at least on a ball of radius $\beta \underline{\tau}$ around θ_t . This implies that the chain is aperiodic since $P_{\beta,p}(\theta_t, W_{\theta_t}) > 0$ for any neighborhood W_{θ_t} contained in the previous ball. Moreover, by induction, the distribution of θ_{t+m} has positive density at least on a ball of radius $m\beta \underline{\tau}$ around θ_t . Thus for m high enough we have $P_{\beta,p}^m(\theta_t, A) > 0$ for any set A with non zero Lebesgue measure. It follows that the Markov chain is irreducible w.r.t. Lebesgue's measure and is thus ψ -irreducible (see (Meyn & Tweedie, 1993, Chapter 4)).

For fixed θ , condition (9) implies $\beta < \frac{2}{\mu+L}$ and using Lemma 1 and denoting α_{θ} and $\bar{\alpha}_{\theta}$ as in Lemma 2, we find:

$$\begin{aligned} P_{\beta,p} \|\theta - \theta^*\|^2 &= \mathbb{E} \|\theta - \beta \alpha_{\theta} G(\theta) - \theta^*\|^2 \\ &\leq \eta \mathbb{E} (\|\theta - \theta^*\| + \beta \tau_{\theta})^2 + (1 - \eta) \mathbb{E} [\|\theta - \alpha_{\theta} \beta \tilde{G}(\theta) - \theta^*\|^2] \\ &\leq \eta \mathbb{E} (\|\theta - \theta^*\| + \beta \tau_{\theta})^2 + (1 - \eta) \mathbb{E} [\|\theta - \bar{\alpha}_{\theta} \beta \nabla \mathcal{L}(\theta) - \theta^*\|^2 - \\ &\quad 2\beta \langle \theta - \bar{\alpha}_{\theta} \beta \nabla \mathcal{L}(\theta) - \theta^*, \alpha_{\theta} \tilde{G}(\theta) - \bar{\alpha}_{\theta} \nabla \mathcal{L}(\theta) \rangle + \beta^2 \|\alpha_{\theta} \tilde{G}(\theta) - \bar{\alpha}_{\theta} \nabla \mathcal{L}(\theta)\|^2] \\ &\leq \eta \mathbb{E} (\|\theta - \theta^*\| + \beta \tau_{\theta})^2 + (1 - \eta) \mathbb{E} [\|\theta - \bar{\alpha}_{\theta} \beta \nabla \mathcal{L}(\theta) - \theta^*\|^2 \\ &\quad + \beta \|\mathbb{E}[\alpha_{\theta} \tilde{G}(\theta)] - \bar{\alpha}_{\theta} \nabla \mathcal{L}(\theta)\|^2] + \beta^2 (A_q \|\theta - \theta^*\| + B_q)^2 + 5\beta^2 (1 - p) \tau_{\theta}^2, \end{aligned}$$

where the last step uses that $\mathbb{E} \|\alpha_{\theta} \tilde{G}(\theta) - \bar{\alpha}_{\theta} \nabla \mathcal{L}(\theta)\|^2 = \mathbb{E} \|\alpha_{\theta} \tilde{G}(\theta) - \mathbb{E}[\alpha_{\theta} \tilde{G}(\theta)]\|^2 + \|\mathbb{E}[\alpha_{\theta} \tilde{G}(\theta)] - \bar{\alpha}_{\theta} \nabla \mathcal{L}(\theta)\|^2$ and Lemma 2.

Using Lemmas 1 and 2 and Assumption 5 and grouping the terms by powers of $\|\theta - \theta^*\|$, we arrive at

$$\begin{aligned}
P_{\beta,p}\|\theta - \theta^*\|^2 &\leq (\eta + (1 - \eta)(1 - \bar{\alpha}_\theta\beta\mu)^2)\mathbb{E}\|\theta - \theta^*\|^2 + \\
&\quad 2\beta\mathbb{E}[\|\theta - \theta^*\|(\eta\tau_\theta + (1 - \eta)(1 - \bar{\alpha}_\theta\beta\mu)\|\mathbb{E}[\alpha_\theta\tilde{G}(\theta)] - \bar{\alpha}_\theta\nabla\mathcal{L}(\theta)\|)] + \\
&\quad \beta^2(\eta\tau_\theta^2 + (1 - \eta)(\|\mathbb{E}[\alpha_\theta\tilde{G}(\theta)] - \bar{\alpha}_\theta\nabla\mathcal{L}(\theta)\|^2 + (A_q\|\theta - \theta^*\| + B_q)^2 + 5(1 - p)\tau_\theta^2)) \\
&\leq \mathbb{E}[\mathfrak{A}\|\theta - \theta^*\|^2 + 2\beta\mathfrak{B}\|\theta - \theta^*\|] + \beta^2\mathfrak{C} \\
&\leq (\mathfrak{A} + \beta\kappa)\mathbb{E}[\|\theta - \theta^*\|^2] + \frac{\beta\mathfrak{B}^2}{\kappa} + \beta^2\mathfrak{C},
\end{aligned} \tag{24}$$

where we used Fact 1 and defined the quantities \mathfrak{A} , \mathfrak{B} , \mathfrak{C} which may be bounded as follows

$$\begin{aligned}
\mathfrak{A} &= \eta + (1 - \eta)(1 - \bar{\alpha}_\theta\beta\mu)^2 + 2\beta(\eta(L + A_q(1 - p)^{-1/q} + (1 - \eta)(1 - \bar{\alpha}_\theta\beta\mu)(1 - p)^{-1/q}A_q) \\
&\quad + \beta^2((\eta + 5(1 - p))(L + (1 - p)^{-1/q}A_q)^2 + 4A_q^2) \\
&\leq 1 - 2\beta((1 - \eta)\bar{\alpha}_\theta\mu - \eta(L + A_q(1 - p)^{-1/q} - (1 - \eta)(1 - \bar{\alpha}_\theta\beta\mu)(1 - p)^{-1/q}A_q) + \\
&\quad \beta^2((1 - \eta)(\bar{\alpha}_\theta\mu)^2 + 2(\eta + 5(1 - p))(L + (1 - p)^{-1/q}A_q)^2 + 4A_q^2) \\
&\leq 1 - 2\beta((1 - \eta)p\mu - \eta L - (1 - p)^{-1/q}A_q(1 - p(1 - \eta))) + \beta^2(\mu^2 + 24\eta L^2 + 28A_q^2) \\
&= 1 - 2\beta\kappa + \beta^2(\mu^2 + 24\eta L^2 + 28A_q^2), \\
\mathfrak{B} &= (\eta(1 - p)^{-1/q} + (1 - \eta)(1 - \bar{\alpha}_\theta\beta\mu)(1 - p)^{1-1/q})B_q \\
&\leq (1 - p)^{-1/q}B_q(\eta + (1 - p)) \\
\mathfrak{C} &= (4 + 2(\eta + 5(1 - p))(1 - p)^{-2/q})B_q^2
\end{aligned}$$

The above bounds use the simple properties $p \leq \bar{\alpha}_\theta \leq 1$, $0 \leq \eta \leq 1$, $1 - \bar{\alpha}_\theta\beta\mu \leq 1$ and $(a + b)^2 \leq 2a^2 + 2b^2$ for all a, b .

Thanks to our choice of β , we get that $\mathfrak{A} + \beta\kappa < 1$. It is now easy to check that $V(\theta) = 1 + \|\theta - \theta^*\|^2$ satisfies the contraction

$$P_{\beta,p}V(\theta) \leq \underbrace{(\mathfrak{A} + \beta\kappa)}_{=: \tilde{\lambda} < 1} V(\theta) + \underbrace{1 - (\mathfrak{A} + \beta\kappa) + \frac{\beta\mathfrak{B}^2}{\kappa} + \mathfrak{C}}_{=: \tilde{b}}$$

We now define the set $\mathcal{C} = \{\theta \in \mathbb{R}^d, V(\theta) \leq 2\tilde{b}/(1 - \tilde{\lambda})\}$ for which we have:

$$\Delta V(\theta) \leq -\frac{1 - \tilde{\lambda}}{2}V(\theta) + \tilde{b}\mathbf{1}_{\theta \in \mathcal{C}}. \tag{25}$$

For such \mathcal{C} , let $\Delta_{\mathcal{C}} = \text{diam}(\mathcal{C})$ be its diameter and set $m_{\mathcal{C}} = \left\lceil \frac{\Delta_{\mathcal{C}}}{\beta\tau} \right\rceil$. As previously mentioned in the beginning of the proof, conditioning on $\theta_t = \theta$, the distribution of θ_{t+m} admits a positive density at least over a ball of radius $m\beta\tau$ around θ i.e. there exists $h_\theta^{+m}(\omega) \geq 0$ satisfying

$$h_\theta^{+m}(\omega) > 0 \quad \text{for } \|\theta - \omega\| < m\beta\tau \quad \text{and} \quad P_{\beta,p}^m(\theta, A) \geq \int_A h_\theta^{+m}(\omega)d\omega \quad \text{for } A \in \mathcal{B}(\mathbb{R}^d).$$

We then let $\underline{h}_{\mathcal{C}}(\omega) = \inf_{\theta \in \mathcal{C}} h_\theta^{+m_{\mathcal{C}}}(\omega)$ and define the measure $\nu_{\mathcal{C}}$ by:

$$\nu_{\mathcal{C}}(A) = \int_{A \cap \mathcal{C}} \underline{h}_{\mathcal{C}}(\theta)d\theta.$$

The above measure is non trivial since our choice of $m_{\mathcal{C}}$ ensures that $\underline{h}_{\mathcal{C}}$ defines a non zero density at least on \mathcal{C} . It follows that for all $\theta_0 \in \mathcal{C}$, we have the following minorization property:

$$P_{\beta,p}^{m_{\mathcal{C}}}(\theta_0, A) \geq \nu_{\mathcal{C}}(A) \quad \text{for all } A \in \mathcal{B}(\mathbb{R}^d), \tag{26}$$

which implies that the set \mathcal{C} is $(m_{\mathcal{C}}, \nu_{\mathcal{C}})$ -small and also $(\delta_{m_{\mathcal{C}}, \nu_{\mathcal{C}}})$ -petite as a result. We have previously shown that the Markov chain (θ_t) is irreducible and aperiodic. Moreover, we have exhibited a petite set \mathcal{C} (via (26)) towards which the geometric drift property (25) holds. Thus, condition (iii) of (Meyn & Tweedie, 1993, Theorem 15.0.1) is fulfilled.

By the latter result, it follows that the chain (θ_t) admits a unique invariant probability measure $\pi_{\beta,p}$ and there exist $r > 1$ and $M < \infty$ such that:

$$\sum_{t \geq 0} r^t \|P_{\beta,p}(\theta_0, \cdot) - \pi_{\beta,p}\|_V \leq MV(\theta_0). \quad (27)$$

Taking $\rho = r^{-1}$ and using that $\|\nu\|_V \geq \|\nu\|_{\text{TV}}$ for any signed measure ν concludes the proof.

D.4 Proof of Proposition 1

We consider the case $q \geq 2$. Since the distribution $\pi_{\beta,p}$ is invariant by the transition kernel $P_{\beta,p}$, we can deduce that for $\theta \sim \pi_{\beta,p}$, we have

$$\mathbb{E}\|\theta - \theta^*\|^2 = \mathbb{E}\|\theta - \beta\alpha_{\theta}G(\theta) - \theta^*\|^2.$$

From here, we can follow similar computations to those leading to (24) in the proof of Theorem 1. By setting $p = 1 - \eta$ and using that $q \geq 2$, we additionally have the bounds $\mathfrak{B} \leq 3\eta^{1-1/q}B_q$ and $\mathfrak{C} \leq 16B_q^2$.

Hence, we have that:

$$\begin{aligned} \mathbb{E}\|\theta - \theta^*\|^2 &\leq (\mathfrak{A} + \beta\kappa)\mathbb{E}\|\theta - \theta^*\|^2 + \frac{\beta\mathfrak{B}^2}{\kappa} + \beta^2\mathfrak{C} \\ \implies \mathbb{E}\|\theta - \theta^*\|^2 &\leq \frac{\beta\mathfrak{B}^2}{\kappa(1 - \mathfrak{A} - \beta\kappa)} + \frac{\beta^2\mathfrak{C}}{1 - \mathfrak{A} - \beta\kappa}. \end{aligned} \quad (28)$$

Note that (9) entails

$$1 - \mathfrak{A} - \beta\kappa \geq \beta\kappa - \beta^2(\mu^2 + 24\eta L^2 + 28A_q^2) \geq 3\beta\kappa/4. \quad (29)$$

Plugging this into (28) and using (10), we find

$$\mathbb{E}\|\theta - \theta^*\|^2 \leq \frac{12\eta^{2-2/q}B_q^2}{\kappa^2} + \frac{64\beta B_q^2}{3\kappa} \leq \frac{34\eta^{2-2/q}B_q^2}{\kappa^2}$$

which implies the result.

D.5 The case $q \in (1, 2)$

Similar results hold for the case $q \in (1, 2)$ but require a different proof given below.

Proposition 4. *Let Assumptions 1-5 hold with $q \in (1, 2)$ and let QC-SGD be run with quantile $p \in [\eta, 1 - \eta]$. Assume that*

$$\kappa' := (1 - \eta)p\mu - q\eta L - q(\eta(1 - p)^{-1/q} + (1 - p)^{1-1/q})A_q > 0, \quad (30)$$

and take a step-size satisfying

$$\beta \leq \left(\frac{\kappa'}{86(L + A_q)^q} \right)^{\frac{1}{q-1}}. \quad (31)$$

then the generated Markov chain $(\theta_t)_t$ converges geometrically to a unique invariant measure $\pi_{\beta,1-\eta}$ as in Theorem 1. In addition, for $p = 1 - \eta$, $\beta \leq \eta/\kappa'$ and $\theta \sim \pi_{\beta,1-\eta}$, we have

$$\mathbb{E}\|\theta - \theta^*\|^q \leq 128 \left(\frac{\eta^{1-1/q}B_q}{\kappa'} \right)^q.$$

Proof. We first need to prove that convergence holds as stated in Theorem 1. We establish the properties of irreducibility, aperiodicity and exhibit a petite similarly as before. However, the demonstration of a geometric drift property such as (25) requires a different approach.

In this proof, we will use the following inequalities valid for all positive x, y and ε and $q \in (1, 2)$

$$(x + y)^q \leq 2^{q-1}(x^q + y^q), \quad (32)$$

$$(x + y)^q \leq x^q + qx^{q-1}y + y^q, \quad (33)$$

(a consequence of the inequality $(1 + a)^q \leq 1 + qa + a^q$ for $a > 0$),

$$(x + y)^{q-1} \leq x^{q-1} + y^{q-1}, \quad (34)$$

and

$$xy \leq \frac{(x\varepsilon)^q}{q} + \frac{q-1}{q}(y/\varepsilon)^{q/(q-1)} \quad (35)$$

which is a consequence of Young's inequality applied to the pair $x\varepsilon$ and y/ε with exponent q and its conjugate. We first write

$$P_{\beta,p}\|\theta - \theta^*\|^q = \mathbb{E}\|\theta - \beta\alpha_\theta G(\theta) - \theta^*\|^q \leq \eta\mathbb{E}(\|\theta - \theta^*\| + \beta\tau_\theta)^q + (1 - \eta)\mathbb{E}\|\theta - \alpha_\theta\beta\tilde{G}(\theta) - \theta^*\|^q$$

Defining the notation $\mathbb{E}_\theta[\cdot] = \mathbb{E}[\cdot|\theta]$, we have

$$\begin{aligned} \mathbb{E}\|\theta - \alpha_\theta\beta\tilde{G}(\theta) - \theta^*\|^q &= \mathbb{E}\left[\left(\|\theta - \bar{\alpha}_\theta\beta\nabla\mathcal{L}(\theta) - \theta^*\|^2 - \right. \right. \\ &\quad \left. \left. 2\beta\langle\theta - \bar{\alpha}_\theta\beta\nabla\mathcal{L}(\theta) - \theta^*, \alpha_\theta\tilde{G}(\theta) - \bar{\alpha}_\theta\nabla\mathcal{L}(\theta)\rangle + \beta^2\|\alpha_\theta\tilde{G}(\theta) - \bar{\alpha}_\theta\nabla\mathcal{L}(\theta)\|^2\right)^{q/2}\right] \\ &\leq \mathbb{E}\left[\left(\|\theta - \bar{\alpha}_\theta\beta\nabla\mathcal{L}(\theta) - \theta^*\|^2 - 2\beta\langle\theta - \bar{\alpha}_\theta\beta\nabla\mathcal{L}(\theta) - \theta^*, \alpha_\theta\tilde{G}(\theta) - \bar{\alpha}_\theta\nabla\mathcal{L}(\theta)\rangle + \right. \right. \\ &\quad \left. \left. \beta^2(\|\alpha_\theta\tilde{G}(\theta) - \mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)]\|^2 + 2\langle\alpha_\theta\tilde{G}(\theta) - \mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)], \mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)] - \bar{\alpha}_\theta\nabla\mathcal{L}(\theta)\rangle + \right. \right. \\ &\quad \left. \left. \|\mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)] - \bar{\alpha}_\theta\nabla\mathcal{L}(\theta)\|^2\right)^{q/2}\right] \\ &\leq \mathbb{E}\left[\left(\|\theta - \bar{\alpha}_\theta\beta\nabla\mathcal{L}(\theta) - \theta^*\|^2 - 2\beta\langle\theta - \bar{\alpha}_\theta\beta\nabla\mathcal{L}(\theta) - \theta^*, \alpha_\theta\tilde{G}(\theta) - \bar{\alpha}_\theta\nabla\mathcal{L}(\theta)\rangle + \right. \right. \\ &\quad \left. \left. \beta^2(2\langle\alpha_\theta\tilde{G}(\theta) - \mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)], \mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)] - \bar{\alpha}_\theta\nabla\mathcal{L}(\theta)\rangle + \|\mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)] - \bar{\alpha}_\theta\nabla\mathcal{L}(\theta)\|^2\right)^{q/2}\right] \\ &\quad + \beta^q\mathbb{E}\|\alpha_\theta\tilde{G}(\theta) - \mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)]\|^q \\ &\leq \mathbb{E}\left[\left(\|\theta - \bar{\alpha}_\theta\beta\nabla\mathcal{L}(\theta) - \theta^*\|^2 - 2\beta\langle\theta - \bar{\alpha}_\theta\beta\nabla\mathcal{L}(\theta) - \theta^*, \mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)] - \bar{\alpha}_\theta\nabla\mathcal{L}(\theta)\rangle + \right. \right. \\ &\quad \left. \left. \beta^2\|\mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)] - \bar{\alpha}_\theta\nabla\mathcal{L}(\theta)\|^2\right)^{q/2}\right] + \beta^q\mathbb{E}\|\alpha_\theta\tilde{G}(\theta) - \mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)]\|^q \\ &\leq \mathbb{E}\left(\|\theta - \bar{\alpha}_\theta\beta\nabla\mathcal{L}(\theta) - \theta^*\| + \beta\|\mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)] - \bar{\alpha}_\theta\nabla\mathcal{L}(\theta)\right)^q + \beta^q\mathbb{E}\|\alpha_\theta\tilde{G}(\theta) - \mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)]\|^q, \end{aligned}$$

where we used (34), conditioned on θ , then used Jensen's inequality for \mathbb{E}_θ and a Cauchy-Schwarz inequality. We focus on the last term. Defining the event $\mathcal{E} = \{\|\tilde{G}(\theta)\| \leq \tau_\theta\}$ such that $\mathbb{P}(\mathcal{E}) = p$, we have

$$\begin{aligned} \mathbb{E}\|\alpha_\theta\tilde{G}(\theta) - \mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)]\|^q &= (1 - p)\mathbb{E}[\|\alpha_\theta\tilde{G}(\theta) - \mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)]\|^q|\bar{\mathcal{E}}] \\ &\quad + p\mathbb{E}[\|\alpha_\theta\tilde{G}(\theta) - \mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)]\|^q|\mathcal{E}] \\ &\leq (1 - p)(2\tau_\theta)^q + p\mathbb{E}[\|\alpha_\theta\tilde{G}(\theta) - \mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)]\|^q|\mathcal{E}]. \end{aligned}$$

In addition, using (34) twice, we find

$$\begin{aligned} \mathbb{E}[\|\alpha_\theta\tilde{G}(\theta) - \mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)]\|^q|\mathcal{E}] &\leq 2^{q-1}(\mathbb{E}[\|\alpha_\theta\tilde{G}(\theta) - \bar{\alpha}_\theta\nabla\mathcal{L}(\theta)\|^q|\mathcal{E}] + \\ &\quad \|\bar{\alpha}_\theta\nabla\mathcal{L}(\theta) - \mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)]\|^q) \\ &\leq 2^{q-1}(2^{q-1}\mathbb{E}[\|\tilde{G}(\theta) - \nabla\mathcal{L}(\theta)\|^q|\mathcal{E}] + 2^{q-1}(1 - \bar{\alpha}_\theta)^q\|\nabla\mathcal{L}(\theta)\|^q + \\ &\quad \|\bar{\alpha}_\theta\nabla\mathcal{L}(\theta) - \mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)]\|^q). \end{aligned}$$

Therefore, from the two previous displays and using Assumption 1, Lemma 2 and the inequality $p\mathbb{E}[\|\tilde{G}(\theta) - \nabla\mathcal{L}(\theta)\|^q|\mathcal{E}] \leq \mathbb{E}[\|\tilde{G}(\theta) - \nabla\mathcal{L}(\theta)\|^q]$, we get

$$\begin{aligned} \mathbb{E}\|\alpha_\theta\tilde{G}(\theta) - \mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)]\|^q &\leq (1-p)(2\tau_\theta)^q + 2^{2q-2}p(1-\bar{\alpha}_\theta)^qL^q\|\theta - \theta^*\|^q \\ &\quad + (2^{2q-2} + 2^{q-1}p(1-p)^{q-1})(A_q\|\theta - \theta^*\| + B_q)^q. \end{aligned}$$

We also have

$$\begin{aligned} (1-p)\tau_\theta^q &\leq (1-p)((L + (1-p)^{-1/q}A_q)\|\theta - \theta^*\| + (1-p)^{-1/q}B_q)^q \\ &\leq ((L + A_q)\|\theta - \theta^*\| + B_q)^q. \end{aligned}$$

Plugging into the previous display and simplifying leads to

$$\begin{aligned} \mathbb{E}\|\alpha_\theta\tilde{G}(\theta) - \mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)]\|^q &\leq 2^{2q-2}((L + A_q)\|\theta - \theta^*\| + B_q)^q + 2^{2q-2}p(1-\bar{\alpha}_\theta)^q(L\|\theta - \theta^*\|)^q \\ &\quad + (2^{2q-2} + 2^{q-1}p(1-p)^{q-1})(A_q\|\theta - \theta^*\| + B_q)^q \\ &\leq 2^{2q-1}((L + A_q)\|\theta - \theta^*\| + B_q)^q + (2^{2q-2} + 2^{q-1}p(1-p)^{q-1})(A_q\|\theta - \theta^*\| + B_q)^q \\ &\leq 14((L + A_q)\|\theta - \theta^*\| + B_q)^q. \end{aligned} \tag{36}$$

Using these inequalities along with Lemma 2 to bound $\mathbb{E}\|\theta - \theta^*\|^q$, we find that

$$\begin{aligned} P_{\beta,p}\|\theta - \theta^*\|^q &\leq \eta\mathbb{E}(\|\theta - \theta^*\| + \beta\tau_\theta)^q + (1-\eta)\beta^q\mathbb{E}\|\alpha_\theta\tilde{G}(\theta) - \mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)]\|^q \\ &\quad + (1-\eta)\mathbb{E}(\|\theta - \bar{\alpha}_\theta\beta\nabla\mathcal{L}(\theta) - \theta^*\| + \beta\|\mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)] - \bar{\alpha}_\theta\nabla\mathcal{L}(\theta)\|)^q \\ &\leq \eta\mathbb{E}(\|\theta - \theta^*\|^q + \beta^q\tau_\theta^q + q\beta\|\theta - \theta^*\|^{q-1}\tau_\theta) + (1-\eta)\mathbb{E}((1-\bar{\alpha}_\theta\beta\mu)^q\|\theta - \theta^*\|^q \\ &\quad + \beta^q\|\mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)] - \bar{\alpha}_\theta\nabla\mathcal{L}(\theta)\|^q + q\beta\|\theta - \theta^*\|^{q-1}\|\mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)] - \bar{\alpha}_\theta\nabla\mathcal{L}(\theta)\|) \\ &\quad + (1-\eta)\beta^q\mathbb{E}\|\alpha_\theta\tilde{G}(\theta) - \mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)]\|^q \\ &\leq \mathbb{E}[(1 - (1-\eta)\bar{\alpha}_\theta\beta\mu)\|\theta - \theta^*\|^q] + q\beta\mathbb{E}[\|\theta - \theta^*\|^{q-1}(\eta\tau_\theta \\ &\quad + (1-\eta)\|\mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)] - \bar{\alpha}_\theta\nabla\mathcal{L}(\theta)\|) + \beta^q(\eta\tau_\theta^q \\ &\quad + (1-\eta)(\|\mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)] - \bar{\alpha}_\theta\nabla\mathcal{L}(\theta)\|^q + \|\alpha_\theta\tilde{G}(\theta) - \mathbb{E}_\theta[\alpha_\theta\tilde{G}(\theta)]\|^q))] \\ &\leq (1-\beta((1-\eta)p\mu - q\eta(L + (1-p)^{-1/q}A_q) - q(1-\eta)(1-p)^{1-1/q}A_q) \\ &\quad + \beta^q2^{q-1}(\eta(L + (1-p)^{-1/q}A_q)^q + (1-\eta)14(L + A_q)^q \\ &\quad + (1-\eta)(1-p)^{q-1}A_q^q)\mathbb{E}\|\theta - \theta^*\|^q + q\beta\mathbb{E}\|\theta - \theta^*\|^{q-1}((1-\eta)(1-p)^{1-1/q}B_q \\ &\quad + \eta(1-p)^{-1/q}B_q) + \beta^q2^{q-1}(\eta(1-p)^{-1}B_q^q + (1-\eta)(1-p)^{q-1}B_q^q + 14B_q^q) \\ &\leq \mathfrak{A}\mathbb{E}\|\theta - \theta^*\|^q + q\beta\mathfrak{B}\mathbb{E}\|\theta - \theta^*\|^{q-1} + \beta^q2^{q-1}\mathfrak{C}^q \end{aligned} \tag{37}$$

where the second inequality uses Lemma 1 and (33) twice, the third inequality rearranges according to powers of $\|\theta - \theta^*\|$, the fourth one applies Lemma 2 and (36) and the last inequality defines the factors \mathfrak{A} , \mathfrak{B} and \mathfrak{C} . Since $p \leq 1 - \eta$, these satisfy:

$$\mathfrak{A} \leq 1 - \beta\kappa' + \beta^q2^{q-1}(16)(L + A_q)^q, \quad \mathfrak{B} \leq 2B_q \quad \text{and} \quad \mathfrak{C} \leq 2^{4/q}B_q$$

Applying (35), we have for all $\varepsilon > 0$ that

$$\mathfrak{B}\mathbb{E}\|\theta - \theta^*\|^{q-1} \leq \frac{\varepsilon^{\frac{q}{q-1}}\mathbb{E}\|\theta - \theta^*\|^q}{q/(q-1)} + \frac{\mathfrak{B}/\varepsilon)^q}{q} \leq \frac{\kappa'\mathbb{E}\|\theta - \theta^*\|^q}{8q} + \frac{\mathfrak{B}^q}{q} \left(\frac{\kappa'}{8(q-1)} \right)^{1-q}, \tag{38}$$

where the second inequality corresponds to the choice $\varepsilon = \left(\frac{\kappa'}{8(q-1)} \right)^{\frac{q-1}{q}}$. Plugging back into (37) and using condition (31) ensures that $\mathfrak{A} + \beta\kappa'/8 < 1$. Defining $V(\theta) = 1 + \|\theta - \theta^*\|^q$, we get that satisfies the contraction

$$P_{\beta,p}V(\theta) \leq \underbrace{(\mathfrak{A} + \beta\kappa'/8)}_{=: \tilde{\lambda} < 1} V(\theta) + \underbrace{1 - (\mathfrak{A} + \beta\kappa'/8) + \beta\mathfrak{B}^q \left(\frac{\kappa'}{8(q-1)} \right)^{1-q} + \beta^q2^{q-1}\mathfrak{C}^q}_{=: \tilde{b}}$$

From here, we use the same argument as the previous proof of Theorem 1 to deduce that the Markov chain converges to a unique invariant distribution $\pi_{\beta,p}$. Now, since convergence occurs to a limit distribution $\pi_{\beta,p}$, we have, as before, for $\theta \sim \pi_{\beta,p}$ the identity

$$\mathbb{E}\|\theta - \theta^*\|^q = \mathbb{E}\|\theta - \beta\alpha_\theta G(\theta) - \theta^*\|^q.$$

By setting $p = 1 - \eta$, we can show that in fact $\mathfrak{B} \leq 2\eta^{1-1/q}B_q$. Plugging back into (37) and using (38) and (31), we find

$$\begin{aligned} \mathbb{E}\|\theta - \theta^*\|^q &\leq (1 - (7/8)\beta\kappa' + 32\beta^q(L + A_q)^q)\mathbb{E}\|\theta - \theta^*\|^q + \beta(2\eta^{1-1/q}B_q)^q \left(\frac{\kappa'}{8(q-1)}\right)^{1-q} + 32\beta^q B_q^q \\ &\leq (1 - (3/8)\beta\kappa')\mathbb{E}\|\theta - \theta^*\|^q + \beta(2\eta^{1-1/q}B_q)^q \left(\frac{\kappa'}{8(q-1)}\right)^{1-q} + 32\beta^q B_q^q \end{aligned}$$

Now using the condition $\beta \leq \eta/\kappa'$ and rearranging the inequality, we finally arrive at

$$\begin{aligned} \mathbb{E}\|\theta - \theta^*\|^q &\leq \frac{8}{3} \left(\frac{\eta^{1-1/q}B_q}{\kappa'}\right)^q \left(2(8(q-1))^{q-1} + 32\right) \\ &\leq 128 \left(\frac{\eta^{1-1/q}B_q}{\kappa'}\right)^q, \end{aligned}$$

which is the desired result. \square

D.6 Proof of Proposition 2

We use the invariance of $\pi_{\beta,p}$ by the transition kernel $P_{\beta,p}$. For real λ , this implies the equality

$$\mathbb{E} \exp(\lambda^2 \|\theta - \theta^*\|^2) = \mathbb{E} \exp(\lambda^2 \|\theta - \alpha_\theta \beta G(\theta) - \theta^*\|^2). \quad (39)$$

We then write:

$$\begin{aligned} \|\theta - \alpha_\theta \beta G(\theta) - \theta^*\|^2 &= \|\theta - \bar{\alpha}_\theta \beta \nabla \mathcal{L}(\theta) - \theta^*\|^2 + \beta^2 \|\alpha_\theta G(\theta) - \bar{\alpha}_\theta \nabla \mathcal{L}(\theta)\|^2 \\ &\quad - 2\beta \langle \theta - \bar{\alpha}_\theta \beta \nabla \mathcal{L}(\theta) - \theta^*, \alpha_\theta G(\theta) - \mathbb{E}[\alpha_\theta G(\theta)] + \mathbb{E}[\alpha_\theta G(\theta)] - \bar{\alpha}_\theta \nabla \mathcal{L}(\theta) \rangle. \end{aligned}$$

We also have the inequality

$$\|\alpha_\theta G(\theta) - \bar{\alpha}_\theta \nabla \mathcal{L}(\theta)\|^2 \leq 2\|\alpha_\theta G(\theta) - \mathbb{E}[\alpha_\theta G(\theta)]\|^2 + 2\|\mathbb{E}[\alpha_\theta G(\theta)] - \bar{\alpha}_\theta \nabla \mathcal{L}(\theta)\|^2.$$

Conditioning upon θ , we have that $\|\alpha_\theta G(\theta) - \mathbb{E}[\alpha_\theta G(\theta)]\| \leq 2\tau_\theta \leq 2\bar{\tau}$. Moreover, the vector $\alpha_\theta G(\theta) - \mathbb{E}[\alpha_\theta G(\theta)]$ is centered, therefore it is sub-Gaussian with constant $2\bar{\tau}$ (see (Vershynin, 2018, Proposition 2.5.2)). Still conditioning on θ , using these two properties and a Cauchy-Schwarz inequality, we find:

$$\begin{aligned} \mathbb{E} \exp(\lambda^2 (2\beta^2 \|\alpha_\theta G(\theta) - \mathbb{E}[\alpha_\theta G(\theta)]\|^2 - 2\beta \langle \theta - \bar{\alpha}_\theta \beta \nabla \mathcal{L}(\theta) - \theta^*, \alpha_\theta G(\theta) - \mathbb{E}[\alpha_\theta G(\theta)] \rangle)) \\ \leq \exp(8\lambda^2 \beta^2 \bar{\tau}^2 + 16\lambda^4 \beta^2 \bar{\tau}^2 \|\theta - \bar{\alpha}_\theta \beta \nabla \mathcal{L}(\theta) - \theta^*\|^2). \end{aligned}$$

Putting everything together in (39), we get:

$$\begin{aligned} \mathbb{E} \exp(\lambda^2 \|\theta - \theta^*\|^2) &\leq \mathbb{E} \exp(\lambda^2 ((1 + 16\lambda^2 \beta^2 \bar{\tau}^2) \|\theta - \bar{\alpha}_\theta \beta \nabla \mathcal{L}(\theta) - \theta^*\|^2 + 8\beta^2 \bar{\tau}^2 \\ &\quad + 2\beta^2 \|\mathbb{E}[\alpha_\theta G(\theta)] - \bar{\alpha}_\theta \nabla \mathcal{L}(\theta)\|^2 + 2\beta \|\theta - \bar{\alpha}_\theta \beta \nabla \mathcal{L}(\theta) - \theta^*\| \|\mathbb{E}[\alpha_\theta G(\theta)] - \bar{\alpha}_\theta \nabla \mathcal{L}(\theta)\|)) \\ &\leq \mathbb{E} \exp(\lambda^2 ((1 + 16\lambda^2 \beta^2 \bar{\tau}^2 + \epsilon) \|\theta - \bar{\alpha}_\theta \beta \nabla \mathcal{L}(\theta) - \theta^*\|^2 + 8\beta^2 \bar{\tau}^2 \\ &\quad + 2\beta^2 (1 + 1/(2\epsilon)) \|\mathbb{E}[\alpha_\theta G(\theta)] - \bar{\alpha}_\theta \nabla \mathcal{L}(\theta)\|^2)), \end{aligned}$$

where we used Fact 1. Now, recalling that $\bar{\alpha}_\theta \geq p$, we set $\epsilon = \bar{\alpha}_\theta \beta \mu / 2$ and restrict λ to $\lambda \leq (4\bar{\tau} \sqrt{2\beta/p\mu})^{-1}$ to find:

$$\begin{aligned} \mathbb{E} \exp(\lambda^2 \|\theta - \theta^*\|^2) &\leq \mathbb{E} \exp(\lambda^2 ((1 + 16\lambda^2 \beta^2 \bar{\tau}^2 + \epsilon) \|\theta - \bar{\alpha}_\theta \beta \nabla \mathcal{L}(\theta) - \theta^*\|^2 + 8\beta^2 \bar{\tau}^2 \\ &\quad + 2\beta^2(1 + 1/(2\epsilon)) \|\mathbb{E}[\alpha_\theta G(\theta)] - \bar{\alpha}_\theta \nabla \mathcal{L}(\theta)\|^2)) \\ &\leq \mathbb{E} \exp\left(\lambda^2 \left((1 + \bar{\alpha}_\theta \beta \mu) \|\theta - \bar{\alpha}_\theta \beta \nabla \mathcal{L}(\theta) - \theta^*\|^2 + 8\beta^2 \bar{\tau}^2 + \frac{4\beta}{\bar{\alpha}_\theta \mu} \|\mathbb{E}[\alpha_\theta G(\theta)] - \bar{\alpha}_\theta \nabla \mathcal{L}(\theta)\|^2 \right)\right) \\ &\leq \mathbb{E} \exp\left(\lambda^2 \left((1 + \bar{\alpha}_\theta \beta \mu)(1 - \bar{\alpha}_\theta \beta \mu)^2 \|\theta - \theta^*\|^2 + 8\beta^2 \bar{\tau}^2 + \frac{4\beta}{\bar{\alpha}_\theta \mu} ((1-p)^{1-1/q} \bar{B}_q)^2 \right)\right) \\ &\leq \mathbb{E} \exp(\lambda^2 ((1 - \bar{\alpha}_\theta \beta \mu)(1 - (\bar{\alpha}_\theta \beta \mu)^2) \|\theta - \theta^*\|^2 + 4\beta^2 (2\bar{\tau}^2 + \bar{B}_q^2/p))), \end{aligned}$$

where we used Lemma 1, inequality (21) from Lemma 2 (recall that $\eta = 0$ in this context) and the imposed bound relating β and p .

Finally, using Jensen's inequality and the fact that $\bar{\alpha}_\theta \geq p \geq 1/2$ we get:

$$\mathbb{E} \exp(\lambda^2 \|\theta - \theta^*\|^2) \leq \exp\left(\frac{8\lambda^2 \beta^2 (\bar{\tau}^2 + \bar{B}_q^2)}{\bar{\alpha}_\theta \beta \mu + (\bar{\alpha}_\theta \beta \mu)^2 - (\bar{\alpha}_\theta \beta \mu)^3}\right) \leq \exp\left(\frac{8\lambda^2 \beta (\bar{\tau}^2 + \bar{B}_q^2)}{p\mu}\right),$$

which concludes the first part of the proof. We now consider the corrupted case $\eta > 0$. Let $\lambda > 0$ and write :

$$\begin{aligned} \mathbb{E}[\exp(\lambda \|\theta - \theta^*\|)] &= \mathbb{E}[\exp(\lambda \|\theta - \alpha_\theta \beta G(\theta) - \theta^*\|)] \\ &\leq \eta \mathbb{E}[\exp(\lambda \|\theta - \alpha_\theta \beta \tilde{G}(\theta) - \theta^*\|)] + (1-\eta) \mathbb{E}[\exp(\lambda \|\theta - \alpha_\theta \beta \tilde{G}(\theta) - \theta^*\|)] \\ &\leq \eta \mathbb{E}[\exp(\lambda (\|\theta - \theta^*\| + \beta \tau_\theta))] + (1-\eta) \mathbb{E}[\exp(\lambda \|\theta - \bar{\alpha}_\theta \beta \nabla \mathcal{L}(\theta) - \theta^*\| \\ &\quad + \lambda \beta (\|\alpha_\theta \tilde{G}(\theta) - \mathbb{E}[\alpha_\theta \tilde{G}(\theta)]\| + \|\mathbb{E}[\alpha_\theta \tilde{G}(\theta)] - \bar{\alpha}_\theta \nabla \mathcal{L}(\theta)\|))] \\ &\leq \eta \mathbb{E}[\exp(\lambda \|\theta - \theta^*\|)] e^{\lambda \beta \tau} \\ &\quad + (1-\eta) \mathbb{E}[\exp(\lambda (1-p\beta\mu) \|\theta - \theta^*\|)] \exp(\lambda \beta (2\bar{\tau} + (1-p)^{1-1/q} \bar{B}_q)), \end{aligned}$$

where we used Lemma 1, the inequality $\alpha_\theta \geq p$, the inequality $\|\alpha_\theta \tilde{G}(\theta) - \mathbb{E}[\alpha_\theta \tilde{G}(\theta)]\| \leq 2\tau_\theta \leq 2\bar{\tau}$ and (21) from Lemma 2. Using Hölder's inequality, this leads to :

$$\mathbb{E}[\exp(\lambda \|\theta - \theta^*\|)] \leq \left(\frac{1-\eta}{1-\eta e^{\lambda \beta \bar{\tau}}}\right)^{1/(p\beta\mu)} \exp(\lambda (2\bar{\tau} + (1-p)^{1-1/q} \bar{B}_q)/(p\mu))$$

Now we use the inequality $\log\left(\frac{1-\eta}{1-\eta e^{\lambda \beta \bar{\tau}}}\right) \leq \frac{\beta \bar{\tau} \lambda / \log(1/\eta)^2}{1-\beta \bar{\tau} \lambda / \log(1/\eta)}$ valid for $\lambda \geq 0$ which leads to:

$$\left(\frac{1-\eta}{1-\eta e^{\lambda \beta \bar{\tau}}}\right)^{1/(p\beta\mu)} \leq \exp\left(\frac{2\lambda \bar{\tau}}{p\mu \log(1/\eta)^2}\right) \quad \text{for } 0 \leq \lambda \leq \frac{\log(1/\eta)}{2\beta \bar{\tau}}.$$

Using that $\eta < 1/2$, we find that for $0 \leq \lambda \leq \frac{\log(1/\eta)}{2\beta \bar{\tau}}$, the following inequality holds :

$$\begin{aligned} \mathbb{E}[\exp(\lambda \|\theta - \theta^*\|)] &\leq \exp\left(\frac{\lambda}{p\mu} \left(2\bar{\tau} \left(1 + \frac{1}{\log(1/\eta)^2}\right) + (1-p)^{1-1/q} \bar{B}_q\right)\right) \\ &\leq \exp\left(\frac{\lambda}{p\mu} (7\bar{\tau} + (1-p)^{1-1/q} \bar{B}_q)\right). \end{aligned}$$

Noticing that $\beta \leq \frac{1}{\mu}$ allows to finish the proof.

D.7 Unimprovability of the sub-exponential property for $\eta > 0$

We consider the Markov chain

$$X_{t+1} = \begin{cases} \alpha X_t + \xi & \text{w.p. } 1-\eta \\ X_t + \tau & \text{w.p. } \eta \end{cases}$$

Assuming that the distribution of the noise ξ has a density, one can show that the chain is aperiodic and satisfies a minorization property as in the proof of Theorem 1 (see Appendix D.3).

Defining $V(x) = 1 + x$, we can show that V verifies a geometric drift property similar to (25). Consequently, Theorem 15.0.1 of (Meyn & Tweedie, 1993) applies to the chain $(X_t)_{t \geq 0}$ and implies that it converges geometrically to a limit distribution π analogously to the claim of Theorem 1.

We denote $M_k^k = \mathbb{E}|X|^k$ the absolute moments of X for $k \geq 1$ and show that $M_k = \Omega(k)$ (we merely provide a sketch and do not attempt to explicitly compute the involved constants). For $X \sim \pi$ following the invariant measure, using the recursion defining X_t and the positivity of ξ , it is easy to establish the inequality for $k \geq 1$

$$(1 - \eta)(1 - \alpha^k)M_k^k \geq \eta \sum_{j=1}^k \binom{k}{j} \tau^j M_{k-j}^{k-j}$$

where one may use the convention that $M_0 = 1$. We now postulate the induction hypothesis $M_j \geq Cj$ up to $j = k - 1$ for some $k > 1$ and $C > 0$. Using Stirling's formula, we find:

$$\begin{aligned} \frac{(1 - \eta)(1 - \alpha^k)}{\eta} M_k^k &\geq \sum_{j=1}^k \binom{k}{j} \tau^j (C(k - j))^{k-j} = \sum_{j=1}^k \frac{k!}{j!(k - j)!} \tau^j (C(k - j))^{k-j} \\ &\gtrsim \sum_{j=1}^k \sqrt{\frac{k}{j(k - j)}} \frac{k^k \tau^j (C(k - j))^{k-j}}{j^j (k - j)^{k-j}} = \sum_{j=1}^k \sqrt{\frac{k}{j(k - j)}} \left(\frac{\tau}{j}\right)^j k^k C^{k-j} \\ &\geq \frac{\tau}{C} (Ck)^k \end{aligned}$$

where \gtrsim denotes an inequality up to a universal constant and we took the term $j = 1$ in the last step. It is only left to set C small enough such that $\frac{\tau\eta}{C(1-\eta)(1-\alpha)} \geq 1$ in order to finish the induction. It follows that $M_k = \Omega(k)$ implying that π may be sub-exponential but cannot be sub-Gaussian since that would require $M_k = \mathcal{O}(\sqrt{k})$ (see (Vershynin, 2018, Chapter 2) for a reference).

D.8 Proof of Corollary 1

We need the following lemma.

Lemma 4. *Let X be a real sub-Gaussian random variable with constant K then, with probability at least δ , we have :*

$$|X| \leq K \sqrt{\log(e/\delta)}$$

Proof. Using Chernoff's method, we find for $t > 0$ and $\lambda > 0$:

$$\begin{aligned} \mathbb{P}(|X| > t) &= \mathbb{P}(\lambda^2 X^2 > \lambda^2 t^2) = \mathbb{P}(\exp(\lambda^2 X^2) > \exp(\lambda^2 t^2)) \\ &\leq \mathbb{E} \exp(\lambda^2 X^2) e^{-\lambda^2 t^2} \leq \exp(\lambda^2 (K^2 - t^2)). \end{aligned}$$

Choosing $\lambda = 1/K$, we have $\exp(1 - (t/K)^2) \leq \delta \iff t \geq K \sqrt{\log(e/\delta)}$ and the result follows. \square

By Theorem 1, the Markov chain $(\theta_t)_{t \geq 0}$ is geometrically converging to the invariant distribution $\pi_{\beta,p}$ w.r.t. the Total Variation distance so that for any event $\mathcal{E} \in \mathcal{B}(\mathbb{R}^d)$, we have:

$$|\mathbb{P}(\theta_T \in \mathcal{E}) - \mathbb{P}_{\theta \sim \pi_{\beta,p}}(\theta \in \mathcal{E})| \leq M \rho^T V(\theta_0). \quad (40)$$

Proposition 2 states that, in the absence of corruption, for $\theta \sim \pi_{\beta,p}$, the variable $\|\theta - \theta^*\|$ is sub-Gaussian with constant $K = 4\sqrt{\frac{2\beta(B_\mu^2 + \bar{\tau}^2)}{p\mu}}$. It is only left to combine this conclusion with Lemma 4 in order to obtain the claimed bound.

D.9 Proof of Corollary 2

We assume without loss of generality that T is a multiple of N . Note that according to the assumptions, the estimators $\theta_T^{(n)}$ for $n \in \llbracket N \rrbracket$ are independent and for each n . For positive $\epsilon < 1$ define the events $E_n := \{\|\theta_T^{(n)} - \theta^\star\| \leq \frac{\eta^{1-\frac{1}{q}} B_q \sqrt{20}}{\kappa \epsilon}\}$. We first assume that $\sum_{n=1}^N \mathbf{1}_{E_n} > N/2$ then there exists $i' \in \llbracket N \rrbracket$ such that

$$r_{N/2}^{(\hat{i})} = \|\theta_T^{(\hat{i})} - \theta_T^{(i')}\| \leq \|\theta_T^{(\hat{i})} - \theta^\star\| + \|\theta_T^{(i')} - \theta^\star\| \leq 2 \frac{\eta^{1-\frac{1}{q}} B_q \sqrt{20}}{\kappa \epsilon}.$$

Moreover, among the $N/2$ estimators closest to $\theta_T^{(\hat{i})}$, at least one of them $\theta_T^{(i')}$ satisfies $\|\theta_T^{(i')} - \theta^\star\| \leq \frac{\eta^{1-\frac{1}{q}} B_q \sqrt{20}}{\kappa \epsilon}$ thus we find :

$$\begin{aligned} \|\hat{\theta} - \theta^\star\| &= \|\theta_T^{(\hat{i})} - \theta^\star\| \leq \|\theta_T^{(\hat{i})} - \theta_T^{(i')}\| + \|\theta_T^{(i')} - \theta^\star\| \\ &\leq r_{N/2}^{(\hat{i})} + \frac{\eta^{1-\frac{1}{q}} B_q \sqrt{20}}{\kappa \epsilon} \leq 3 \frac{\eta^{1-\frac{1}{q}} B_q \sqrt{20}}{\kappa \epsilon}. \end{aligned} \quad (41)$$

Notice that setting $\epsilon = 1/2$ immediately yields (13). We now show that $\sum_{n=1}^N \mathbf{1}_{E_n} > N/2$ happens with high probability. Thanks to Theorem 1 and Proposition 1, we have:

$$\mathbb{P}(\overline{E}_n) \leq \epsilon^2 + \underbrace{M \rho^{T/N}}_{\epsilon'} (1 + \|\theta_0 - \theta^\star\|^2).$$

Consequently, the variables $\mathbf{1}_{\overline{E}_n}$ are stochastically dominated by Bernoulli variables with parameter $\epsilon^2 + \epsilon'$ so that their sum is stochastically dominated by a Binomial random variable $S := \text{Bin}(N, \epsilon^2 + \epsilon')$. We compute :

$$\begin{aligned} \mathbb{P}\left(\sum_{n=1}^N \mathbf{1}_{E_n} < N/2\right) &= \mathbb{P}\left(\sum_{n=1}^N \mathbf{1}_{\overline{E}_n} > N/2\right) \leq \mathbb{P}(S - \mathbb{E}S > N/2 - (\epsilon^2 + \epsilon')N) \\ &\leq \exp(-2N(1/2 - \epsilon^2 - \epsilon')^2) \\ &\leq \exp(-2N(1/2 - 1/4 - M \rho^{T/N}(1 + \|\theta_0 - \theta^\star\|^2))^2) \\ &\leq \exp(-2N(1/4 - 1/15)^2) \leq \exp(-\log(1/\delta)) = \delta \end{aligned}$$

where we used Hoeffding's inequality, the choice $\epsilon = 1/2$ and the fact that our condition on T implies $M \rho^{T/N}(1 + \|\theta_0 - \theta^\star\|^2)^2 \leq 1/15$. The last inequalities result from our condition on N .

D.10 Proof of Theorem 2

As previously done in the proof of Theorem 1, we show that the Markov chain is aperiodic. Note that since \mathcal{L} has finite lower bound $\inf_\theta \mathcal{L}$, we can replace it with $1 + \mathcal{L}(\theta) - \inf_\theta \mathcal{L}$ and assume it is positive in the rest of the proof without loss of generality. We will now show that it satisfies a drift property. Let $\theta \in \Theta$ be fixed, using Assumptions 1 and 3, we have :

$$\begin{aligned} \mathbb{E}[\mathcal{L}(\theta - \alpha_\theta \beta G(\theta))] - \mathcal{L}(\theta) &\leq \mathbb{E}\left[-\beta \langle \nabla \mathcal{L}(\theta), \alpha_\theta G(\theta) \rangle + \frac{L\beta^2}{2} \|\alpha_\theta G(\theta)\|^2\right] \\ &\leq \mathbb{E}\left[\eta \left(-\beta \langle \nabla \mathcal{L}(\theta), \alpha_\theta \check{G}(\theta) \rangle + \frac{L\beta^2}{2} \|\alpha_\theta \check{G}(\theta)\|^2\right) + \right. \\ &\quad \left. (1 - \eta) \left(-\beta \langle \nabla \mathcal{L}(\theta), \alpha_\theta \tilde{G}(\theta) \rangle + \frac{L\beta^2}{2} \|\alpha_\theta \tilde{G}(\theta)\|^2\right)\right] \end{aligned}$$

Note that we have $\|\alpha_\theta \tilde{G}(\theta)\|, \|\alpha_\theta \check{G}(\theta)\| \leq \tau_\theta$, therefore

$$\begin{aligned} \mathbb{E}[\mathcal{L}(\theta - \alpha_\theta \beta G(\theta))] - \mathcal{L}(\theta) &\leq \mathbb{E}\left[\eta \left(-\beta \langle \nabla \mathcal{L}(\theta), \alpha_\theta \check{G}(\theta) \rangle \right) + \right. \\ &\quad \left. (1 - \eta) \left(-\beta \langle \nabla \mathcal{L}(\theta), \alpha_\theta \tilde{G}(\theta) \rangle \right) \right] + \frac{L\beta^2 \tau_\theta^2}{2}. \end{aligned}$$

Further, we write

$$\begin{aligned} \mathbb{E}[\langle \nabla \mathcal{L}(\theta), \alpha_\theta \tilde{G}(\theta) \rangle] &= \langle \nabla \mathcal{L}(\theta), \mathbb{E}[\alpha_\theta \tilde{G}(\theta)] - \bar{\alpha}_\theta \nabla \mathcal{L}(\theta) + \bar{\alpha}_\theta \nabla \mathcal{L}(\theta) \rangle \\ &= \bar{\alpha}_\theta \|\nabla \mathcal{L}(\theta)\|^2 + \langle \nabla \mathcal{L}(\theta), \mathbb{E}[\alpha_\theta \tilde{G}(\theta)] - \bar{\alpha}_\theta \nabla \mathcal{L}(\theta) \rangle. \end{aligned}$$

Plugging back and using Cauchy-Schwartz and the inequality $\|\alpha_\theta \check{G}(\theta)\| \leq \tau_\theta$ leads to

$$\begin{aligned} \mathbb{E}[\mathcal{L}(\theta - \alpha_\theta \beta G(\theta))] - \mathcal{L}(\theta) &\leq -\eta \beta \langle \nabla \mathcal{L}(\theta), \mathbb{E}[\alpha_\theta \check{G}(\theta)] \rangle - \beta(1 - \eta) \bar{\alpha}_\theta \|\nabla \mathcal{L}(\theta)\|^2 - \\ &\quad \beta(1 - \eta) \langle \nabla \mathcal{L}(\theta), \mathbb{E}[\alpha_\theta \tilde{G}(\theta)] - \bar{\alpha}_\theta \nabla \mathcal{L}(\theta) \rangle + \frac{L\beta^2 \tau_\theta^2}{2} \\ &\leq \eta \beta \tau_\theta \|\nabla \mathcal{L}(\theta)\| - \beta(1 - \eta) \bar{\alpha}_\theta \|\nabla \mathcal{L}(\theta)\|^2 + \\ &\quad \beta(1 - \eta) \|\nabla \mathcal{L}(\theta)\| \|\mathbb{E}[\alpha_\theta \tilde{G}(\theta)] - \bar{\alpha}_\theta \nabla \mathcal{L}(\theta)\| + \frac{L\beta^2 \tau_\theta^2}{2}. \end{aligned}$$

We now use the inequalities $\tau_\theta \leq \|\nabla \mathcal{L}(\theta)\| + Q_p(\|\varepsilon_\theta\|)$ (see Lemma 2) and $(a + b)^2 \leq 2a^2 + 2b^2$, to find that

$$\begin{aligned} \mathbb{E}[\mathcal{L}(\theta - \alpha_\theta \beta G(\theta))] - \mathcal{L}(\theta) &\leq -\beta \|\nabla \mathcal{L}(\theta)\|^2 ((1 - \eta) \bar{\alpha}_\theta - L\beta - \eta) + L\beta^2 Q_p(\|\varepsilon_\theta\|)^2 + \\ &\quad \beta \|\nabla \mathcal{L}(\theta)\| (\eta Q_p(\|\varepsilon_\theta\|) + (1 - \eta) \|\mathbb{E}[\alpha_\theta \tilde{G}(\theta)] - \bar{\alpha}_\theta \nabla \mathcal{L}(\theta)\|). \end{aligned}$$

Next, letting $\epsilon > 0$, and using Fact 1 gives

$$\begin{aligned} \|\nabla \mathcal{L}(\theta)\| (\eta Q_p(\|\varepsilon_\theta\|) + (1 - \eta) \|\mathbb{E}[\alpha_\theta \tilde{G}(\theta)] - \bar{\alpha}_\theta \nabla \mathcal{L}(\theta)\|) &\leq \epsilon \|\nabla \mathcal{L}(\theta)\|^2 / 2 + \\ &\quad (\eta Q_p(\|\varepsilon_\theta\|) + (1 - \eta) \|\mathbb{E}[\alpha_\theta \tilde{G}(\theta)] - \bar{\alpha}_\theta \nabla \mathcal{L}(\theta)\|)^2 / (2\epsilon). \end{aligned}$$

We now plug back with the choice $\epsilon = p(1 - \eta)/2$ and use that $\bar{\alpha}_\theta \geq p$ to find

$$\begin{aligned} \mathbb{E}[\mathcal{L}(\theta - \alpha_\theta \beta G(\theta))] - \mathcal{L}(\theta) &\leq -\beta \|\nabla \mathcal{L}(\theta)\|^2 (3p(1 - \eta)/4 - L\beta - \eta) + L\beta^2 Q_p(\|\varepsilon_\theta\|)^2 + \\ &\quad \frac{\beta (\eta Q_p(\|\varepsilon_\theta\|) + (1 - \eta) \|\mathbb{E}[\alpha_\theta \tilde{G}(\theta)] - \bar{\alpha}_\theta \nabla \mathcal{L}(\theta)\|)^2}{p(1 - \eta)}. \end{aligned}$$

Finally, we use Lemma 2 to bound the terms $Q_p(\|\varepsilon_\theta\|)$ and $\|\mathbb{E}[\alpha_\theta \tilde{G}(\theta)] - \bar{\alpha}_\theta \nabla \mathcal{L}(\theta)\|$ leading to

$$\begin{aligned} \mathbb{E}[\mathcal{L}(\theta - \alpha_\theta \beta G(\theta))] - \mathcal{L}(\theta) &\leq -\beta \|\nabla \mathcal{L}(\theta)\|^2 (3p(1 - \eta)/4 - L\beta - \eta) + L\beta^2 (1 - p)^{-\frac{2}{q}} B_q^2 + \\ &\quad \frac{\beta B_q^2 (\eta(1 - p)^{-\frac{1}{q}} + (1 - \eta) \eta^{1 - \frac{1}{q}})^2}{p(1 - \eta)} \\ &\leq -\beta \|\nabla \mathcal{L}(\theta)\|^2 (3p(1 - \eta)/4 - L\beta - \eta) + \\ &\quad \frac{\beta B_q^2 ((1 - p)^{-\frac{2}{q}} (L\beta + 2\eta^2) + 2\eta^{2 - \frac{2}{q}})}{p(1 - \eta)}. \end{aligned} \tag{42}$$

By assumption, we have $3p(1 - \eta)/4 - L\beta - \eta > 0$. Define the quantity $\xi = \frac{B_q^2 ((1 - p)^{-\frac{2}{q}} (L\beta + 2\eta^2) + 2\eta^{2 - \frac{2}{q}})}{p(1 - \eta)}$ and the set $\mathcal{C} := \left\{ \theta \in \mathbb{R}^d : \frac{1}{2} \|\nabla \mathcal{L}(\theta)\|^2 \leq \frac{\xi}{3p(1 - \eta)/4 - L\beta - \eta} \right\}$. By assumption, \mathcal{C} is bounded and it is clear that the right hand side in (42) is negative outside \mathcal{C} . Define the function $V(\theta) = \mathcal{L}(\theta)/(\beta\xi)$, which is positive and satisfies:

$$\Delta V(\theta) \leq -1 + 2 \cdot \mathbf{1}_{\theta \in \mathcal{C}}. \tag{43}$$

In addition, we show similarly to Theorem 1 that the set \mathcal{C} is *small* and, therefore, also *petite* according to the definitions of (Meyn & Tweedie, 1993, Chapter 5). Since V is everywhere finite and bounded on \mathcal{C} (because the latter is compact), the conditions of (Meyn & Tweedie, 1993, Theorem 11.3.4) are fulfilled implying that the chain is Harris recurrent.

We have shown that the Markov chain verifies the fourth condition of (Meyn & Tweedie, 1993, Theorem 13.0.1). This allows us to conclude that the Markov chain is ergodic i.e. we have for any initial measure λ that $\|\lambda P^t - \pi_{\beta,p}\|_{\text{TV}} \rightarrow 0$ and the following sum is finite

$$\sum_t \|\lambda P_{\beta,p}^t - \pi_{\beta,p}\|_{\text{TV}} < \infty.$$

In addition, by (Meyn & Tweedie, 1993, Proposition 13.3.2) the terms in the above sum are non-increasing which implies that $\|\lambda P_{\beta,p}^t - \pi_{\beta,p}\|_{\text{TV}} = \mathcal{O}(t^{-1})$ and the result follows.

D.11 Proof of Proposition 3

By Theorem 2, the assumptions imply that the Markov chain $(\theta_t)_{t \geq 0}$ converges to an invariant distribution $\pi_{\beta,p}$. For $\theta \sim \pi_{\beta,1-\eta}$, by invariance of $\pi_{\beta,1-\eta}$, we have that the variables $\mathcal{L}(\theta - \alpha_\theta \beta G(\theta))$ and $\mathcal{L}(\theta)$ are identically distributed. Taking the expectation w.r.t. θ , this implies the identity $\mathbb{E}[\mathcal{L}(\theta - \alpha_\theta \beta G(\theta))] = \mathbb{E}[\mathcal{L}(\theta)]$. Plugging into Inequality (42) from the proof of Theorem 2, we find

$$\begin{aligned} \mathbb{E}[\|\nabla \mathcal{L}(\theta)\|^2] &\leq \frac{B_q^2((1-p)^{-\frac{2}{q}}(L\beta + 2\eta^2) + 2\eta^{2-\frac{2}{q}})}{p(1-\eta)(3p(1-\eta)/4 - L\beta - \eta)} \\ &\leq \frac{B_q^2(3(1-p)^{-\frac{2}{q}}\eta^2 + 2\eta^{2-\frac{2}{q}})}{p(1-\eta)(3p(1-\eta)/4 - L\beta - \eta)} \\ &\leq \frac{5B_q^2\eta^{2-\frac{2}{q}}}{p(1-\eta)(3p(1-\eta)/4 - L\beta - \eta)} \end{aligned}$$

where we used the choices $\beta \leq \frac{\eta^2}{L}$ and $p = 1 - \eta$.