
Minimax M-estimation under Adversarial Corruption

Sujay Bhatt, Guanhua Fang, Ping Li

Cognitive Computing Lab
Baidu Research
10900 NE 8th St. Bellevue, WA 98004, USA
{sujaybhatt.hr, fanggh2018, pingli98}@gmail.com

Gennady Samorodnitsky

School of ORIE
Cornell University
220 Frank T Rhodes Hall, Ithaca, NY 14853, USA
gs18@cornell.edu

Abstract

¹We present a new finite-sample analysis of Catoni’s M-estimator under adversarial contamination, where an adversary is allowed to corrupt a fraction of the samples arbitrarily. We make minimal assumptions on the distribution of the *uncorrupted* random variables, namely, we only assume the existence of a known upper bound on the $(1 + \varepsilon)^{th}$ central moment. We provide a lower bound on the minimax error rate for the mean estimation problem under adversarial corruption under this weak assumption, and establish that the proposed M-estimator achieves this lower bound (up to multiplicative constants). When variance is infinite, the tolerance to contamination of any estimator reduces as $\varepsilon \downarrow 0$. We establish a tight upper bound that characterizes this bargain. To illustrate the usefulness of the derived robust M-estimator in an online setting, we present a bandit algorithm for the partially identifiable best arm identification problem that improves upon the sample complexity of the state of the art algorithms.

1. Introduction

Univariate mean estimation plays an important role in many statistical learning problems, ranging from classification and regression (James et al., 2013) to online learning (Lattimore and Szepesvári, 2020; Agarwal et al., 2019). A fundamental challenge in machine learning, using diverse data drawn from heterogeneous sources, is that outliers and adversarial contamination is unavoidable. For example, the presence of outliers is the primary source of invalid inferences in fMRI (Eklund et al., 2016). In addition to dealing with

heavy-tails, modern machine learning also has to deal with malicious noise (Auer and Cesa-Bianchi, 1998; Diakonikolas et al., 2018). Here an adversary can arbitrarily corrupt a fraction of the data to disrupt inference. These challenges motivate the study of robust mean estimation in this paper, with a focus on providing estimators that can *tolerate* large amount of adversarial corruption, possibly in the presence of heavy-tailed data.

Standard mean estimation based on mean squared error is insufficient to deal with such diverse data (Catoni, 2012); and mean estimation based on deviations, namely, that of estimation of confidence intervals offers a much better alternative (Catoni, 2012; Brownlees et al., 2015). Here the fundamental problem is of designing an estimator $\hat{\mu}_n = \hat{\mu}_n(\{X_i\}_{i=1}^n)$ that for a given *confidence* $\delta \in (0, 1)$ has the smallest possible $\varrho = \varrho(n, \delta)$ such that

$$\mathbb{P}\left\{\left|\hat{\mu}_n - \mu\right| > \varrho\right\} \leq \delta.$$

In next two subsections, we briefly review the standard results (and key challenges) for robust mean estimation under finite and infinite variance settings, separately.

1.1. Finite Variance, $\varepsilon = 1$

In the case of finite variance, the classical empirical mean is a poor estimator of mean due to the presence of outliers (Catoni, 2012; Devroye et al., 2016; Oliveira and Orenstein, 2019). In fact, it was shown in Catoni (2012) that, if v is the variance of the observations, then the best possible error $\varrho(n, \delta)$ for empirical mean of $\{X_i\}_{i=1}^n$ is of the order of $\frac{v}{\sqrt{\delta n}}$, which is far from the best possible, as informed by the central limit theorem (Lugosi and Mendelson, 2019a). Devroye et al. (2016) exhibit several sub-Gaussian estimators² of the mean such as median-of-means, trimmed mean,

²An estimator $\hat{\mu}_n$ is L -sub-Gaussian for a constant $L > 0$ if for random variables with variance v and (any) sample size n , with probability at least $1 - \delta$,

$$\left|\hat{\mu}_n - \mu\right| \leq \frac{L\sqrt{v \log(2/\delta)}}{\sqrt{n}}.$$

and Catoni’s M-estimator.

In the presence of malicious noise, however, every robust estimator has an asymptotic bias of $O(\sqrt{v\eta})$, where η is the fraction of samples corrupted by the adversary (Lai et al., 2016; Hopkins and Li, 2018; Lugosi and Mendelson, 2021). In the univariate setting, such estimators have been derived using the well known estimators for *finite* variance:

(i) Using *empirical mean*: The main idea here is based on the robust estimator proposed in Lai et al. (2016). Essentially, the data is split into two halves— one half is used to construct bounded intervals that ‘trap’ many samples from the uncorrupted distribution with a high confidence, while an empirical estimate of the other half contained in the constructed interval is returned as the robust estimate of the mean. It is shown this estimator has the (minimax) asymptotic bias of $O(\sqrt{v\eta})$ (Prasad et al., 2020a, Lemma 3).

(ii) Using *trimmed mean*: The main idea is similar to that in Lai et al. (2016), except that instead of the intervals, order statistics on one half of the data is used to compute the minimum and maximum truncation levels, while a smoothed estimate of the other half data is returned as the robust estimated of the mean. A subtle difference here is that the samples in the second half outside the truncation levels are not discarded, however, are set to the truncation level. It has been shown that this estimator again achieves the (minimax) asymptotic bias of $O(\sqrt{v\eta})$, see Lugosi and Mendelson (2021, Theorem 1).

It is clear that the above methods, while achieving the minimax error up to constants, do not make use of the data effectively— one half of the data is used only to extract periphery information. This, in part, motivates robust estimation using M-estimators that effectively utilize the entire data to return a robust estimate of the mean, while achieving a high tolerance to adversarial contamination. For example, when $\varepsilon = 1$, we establish that due to inherent robustness of the proposed M-estimator, it tolerates upto 36% arbitrary contamination, and outperforms the state of the art estimators based on trimmed mean (Lugosi and Mendelson, 2021).

There is related work using Median-of-Means estimator (Laforgue et al., 2021), where the number of groups is modulated using the fraction of the outliers (η), and the authors obtain L -sub-Gaussian like estimators with very large $L(\eta)$. The asymptotic bias is not characterized and it is unclear whether the proposed estimator achieves the minimax error bound. In contrast, the focus in this paper, is to obtain sharp constants with minimax asymptotic bias.

L -sub-Gaussian estimators are optimal up to constants with $L \leq \sqrt{2} + o(1)$ identified as nearly-optimal (Devroye et al., 2016). Also see Buldygin and Kozachenko (2000) and Li (2007, Chapter 4) for essentially equivalent definitions (and more applications) of the sub-Gaussian estimators.

1.2. Infinite Variance, $\varepsilon < 1$

When $\varepsilon < 1$, Devroye et al. (2016) establish that the achievable $\varrho(n, \delta)$ is no longer sub-Gaussian and is in fact given by the following result:

Theorem (Lower Bound). There exists a distribution with mean μ and $(1 + \varepsilon)^{th}$ central moment v_ε such that for any mean estimator $\hat{\mu}_n$ and $\delta \in (2e^{-n/4}, 1/2)$,

$$\mathbb{P}\left\{|\hat{\mu}_n - \mu| > \left(\frac{v_\varepsilon^{\frac{1}{\varepsilon}} \log(2/\delta)}{n}\right)^{\frac{\varepsilon}{1+\varepsilon}}\right\} \geq \delta.$$

This is established in Devroye et al. (2016, Theorem 3.1). There are estimators that obtain the optimal order such as Median-of-Means (Bubeck et al., 2013; Minsker, 2019; Lecué and Lerasle, 2020) and Trimmed Mean (Oliveira and Orenstein, 2019; Lugosi and Mendelson, 2021), each with their own merits and shortcomings. See Lugosi and Mendelson (2019a) for an excellent summary and related literature.

In the presence of malicious noise, however, there are no simple estimators that achieve minimax error rate. Chrapanajerri et al. (2020) consider a similar problem over random vectors and propose a two-stage algorithm that is inspired by semi-definite relaxations of the results derived in Lugosi and Mendelson (2019b). However, for the univariate mean estimation, the algorithm proposed is neither efficient nor obtains sharp error bounds. The M-estimators derived in this paper alleviate both these shortcomings and achieve the minimax rate in univariate settings.

1.3. Main Results

Catoni’s estimator (Catoni, 2012) is a nearly-optimal estimator ($L = \sqrt{2} + o(1)$) in the *absence* of adversarial contamination and finite variance. This motivates extending the analysis to deal with adversarial contamination and under weaker moment assumptions. The main technical contributions are highlighted below.

1. We provide an M-estimator based on Catoni’s influence functions (Catoni, 2012) that can deal with a fraction η of arbitrary adversarial contamination. The error bounds are order optimal in η , δ and n , and achieve the best minimax error obtainable under finite variance assumption. While the tight non-asymptotic bounds under contamination are of independent interest, an interesting feature of the proposed M-estimator is that the inherent robustness offered by the influence functions facilitates a large (compared to the state of the art) tolerance to adversarial contamination.
2. The finite variance assumption is not always valid in machine learning applications, where the existence of a bound on the $(1 + \varepsilon)^{th}$ central moment is all that is

available. We extend the analysis to deal with adversarial contamination and weak moment assumptions. We first derive the minimax rate achievable under this setting, and establish that the derived Catoni’s estimator achieves this rate. The novel analysis also offers insights into the nature of tolerance to contamination of any minimax estimator as $\varepsilon \downarrow 0$. In line with the intuition, the tolerance reduces, as it becomes increasingly difficult to distinguish contaminated samples from the underlying data. We also explicitly characterize the rate of this gradual decline.

3. As an application of the developed error bounds, we demonstrate their usefulness in a multi-armed bandit application. We propose a novel best arm identification algorithm under adversarial contamination that outperforms the state of the art algorithms in terms of sample complexity and exhibits excellent empirical performance.

1.4. Other Related Work

Mean estimation under adversarial contamination is well researched topic in the robust statistics community. The contamination model introduced by Huber (1964), called the Huber’s contamination model, provides a solid framework design optimal estimations that achieve statistical efficiency and robustness simultaneously (Huber, 2004). It is not surprising that it has been widely used to model adversarial contamination in statistical learning (Chen et al., 2016; 2018; Prasad et al., 2020b; Sun et al., 2020; Laforgue et al., 2021; Bhatt et al., 2022a) and online learning (Chen et al., 2021b; Prasad et al., 2020a; Zhu et al., 2020; Bhatt et al., 2022b).

In the context of bandit applications, best arm identification algorithms under adversarial contamination based on Huber’s contamination model was first considered in Altschuler et al. (2019) and several abstract frameworks were introduced to address this challenging setting. We adopt one such framework called the Partially Identifiable Best Arm Identification (PIBAI) framework to propose algorithms using the developed estimators. While Altschuler et al. (2019) used the median to define and identify the best arm, Mukherjee et al. (2021) consider the traditional setting using the mean instead. In this work, we consider PIBAI problem in the fixed confidence setting, and seek to identify the arm having the largest mean using as few samples as possible.

2. Contamination Model

Let $\{X_i\}_{i \leq n}$ be i.i.d observations with $\mathbb{E}X_1 = \mu$. We assume that for some corruption rate $0 < \eta < 1$, an adversary may change at most ηn of these observations to arbitrary values. The resulting set of observations will be $\tilde{X}_1, \tilde{X}_2, \tilde{X}_n$,

so that

$$\sum_{i=1}^n \mathbf{1}(\tilde{X}_i \neq X_i) \leq \eta n. \quad (1)$$

The task is to estimate the true mean μ based on the observations $\tilde{X}_1, \tilde{X}_2, \dots, \tilde{X}_n$.

This strong adversarial contamination model is well-studied in machine learning; see Charikar et al. (2017); Hopkins and Li (2018); Lugosi and Mendelson (2021) and the related references for existing results. It should be noted that the well known ϵ -Huber contamination model (Huber, 2004), is just a special case of (1); see Lai et al. (2016) and Prasad et al. (2020a) for robust mean estimation procedures using this model. We will develop M-estimators to deal with the contamination model given in Eq. (1).

The contamination model Eq. (1), while allowing arbitrary contamination, is retrospective in that the adversary corrupts a fraction of the sample possibly with the knowledge of the whole data set. This is more than what the adversary can do in an online corruption setting, where the adversary can choose to contaminate based only on the observed data. So hereafter, we assume η for online settings as well without loss of generality (see Appendix B).

3. Robust M-estimator for Finite Variance

We use a special version of the estimator of the mean that was proposed by Catoni (2012). In the context of L -sub-Gaussian estimators, the estimator proposed in Catoni (2012), which is henceforth referred to as *Catoni’s estimator*, is significant owing to the fact that it is a (nearly) *optimal*³ sub-Gaussian estimator of the mean with $L = \sqrt{2} + o(1)$. This motivates extending this estimator to deal with adversarial contamination. In the following, the terminology “robust” is used to collectively describe the stability of the estimator with respect to the tail-behavior of the data, and how well it ‘tolerates’ adversarial corruption.

Start with a non-decreasing function $\psi : \mathbb{R} \rightarrow \mathbb{R}$ such that

$$-\log(1 - x + x^2/2) \leq \psi(x) \leq \log(1 + x + x^2/2) \quad (2)$$

for all $x \in \mathbb{R}$. One can choose such a function that is bounded⁴: specifically, we for some $0 < A < \infty$,

$$|\psi(x)| \leq A \text{ for all } x \in \mathbb{R}. \quad (3)$$

Robust Catoni’s Estimator: We define the estimator of μ , denoted by $\hat{\mu}$, as the solution of the equation in the variable θ

$$\sum_{i=1}^n \psi[\alpha(\tilde{X}_i - \theta)] = 0. \quad (4)$$

³Here optimal is to be understood in the sense that the Catoni’s estimator comes close the best possible $L(= \sqrt{2})$.

⁴see Eq. (67) in Appendix C.1 for an explicit representation.

Clearly, $\psi(x) = x$ corresponds to the empirical mean, though it does not satisfy Eq. (2). [Catoni \(2012\)](#) notes that the empirical mean is unduly influenced by large values when the distribution tails are not themselves sub-Gaussian. When instead $\psi(x)$ satisfies Eq. (2), it is similar to the linear function for small and moderate values of x , but its logarithmic rate of growth reduces the effect of large values. Unlike classical M-estimation ([Huber, 2004](#)) literature, [Catoni \(2012\)](#) uses unbounded influence functions to obtain sharp confidence bounds. However, we will show that when we use a bounded function instead, Catoni’s estimator can deal with adversarial contamination as well.

Theorem 3.1. *Let $\delta \in (0, 1)$ such that $\delta \geq 2e^{-n/4}$. Let $\{X_i\}_{i=1}^n$ be i.i.d random variables with mean μ and $\mathbb{E}|X_1 - \mu|^2 \leq v$. Let the corruption parameter $\eta \in [0, \frac{1}{4A}]$ for $A > 0$ in Eq. (3). Let $\Omega \in (0, \frac{1}{A\eta} - 4)$ be such that*

$$n \geq \frac{4 \log(2/\delta)}{1 - (\Omega + 4)A\eta}.$$

Robust Catoni’s M-estimator $\hat{\mu}$ with parameter

$$\alpha = \frac{1}{v^{1/2}} \left(\Omega A \eta + \frac{2 \log(2/\delta)}{n} \right)^{1/2},$$

satisfies, with probability at least $1 - \delta$,

$$|\hat{\mu} - \mu| < v^{1/2} \frac{\left(\frac{\Omega+4}{2\Omega^{1/2}} A^{1/2} \eta^{1/2} + \left(\frac{2 \log(2/\delta)}{n} \right)^{1/2} \right)}{1 - (\Omega + 4)A\eta/2 - 2 \log(2/\delta)/n}. \quad (5)$$

Moreover, with probability at least $1 - 2 \exp\left(-\frac{\eta A}{4} n\right)$, we have that

$$|\hat{\mu} - \mu| < \tilde{C} \sqrt{v\eta}, \quad (6)$$

for some \tilde{C} .

Theorem 3.1 provides a non-asymptotic error bound for the robust mean estimator in the presence of adversarial contamination. Clearly, the asymptotic bias is $O(\sqrt{v\eta})$, which is information theoretically optimal ([Hopkins and Li, 2018](#)). The error bound obtained in Eq. (5) is not the tightest possible. A slightly sharper error bound that is less explicit in the dependence on the parameters can be similarly obtained.

Corollary 3.2. *Under the same assumptions as in Theorem 3.1, with probability at least $1 - \delta$,*

$$|\hat{\mu} - \mu| < \frac{(\Omega + 4)A\eta + 4(\log(2/\delta)/n)}{2\alpha \left(1 - (\Omega + 4)A\eta/2 - 2 \log(2/\delta)/n \right)}.$$

3.1. Minimax Lower Bound

The minimax error bound for estimators under finite variance bound assumption has been derived in [Diakonikolas et al. \(2017\)](#); [Lugosi and Mendelson \(2021\)](#). In this section,

we will repeat the key ideas that will enable the proof of the minimax lower bound for variables with bounded $(1 + \varepsilon)^{th}$ moments in Sec.4.3.

Since at most η fraction of the samples can be corrupted by the adversary, [Lugosi and Mendelson \(2021\)](#) argue that the adversary can arbitrarily corrupt the values in the tail quantiles⁵ $Q_{1-\eta/2}(X - \mu)$ and $Q_{\eta/2}(X - \mu)$ to introduce the worst possible error. Then with a high probability, no estimator can perform better than

$$\max \left\{ E[|X - \mu - Q_{\eta/2}| \mathbf{1}(X - \mu \leq Q_{\eta/2})], \right. \\ \left. E[|X - \mu - Q_{1-\eta/2}| \mathbf{1}(X - \mu \geq Q_{1-\eta/2})] \right\}.$$

This combined with the sub-Gaussian error bound under no contamination obtains the best minimax error as

$$cv^{1/2} \max \left\{ \eta^{1/2}, \left((\log 2/\delta)/n \right)^{1/2} \right\}, \quad (7)$$

where c is an absolute constant. By considering a threshold on the range of n as a function of δ and η , it follows that Eq. (5) switches between a sub-Gaussian estimator and that having order optimal corruption bias Eq. (6), essentially matching the lower bound.

3.2. Comparison with Lugosi and Mendelson (2021)

[Lugosi and Mendelson \(2021\)](#) develop a mean estimator based on trimmed mean that is robust to adversarial contamination. The corruption parameter η is required to satisfy

$$8\eta + 12 \frac{\log(4/\delta)}{n} \leq \frac{1}{2}.$$

An upper bound on the corruption fraction tolerated by the robust trimmed mean estimator is $\eta < 1/16$ or around **6%**. In contrast, for the robust Catoni’s estimator in Eq. (5), the upper bound on η is $1/4A$. For $A = \log 2$ as in Eq. (67), the robust estimator can tolerate up to **36%** contamination – around *three* times higher!

3.3. Comparison with Robust Empirical Mean

[Prasad et al. \(2020a\)](#) develop an interval estimator of empirical mean that is robust to adversarial contamination. The corruption parameter η is required to satisfy

$$2\eta + \sqrt{\eta \frac{\log(4/\delta)}{n}} + \frac{\log(2/\delta)}{n} \leq \frac{1}{2}.$$

An upper bound on the corruption fraction tolerated by the robust trimmed mean estimator is $\eta < 1/4$ or at most **25%**,

⁵Here the quantile is identified as

$$Q_p(X - \mu) := \sup_{M \in \mathbb{R}} \mathbb{P}\{X - \mu \geq M\} \geq 1 - p$$

compared with the the robust Catoni's estimator in Eq. (5) that can tolerate up to **36%** contamination.

Remark 1. Here we would like to point out that [Lugosi and Mendelson \(2021\)](#) and [Prasad et al. \(2020a\)](#) did not push the requirement of η to the limit. The constants cannot be easily improved, however, using their current proof techniques. So there might be room for improvement in each of the estimators, which could be worthwhile exploring as future work.

4. Robust M-estimator for Infinite Variance

Let $\varepsilon \in (0, 1]$ and let a sequence of i.i.d random variables $\{X_i\}_{i=1}^n$ be such that $\mathbb{E}(X_1) = \mu$ and $\mathbb{E}|X_1 - \mu|^{1+\varepsilon} \leq v_\varepsilon$. For $C_\varepsilon > 0$, let $\psi : \mathbb{R} \rightarrow \mathbb{R}$ be a non-decreasing influence function such that for all $x \in \mathbb{R}$

$$-\log(1 - x + C_\varepsilon|x|^{1+\varepsilon}) \leq \psi(x) \leq \log(1 + x + C_\varepsilon|x|^{1+\varepsilon}). \quad (8)$$

[Chen et al. \(2021a\)](#) choose $C_\varepsilon = \frac{1}{\varepsilon}$ inspired by Taylor-like expansions. [Minsker \(2018, Section 3.4\)](#) choose $C_\varepsilon = \frac{\varepsilon}{1+\varepsilon} \vee \sqrt{\frac{1-\varepsilon}{1+\varepsilon}}$. Motivated by the choice of the co-efficient in Eq. (2), we choose a value that satisfies

$$(1 - x + C_\varepsilon x^{1+\varepsilon})(1 + x + C_\varepsilon x^{1+\varepsilon}) \geq 1, \quad \forall x \geq 0,$$

and is chosen as

$$C_\varepsilon = \left(\frac{\varepsilon}{1+\varepsilon}\right)^{\frac{1+\varepsilon}{2}} \left(\frac{1-\varepsilon}{\varepsilon}\right)^{\frac{1-\varepsilon}{2}}.$$

When $\varepsilon = 1$, we also recover the coefficient in [Catoni \(2012\)](#), namely $C_1 = 1/2$.

One can choose a bounded function⁶ satisfying Eq. (8): specifically, we assume that for some $0 < A_\varepsilon < \infty$,

$$|\psi(x)| \leq A_\varepsilon \text{ for all } x \in \mathbb{R}. \quad (9)$$

Robust Catoni's Estimator for $\varepsilon < 1$: As before, we define the Catoni's M-estimator $\hat{\mu}_\varepsilon$ as a solution of the equation in the variable θ , namely, $\sum_{i=1}^n \psi(\alpha(\tilde{X}_i - \theta)) = 0$ using an influence function ψ satisfying Eq. (8). If the solution is not unique, choose $\hat{\mu}_\varepsilon$ to be the median solution. Again the motivation is to reduce the effect of large values using a logarithmic function.

Theorem 4.1. *Let $\{X_i\}_{i=1}^n$ be i.i.d random variables with mean μ and $\mathbb{E}|X_1 - \mu|^{1+\varepsilon} \leq v_\varepsilon$. Fix $\delta \in (0, 1)$, $\tau > 0$, $0 < h < 1$ and $B > 0$, and let the corruption parameter $\eta \in [0, \frac{\tau^{1/\varepsilon}}{2A_\varepsilon(1+\tau)^{(1+\varepsilon)/\varepsilon}} C_\varepsilon^{-1/\varepsilon}]$ for $\varepsilon \in (0, 1]$. Let n be such that*

$$n \geq \frac{\log(2/\delta)(1 + h^{-\varepsilon} B^{1+\varepsilon} C_\varepsilon)}{\frac{\tau^{1/\varepsilon}}{(1+\tau)^{(1+\varepsilon)/\varepsilon}} (1-h) C_\varepsilon^{-1/\varepsilon} - 2A_\varepsilon \eta (1 + h^{-\varepsilon} B^{1+\varepsilon} C_\varepsilon)}.$$

⁶see Eq. (69) in Appendix C.2 for an explicit function representation.

Catoni's M-estimator $\hat{\mu}_\varepsilon$ with parameter

$$\alpha = Bv_\varepsilon^{-1/(1+\varepsilon)} \left(2A_\varepsilon \eta + \frac{\log 2/\delta}{n}\right)^{1/(1+\varepsilon)}.$$

satisfies, with probability at least $1 - \delta$,

$$|\hat{\mu}_\varepsilon - \mu| < (1 + \tau)v_\varepsilon^{1/(1+\varepsilon)} \left(2A_\varepsilon \eta + \frac{\log 2/\delta}{n}\right)^{\varepsilon/(1+\varepsilon)} \times (h^{-\varepsilon} B^\varepsilon C_\varepsilon + 1/B). \quad (10)$$

Note that the error we have obtained in Eq. (10) has the $\eta^{\varepsilon/(1+\varepsilon)}$ dependence on the corruption rate, $((\log 2/\delta)/n)^{\varepsilon/(1+\varepsilon)}$ dependence on the number of observations, and has a $v_\varepsilon^{1/(1+\varepsilon)}$ dependence on the centered moment. We will see below that this order of magnitude is optimal.

4.1. Choice of τ , h and B in Eq. (10)

Choose $\tau_{\varepsilon, \eta} > 0$ and $1 > h_{\varepsilon, \eta} > 0$, and $B > 0$ such that

$$(1 - h_{\varepsilon, \eta}) \frac{\tau_{\varepsilon, \eta}^{1/\varepsilon}}{(1 + \tau_{\varepsilon, \eta})^{(1+\varepsilon)/\varepsilon}} > (1 + h_{\varepsilon, \eta}^{-\varepsilon} B^{1+\varepsilon} C_\varepsilon) C_\varepsilon^{1/\varepsilon} 2A_\varepsilon \eta.$$

As $\eta \rightarrow 0$, we can choose $\tau_{\varepsilon, \eta}$ arbitrarily close to 0, $h_{\varepsilon, \eta}$ arbitrarily close to 1, and then choose B that minimizes the expression in the right hand side of Eq. (10) for h arbitrarily close to 1, i.e. $B = (\varepsilon C_\varepsilon)^{1/(1+\varepsilon)}$. This gives us the following result.

Corollary 4.2. *For large n and small η , with probability at least $1 - \delta$,*

$$|\hat{\mu}_\varepsilon - \mu| \leq (1 + o(1))v_\varepsilon^{1/(1+\varepsilon)}(1 + \varepsilon)\varepsilon^{-\varepsilon/(1+\varepsilon)} C_\varepsilon^{1/(1+\varepsilon)} \times \left(2A_\varepsilon \eta + \frac{\log 2/\delta}{n}\right)^{\varepsilon/(1+\varepsilon)}.$$

Note that for $\varepsilon = 1$, the result recovers the near-optimal asymptotic constant for $\eta \approx 0$ as in [Catoni \(2012\)](#) (see also [Theorem 3.1](#)) as $C_1^{1/2} = 1/\sqrt{2}$. This indicates that error bounds are nearly-tight for $\varepsilon < 1$ as well in the absence of contamination. This is not the case, for example, with the estimator proposed in [Chen et al. \(2021a\)](#) for $\eta = 0$.

4.2. Corruption tolerance under $\varepsilon < 1$

Below we present a tight upper bound on the tolerance to contamination when $\varepsilon \in (0, 1]$.

Corollary 4.3. *For any $\varepsilon \in (0, 1]$, the robust M-estimator for infinite variance can tolerate a corruption level $\eta \in [0, \Lambda(\varepsilon))$ with*

$$\Lambda(\varepsilon) := \frac{\Omega(\varepsilon)}{-2\log(1 - \Omega(\varepsilon))},$$

where $\Omega(\varepsilon) := ((1 + \varepsilon)^{(1+\varepsilon)/(2\varepsilon)}(1 - \varepsilon)^{(1-\varepsilon)/(2\varepsilon)})^{-1}$.

It turns out that, for a fixed value of the parameter $\tau > 0$, an upper bound on the contamination rate that our estimator can handle is:

$$2A_\varepsilon\eta < \frac{\tau^{1/\varepsilon}}{(1 + \tau)^{(1+\varepsilon)/\varepsilon}} C_\varepsilon^{-1/\varepsilon}.$$

The value $\tau = 1/\varepsilon$ maximizes the expression in the right hand side above, so assuming that A_ε is given in Eq. (68), the upper bound on the contamination rate $\Lambda(\varepsilon)$ given in the corollary is obtained.

ε	0.001	0.1	0.3	0.5	0.9	1
$\Lambda(\varepsilon)$ (as %)	7%	16%	22%	26%	34%	36%

Table 1. Percentage corruption tolerance

In Table 1, as ε decreases from 1 to 0, $\Omega(\varepsilon)$ increases from 1/2 to 1, and the upper bound on the contamination rate decreases from $1/(4 \log 2)$ to 0. Remarkably, even in the challenging setting when $\varepsilon = 0.001$, the proposed M-estimator tolerates around **7%** contamination.

4.3. Minimax Lower Bound for $\varepsilon < 1$

We establish now a lower bound on the best minimax error rate achievable under adversarial contamination in the weak moment setting.

Theorem 4.4. *Let $n > 0$ and $\delta \in (0, 1)$ be such that $\delta \geq 2e^{-n/4}$. Let $\varepsilon \in (0, 1]$ and distribution \mathcal{D} be such that $\mathbb{E}_{X_1 \sim \mathcal{D}} |X_1 - \mu|^{1+\varepsilon} = v_\varepsilon$. The error bound of any estimator $\hat{\mu}_n(\{X_i\}_{i \leq n})$ in the adversarial contamination setting with corruption parameter $0 < \eta < 1$ is at least*

$$qv_\varepsilon^{1/(1+\varepsilon)} \max \left\{ \eta^{\frac{\varepsilon}{1+\varepsilon}}, \left(\frac{\log(2/\delta)}{n} \right)^{\frac{\varepsilon}{1+\varepsilon}} \right\},$$

for a suitable absolute constant q .

Clearly, the error we obtain in Eq. (10) is minimax optimal in the sense of Theorem 4.4 using arguments similar to those in Sec 3.1. Essentially, there is threshold for n depending on ε, δ & η beyond which the bias dominates and below which the natural variability of the data dominates.

Remark 2. All results derived in the previous sections require the knowledge of v_ε . The extension to the unknown moment parameter v_ε can be achieved using Lepskii's method (Lepskii, 1992), which adapts to any unknown moment of the problem, by compromising on the tightness of the deviations.

5. Application: Best Arm Identification

Consider the best arm identification problem on a multi-armed bandit with $[K] := \{1, 2, \dots, K\}$ arms in a fixed confidence setting. The goal is to identify the best arm with a high probability, while providing qualitative guarantees. A primary difficulty in the contaminated setting, as opposed to the classical setting (Even-Dar et al., 2006), is that the true parameters can only be *partially identified* (Altschuler et al., 2019), even with infinite samples. That is, there is an inherent asymptotic bias associated, which needs to be taken into account while identifying the best arm. This motivates the partially identifiable best arm identification framework (PIBAI) of Altschuler et al. (2019), summarized as follows.

With each arm $i \in [K]$, associate a family of reward distributions $\mathcal{D}_i = \{D_i(\mu_i, \eta)\}_{\eta \in \mathcal{E}}$, where η represents a corruption, and \mathcal{E} is some space. The uncorrupted reward associated with arm i has mean μ_i and a centered $1 + \varepsilon$ -moment not exceeding v_ε , $0 < \varepsilon \leq 1$. Let $* = \operatorname{argmax}_{i \in [K]} \mu_i$ denote the best arm.

To simplify notation, rewrite Eq. (5) and Eq. (10) as follows. For any arm $i \in [K]$,

$$|\hat{\mu}_i(t) - \mu| \leq H_i(\eta) + G_{\varepsilon, \eta}(\delta) \left(\frac{1}{t} \right)^{\frac{\varepsilon}{1+\varepsilon}}, \quad (11)$$

where $H_i(\eta)$ and $G_{\varepsilon, \eta}(\delta)$ absorb the missing terms (see Appendix C.3).

PIBAI Model Assumptions (Altschuler et al., 2019):

- A1. Even with infinite samples $X_t(i) \sim D_i(\mu_i, \eta_t)$ for unknown $\eta_t, t = 1, 2, \dots$, it is impossible to estimate μ_i more precisely than the region $[\mu_i \pm H_i]$.
- A2. The unavoidable biases $\{H_i\}_{i \in [K]}$ are such that the *effective gaps*

$$\Delta_i := (\mu_* - H_*) - (\mu_i + H_i) \quad (12)$$

are strictly positive for each sub-optimal arm $i \neq *$.

- A3. There exists an algorithm that is δ -PAC⁷ for the given contaminated bandit instance (\mathcal{D}) .

The above assumptions are motivated by the fact that, even if any estimator computes the means of the individual arms with large tolerance in the presence of contamination, it is not guaranteed that the relative ordering between the

⁷Any PIBAI algorithm is said to be δ -PAC if it outputs an arm \hat{I} that satisfies the following with probability at least $1 - \delta$,

$$\mathbb{P} \left\{ \mu_{\hat{I}} + H_{\hat{I}} < \mu_* - H_* \right\} \leq \delta.$$

estimated means remains the same. In fact, if Eq. (12) does not hold, it can be shown that no algorithm can distinguish between the best and the second best arms, even with access to infinitely many samples (Altschuler et al., 2019).

5.1. Adversarial Elimination with Catoni

A naïve approach to identify the best arm that attains optimal order of sample complexity up to logarithmic terms is based on successive elimination. However, the standard algorithms achieve the optimal order with doubly-logarithmic terms. When $\varepsilon = 1$, Altschuler et al. (2019) argue that the most common approaches (Karnin et al., 2013; Jamieson et al., 2013) to improve order to $O(\log(\frac{K \log(1/\Delta_i)}{\delta}) \frac{1}{\Delta_i^2})$ are unsuccessful in the presence of contamination. The key issue that contributes to this shortcoming – unlike a classical successive elimination algorithm (Mukherjee et al., 2021) – is that these algorithms heavily rely on the “additive property of suboptimality” for successful identification: When the biases $\forall i, H_i = 0, \Delta_i = \Delta_j + \mu_j - \mu_i$. Clearly this is not true when $H_i \neq 0$. We next provide an improvement of successive elimination algorithm that alleviates this shortcoming and achieves better sample complexity in terms of Δ_i for $\varepsilon \in (0, 1]$, in the contamination setting.

For $i \in S$, let the *elimination criterion* for Algorithm 1 be specified as

$$\left\{ \hat{\mu}_i(t_m) + G_{\varepsilon, \eta} \left(\frac{\delta}{2K(m+1)^2} \right) \left(\frac{1}{t_m} \right)^{\frac{\varepsilon}{1+\varepsilon}} \right\} \quad (13)$$

$$< \max_{j \in S} \left\{ \hat{\mu}_j(t_m) - G_{\varepsilon, \eta} \left(\frac{\delta}{2K(m+1)^2} \right) \left(\frac{1}{t_m} \right)^{\frac{\varepsilon}{1+\varepsilon}} \right\},$$

where m is the phase index.

Theorem 5.1 (Sample Complexity). *Let $\varepsilon \in (0, 1]$. Suppose Assumptions A1-A3 hold. With probability at least $1 - \delta$, Algorithm 1 outputs $S = \{*\}$, after pulling at most*

$$\max \left\{ O \left(\sum_{i \in [K]} \log \left(\frac{K \log(1/\Delta_i)}{\delta} \right) \frac{1}{\Delta_i^{\frac{1+\varepsilon}{\varepsilon}}} \right), K t_{\text{init}} \right\},$$

samples, where Δ_i is defined as Eq. (12) using $\eta = \eta_{\text{init}}$ and t_{init} as defined in Algorithm 1.

The key intuition for reduction in sample complexity is as follows: various successive elimination algorithms were initially proposed in Even-Dar et al. (2006), and later modified to regret minimization setting in Auer and Ortner (2010). Here the length of the phase is modulated by a parameter representing unknown sub-optimality gaps ($\tilde{\Delta}$), and elimination of arm i is identified when $\tilde{\Delta} < \Delta_i/2$. While this is sufficient for the purposes of regret minimization, it falls short in terms of the reducing samples for best arm identification. The elimination parameter γ , barely greater than 1, in Algorithm 1 balances this trade-off between elimination

Algorithm 1 Adversarial Elimination with Catoni (AECat)

- 1: **Input:** $\delta, K, \sigma, \varepsilon, v_\varepsilon, \gamma (> 1)$
 - 2: **Initialization:** Set $S := [K]$, phase index $m = 0$
 - 3: Set $t_0 = 0$ and

$$t_1 := t_{\text{init}} = \max \left\{ \left(\gamma G_{\varepsilon, \eta}(\delta/2K) \right)^{\frac{1+\varepsilon}{\varepsilon}}, T_\varepsilon(\sigma, \frac{\delta}{2K}) \right\}$$
 - 4: **while** $|S| > 1$ **do**
 - 5: Increase phase index m by 1
 - 6: Sample every arm in S for $\max\{t_m - t_{m-1}, 0\}$ times
 - 7: Compute $\hat{\mu}_i(t)$ with appropriate α
 - 8: Remove all arms i from S which satisfy Eq. (13)
 - 9: Update S as the remaining arms and set $m = m + 1$
 - 10: Set $t_m = \left(\gamma^m G_{\varepsilon, \eta}(\delta/2K(m+1)^2) \right)^{\frac{1+\varepsilon}{\varepsilon}}$
 - 11: **end while**
 - 12: **Output:** S
-

and sample complexity. We establish that once the phase index is such that

$$\left(\gamma^m G_{\varepsilon, \eta}(\delta/2K(m+1)^2) \right)^{\frac{1+\varepsilon}{\varepsilon}} \geq \left\lceil \left(\frac{4G_{\varepsilon, \eta}(\frac{\delta}{4Km^2})}{\Delta_i} \right)^{\frac{1+\varepsilon}{\varepsilon}} \right\rceil,$$

then arm i will be eliminated before phase m , whence we obtain $m = O(\log_\gamma(4/\Delta_i))$.

In Algorithm 1, the Catoni’s estimate is computed at the end of each phase. To simplify algorithmic notation, let $T_\varepsilon(\sigma, \delta) = \sigma \log \frac{2}{\delta}$ to denote the minimum number of samples required in Theorem 3.1 and Theorem 4.1, where

$$\sigma = \begin{cases} \frac{4}{1 - (\Omega + 4)A\eta}, & \varepsilon = 1, \\ \frac{(1+h^{-\varepsilon}B^{1+\varepsilon}C_\varepsilon)}{\frac{\tau^{1/\varepsilon}}{(1+\tau)^{(1+\varepsilon)/\varepsilon}}(1-h)C_\varepsilon^{-1/\varepsilon} - 2A_\varepsilon\eta(1+h^{-\varepsilon}B^{1+\varepsilon}C_\varepsilon)}, & \varepsilon < 1. \end{cases}$$

A few parameters appearing in $T(\sigma, \delta)$ are suppressed in the algorithm inputs for simplicity. The choices of the missing parameters are guided by the discussion in Sec. 4.1 for implementation purposes. The choice of α in Step 9 is further guided by Theorem 3.1 and Theorem 4.1 for $\varepsilon = 1$ and $\varepsilon < 1$ cases, respectively.

There are two main drawbacks of the traditional *successive elimination-type* algorithm when using Catoni’s estimator: (i) Unlike the proposed AECat method, which is phase-based, the mean estimation by root finding needs to be performed at every time for all the non-eliminated arms. We find it extremely computationally inefficient in our experiments. (ii) The sample complexity is at best $O(\log(K/\delta\Delta_i)/\Delta_i^{1+\varepsilon/\varepsilon})$ as opposed to $O(\log(\frac{K \log(1/\Delta_i)}{\delta})/\Delta_i^{1+\varepsilon/\varepsilon})$ using AECat, and this translates to better empirical performance. When Δ_i is sufficiently small, the successive elimination-type method exhibits a very bad empirical performance.

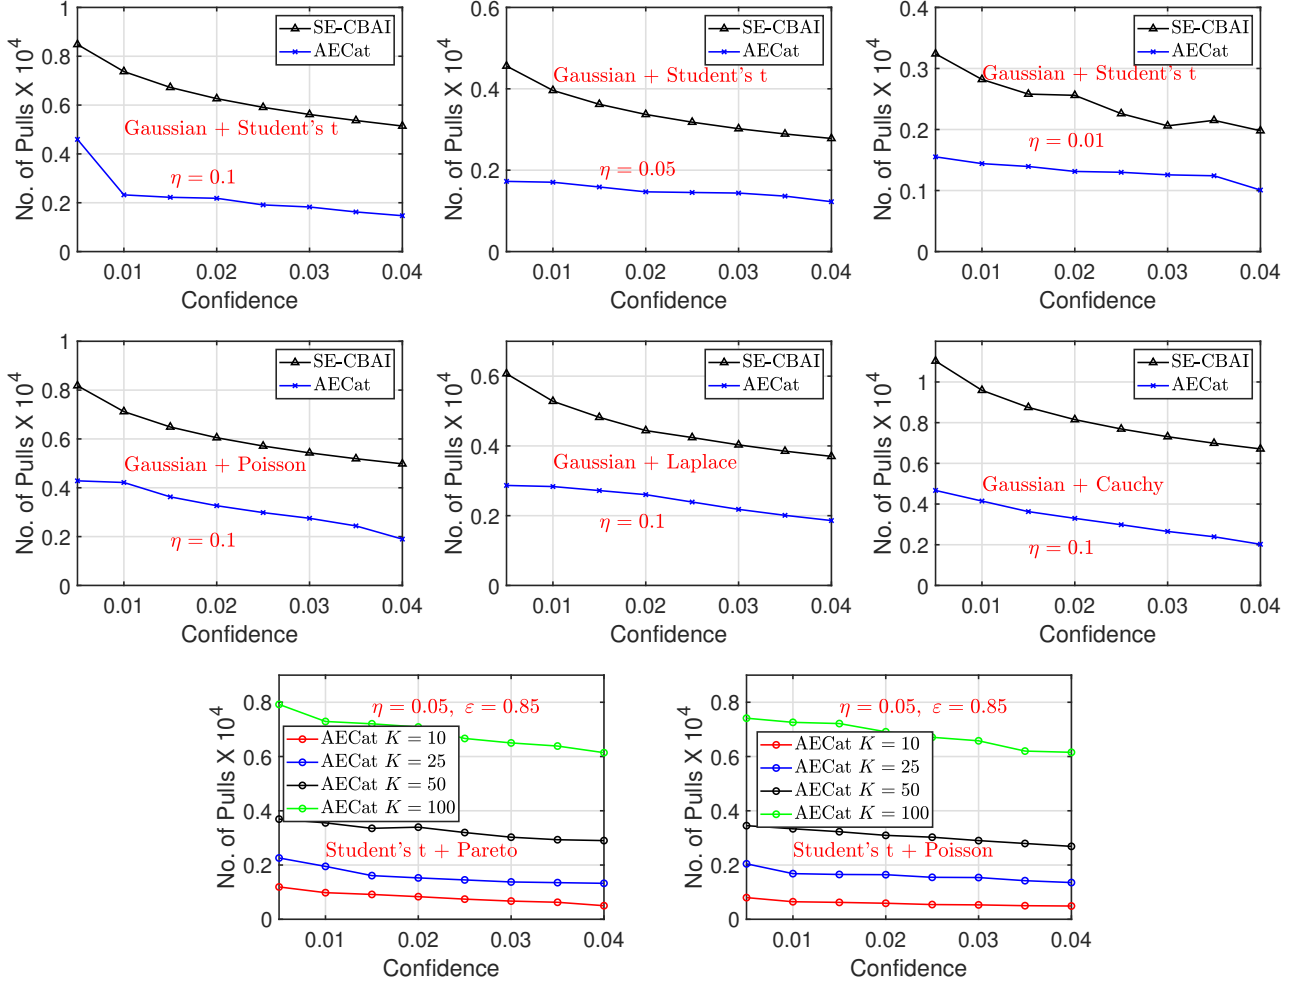


Figure 1. Average sample complexity over 50 iterations for the two algorithms: SE-CBAI as proposed in Mukherjee et al. (2021) and Algorithm 1 shortened as AECat. The true/ uncontaminated distribution is taken to be Gaussian, and the sensitivity w.r.t to contamination distribution and fraction of contamination is shown. As expected, the sample complexity decreases with lower contamination levels. True mean values are chosen as $\mu_k = 2 - (i/K)^{0.1}$ for $0 \leq i < K$. The top six figures correspond to $\varepsilon = 1$. The last two figures correspond to best arm identification in heavy-tail settings under contamination, and use $\varepsilon = 0.85$ for a bound of $v_\varepsilon = 50$. These indicate how the number of arms affect the identification as a function of δ .

5.2. Experimental Results

In this section, we describe the experimental setup designed to evaluate the performance of the proposed algorithm against the existing baseline.

Mukherjee et al. (2021) propose a successive elimination algorithm (SE-CBAI) for contaminated best arm identification *sub-Gaussian* setting, using a suitable trimmed mean estimator for robustness and confidence bounds adjusted to provide good sample complexity. Since SE-CBAI is proposed in the sub-Gaussian setting, we provide the performance comparison as shown in this setting, while allowing the contamination distribution to be selected from common models of noise. As the implementation details of SE-CBAI and technical details of another gap based al-

gorithm (G-CBAI) – for the asymptotic setting ($\delta \downarrow 0$) – are not clear, we tuned the parameters that reflect the obtained performance in the paper and use that throughout for comparison. The hyper-parameters ($\Omega, \tau, B, h, \gamma$) in Algorithm 1 are tuned as follows: To compute the Catoni’s estimator Eq. (4), for $\varepsilon = 1$, we need to calculate α , which depends on $\Omega \in (0, 1/A\eta - 4)$. Smaller Ω results in smaller initial exploration, while increasing the magnitude of $H(\eta)$ in Eq. (5) – a quantity of interest from assumption A2. We choose $\Omega = 0.25 \cdot (1/A\eta - 4)$. The factor $\gamma > 1$ in Algorithm 1 that controls the exploration should be chosen barely greater than 1 for good performance, and we choose $\gamma = 1.01$. From Eq. (10), for $\varepsilon < 1$, we can choose $h \geq 0.5$, and a $\tau \in (0, 2)$ that obtains a valid but large B (note that this affects the error bound in Eq. (10)).

We take $h = 0.5$, $\tau = 1.2$ and use $B = 0.8$. While the algorithm has more inputs than a typical successive elimination algorithm, it should be noted that tuning is straightforward here as we know the trade-offs. Results under different settings are summarized in Figure 1. Our method is uniformly better than SE-CBAI under various scenarios.

6. Generalization to Unknown η

Suppose, the true contamination level η_{true} is significantly smaller than the upper bound η , then the upper bounds given in the main theorems may be larger than desirable. We will use the idea proposed in Jain et al. (2022) to deal with an unknown or a loose η . Jain et al. (2022) do not address probabilistic statements, and we contribute to this literature. In this section, we only consider the finite variance case, i.e., $\varepsilon = 1$. The extension to the infinite variance case is easily derived using similar constructions.

Let $0 < \eta_{\min} < \eta < \frac{1}{4A}$ be a threshold to be decided upon in the sequel. Choose a number $\theta \in (0, 1)$ and consider the following sequential contamination bounds

$$\eta_k = \eta\theta^k, \quad k = 0, 1, \dots, J, \quad (14)$$

where

$$J = \min\{k \geq 1 : \eta\theta^k \leq \eta_{\min}\}. \quad (15)$$

Let δ_k for $k = 0, 1, \dots, J$ be confidence levels. For every k , we compute the estimator $\hat{m}(\eta_k, \delta_k)$ in the same way as derived in Theorem 3.1 to obtain a confidence interval as

$$I_k = (\hat{m}(\eta_k, \delta_k) - B(\eta_k, \delta_k), \hat{m}(\eta_k, \delta_k) + B(\eta_k, \delta_k)),$$

where

$$B(\eta_k, \delta_k) = v^{1/2} \frac{(K+4)/(2K^{1/2})A^{1/2}\eta^{1/2} + ((2\log 2/\delta)/n)^{1/2}}{1 - (K+4)A\eta/2 - 2(\log 2/\delta)/n}. \quad (16)$$

By using interval sequences $\{I_k\}$, we can construct a tighter estimator as follows. We first define index J_0 to be

$$J_0 := \max\{k = 0, \dots, J : \cap_{j=0}^k I_j \neq \emptyset\}. \quad (17)$$

The desired estimator is defined as

$$\hat{\mu} := \hat{m}(\eta_{J_0}, \delta_{J_0}). \quad (18)$$

In particular, we can choose the confidence sequence $\{\delta_k\} = \{2^{-k}\delta\}$ for convenience. For θ , we can set it as $1/4$. In terms of η_{\min} , we can set

$$\eta_{\min} = \min \left\{ \frac{2\log(2/\delta)/n}{(\Omega+4)^2 A/(4\Omega)}, \eta \right\}.$$

The choice of such η_{\min} has the following advantages.

1. If $\eta_{\text{true}} \geq \eta_{\min}$, we can guarantee that $\eta_{J_0} \leq 4\eta_{\text{true}}$ with a high probability. This implies an error order of $O(\eta_{\text{true}}^{1/2})$.
2. If $\eta_{\text{true}} < \eta_{\min}$, we know that $\hat{\mu}$ will have error of order $O(\eta_{\min}^{1/2})$. Thanks to the choice of η_{\min} , the error will never exceed $O((\frac{\log(2/\delta)}{n})^{1/2})$.

The above arguments imply that $\hat{\mu}$ in (18) enjoys an error bound of $O(\max\{\eta_{\text{true}}^{1/2}, (\frac{\log(2/\delta)}{n})^{1/2}\})$. Formally, we have the following result.

Theorem 6.1. *Let $\delta \in (0, 1)$ such that $\delta \geq 2e^{-n/4}$. Let $\{X_i\}_{i=1}^n$ be i.i.d. random variables with mean μ and $\mathbb{E}|X_1 - \mu|^2 \leq v$. Suppose the true corruption parameter $\eta_{\text{true}} \in [0, \frac{1}{4A}]$. Let Ω satisfy*

$$(\Omega + 4)A\eta + 4(\log 2/\delta)/n + \frac{4\log 2}{n\log 4} \left(\log [(\Omega + 4)^2 A\eta/(4\Omega)] + \log n \right) \leq 1. \quad (19)$$

Then the estimator defined in (18) satisfies

$$|\hat{\mu} - \mu| \leq v^{1/2} \frac{\frac{(\Omega+4)}{2\Omega^{1/2}} A^{1/2} \eta_{\text{true}}^{1/2} + (\frac{2J_0 \log 2}{n})^{1/2} + 2(\frac{2\log(2/\delta)}{n})^{1/2}}{1 - (\Omega + 4)A\eta_{\text{true}}/2 - 2J_0 \log(2)/n - 2\log(2/\delta)/n}. \quad (20)$$

with probability at least $1 - 2\delta$.

The error bound in (20) is free of choice of η , which makes our theoretical results more appealing in the practical problems.

7. Conclusion

We provided a minimax M-estimator based on influence functions inspired by Catoni (2012), which is known to be nearly-optimal in the absence of contamination. In the adversarial contamination setting, the proposed M-estimator tolerates more corruption than the state of the art estimators and achieves the minimax error rate both in the finite ($\varepsilon = 1$) and infinite variance ($\varepsilon < 1$) setting. We also explicitly characterized the maximum tolerance as a function of ε for the proposed estimators. We then proposed a novel best arm identification algorithm in the contaminated setting, that works in both bounded and heavy-tailed settings, and achieves better theoretical sample complexity and empirical performance than the state of the art. Finally, we extend the minimax estimation procedure to incorporate the challenging setting where a tight upper bound on the corruption level η is unknown, greatly improving the applicability of the proposed minimax M-estimators in applications.

Currently, we require the knowledge of ε for obtaining the bounds. A recent work (Ashutosh et al., 2021) discusses adhoc ways of choosing ε that obtains a decent compromise. Extending these ideas for minimax M-estimation might be a worthwhile avenue for further exploration.

References

- Alekh Agarwal, Nan Jiang, Sham M Kakade, and Wen Sun. Reinforcement learning: Theory and algorithms. 2019.
- Jason Altschuler, Victor-Emmanuel Brunel, and Alan Malek. Best arm identification for contaminated bandits. *J. Mach. Learn. Res.*, 20(91):1–39, 2019.
- Kumar Ashutosh, Jayakrishnan Nair, Anmol Kagrecha, and Krishna P. Jagannathan. Bandit algorithms: Letting go of logarithmic regret for statistical robustness. In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 622–630, Virtual Event, 2021.
- Peter Auer and Nicolo Cesa-Bianchi. On-line learning with malicious noise and the closure algorithm. *Annals of mathematics and artificial intelligence*, 23(1):83–99, 1998.
- Peter Auer and Ronald Ortner. Ucb revisited: Improved regret bounds for the stochastic multi-armed bandit problem. *Periodica Mathematica Hungarica*, 61(1-2):55–65, 2010.
- Sujay Bhatt, Guanhua Fang, and Ping Li. Offline change detection under contamination. In *Proceedings of the Thirty-Eighth Conference on Uncertainty in Artificial Intelligence (UAI)*, Eindhoven, The Netherlands, 2022a.
- Sujay Bhatt, Guanhua Fang, Ping Li, and Gennady Samorodnitsky. Nearly optimal catoni’s M-estimator for infinite variance. In *Proceedings of the 39th International Conference on Machine Learning (ICML)*, Bartimore, MD, 2022b.
- Christian Brownlees, Emilien Joly, and Gábor Lugosi. Empirical risk minimization for heavy-tailed losses. *The Annals of Statistics*, 43(6):2507–2536, 2015.
- Sébastien Bubeck, Nicolò Cesa-Bianchi, and Gábor Lugosi. Bandits with heavy tail. *IEEE Trans. Inf. Theory*, 59(11):7711–7717, 2013.
- Valerii Vladimirovich Buldygin and YU V Kozachenko. *Metric characterization of random variables and random processes*, volume 188. American Mathematical Soc., 2000.
- Olivier Catoni. Challenging the empirical mean and empirical variance: a deviation study. In *Annales de l’IHP Probabilités et statistiques*, volume 48, pages 1148–1185, 2012.
- Moses Charikar, Jacob Steinhardt, and Gregory Valiant. Learning from untrusted data. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 47–60, Montreal, Canada, 2017.
- Mengjie Chen, Chao Gao, and Zhao Ren. A general decision theory for huber’s ϵ -contamination model. *Electronic Journal of Statistics*, 10(2):3752–3774, 2016.
- Mengjie Chen, Chao Gao, and Zhao Ren. Robust covariance and scatter matrix estimation under huber’s contamination model. *The Annals of Statistics*, 46(5):1932–1960, 2018.
- Peng Chen, Xinghu Jin, Xiang Li, and Lihu Xu. A generalized catoni’s m-estimator under finite α -th moment assumption with $\alpha \in (1, 2)$. *Electronic Journal of Statistics*, 15(2):5523–5544, 2021a.
- Sitan Chen, Frederic Koehler, Ankur Moitra, and Morris Yau. Online and distribution-free robustness: Regression and contextual bandits with huber contamination. In *Proceedings of the 62nd IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 684–695, Denver, CO, 2021b.
- Yeshwanth Cherapanamjeri, Nilesh Tripuraneni, Peter L Bartlett, and Michael I Jordan. Optimal mean estimation without a variance. *arXiv preprint arXiv:2011.12433*, 2020.
- Luc Devroye, Matthieu Lerasle, Gábor Lugosi, and Roberto I Oliveira. Sub-gaussian mean estimators. *The Annals of Statistics*, 44(6):2695–2725, 2016.
- Ilias Diakonikolas, Gautam Kamath, Daniel M. Kane, Jerry Li, Ankur Moitra, and Alistair Stewart. Being robust (in high dimensions) can be practical. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, pages 999–1008, Sydney, Australia, 2017.
- Ilias Diakonikolas, Daniel M. Kane, and Alistair Stewart. Learning geometric concepts with nasty noise. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 1061–1073, Los Angeles, CA, 2018.
- Anders Eklund, Thomas E Nichols, and Hans Knutsson. Cluster failure: Why fmri inferences for spatial extent have inflated false-positive rates. *Proceedings of the national academy of sciences*, 113(28):7900–7905, 2016.
- Eyal Even-Dar, Shie Mannor, and Yishay Mansour. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *J. Mach. Learn. Res.*, 7:1079–1105, 2006.
- 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 (STOC)*, pages 1021–1034, Los Angeles, CA, 2018.

- Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 49(2):1055–1080, 2021.
- Peter J Huber. Robust estimation of a location parameter. *The Annals of Mathematical Statistics*, 35(1):73–101, 1964.
- Peter J Huber. *Robust statistics*, volume 523. John Wiley & Sons, 2004.
- Ayush Jain, Alon Orlitsky, and Vaishakh Ravindrakumar. Robust estimation algorithms don’t need to know the corruption level. *arXiv preprint arXiv:2202.05453*, 2022.
- Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani. *An introduction to statistical learning*, volume 112. Springer, 2013.
- Kevin Jamieson, Matthew Malloy, Robert Nowak, and Sebastien Bubeck. On finding the largest mean among many. *arXiv preprint arXiv:1306.3917*, 2013.
- Zohar Shay Karnin, Tomer Koren, and Oren Somekh. Almost optimal exploration in multi-armed bandits. In *Proceedings of the 30th International Conference on Machine Learning (ICML)*, pages 1238–1246, Atlanta, GA, 2013.
- Pierre Laforgue, Guillaume Staerman, and Stéphan Cléménçon. Generalization bounds in the presence of outliers: a median-of-means study. In *Proceedings of the 38th International Conference on Machine Learning (ICML)*, pages 5937–5947, Virtual Event, 2021.
- Kevin A. Lai, Anup B. Rao, and Santosh S. Vempala. Agnostic estimation of mean and covariance. In *Proceedings of the IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 665–674, New Brunswick, NJ, 2016.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Guillaume Lecué and Matthieu Lerasle. Robust machine learning by median-of-means: theory and practice. *The Annals of Statistics*, 48(2):906–931, 2020.
- OV Lepskii. Asymptotically minimax adaptive estimation. I: Upper bounds. optimally adaptive estimates. *Theory of Probability & Its Applications*, 36(4):682–697, 1992.
- Ping Li. Stable random projections and conditional random sampling, two sampling techniques for modern massive datasets, PhD Thesis. 2007. URL https://hastie.su.domains/THESES/pingli_thesis.pdf.
- Gábor Lugosi and Shahar Mendelson. Mean estimation and regression under heavy-tailed distributions: A survey. *Found. Comput. Math.*, 19(5):1145–1190, 2019a.
- Gábor Lugosi and Shahar Mendelson. Sub-gaussian estimators of the mean of a random vector. *The Annals of Statistics*, 47(2):783–794, 2019b.
- Gábor Lugosi and Shahar Mendelson. Robust multivariate mean estimation: the optimality of trimmed mean. *The Annals of Statistics*, 49(1):393–410, 2021.
- Stanislav Minsker. Sub-gaussian estimators of the mean of a random matrix with heavy-tailed entries. *The Annals of Statistics*, 46(6A):2871–2903, 2018.
- Stanislav Minsker. Distributed statistical estimation and rates of convergence in normal approximation. *Electronic Journal of Statistics*, 13(2):5213–5252, 2019.
- Arpan Mukherjee, Ali Tajer, Pin-Yu Chen, and Payel Das. Best arm identification in contaminated stochastic bandits. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 9651–9662, virtual, 2021.
- Roberto I Oliveira and Paulo Orenstein. The sub-gaussian property of trimmed means estimators. Technical report, Technical report, IMPA, 2019.
- Adarsh Prasad, Sivaraman Balakrishnan, and Pradeep Ravikumar. A robust univariate mean estimator is all you need. In *Proceedings of the 23rd International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 4034–4044, Online [Palermo, Sicily, Italy], 2020a.
- 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(3):601–627, 2020b.
- Qiang Sun, Wen-Xin Zhou, and Jianqing Fan. Adaptive huber regression. *Journal of the American Statistical Association*, 115(529):254–265, 2020.
- Yinglun Zhu, Sumeet Katariya, and Robert D. Nowak. Robust outlier arm identification. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*, pages 11566–11575, Virtual Event, 2020.

A. Proofs of Main Results

A.1. Proof of Theorem 3.1

For $\alpha > 0$ define two functions of $\theta \in \mathbb{R}$:

$$r(\theta) = \frac{1}{\alpha n} \sum_{i=1}^n \psi(\alpha(X_i - \theta)), \quad \tilde{r}(\theta) = \frac{1}{\alpha n} \sum_{i=1}^n \psi(\alpha(\tilde{X}_i - \theta)). \quad (21)$$

Note that by Eq. (1) and Eq. (3),

$$|r(\theta) - \tilde{r}(\theta)| \leq 2A\eta/\alpha \quad (22)$$

for any $\theta \in \mathbb{R}$. As in Catoni (2012), for $\theta \in \mathbb{R}$, we define

$$B_+(\theta) = \mu - \theta + \frac{\alpha}{2} [v + (\mu - \theta)^2] + \frac{\log 2/\delta}{n\alpha},$$

$$B_-(\theta) = \mu - \theta - \frac{\alpha}{2} [v + (\mu - \theta)^2] - \frac{\log 2/\delta}{n\alpha}.$$

Let P denote a probability measure that encompasses the randomness in the observations X_1, \dots, X_n as well as any randomness that may be used by the adversary in corrupting the sample. It follows from Eq. (22) that for any θ

$$P(\tilde{r}(\theta) \geq B_+(\theta) + 2A\eta/\alpha) \leq P(r(\theta) \geq B_+(\theta)) \leq \delta/2,$$

$$P(\tilde{r}(\theta) \leq B_-(\theta) - 2A\eta/\alpha) \leq P(r(\theta) \leq B_-(\theta)) \leq \delta/2. \quad (23)$$

Let θ_+ be the smallest solution of the equation

$$0 = B_+(\theta) + 2A\eta/\alpha = \mu - \theta + \frac{\alpha}{2} [v + (\mu - \theta)^2] + \frac{\log 2/\delta}{n\alpha} + \frac{2A\eta}{\alpha} \quad (24)$$

and let θ_- be the largest solution of the equation

$$0 = B_-(\theta) - 2A\eta/\alpha = \mu - \theta - \frac{\alpha}{2} [v + (\mu - \theta)^2] - \frac{\log 2/\delta}{n\alpha} - \frac{2A\eta}{\alpha}, \quad (25)$$

provided, of course, that these solutions exist.

Let us concentrate first at θ_+ . Denoting $x = \theta - \mu$, the equation Eq. (24) becomes

$$\frac{\alpha}{2} x^2 - x + \left(\frac{\alpha}{2} v + \frac{\log 2/\delta}{n\alpha} + \frac{2A\eta}{\alpha} \right) = 0, \quad (26)$$

and, if real solutions exist, then the smallest such solution is given by

$$x_+ = \frac{1 - \sqrt{1 - \alpha^2 v - 4A\eta - 2(\log 2/\delta)/n}}{\alpha}. \quad (27)$$

Of course, for a real solution to exist, the expression under the square root must be non-negative. This immediately says that our approach may work only if the corruption level η satisfies

$$\eta < 1/(4A). \quad (28)$$

With $A = \log 2$, the limit on the corruption level becomes $\eta < 0.36$.

Assuming that Eq. (28) holds, a real solution to the equation Eq. (26) exists whenever

$$\alpha^2 v + 2(\log 2/\delta)/n \leq 1 - 4A\eta. \quad (29)$$

In that case x_+ in Eq. (27) is real, and

$$\begin{aligned} 0 \leq x_+ &= \frac{\alpha v + 4A\eta/\alpha + 2(\log 2/\delta)/(\alpha n)}{1 + \sqrt{1 - \alpha^2 v - 4A\eta - 2(\log 2/\delta)/n}} \\ &\leq \frac{\alpha v/2 + 2A\eta/\alpha + (\log 2/\delta)/(\alpha n)}{1 - \alpha^2 v/2 - 2A\eta - (\log 2/\delta)/n}. \end{aligned}$$

That is,

$$\theta_+ \leq \mu + \frac{\alpha v/2 + 2A\eta/\alpha + (\log 2/\delta)/(\alpha n)}{1 - \alpha^2 v/2 - 2A\eta - (\log 2/\delta)/n}. \quad (30)$$

We need to choose α . Choose Ω so that

$$0 < \Omega < 1/(A\eta) - 4,$$

and let

$$\alpha = \frac{1}{v^{1/2}} \left(\Omega A\eta + \frac{2 \log 2/\delta}{n} \right)^{1/2}.$$

Then Eq. (29) holds whenever

$$(\Omega + 4)A\eta + 4(\log 2/\delta)/n \leq 1. \quad (31)$$

It follows from Eq. (24) that

$$\begin{aligned} \theta_+ &\leq \mu + \frac{(v^{1/2}/2)((\Omega + 4)A\eta + 4(\log 2/\delta)/n)/(\Omega A\eta + \frac{2 \log 2/\delta}{n})^{1/2}}{1 - (\Omega + 4)A\eta/2 - 2(\log 2/\delta)/n} \\ &\leq \mu + v^{1/2} \frac{(\Omega + 4)/(2\Omega^{1/2})A^{1/2}\eta^{1/2} + ((2 \log 2/\delta)/n)^{1/2}}{1 - (\Omega + 4)A\eta/2 - 2(\log 2/\delta)/n}. \end{aligned} \quad (32)$$

It follows by the monotonicity of the function ψ that, with probability at least $1 - \delta/2$,

$$\hat{\mu} - \mu \leq v^{1/2} \frac{(\Omega + 4)/(2\Omega^{1/2})A^{1/2}\eta^{1/2} + ((2 \log 2/\delta)/n)^{1/2}}{1 - (\Omega + 4)A\eta/2 - 2(\log 2/\delta)/n}. \quad (33)$$

Performing the same analysis with θ_- , we conclude that, with probability at least $1 - \delta/2$,

$$\mu - \hat{\mu} \leq v^{1/2} \frac{(\Omega + 4)/(2K^{1/2})A^{1/2}\eta^{1/2} + ((2 \log 2/\delta)/n)^{1/2}}{1 - (\Omega + 4)A\eta/2 - 2(\log 2/\delta)/n}. \quad (34)$$

Combining Eq. (33) with Eq. (34), we conclude that, with probability at least $1 - \delta$,

$$|\hat{\mu} - \mu| \leq v^{1/2} \frac{(\Omega + 4)/(2K^{1/2})A^{1/2}\eta^{1/2} + ((2 \log 2/\delta)/n)^{1/2}}{1 - (\Omega + 4)A\eta/2 - 2(\log 2/\delta)/n}. \quad (35)$$

The second part of the theorem follows from using $n \geq \frac{4}{A\eta} \log(2/\delta)$. By routine algebraic manipulations, we obtain

$$|\hat{\mu} - \mu| < \tilde{C} \sqrt{\eta v},$$

where $\tilde{C} := \frac{(\Omega+4)/(2K^{1/2})A^{1/2}\eta^{1/2} + \frac{1}{\sqrt{2}}}{1 - (\Omega+4)A\eta/2 - \frac{A\eta}{2}}$.

A.2. Proof of Corollary 3.2

The result follows from Eq. (32) in Theorem 3.1 and the definition of α .

A.3. Proof of Theorem 4.1

Initial analysis is similar to Theorem 3.1. Consider the following convexity upper bound as follows. For $a, b \geq 0$ and $0 < h < 1$,

$$\begin{aligned} (a+b)^{1+\varepsilon} &= \left(h \frac{a}{h} + (1-h) \frac{b}{1-h} \right)^{1+\varepsilon}, \\ &\leq h \left(\frac{a}{h} \right)^{1+\varepsilon} + (1-h) \left(\frac{b}{1-h} \right)^{1+\varepsilon} = \frac{a^{1+\varepsilon}}{h^\varepsilon} + \frac{b^{1+\varepsilon}}{(1-h)^\varepsilon}. \end{aligned} \quad (36)$$

Therefore, for any $0 < h < 1$,

$$\mathbb{E}|X_1 - \theta|^{1+\varepsilon} \leq h^{-\varepsilon} \mathbb{E}|X_1 - \mu|^{1+\varepsilon} + (1-h)^{-\varepsilon} |\mu - \theta|^{1+\varepsilon}. \quad (37)$$

For $0 < h < 1$, define

$$\begin{aligned} B_+(\theta) &= (\mu - \theta) + h^{-\varepsilon} 2^\varepsilon C_\varepsilon \alpha^\varepsilon v_\varepsilon + (1-h)^{-\varepsilon} 2^\varepsilon C_\varepsilon \alpha^\varepsilon |\mu - \theta|^{1+\varepsilon} + \frac{\log 2/\delta}{\alpha n}, \\ B_-(\theta) &= (\mu - \theta) - h^{-\varepsilon} 2^\varepsilon C_\varepsilon \alpha^\varepsilon v_\varepsilon - (1-h)^{-\varepsilon} 2^\varepsilon C_\varepsilon \alpha^\varepsilon |\mu - \theta|^{1+\varepsilon} - \frac{\log 2/\delta}{\alpha n}. \end{aligned} \quad (38)$$

Observe that the bounds Eq. (23) still hold.

Now let θ_+ be the smallest solution of the equation

$$\begin{aligned} 0 &= B_+(\theta) + 2A_\varepsilon \eta / \alpha \\ &= \mu - \theta + h^{-\varepsilon} 2^\varepsilon C_\varepsilon \alpha^\varepsilon v_\varepsilon + (1-h)^{-\varepsilon} 2^\varepsilon C_\varepsilon \alpha^\varepsilon |\mu - \theta|^{1+\varepsilon} + \frac{\log 2/\delta}{\alpha n} + \frac{2A_\varepsilon \eta}{\alpha} \end{aligned} \quad (39)$$

and let θ_- be the largest solution of the equation

$$\begin{aligned} 0 &= B_-(\theta) - 2A_\varepsilon \eta / \alpha \\ &= \mu - \theta - h^{-\varepsilon} 2^\varepsilon C_\varepsilon \alpha^\varepsilon v_\varepsilon - (1-h)^{-\varepsilon} 2^\varepsilon C_\varepsilon \alpha^\varepsilon |\mu - \theta|^{1+\varepsilon} - \frac{\log 2/\delta}{\alpha n} - \frac{2A_\varepsilon \eta}{\alpha}, \end{aligned} \quad (40)$$

assuming such solutions exist.

Concentrating first on θ_+ , we note that, if the equation in Eq. (39) has real roots, they are larger than μ . Setting $x = \theta - \mu$, $\theta \geq \mu$, the equation in Eq. (39) can be written in the form

$$Kx^{1+\varepsilon} - x + M = 0 \quad (41)$$

with

$$K = (1-h)^{-\varepsilon} \alpha^\varepsilon C_\varepsilon, \quad M = h^{-\varepsilon} \alpha^\varepsilon C_\varepsilon v_\varepsilon + \frac{\log 2/\delta}{\alpha n} + \frac{2A_\varepsilon \eta}{\alpha}.$$

As in the paper, we denote

$$D = K^{1/\varepsilon} M,$$

and $y = K^{1/\varepsilon} x$, so the equation Eq. (41) becomes

$$y^{1+\varepsilon} - y + D = 0, \quad y \geq 0.$$

Assuming that for some $\tau > 0$ the condition

$$D \leq \frac{\tau^{1/\varepsilon}}{(1+\tau)^{(1+\varepsilon)/\varepsilon}} \quad (42)$$

holds, we know that

$$\theta_+ - \mu \leq (1 + \tau)M,$$

so that

$$\theta_+ - \mu \leq (1 + \tau) \left(h^{-\varepsilon} \alpha^\varepsilon C_\varepsilon v_\varepsilon + \frac{\log 2/\delta}{\alpha n} + \frac{2A_\varepsilon \eta}{\alpha} \right) \quad (43)$$

and the condition Eq. (42) has the following explicit form;

$$h^{-\varepsilon} \alpha^{1+\varepsilon} C_\varepsilon v_\varepsilon + 2A_\varepsilon \eta + \frac{\log 2/\delta}{n} \leq \frac{\tau^{1/\varepsilon}}{(1 + \tau)^{(1+\varepsilon)/\varepsilon}} (1 - h) C_\varepsilon^{-1/\varepsilon}. \quad (44)$$

Note that this puts an upper bound on the contamination rate our estimator can tolerate:

$$2A_\varepsilon \eta < \frac{\tau^{1/\varepsilon}}{(1 + \tau)^{(1+\varepsilon)/\varepsilon}} C_\varepsilon^{-1/\varepsilon}.$$

The value $\tau = 1/\varepsilon$ maximizes the expression in the right hand side above, so assuming that A_ε is given in Eq. (68), the bound on the contamination rate is

$$\begin{aligned} \eta &< \frac{\varepsilon/(1 + \varepsilon)^{(1+\varepsilon)/\varepsilon}}{2A_\varepsilon C_\varepsilon^{1/\varepsilon}} \\ &= \left\{ -2 \log \left[1 - ((1 + \varepsilon)^{(1+\varepsilon)/(2\varepsilon)} (1 - \varepsilon)^{(1-\varepsilon)/(2\varepsilon)})^{-1} \right] (1 + \varepsilon)^{(1+\varepsilon)/(2\varepsilon)} (1 - \varepsilon)^{(1-\varepsilon)/(2\varepsilon)} \right\}^{-1} \\ &= \frac{K(\varepsilon)}{-2 \log(1 - K(\varepsilon))}, \end{aligned} \quad (45)$$

where

$$K(\varepsilon) = ((1 + \varepsilon)^{(1+\varepsilon)/(2\varepsilon)} (1 - \varepsilon)^{(1-\varepsilon)/(2\varepsilon)})^{-1}. \quad (46)$$

Assuming that the contamination rate η satisfies Eq. (45), condition Eq. (44) holds for some choices of α , h and τ and for n large enough. In that case, as in the paper we conclude that, with probability at least $1 - \delta/2$,

$$\hat{\mu} - \mu \leq (1 + \tau) \left(h^{-\varepsilon} \alpha^\varepsilon C_\varepsilon v_\varepsilon + \frac{\log 2/\delta}{\alpha n} + \frac{2A_\varepsilon \eta}{\alpha} \right). \quad (47)$$

Performing the same analysis with θ_- , we conclude that, with probability at least $1 - \delta/2$,

$$\mu - \hat{\mu} \leq (1 + \tau) \left(h^{-\varepsilon} \alpha^\varepsilon C_\varepsilon v_\varepsilon + \frac{\log 2/\delta}{\alpha n} + \frac{2A_\varepsilon \eta}{\alpha} \right). \quad (48)$$

It follows from Eq. (47) and Eq. (48) that, with probability at least $1 - \delta$,

$$|\hat{\mu} - \mu| \leq (1 + \tau) \left(h^{-\varepsilon} \alpha^\varepsilon C_\varepsilon v_\varepsilon + \frac{\log 2/\delta}{\alpha n} + \frac{2A_\varepsilon \eta}{\alpha} \right) \quad (49)$$

whenever Eq. (44) holds.

Next, we address the choice of α , τ , h . For $B > 0$ can choose

$$\alpha = B v_\varepsilon^{-1/(1+\varepsilon)} \left(2A_\varepsilon \eta + \frac{\log 2/\delta}{n} \right)^{1/(1+\varepsilon)}. \quad (50)$$

Then the bound Eq. (49) becomes

$$|\hat{\mu} - \mu| \leq (1 + \tau) v_\varepsilon^{1/(1+\varepsilon)} \left(2A_\varepsilon \eta + \frac{\log 2/\delta}{n} \right)^{\varepsilon/(1+\varepsilon)} (h^{-\varepsilon} B^\varepsilon C_\varepsilon + 1/B), \quad (51)$$

while the constraint Eq. (44) takes the form

$$(1 + h^{-\varepsilon} B^{1+\varepsilon} C_\varepsilon) \left(2A_\varepsilon \eta + \frac{\log 2/\delta}{n} \right) \leq \frac{\tau^{1/\varepsilon}}{(1 + \tau)^{(1+\varepsilon)/\varepsilon}} (1 - h) C_\varepsilon^{-1/\varepsilon}. \quad (52)$$

A.4. Proof of Corollary 4.3

The proof follows from Eq. (45) in Theorem 4.1.

A.5. Proof of Theorem 4.4

Arguing as in Lugosi and Mendelson (2021), the lower bound on the error one can get with probability at least $1 - \delta$ is

$$\text{const.} \max \left\{ E[|X - \mu - Q_{\eta/2}| \mathbf{1}(X - \mu \leq Q_{\eta/2})], E[|X - \mu - Q_{1-\eta/2}| \mathbf{1}(X - \mu \geq Q_{1-\eta/2})], \right. \\ \left. v_\varepsilon^{1/(1+\varepsilon)} ((\log 2/\delta)/n)^{\varepsilon/(1+\varepsilon)} \right\}, \quad (53)$$

where for $0 < p < 1$, Q_p is a p th quantile of the distribution of $X - \mu$. Therefore, we only need to show that there is a constant c and a random variable X with mean μ , $E|X - \mu|^{1+\varepsilon} \leq v_\varepsilon$ such that the maximum of the first two terms in the right hand side of Eq. (53) is at least $c v_\varepsilon^{1/(1+\varepsilon)} \eta^{\varepsilon/(1+\varepsilon)}$.

We can simply choose $\mu = 0$ and set

$$X = \begin{cases} 0 & \text{with probability } 1 - \eta, \\ \pm v_\varepsilon^{1/(1+\varepsilon)} \eta^{-1/(1+\varepsilon)} / 2 & \text{with probability } \eta/4 \text{ each} \\ \pm v_\varepsilon^{1/(1+\varepsilon)} \eta^{-1/(1+\varepsilon)} & \text{with probability } \eta/4 \text{ each.} \end{cases}$$

Then $EX = 0$, $E|X|^{1+\varepsilon} \leq v_\varepsilon$ and the quantile $Q_{\eta/2} = -v_\varepsilon^{1/(1+\varepsilon)} \eta^{-1/(1+\varepsilon)} / 2$. Therefore,

$$E[|X - Q_{\eta/2}| \mathbf{1}(X - m \leq Q_{\eta/2})] = \frac{1}{8} v_\varepsilon^{1/(1+\varepsilon)} \eta^{\varepsilon/(1+\varepsilon)},$$

as required.

A.6. Proof of Theorem 5.1

Rewriting the error bound in Theorem 4.1 as in Eq. (11), with probability at least $1 - \frac{\delta}{Km^2}$, we have each of the following events

$$\hat{\mu}_i(t) \leq \mu(i) + H_i(\eta) + G_{\varepsilon,\eta} \left(\frac{\delta}{2Km^2} \right) \left(\frac{1}{t} \right)^{\frac{\varepsilon}{1+\varepsilon}}, \quad (54)$$

and

$$\hat{\mu}_*(t) \geq \mu(*) - H_*(\eta) - G_{\varepsilon,\eta} \left(\frac{\delta}{2Km^2} \right) \left(\frac{1}{t} \right)^{\frac{\varepsilon}{1+\varepsilon}}. \quad (55)$$

We'll first establish that if the number of pulls of sub-optimal arm i is larger than

$$n_{m_i} := \left\lceil \left(\frac{4G_{\varepsilon,\eta} \left(\frac{\delta}{2Km^2} \right)}{\Delta_i} \right)^{\frac{1+\varepsilon}{\varepsilon}} \right\rceil,$$

at phase m (m will be determined later, i.e. Eq. (56)), then arm i will be eliminated. Essentially, this means that

$$\Delta_i \geq 4G_{\varepsilon,\eta} \left(\frac{\delta}{2Km^2} \right) \left(\frac{1}{n_{m_i}} \right)^{\frac{\varepsilon}{1+\varepsilon}}.$$

Therefore, by (54) and (55), we have

$$\begin{aligned} \hat{\mu}_i(n_{m_i}) + G_{\varepsilon,\eta} \left(\frac{\delta}{2Km^2} \right) \left(\frac{1}{n_{m_i}} \right)^{\frac{\varepsilon}{1+\varepsilon}} + H_i(\eta) &\leq \mu(i) + 2G_{\varepsilon,\eta} \left(\frac{\delta}{2Km^2} \right) \left(\frac{1}{n_{m_i}} \right)^{\frac{\varepsilon}{1+\varepsilon}} + 2H_i(\eta) \\ &< \mu(i) + \Delta_i - 2G_{\varepsilon,\eta} \left(\frac{\delta}{2Km^2} \right) \left(\frac{1}{n_{m_i}} \right)^{\frac{\varepsilon}{1+\varepsilon}} + 2H_i(\eta) \\ &= \mu(*) - H_*(\eta) - H_i(\eta) - 2G_{\varepsilon,\eta} \left(\frac{\delta}{2Km^2} \right) \left(\frac{1}{n_{m_i}} \right)^{\frac{\varepsilon}{1+\varepsilon}} + 2H_i(\eta) \\ &< \hat{\mu}_*(n_{m_i}) - G_{\varepsilon,\eta} \left(\frac{\delta}{2Km^2} \right) \left(\frac{1}{n_{m_i}} \right)^{\frac{\varepsilon}{1+\varepsilon}} + H_i(\eta). \end{aligned}$$

Note that the optimal arm will not be eliminated if (54) and (55) hold, as then the required relation

$$\widehat{\mu}_i(n_{m_i}) - G_{\varepsilon,\eta}\left(\frac{\delta}{2Km^2}\right)\left(\frac{1}{n_{m_i}}\right)^{\frac{\varepsilon}{1+\varepsilon}} > \widehat{\mu}_*(n_{m_i}) + G_{\varepsilon,\eta}\left(\frac{\delta}{2Km^2}\right)\left(\frac{1}{n_{m_i}}\right)^{\frac{\varepsilon}{1+\varepsilon}}$$

leads to $\mu(i) + H_i(\eta) > \mu(*) - H_*(\eta)$ for a sub-optimal arm i , which violates the assumption (A2) in PIBAI framework Eq. (12). This implies that once the phase index m_i^0 is such that

$$m_i^0 := \min \left\{ m : \left(\gamma^m G_{\varepsilon,\eta}(\delta/2Km^2) \right)^{\frac{1+\varepsilon}{\varepsilon}} > n_{m_i} \right\}, \quad (56)$$

then arm i will be eliminated before phase m_i^0 with probability at least $1 - \sum_{m=2}^{m_i^0} \frac{\delta}{Km^2} \geq 1 - \frac{\delta}{K}$.

Solving Eq. (56), we have $\left(\gamma^m G_{\varepsilon,\eta}(\delta/2Km^2) \right)^{\frac{1+\varepsilon}{\varepsilon}} > \left[\left(\frac{4G_{\varepsilon,\eta}(\frac{\delta}{2Km^2})}{\Delta_i} \right)^{\frac{1+\varepsilon}{\varepsilon}} \right]$. It suffices to have $m > \log_\gamma(4/\Delta_i)$. We then know m_i^0 is bounded by $\lceil \log_\gamma(4/\Delta_i) \rceil$. Therefore, the number of times pulling arm i is bounded by

$$\begin{aligned} & \left(\gamma^{m_i^0} G_{\varepsilon,\eta}\left(\frac{\delta}{2Km_i^2}\right) \right)^{\frac{1+\varepsilon}{\varepsilon}} \\ &= \left(G_{\varepsilon,\eta}\left(\frac{\delta}{2K(m_i^0)^2}\right) \right)^{\frac{1+\varepsilon}{\varepsilon}} (\gamma^{m_i^0})^{\frac{1+\varepsilon}{\varepsilon}} \\ &\leq 2^{\frac{1+\varepsilon}{\varepsilon}} \left(G_{\varepsilon,\eta}\left(\frac{\delta}{2K(m_i^0)^2}\right) \right)^{\frac{1+\varepsilon}{\varepsilon}} \left(\frac{4}{\Delta_i} \right)^{\frac{1+\varepsilon}{\varepsilon}} \\ &\leq O\left(\log\left(\frac{K \log_\gamma(4/\Delta_i)}{\delta}\right) \left(\frac{4}{\Delta_i} \right)^{\frac{1+\varepsilon}{\varepsilon}} \right). \end{aligned} \quad (57)$$

The result follows by substituting the expression for $G_{\varepsilon,\eta}(\cdot)$.

A.7. Proof of Theorem 6.1

According to Theorem 3.1, we know, for any Ω satisfying $0 < \Omega < 1/(A\eta) - 4$ and $(\Omega + 4)A\eta + 4(\log 2/\delta)/n \leq 1$, the estimator $\widehat{m}(\eta, \delta)$ produced by our algorithm holds

$$|\widehat{m}(\eta, \delta) - \mu| \leq v^{1/2} \frac{(\Omega + 4)/(2\Omega^{1/2})A^{1/2}\eta^{1/2} + ((2 \log 2/\delta)/n)^{1/2}}{1 - (\Omega + 4)A\eta/2 - 2(\log 2/\delta)/n} \quad (58)$$

with probability at least $1 - \delta$.

By construction, we know

$$P(\mu \in I_k) \geq 1 - \delta_k, \quad (59)$$

for any k such that $\eta_{\text{true}} \leq \eta_k$. Thus, it holds

$$P(\mu \in I_k \text{ for any } k \text{ such that } \eta_{\text{true}} \leq \eta_k) \geq 1 - \sum_{k=0}^J \delta_k. \quad (60)$$

By definition of J_0 , it gives

$$P(m \in I_{J_0}) \geq 1 - \sum_{k=0}^J \delta_k. \quad (61)$$

By the choice that $\delta_k = \delta 2^{-k}$, (61) implies

$$|\widehat{m} - m| \leq v^{1/2} \frac{(\omega + 4)/(2\Omega^{1/2})A^{1/2}\eta_{J_0}^{1/2} + (2J_0(\log 2)/n)^{1/2} + ((2 \log 2/\delta)/n)^{1/2}}{1 - (\Omega + 4)A\eta/2 - 2J_0 \log 2/n - 2(\log 2/\delta)/n} \quad (62)$$

with probability at least $1 - 2\delta$.

If $\eta_{\text{true}} \geq \eta_{\min}$, then on the above event of probability at least $1 - 2\delta$, we have $\eta_{\text{true}} \geq \eta_{J_0}/4$, so that

$$|\hat{\mu} - \mu| \leq v^{1/2} \frac{2(\Omega + 4)/(2\Omega^{1/2})A^{1/2}\eta_{\text{true}}^{1/2} + (2J_0(\log 2)/n)^{1/2} + ((2 \log 2/\delta)/n)^{1/2}}{1 - (\Omega + 4)A\eta/2 - 2J_0 \log 2/n - 2(\log 2/\delta)/n}.$$

On the other hand, if $\eta_{\text{true}} < \eta_{\min}$, then on the same event

$$\begin{aligned} |\hat{\mu} - \mu| &\leq v^{1/2} \frac{(\Omega + 4)/(2\Omega^{1/2})A^{1/2}\eta_{\min}^{1/2} + (2J_0(\log 2)/n)^{1/2} + ((2 \log 2/\delta)/n)^{1/2}}{1 - (\Omega + 4)A\eta/2 - 2J_0 \log 2/n - 2(\log(2/\delta))/n} \\ &\leq v^{1/2} \frac{(2J_0(\log 2)/n)^{1/2} + 2((2 \log(2/\delta))/n)^{1/2}}{1 - (K + 4)A\eta/2 - 2J_0 \log 2/n - 2(\log(2/\delta))/n}. \end{aligned}$$

Therefore, in any case, ‘

$$|\hat{\mu} - \mu| \leq v^{1/2} \frac{2(\Omega + 4)/(2\Omega^{1/2})A^{1/2}\eta_{\text{true}}^{1/2} + (2J_0(\log 2)/n)^{1/2} + 2((2 \log 2/\delta)/n)^{1/2}}{1 - (\Omega + 4)A\eta/2 - 2J_0 \log 2/n - 2(\log 2/\delta)/n} \quad (63)$$

with probability at least $1 - 2\delta$.

This, of course, holds assuming that (31) holds for every $k = 0, 1, \dots, J_0$. Since $J_0 \leq J$, this will hold if

$$(\Omega + 4)A\eta + 4(\log 2/\delta)/n + \frac{4 \log 2}{n} \left(\log[(\Omega + 4)^2 A\eta/(4\Omega)] + \log n \right) / \log 4 \leq 1. \quad (64)$$

B. ϵ -Huber Contamination Model

There is a related contamination model considered in Lai et al. (2016); Prasad et al. (2020a), known as the ϵ -Huber contamination model (Huber, 2004):

$$\tilde{P} = (1 - \epsilon)P^* + \epsilon Q. \quad (65)$$

Here $\{X_i\}_{i \leq n}$ are drawn i.i.d from the mixture model \tilde{P} , with the uncontaminated distribution P^* and arbitrary contamination distribution Q chosen based on Bernoulli(ϵ) flip, possibly in an online fashion. Proposition B.1 provides a high confidence bound on the empirical fraction $\hat{\epsilon}_n = \frac{1}{n} \sum_{i=1}^n \mathbf{1}(\tilde{X}_i \neq X_i)$.

Proposition B.1. *Let $\hat{\epsilon}_n$ denote the empirical fraction of observations drawn from Q up to time n . With probability at least $1 - \beta$, for all $n \geq 1$*

$$\hat{\epsilon}_n \leq \underbrace{\epsilon + 1.7\sqrt{\epsilon(1 - \epsilon)} \sqrt{\frac{\log(\log(2n)) + 0.72 \log \frac{10.4}{\beta}}{n}}}_{:=f(\beta, \epsilon, n)}.$$

Proof of Proposition B.1 In the empirical fraction $\hat{\epsilon}_n = \frac{1}{n} \sum_{i=1}^n \mathbf{1}(\tilde{X}_i \neq X_i)$, the indicator random variable $\mathbf{1}(\tilde{X}_i \neq X_i)$ has a sub-Gaussian distribution. The result follows using Howard et al. (2021, Theorem 1). \square

The transition from Eq. (1) to Eq. (65) is as follows: fix $n_0 > 1$ and set

$$\epsilon + f(\beta, \epsilon, n_0) = \eta. \quad (66)$$

Clearly, for all $n \geq n_0$, we have the corruption fraction $\hat{\epsilon}_n$ to be at most η with a very high probability.

C. Detailed Expressions for Suppressed Notation

C.1. Influence function when $\epsilon < 1$

As explained in Catoni (2012), the narrowest possible ψ that satisfies Eq. (2) has $A = \log 2$, and is given by

$$\psi(x) = \begin{cases} -\log(1 - x + x^2/2), & 0 \leq x \leq 1, \\ \log(2), & x \geq 1, \\ -\psi(-x), & x \leq 0. \end{cases} \quad (67)$$

C.2. Influence function when $\varepsilon < 1$

Influence Function: Note that the function

$$-\log(1 - x + C_\varepsilon x^{1+\varepsilon}), \quad x \geq 0$$

achieves its maximum at the point $((1 + \varepsilon)C_\varepsilon)^{-1/\varepsilon}$ and its maximal value is

$$-\log\left(1 - \frac{\varepsilon}{1 + \varepsilon}((1 + \varepsilon)C_\varepsilon)^{-1/\varepsilon}\right).$$

Similarly, the function

$$\log(1 + x + C_\varepsilon |x|^{1+\varepsilon}), \quad x \leq 0$$

achieves its minimum at the point $-((1 + \varepsilon)C_\varepsilon)^{-1/\varepsilon}$, and its minimal value is

$$\log\left(1 - \frac{\varepsilon}{1 + \varepsilon}((1 + \varepsilon)C_\varepsilon)^{-1/\varepsilon}\right).$$

Therefore, we can choose a specific function $\psi(x)$ as a bounded function satisfying

$$|\psi(x)| \leq A_\varepsilon =: -\log\left(1 - \frac{\varepsilon}{1 + \varepsilon}((1 + \varepsilon)C_\varepsilon)^{-1/\varepsilon}\right). \quad (68)$$

Note that $A_1 = \log 2$, and $A_\varepsilon \rightarrow \infty$ as $\varepsilon \downarrow 0$. We assume, therefore, that the function ψ is bounded by some A_ε , and we know that A_ε can be chosen as in Eq. (68). A bounded influence function $\psi(x)$ that satisfies Eq. (68) is given as follows:

$$\begin{cases} \log\left(1 - \frac{\varepsilon}{1 + \varepsilon}((1 + \varepsilon)C_\varepsilon)^{-1/\varepsilon}\right) & \text{if } x \leq -((1 + \varepsilon)C_\varepsilon)^{-1/\varepsilon}, \\ \log(1 + x + C_\varepsilon |x|^{1+\varepsilon}) & \text{if } -((1 + \varepsilon)C_\varepsilon)^{-1/\varepsilon} \leq x \leq 0, \\ -\log(1 - x + C_\varepsilon x^{1+\varepsilon}) & \text{if } 0 \leq x \leq ((1 + \varepsilon)C_\varepsilon)^{-1/\varepsilon}, \\ -\log\left(1 - \frac{\varepsilon}{1 + \varepsilon}((1 + \varepsilon)C_\varepsilon)^{-1/\varepsilon}\right) & \text{if } x \geq ((1 + \varepsilon)C_\varepsilon)^{-1/\varepsilon}. \end{cases} \quad (69)$$

C.3. Compact error bounds

To simplify notation, we rewrote Eq. (5) and Eq. (10) using Eq. (11).

When $\varepsilon = 1$, we can compactly represent Eq. (5) as follows: For any arm $i \in K$,

$$\begin{aligned} |\widehat{\mu}_i(t) - \mu| &\leq H_i(\eta) + G_{\varepsilon, \eta}(\delta) \left(\frac{1}{t}\right)^{\frac{\varepsilon}{1+\varepsilon}}, \quad \text{where,} \\ H_i(\eta) &:= \frac{v^{1/2}(\Omega + 4)\sqrt{\eta A}}{2\sqrt{\Omega}(1 - (\Omega + 4)A\eta/2 - 2\log(2/\delta)/n)}, \quad \text{and} \\ G_{\varepsilon, \eta}(\delta) &:= \frac{v^{1/2}\sqrt{2\log(2/\delta)/n}}{(1 - (\Omega + 4)A\eta/2 - 2\log(2/\delta)/n)}. \end{aligned}$$

Similarly, when $\varepsilon < 1$, we can compactly represent Eq. (10) (using concavity) as follows: For any arm $i \in K$,

$$\begin{aligned} |\widehat{\mu}_i(t) - \mu| &\leq H_i(\eta) + G_{\varepsilon, \eta}(\delta) \left(\frac{1}{t}\right)^{\frac{\varepsilon}{1+\varepsilon}}, \quad \text{where,} \\ H_i(\eta) &:= (1 + \tau)v_\varepsilon^{1/1+\varepsilon}(h^{-\varepsilon}B^\varepsilon C_\varepsilon + 1/B)(2A_\varepsilon\eta)^{\frac{\varepsilon}{1+\varepsilon}}, \quad \text{and} \\ G_{\varepsilon, \eta}(\delta) &:= (1 + \tau)v_\varepsilon^{1/1+\varepsilon}(h^{-\varepsilon}B^\varepsilon C_\varepsilon + 1/B) \left(\frac{\log(2/\delta)}{n}\right)^{\frac{\varepsilon}{1+\varepsilon}}. \end{aligned}$$