
Top K Ranking for Multi-Armed Bandit with Noisy Evaluations

Evrard Garcelon
Meta AI & CREST, ENSAE

Vashist Avadhanula
Core Data Science, Meta

Alessandro Lazaric
Meta AI

Matteo Pirotta
Meta AI

Abstract

We consider a multi-armed bandit setting where, at the beginning of each round, the learner receives noisy independent, and possibly biased, *evaluations* of the true reward of each arm and it selects K arms with the objective of accumulating as much reward as possible over T rounds. Under the assumption that at each round the true reward of each arm is drawn from a fixed distribution, we derive different algorithmic approaches and theoretical guarantees depending on how the evaluations are generated. First, we show a $\tilde{O}(T^{2/3})$ regret in the general case when the observation functions are a generalized linear function of the true rewards. On the other hand, we show that an improved $\tilde{O}(\sqrt{T})$ regret can be derived when the observation functions are noisy linear functions of the true rewards. Finally, we report an empirical validation that confirms our theoretical findings, provides a thorough comparison to alternative approaches, and further supports the interest of this setting in practice.

1 INTRODUCTION

Consider an idealized content reviewing task in a large social media firm, where the objective is to identify harmful content that violates the platforms’ community standards. Given the large volume of content generated on a daily basis, it may not be possible to ask human reviewers to provide a thorough assessment of each piece of content. For this reason, the platform may automatically assign a badness score for each piece of content depending on their estimated level of severity. For example, a hate speech related post

may be assigned a higher badness score in comparison to a click bait post. The content with higher badness score may then be prioritized for human review, which eventually leads to what we can consider as a “ground-truth” evaluation of the severity of the content. The more accurate the badness score is in predicting the actual severity, the higher the chance that harmful content is passed for human review and properly identified. In practice, the badness score may be obtained by aggregating predictions returned by different automatic systems (e.g., rule-based, ML-based systems). For instance, the platform could rely on NLP-based classifiers for hostile speech detection, or CV-based classifiers for graphic images. As such, it is crucial to properly calibrate the predictions returned by each of these classifiers to ensure that the scores can be compared meaningfully and then return an aggregate and reliable badness score that correctly prioritizes the most harmful content for human review.

The problem sketched before¹ can be seen as an instance of the multi-armed bandit (MAB) framework, where each piece of content is an arm and the objective of the bandit algorithm is to select arms/content with the higher reward (e.g., severity). The algorithm can rely on the estimations returned by a set of *evaluators* (e.g., a set of classifiers) to decide which arm to pull at each step (e.g., content to pass to human review). This setting can be formalized using a number of existing frameworks, such as MAB with expert advice, contextual bandit, bandit with side observation, and contextual bandit with noisy context. We postpone a thorough discussion about these models to Sect. 3.

In this paper, we consider the case where evaluators are noisy and possibly biased functions of the true reward of each arm. In particular, we consider two alternative settings, where evaluations are generated according to: **1)** a noisy generalized linear function of the true reward; **2)** a noisy linear function of the true reward. In both cases, we first define an “oracle” strategy that have prior knowledge of the evaluation

Proceedings of the 25th International Conference on Artificial Intelligence and Statistics (AISTATS) 2022, Valencia, Spain. PMLR: Volume 151. Copyright 2022 by the author(s).

¹Similar problems are patient prioritization in hospitals (Déry et al., 2019), credit scoring (Provenzano et al., 2020), and resume review (Li et al., 2020).

function, including the noise distribution, and it is designed to maximize the rewards of the arms chosen at each round given the evaluations provided as input. We then devise the most suitable MAB strategies to approach the oracle’s performance over time. In the first case, we show that one has to rely on an ϵ -greedy strategy to avoid dependencies between the evaluations observed over time and the decisions taken by the algorithm. This eventually leads to a regret w.r.t. the oracle of order $\tilde{O}(T^{2/3})$ over T rounds. On the other hand, if the evaluation functions are linear and the variance of the additive noise is known, we show that a simple greedy strategy leveraging the specific structure of the problem is able to recover a $\tilde{O}(\sqrt{T})$ regret. We then validate these results in a number of experiments. We first consider synthetic problems where we carefully design the MAB instances to support our theoretical findings and to compare to alternative approaches. Then we move to problems based on real data, where our assumptions may not be verified, to provide a more thorough evaluation of performance and robustness of our approach. Notably, we study a problem related to content review prioritization for integrity in social media platforms.

2 PRELIMINARIES

We consider a multi-armed bandit problem where, at each round t , the learner is provided with a set of $K_t > K \geq 1$ arms (e.g., the content to be reviewed at time t) characterized by a reward $r_{i,t} \in \mathbb{R}$ for each $i = 1, \dots, K_t$ (e.g., the badness score). While the true reward is unknown to the learner, J evaluators return noisy, possibly biased, evaluations $f_{i,t,j}$ for each arm (e.g., different rule-based and/or ML-classifiers). At the beginning of each round, the learner receives the evaluations $\{f_{i,t,j}\}$, it returns a set $\mathcal{A}_t \subseteq \{1, \dots, K_t\}$ of K arms (i.e., $|\mathcal{A}_t| = K$), and it observes their associated rewards $r_{i,t}$ for $i \in \mathcal{A}_t$.² The learner’s objective is to accumulate as much reward as possible over T rounds by selecting the K arms with larger rewards.

Without any further assumption, this problem is not tractable since the rewards may change arbitrarily over time and the evaluations may not be predictive of the true rewards, thus making it impossible for any learner to achieve a satisfactory performance. Throughout the paper, we make a series of assumptions to make the problem solvable. We start from the rewards.

Assumption 1. *The rewards $r_{i,t}$ of each arm $i = 1, \dots, K_t$ at round $t = 1, \dots, T$ are drawn i.i.d. from a common distribution ν supported on $[0, C]$ with C*

²For the sake of simplicity, we consider that the learner receives the exact reward $r_{i,t}$, but all our results can be adapted to the case of noisy feedback.

a positive constant. The number of arms K_t at each round t is arbitrary and $K < K_t \leq K_{\max} < \infty$.

While this assumption simplifies the treatment of the problem, it does not affect the objective of the learner, which is to select the top- K arms *at each round*, i.e., for the specific realizations $\{r_{i,t}\}$. For instance, for $K = 1$, the objective is to return the arm $i_t^* = \arg \max_{i=1, \dots, K_t} r_{i,t}$. We do not assume that the learner has any prior knowledge of the distribution ν .

In general, the evaluators may rely on some contextual information $x_{i,t}$ associated to each arm i (e.g., texts, images, meta-data related to the piece of content) to return their evaluation $f_{i,t,j}$ and their accuracy in predicting the true reward $r_{i,t}$ may vary depending on the evaluator and the specific context $x_{i,t}$. Nonetheless, we assume that the learner has no access to the context or the actual mechanism that generates the evaluations (e.g., the evaluators may be external services) and we rather rely on the following model to describe how the evaluations are generated

$$f_{i,t,j} = f_j(r_{i,t}) + \epsilon_{i,t,j}, \quad j = 1, \dots, J, \quad (1)$$

where $f_j : \mathbb{R} \rightarrow \mathbb{R}$ is the evaluation function, and $\epsilon_{i,t,j}$ is a stochastic error. This general formulation can be seen as the inverse of a calibration function, as it describes the intrinsic bias of each evaluator j and the noise associated to the evaluations of the true reward. See Fig. 1 for a qualitative illustration of this model. We assume the noise in the evaluations satisfy a rather mild assumption.³

Assumption 2. *Each error $\epsilon_{i,t,j}$ is generated i.i.d. from a sub-Gaussian distribution with zero mean and parameter σ_j , assumed to be known to the learner.*

As far as the evaluation function is concerned, we distinguish two settings.

Assumption 3 (Generalized linear setting). *We assume that each evaluation function is a generalized linear model w.r.t. the true reward, i.e., $f_j(r) = g(\alpha_j \cdot r)$, for all $j \leq J$, where $\alpha \in \mathbb{R}$ and g is a strictly increasing function and twice-differentiable with $\|g'\|_\infty \leq L_g$ and $\|g''\|_\infty \leq M_g$, and $c_g := \inf_{x,\theta} g'(x \cdot \theta) > 0$. The function g is known to the learner, while the evaluator-specific parameters α_j are unknown.*

Assumption 4 (Linear setting). *We assume that each evaluation function is linear w.r.t. the true reward, i.e., $f_j(r) = \alpha_j r$, for all $j \leq J$. While the shape of the function is known to the learner, the evaluator-specific parameters α_j are unknown.*

³A similar assumption is used in (Yun et al., 2017), where the covariance matrix of the distribution generating the noisy features is assumed to be known to the learning.




Figure 1: As an illustrative example, consider the case where at each round t each arm is associated to a context $x_{i,t}$ drawn from a context distribution ρ and there exists a function generating the true rewards as $r_{i,t} = h^*(x_{i,t})$. The distribution ν is then defined by the distribution on rewards $r = h^*(x)$ induced by $x \sim \rho$. We also denote by $\rho|h^*(x) = r$ the conditional distribution over contexts associated with reward r . Consider then a neural network h , trained on past context-reward pairs, that returns an evaluation for arm i characterized by a context $x_{i,t}$ as $f_{i,t} = h(x_{i,t})$. The black crosses in the plot are the pairs $(h^*(x_{i,t}), h(x_{i,t}))$. The evaluation function associated to the neural network h is then defined as $f_h(r) = \mathbb{E}_{x \sim \rho|h^*(x)=r}[h(x)]$ (green line) and the noise ϵ is the deviation from $h(x)$ and $f(r)$ depending on the specific realization of x , i.e., $\epsilon_{i,t} = h(x_{i,t}) - f_h(r)$ for $r = h^*(x_{i,t})$. The blue line illustrates the perfectly calibrated case, where h^* itself is used for prediction, in this case $f_{h^*}(r) = r$.

In the following we use $\alpha = (\alpha_1, \dots, \alpha_J) \in \mathbb{R}^J$ and $\sigma = (\sigma_1, \dots, \sigma_J) \in \mathbb{R}^J$. We use the standard notation $\|\cdot\|$ and $\|\cdot\|_\infty$ for the ℓ_2 and the maximum norm respectively, while for any two vectors $x, y \in \mathbb{R}^J$, $x \cdot y \in \mathbb{R}^J$ denotes the component-wise product.

We consider a setting where the bandit algorithm has only access to the predictors’ evaluations of the true reward of an arm. This is a very generic scenario that encompasses the case where evaluations are a function of a context characterizing an arm. The generality of our framework allows us to deal with problems where the context is not directly observable (e.g., because it is kept private) or where it differs across evaluators. For example, similarly to bandits with expert advice, evaluators may use very different context sources (e.g., visual information, text, meta data) to build their predictions, but these are unknown to the bandit algorithm (e.g., because evaluators are external services). The resulting model in Eq. 1 is then a calibration function which, in the case evaluations are function of a context, can be understood as explaining the connection between the true reward and the evaluations after averaging over the stochasticity in the (non-observable) context information (Fig. 1). Notice that if the context was available, the evaluations could be disregarded as the bandit could directly rely on the context to predict the rewards in the first place, as in standard contextual bandit.

We conclude by noticing, despite these additional assumptions, no learner can retrieve the best choice of top- K arms at each round (i.e., $\max_{i_1, \dots, i_K} \sum_{l=1}^K r_{i_l, t}$), since the only information available to the learner is from biased and noisy evaluators (see Lemma 3 in App. A). As a result, instead of targeting the top- K arms, in the following we introduce *oracle* strategies that leverage the full knowledge of the problem (i.e., the evaluation function and the noise distribution) and use their performance as reference for the learner.

3 RELATED WORK

Before diving into how to solve the problem introduced in the previous section, we review alternative models that are related to our setting. Let consider the case with $K = 1$ (i.e., the learner returns one arm at each round). The most direct way to model our setting is MAB with expert advice (Auer et al., 2003), where the evaluators are *experts* and evaluations $\{f_{i,t,j}\}$ are the experts feedback. In this case, it is possible to derive algorithms with sublinear regret w.r.t. the best expert in hindsight (Beygelzimer et al., 2011). While this is a very general model, where no assumption is imposed either on the rewards or on the experts feedback (they could even be generated adversarially), algorithms designed for this setting tend to be over conservative in practice, as they have to be robust to any sort of data process. Furthermore, none of the evaluators may be very accurate (e.g., they all have very large variance) and targeting the performance of the best among them may not correspond to a satisfactory performance.

Alternatively, we can frame our problem as a contextual MAB problem (Agrawal and Goyal, 2013; Agarwal et al., 2014). We could aggregate all J evaluations for arm i into a context representation

$$\phi_{i,t} = (f_{i,t,1}, \dots, f_{i,t,j}, \dots, f_{i,t,J}) \in \mathbb{R}^J. \quad (2)$$

Unfortunately, there are two major issues using this model: **1)** in general, the reward $r_{i,t}$ may not be a simple function (e.g., linear) of $\phi_{i,t}$; **2)** the contextual features may be noisy realizations of some “true” features (e.g., due to the noise factor $\epsilon_{i,t,j}$ in Eq. 1). In order to deal with the first issue, we could rely on Asm. 4, which would lead to a linear contextual problem. The second issue could be dealt by using the approach proposed by Yun et al. (2017) for linear contextual bandit with noisy features. While the setting in (Yun et al., 2017) bears some similarities (e.g., the noise distribution is assumed to be known, they consider a similar notion of relative regret and study a greedy algorithm), there remain some crucial differences: **1)** they consider unbiased features, which corresponds to a very specific instance of Asm. 4 with $\alpha_j = 1$; **2)** they provide guarantees only for Gaussian

noise, while the algorithm designed to handle the general case has no regret guarantee.

Finally, alternative models of contextual bandit with non-deterministic features considered the case where the full distribution of the features is known (Yang et al., 2020) or part of the features are corrupted (Gajane et al., 2016; Bouneffouf, 2021). These settings do not match the use cases studied in this paper.

4 GENERALIZED LINEAR CASE

We consider the case where the evaluator functions satisfy the generalized linear model in Asm. 3.

4.1 The Oracle Strategy

We first define an *oracle* strategy that, beside σ_j and g , has prior knowledge about the parameters α_j . At each round, the oracle receives as input the evaluations $\{f_{i,t,j}\}$ and has to select K arms. We focus on oracle strategies \mathfrak{D} of the following form

1. The oracle \mathfrak{D} first aggregates the evaluations into a reward estimation $\widehat{r}_{i,t}^{\mathfrak{D}}$ for each arm i using a weighted average scheme. Let $\phi_{i,t} \in \mathbb{R}^J$ the vector collecting all evaluations as in Eq. 2 and $w \in \mathbb{R}^J$ a weight vector, then we define⁴

$$\widehat{r}_{i,t}^{\mathfrak{D}} = \langle w, g^{-1}(\phi_{i,t}) \rangle, \quad (3)$$

where g^{-1} is the inverse of the link function applied component-wise to $\phi_{i,t}$. The choice of the weights is fixed and independent from the actual evaluations, but it may depend on the evaluation function and the noise distribution.

2. The oracle \mathfrak{D} then returns the top- K arms according to the estimates $\widehat{r}_{i,t}^{\mathfrak{D}}$, i.e.,

$$\mathcal{A}_t^{\mathfrak{D}} := \arg \max_i^K \langle w, g^{-1}(\phi_{i,t}) \rangle \quad (4)$$

The crucial aspect is then to find the weighting scheme w that guarantees the best performance for the oracle. Let i_1^*, \dots, i_K^* be the true top- K arms and $\mathcal{A}_t^{\mathfrak{D}} = \{i_1^{\mathfrak{D}}, \dots, i_K^{\mathfrak{D}}\}$ be the *estimated top- K arms* according to the estimated rewards $\widehat{r}_{i,t}^{\mathfrak{D}}$. Ideally, at each round t , we would like to find the oracle weights that minimize the suboptimality gap

$$\Delta_t^{\mathfrak{D}} = \sum_{l=1}^K r_{i_l^*,t} - \sum_{l=1}^K r_{i_l^{\mathfrak{D}},t}. \quad (5)$$

⁴In the linear case for $\alpha_j = 1$ (i.e., the link function g is the identity function) and Gaussian noise, (Yun et al., 2017) showed the exact posterior over $r_{i,t}$ given the evaluations $\{f_{i,t,j}\}$ takes a weighted average form as in Eq. 3.

Since the rewards $r_{i,t}$ as well as the evaluations $\{f_{i,t,j}\}$ are random, it is not possible to minimize the previous expression for any possible realization using a fixed set of weights. Thus, we rather focus on minimizing a high-probability upper-bound of Eq. 5.

Lemma 1. *Under Asm. 2 and 3, with α_j being the parameter of the generalized linear model for each evaluator $j = 1, \dots, J$ and σ_j being the sub-Gaussian parameter of the noise $\epsilon_{i,t,j}$, let $\delta \in (0, 1)$ be a desired confidence level, then the oracle strategy designed to minimize a $(1 - \delta)$ -upper bound of Eq. 5 is characterized by the weights solving the optimization problem*

$$\begin{aligned} \min_{w \in \mathbb{R}^J} & 2 \sqrt{K^3 \sum_{j=1}^J (w_j \sigma_j)^2 \ell_\delta} + K \sqrt{J \sum_{j=1}^J (w_j \sigma_j)^2} \\ \text{s.t.} & \sum_{j=1}^J w_j \alpha_j = 1 \end{aligned} \quad (6)$$

where $\ell_\delta = \ln\left(\frac{K_{\max}}{\delta}\right)$. The previous problem has a closed-form solution $w^+ \in \mathbb{R}^J$ such that

$$w_j^+ = \frac{\alpha_j}{\sigma_j^2 \|\alpha \cdot \sigma^{-1}\|}. \quad (7)$$

The resulting oracle has suboptimality gap w.p. $1 - \delta$

$$\Delta_t^+ \leq \frac{2K \sqrt{\ln\left(\frac{K_{\max}}{\delta}\right)} + K\sqrt{J}}{\|\alpha \cdot \sigma^{-1}\|}. \quad (8)$$

We first remark that the previous lemma does not use Asm. 1 and it holds for any realization of the rewards, where the probability $(1 - \delta)$ is w.r.t. to noise in the evaluations. We notice that the optimal weight w_j^+ is proportional to the ratio α_j/σ_j^2 , which describes the amount of “signal” with respect to the noise for evaluator j . Indeed, the oracle gives less weight to evaluators that are noisy (large σ_j), while relying more on evaluators with strong “signal” (large α_j). Indeed, as α_j increases w.r.t. σ_j , the evaluations tend to be near deterministic and thus more reliable. Interestingly, the suboptimal gap in Eq. 8 shows that the oracle improves as the number of evaluators increases (the term $\|\alpha \cdot \sigma^{-1}\|$ grows as \sqrt{J}) but $\Delta_t^{\mathfrak{D}}$ does not tend to zero even when $J \rightarrow \infty$. While this might be counterintuitive (the learner is provided with an infinite number of independent evaluations), the residual gap is due to the nonlinear nature of the generalized linear model, where the zero-mean noise added to the evaluations may be amplified or decreased through g^{-1} while reconstructing the unknown parameters α_j .

Based on the previous lemma, the oracle strategy for the generalized linear case is defined by the weights w^+ and we denote by $\widehat{r}_{i,t}^+$ and \mathcal{A}_t^+ the associated reward estimates and top- K arm selection rule.

Algorithm 1: GLM- ε -GREEDY algorithm

Input: Noise parameters $\{\sigma_j\}_{j \leq J}$, confidence level δ
Parameters: exploration level ε ; number of arms to pull K ; regularization λ

 Set $\mathcal{H}_0 = \emptyset$, $\hat{\alpha} = 0$ and $w_0 = 0$
for $t = 1, \dots, T$ **do**

 Sample $Z_t \sim \text{Ber}(\varepsilon)$

 Observe evaluations for each arm $(\phi_{i,t})_{i \leq K_t}$
if $Z_t = 1$ **then**

 Pull arms in \mathcal{A}_t obtained by sampling K arms uniformly in $\{1, \dots, K_t\}$

 Observe rewards $r_{i,t}$ for all $i \in \mathcal{A}_t$

Add sample to dataset

 $\mathcal{H}_t = \mathcal{H}_{t-1} \cup (\cup_{i \in \mathcal{A}_t} \{(\phi_{i,t}, r_{i,t})\})$

 Update estimators $\hat{\alpha}_{j,t}$ by solving

$$\sum_{\phi, r \in \mathcal{H}_t} r(g(\hat{\alpha}_{t,j} \cdot r) - \phi_j) - \lambda \hat{\alpha}_{t,j} = 0 \quad (9)$$

 Update weights $w_{t,j} = \hat{\alpha}_{t,j} / \|\hat{\alpha}_t\| \cdot \sigma^{-1}$
else

 Select $\mathcal{A}_t = \arg \max_i^K \langle w_t, \phi_{i,t} \rangle$

4.2 The GLM- ε -greedy Algorithm

Building on the oracle strategy defined in the previous section, we now consider the learning problem when the link function g and the noise distribution are known, but the learner has no knowledge of the parameters α_j . As customary in MAB problems, at each round t , the learner observes only the rewards of the selected arms in \mathcal{A}_t . The main challenge in this setting is that a learner leveraging the evaluations $\{f_{i,t,j}\}$ is directly affected in its choices (i.e., the set \mathcal{A}_t) by the noise $\epsilon_{i,t,j}$ generated at the beginning of round t . This is radically different from the standard MAB setting, where the noise (in the reward) follows the arms played by the learner, which are then independent from any noise conditionally on the past (see App. B). A way to circumvent this dependency is to rely on an ε -greedy strategy, where only the samples obtained in exploratory steps are actually used to build an estimator of the unknown parameters α_j .⁵ The resulting algorithm is detailed in Alg. 1. The core of Alg. 1 is a *maximum likelihood estimation* step where the algorithm learn the parameter α . Coherently with the evaluation model in Eq. 1, the true rewards $(r_{i,t})_{i,t}$ serves as the input for the function f_j (in this case the GLM model), while the evaluations $(\phi_{i,t})_{i,t}$ work as the values we need to fit.

We now compare the performance of Alg. 1 to the or-

⁵While we study an ε -greedy type of algorithm, any type of exploration method decorrelating the estimation procedure and the exploration, like an *Explore-Then-Commit* strategy, would achieve a similar regret.

acle strategy and define the notion of *relative* regret

$$R_T = \sum_{t=1}^T \left(\sum_{i \in \mathcal{A}_t^+} \hat{r}_{i,t}^+ - \sum_{i \in \mathcal{A}_t} \hat{r}_{i,t}^+ \right), \quad (10)$$

where $w^+ \in \mathbb{R}^J$ is the oracle weight vector, $\hat{r}_{i,t}^+$ are the associated reward estimates, and \mathcal{A}_t^+ is the oracle set of top- K arms. This notion of regret, first introduced by Yun et al. (2017) in noisy contextual bandit, is comparing the quality of the arms returned by the algorithm and the oracle according to the *estimated rewards*, for which the oracle is optimal (see App. E for further discussion). The following theorem shows that if ε is properly tuned, ε -greedy has sublinear regret.

Theorem 1. *For any $\delta \in (0, 1)$, $T \geq 8$, $\lambda = J^{-1}$ and set $\varepsilon = T^{-1/3}$, let $\eta_{\nu,j}^2 = \mathbb{E}_{r \sim \nu}(g(\alpha_j r)^2)$ and $\eta_{\nu,\min} = \min_j \eta_{\nu,j}$. Then under Asm. 1, 2, and 3 we have that with probability at least $1 - \delta$ the regret of Alg. 1 is bounded as*

$$R_T \leq \tilde{O} \left(T^{2/3} \left(\frac{(1 + \sqrt{J^{-1}} \|\alpha\|) \Phi S \|\sigma\|_\infty}{\sqrt{K} \eta_{\nu,\min}} + \frac{\|g\|_\infty^3}{K^2 \eta_{\nu,\min}^2} \right) \right)$$

where $\|g\|_\infty = \max_{j,x \in [0,C]} g(\alpha_j x)$ and

$$\Phi = 2K \|\sigma\|_\infty \left(2\sqrt{J} + \sqrt{K \ln \left(\frac{eK_{\max}}{K\delta} \right)} \right) + K \|g\|_\infty$$

$$S = (\|\alpha \cdot \sigma^{-2}\|_2 + \|\sigma^{-2}\|_2) \frac{\|\sigma^{-2} \cdot \alpha\|_2}{\|\sigma^{-1} \cdot \alpha\|_2^4} + \left(\frac{\|\sigma^{-1}\|_4}{\|\sigma^{-1} \cdot \alpha\|_2} \right)^2.$$

As expected, the regret of GLM- ε -greedy increases as $\tilde{O}(T^{2/3})$. While this shows that the algorithm is able to approach the performance of the oracle, it also illustrates the difficulty of this setting, where the strict decoupling between explorative and exploitative steps used to guarantee a consistent estimation process translates into a higher regret. Interestingly, this result matches the regret of Yun et al. (2017) for contextual bandit with noisy features.

In order to investigate the main terms appearing in the previous bound, we consider the case where all evaluators share the same parameters $\alpha_j = \alpha_0$ and $\sigma_j = \sigma_0$ for some $(\alpha_0, \sigma_0) \in \mathbb{R}_+^2$ (with $\alpha_0 \geq 1$). In this case, the regret bound of Thm. 1 reduces to

$$R_T \leq \tilde{O} \left(T^{2/3} \sqrt{K} \left(1 + \sqrt{\frac{K}{J}} \right) \frac{\sigma_0^2}{\alpha_0} \right)$$

We first notice that the bound scales as $\max\{\sqrt{K}, K\sqrt{J^{-1}}\}$. Similar to the suboptimality gap for the oracle (see Lemma 1), when $J \rightarrow \infty$, the bound improves but does not decrease down to 0 and rather converges to \sqrt{K} . The dependency on K may be surprising, since for $K = K_{\max}$ the regret is

trivially 0 at each round (i.e., the algorithm returns all arms and it cannot make any error in the ranking). However, this dependency comes from the fact that in the regret analysis we bound the norm of the evaluations for the selected jobs, which scales linearly with K . We believe a more refined analysis could alleviate this dependency. Last, the regret depends inversely on the "signal-to-noise" ratio $\frac{\alpha_0}{\sigma_0^2}$.

5 LINEAR CASE

We consider the special case of linear evaluations and show how this relatively minor change to the problem has a major impact on how to approach the learning problem and the regret.

5.1 The Linear Oracle Strategy

Similar to the GLM case, we define the oracle strategy that defines an estimated reward $\widehat{r}_{i,t}^{\mathcal{D}} = \langle w, \phi_{i,t} \rangle$. Then the oracle ranks arms according to $\widehat{r}_{i,t}^{\mathcal{D}}$ and selects the set of top- K $\mathcal{A}_t^{\mathcal{D}}$ accordingly. We then optimize weights to minimize a high-probability upper bound to the suboptimality gap $\Delta_t^{\mathcal{D}}$.

Lemma 2. *Under Asm. 2 and 4, with α_j , being the parameter for each evaluator $j = 1, \dots, J$, σ_j being the sub-Gaussian parameter of the noise $\epsilon_{i,t,j}$, let $\delta \in (0, 1)$ be a desired confidence level, then the oracle strategy designed to minimize a $(1 - \delta)$ -upper bound of Eq. 5 is characterized by the weights obtained as the solution of the optimization problem*

$$\min_{w \in \mathbb{R}^J} 2 \sqrt{K^3 \sum_{j=1}^J (w_j \sigma_j)^2 \ell_\delta} \text{ s.t. } \sum_{j=1}^J w_j \alpha_j = 1 \quad , \quad (11)$$

where $\ell_\delta = \ln\left(\frac{K_{\max}}{\delta}\right)$. The previous problem has a closed-form solution $w^+ \in \mathbb{R}^J$ such that

$$w_j^+ = \frac{\alpha_j}{\sigma_j^2 \|\alpha \cdot \sigma^{-1}\|^2}. \quad (12)$$

The resulting oracle has suboptimality gap w.p. $1 - \delta$

$$\Delta_t^+ \leq \frac{2K \sqrt{\ln\left(\frac{K_{\max} e}{\delta}\right)}}{\|\alpha \cdot \sigma^{-1}\|} \quad (13)$$

The weights of the oracle have the same expression in the GLM case in Lemma 1. Nonetheless, the suboptimality gap is smaller than in Lemma 1 as the linear structure allows to concentrate the noise of the evaluations, unlike in the GLM setting where the potential non linearity of g forces us to study the worst-case scenario. Notably, in the linear case, we see that as J tends to infinity the suboptimality gap tends to zero.

5.2 Evaluation-Structure-Aware Greedy

Similarly to Sec. 4.2, the learner has no knowledge of the parameters α_j but only knows the noise distribution and the linear structure of the evaluations. The main estimation difficulty of the general case still applies to this setting, i.e., using samples obtained by selecting arms based on the evaluations may introduce a bias in the estimation process.

Instead of introducing explicit exploration steps to gather "unbiased" samples, as done in the GLM case, we exploit a more subtle property for this case. We notice for any evaluator j , the expected evaluation is

$$\mathbb{E}[f_{i,t,j}] = \mathbb{E}[\alpha_j r_{i,t} + \epsilon_{i,t,j}] = \alpha_j \bar{r}, \quad (14)$$

where $\bar{r} = \mathbb{E}[r_{i,t}]$ is the expectation of the reward distribution ν in Asm. 1. Consider an oracle strategy that is fed with the parameters $\alpha_j \bar{r}$, then the optimal weights in Eq. 16 become $\widetilde{w}^+ = w^+ / \bar{r}$. While this leads to estimates $\widehat{r}_{i,t}^+$ that are biased w.r.t. the oracle estimates $\widehat{r}_{i,t}^{\mathcal{D}}$, the factor $1/\bar{r}$ is evaluator- and arm-independent and it does not impact the ranking returned by this *biased* oracle, i.e., $\mathcal{A}_t = \mathcal{A}_t^+$. This in striking contrast with the GLM where it is not possible to easily evaluate a vector proportional to α because of the potential non-linear behavior of g .

Building on this evidence, we define an algorithm, the *Evaluation-Structure-Aware Greedy* (ESAG) in Alg. 2, that avoids the use of the observed rewards altogether, thus removing the statistical dependency between noise and decisions, and rather tries to estimate the expected evaluations in Eq. 14 (see Eq. 15 in Alg. 2) and use them to build estimates, as in the biased oracle. Since ESAG only relies on the evaluations available at round t , it does not need any *explicit* exploration strategy to collect useful information, and it executes greedy actions according to the current weights w_t at each round. We can derive the following regret guarantees.⁶

Theorem 2. *Under Asm. 1, 2, and 4 for any $\delta \in (0, 1)$, $T \geq 1$ with probability at least $1 - \delta$ the regret of Alg. 2 is bounded by:*

$$R_T \leq \mathbb{E}_{r \sim \nu}(r) \Phi' K \left[\left(\frac{\Phi'}{\mathbb{E}_{r \sim \nu}(r) \|\alpha\|_\infty \sqrt{K}} \right)^2 + \frac{8\sqrt{T} \Phi' S}{\sqrt{\bar{K}_T}} \right]$$

where $\bar{K}_T = \frac{T}{\sum_{t=1}^T \frac{t-1}{\sum_{k=1}^{t-1} K_k}} \geq K$ is the harmonic average of the number of arms over T steps, S is the same as in Thm. 1 and

$$\Phi' = 2\|\alpha\| C \ln\left(\frac{4}{\delta}\right) + 2\|\sigma\|_\infty \left(2\sqrt{J} + \sqrt{K \ln\left(\frac{K_{\max}}{K\delta}\right)} \right)$$

⁶While we derive Thm. 2 for a greedy algorithm, similar results hold for an optimistic exploration strategy.

Algorithm 2: The *Evaluation-Structure-Aware Greedy* (ESAG) algorithm for the linear case.

Input: Noise parameters $\{\sigma_j\}_{j \leq J}$, confidence level δ

Parameters: number of arms to pull K

Set $\hat{\alpha} = 0$, $w_0 = 0$ and $N_t = 0$

for $t = 1, \dots, T$ **do**

Observe evaluations for each arm $(\phi_{i,t})_{i \leq K_t}$

Select $\mathcal{A}_t = \arg \max_i^K (w_t, \phi_{i,t})$

Observe rewards $r_{i,t}$ for all $i \in \mathcal{A}_t$

Update $\hat{\alpha}_t$ as:

$$\hat{\alpha}_{t+1} = \hat{\alpha}_t \frac{N_t + K_t - N_t K_t}{(N_t + K_t)N_t} + \frac{\sum_{i=1}^{K_t} \phi_{i,t}}{K_t} \quad (15)$$

Update $N_{t+1} = N_t + K_t$ and w_t as

$$w_{t+1,j} = \frac{\hat{\alpha}_{t+1,j}}{\sigma_j^2 \|\hat{\alpha}_{t+1,j} \cdot \sigma^{-1}\|^2} \quad (16)$$

The most interesting aspect of the previous theorem is that the regret is of order $\tilde{O}(\sqrt{T})$, since ESAG does not pay for the sharp separation between exploration and exploitation steps as for GLM- ε -GREEDY.

6 EMPIRICAL VALIDATION

In order to study different aspects of the settings and algorithms introduced in the previous sections, we focus on both synthetic and real data experiments. We report further results in the supplementary material.

6.1 Synthetic Data

We first validate our algorithms on synthetic data. We consider $K_t = K_{\max} = 20$ arms where for each arm we get $J = 10$ evaluations. The reward distribution ν is a Gaussian distribution centered at 0 and truncated between $[0, 20]$. We consider both the logistic case with $g(x) = (1 + \exp(-x))^{-1}$ and the linear case where for all evaluators $f_j(x) = \alpha_j x$. At the beginning of each experiment, we draw the coefficients α_j and the parameters σ_j from uniform distributions in $[\alpha_0/2; 3\alpha_0/2]$ and $[\sigma_0/2; 3\sigma_0/2]$ respectively. As discussed in Sect. 4, a critical term characterizing the problem structure is the ratio α_j/σ_j . We then set $\alpha_0 = 1$ and consider three different values for σ_0 such that $\alpha_0/\sigma_0 \in \{0.1, 1, 10\}$. Finally, the noise in the evaluations are generated as $\epsilon_{i,t,j} \sim \mathcal{N}(0, \sigma_j^2)$ in the linear case whereas in the GLM case $\epsilon_{i,t,j}$ is drawn from a truncated centered gaussian distribution with variance $2\sigma_j^2$. We average all results over 80 runs and we report 95% confidence intervals. Additional details are reported in the supplement.

Oracle performance. Before investigating the per-

formance of the learning algorithms, we compare the performance of the oracle strategy to a simple average strategy that defines $\hat{r}_{i,t}^{\text{avg}} = 1/J \sum_{j=1}^J f_{i,t,j}$ and ranks arms accordingly. We also study how the suboptimality gap of the oracle changes as J increases for both the GLM and linear case. As shown in Fig. 2a, in both settings that oracle strategies outperform the simple average. Furthermore, as predicted by Lemma 1 and 2, the oracle suboptimality gap decreases as J increases, but it plateaus to a fixed value for GLM, while it tends to zero for the linear case.

6.1.1 The GLM Case

We consider different types of algorithms: GLM- ε -GREEDY (Alg. 1); GLM- ε -GREEDY-ALL, it has same structure as Alg. 1 but uses samples from both explorative and exploitative steps to build the estimator $\hat{\alpha}$; GLM-EVALBASEDUCB, it uses all samples to build an estimator $\hat{\alpha}$ and leverages a high-probability confidence interval on $\hat{\alpha}$ to derive optimistic weights w and rank and select arms accordingly; GLM-LINUCB, the algorithm of Abbasi-Yadkori et al. (2011) using $g^{-1}(\phi_{i,t})$ as features; GLM-ESAG, Alg. 2 adapted for the GLM case; RAND, the fully random strategy; GLM-GREEDY, the greedy strategy using the MLE estimator $\hat{\alpha}$; EXP4.P, the bandit with expert advice algorithm in Beygelzimer et al. (2011).

Estimation Bias. As discussed in Sect. 4.2, one of the critical aspects that motivated the use of an ε -greedy approach and led to the $\tilde{O}(T^{2/3})$ is the fact that whenever the set of arms is chosen according to the noisy evaluations $\{f_{i,t,j}\}$, the dependency between $\epsilon_{i,t,j}$ and \mathcal{A}_t may create a bias when estimating the parameters α_j using the samples $r_{i,t}$ observed after selecting \mathcal{A}_t . We illustrate this effect in Fig. 2b, where we report the error of the estimates $\hat{\alpha}_t$ computed by GLM- ε -GREEDY and GLM- ε -GREEDY-ALL w.r.t. the true parameters α . While the error of the estimator computed by GLM- ε -GREEDY decreases over time, the error for GLM- ε -GREEDY-ALL has a residual bias due to the estimation procedure. Similar results can be shown for all the algorithms (e.g., GLM-ESAG) that rely on samples generated by selecting \mathcal{A}_t based on the evaluations $\{f_{i,t,j}\}$ and they can be reproduced in the linear setting as well.

Regret Performance. We compare the performance of different learning algorithms in terms of relative regret w.r.t. the oracle defined in Sec. 4. We remove from Fig. 2c all algorithms (i.e., RAND, GREEDY, EXP4.P) that suffer large linear regret to avoid losing resolution on the other algorithms. While GLM-ESAG and GLM-LINUCB have better regret, they both have a linear regret that keeps increasing over

Table 1: Cumulative badness over $T = 2000$ steps

Alg.	Badness
RAND	30076.18
GLM- ϵ -GREEDY	53403.6
GLM- ϵ -GREEDY-ALL	53419
GLM-EVALBASEDUCB	53328.3
GLM-LINUCB	30910.5
GLM-ESAG	53393.9
GLM-GREEDY	53371.9
EXP4.P	47766.7
EVALBASEDUCB	90768.2
LINUCB	90332.2
ESAG	90790.5
GREEDY	69956.3

time, while GLM- ϵ -GREEDY has a sublinear regret.⁷

6.1.2 The Linear Case

We also study the linear case and compare the linear versions of the algorithms described above. As illustrated in Fig. 2d, the crucial difference w.r.t. the GLM case is that now ESAG has sublinear regret, whereas other algorithms have linear regret. Interestingly, EVALBASEDUCB, despite the bias introduced by using samples obtained by selecting actions based on the noisy observation, is able to learn a good strategy and it outperforms LINUCB, but still suffers linear regret.

6.2 Content Review Prioritization

We now move to a real-world problem to investigate the performance of our algorithms when their assumptions are no longer verified.

Data description. We consider a small dataset of content shared on a large social media firm that has been reported for violating the platforms’ community standards. In order to ensure that the most harmful content is prioritized for reviewing, the platform assigns badness score for each piece of content which increases with the severity of the content. We consider four different classifiers that provide badness estimates for the sampled content. Each of these classifiers are trained on the real data and follow different modeling architectures to predict the badness score. Our objective is to leverage scores from these four classifiers to identify the most harmful subset of content and flag them for the platform to review and action them.

Performance. We evaluate the cumulative badness of the content selected by the algorithms. The higher

⁷Notice that we have not actively tried to find problem parameters that would make the linear regret of GLM-ESAG and GLM-LINUCB larger than GLM- ϵ -GREEDY in a shorter time. The linear regret of GLM-ESAG is due to the residual estimation bias illustrated in Fig. 2b.

the score the better. The results are reported in Tab. 2. We first notice that all learning algorithms perform significantly better than the random strategy, thus indicating that the GLM and linear assumptions are accurate enough to return meaningful rankings. Nonetheless, we notice that GLM-based algorithms do not perform as well as linear ones, probably due to the choice of the logistic function, which, in this case, does not fit data accurately. On the other hand, ESAG is the algorithm that performs best, followed by EVALBASEDUCB and LINUCB. Notice that in the case of real data, we may even expect LINUCB to perform better, since it relies on somewhat less restrictive assumptions (no assumption is made on how the features are generated) and it relies on the true rewards to estimate parameters (this is also the case for EVALBASEDUCB). This is in contrast with ESAG that exclusively builds on the evaluations, which clearly do not respect an exact linear model, to estimate the unknown parameters. This shows that, even in problems where the assumptions do not hold, ESAG is robust enough and it is competitive w.r.t. a large variety of algorithms.

7 DISCUSSION

We studied a MAB problem where the learner is provided with noisy and biased evaluations of the true reward for each arm. We showed that under specific assumptions it is possible to design learning algorithms that are able to compete to oracle strategies both in theory and in practice. The empirical validation on real data also shows that this model and the associated algorithms are promising for solving challenging real-world problems.

Extensions. There is a number of directions that could be pursued to extend our current results.

- *Evaluation functions.* The GLM and linear assumptions are relatively strong. A natural venue of improvement is to generalize our results to richer function spaces, such as Gaussian processes.
- *Persistent arms.* While we assume that the set of arms is “refreshed” at each round, we can easily extend our setting to the case where all the arms that are not selected in \mathcal{A}_t remain in the pool of arms available at the next round (e.g., content that has not been reviewed). All our results naturally extend to this case, except for ESAG, which would not be able to use K_t samples at each round, but would rather get K new samples corresponding to the arms replaced at each round.
- *Heteroschedastic noise.* In the model illustrated in Fig. 1, the noise $\epsilon_{i,t,j}$ is heteroschedastic, where the variance may depend on the reward value r . While our model leverages an upper-bound σ_j on




Figure 2: **a)** Average suboptimality gap Δ_t^Δ for the oracle strategy in the GLM and linear cases as a function of the number of evaluators J . **b)** Estimation error $\|\hat{\alpha}_t - \alpha\|$. **c)** Regret w.r.t. to the oracle as defined in Sect. 4 for the GLM case. **d)** Regret w.r.t. to the oracle as defined in Sect. 4 for the linear case in the high noise regime (i.e., $\alpha_j = 0.1\sigma_j$).

the actual variance, better adapting to the reward-dependent variance may improve the performance.

References

- Abbasi-Yadkori, Y., Pál, D., and Szepesvári, C. (2011). Improved algorithms for linear stochastic bandits. In *NIPS*, pages 2312–2320.
- Agarwal, A., Hsu, D. J., Kale, S., Langford, J., Li, L., and Schapire, R. E. (2014). Taming the monster: A fast and simple algorithm for contextual bandits. In *ICML*, volume 32 of *JMLR Workshop and Conference Proceedings*, pages 1638–1646. JMLR.org.
- Agrawal, S. and Goyal, N. (2013). Thompson sampling for contextual bandits with linear payoffs. In *ICML (3)*, volume 28 of *JMLR Workshop and Conference Proceedings*, pages 127–135. JMLR.org.
- Auer, P., Cesa-Bianchi, N., Freund, Y., and Schapire, R. E. (2003). The nonstochastic multiarmed bandit problem. *SIAM J. Comput.*, 32(1):48–77.
- Beygelzimer, A., Langford, J., Li, L., Reyzin, L., and Schapire, R. E. (2011). Contextual bandit algorithms with supervised learning guarantees. In *AISTATS*, volume 15 of *JMLR Proceedings*, pages 19–26. JMLR.org.
- Bouneffouf, D. (2021). Corrupted contextual bandits: Online learning with corrupted context. In *ICASSP*, pages 3145–3149. IEEE.
- Déry, J., Ruiz, A. B., Routhier, F., Gagnon, M.-P., Côté, A., Ait-kadi, D., Bélanger, V., Deslauriers, S., and Lamontagne, M.-È. (2019). Patient prioritization tools and their effectiveness in non-emergency healthcare services: a systematic review protocol. *Systematic Reviews*, 8.
- Filippi, S., Cappe, O., Garivier, A., and Szepesvári, C. (2010). Parametric bandits: The generalized linear case. In Lafferty, J., Williams, C., Shawe-Taylor, J., Zemel, R., and Culotta, A., editors, *Advances in Neural Information Processing Systems*, volume 23. Curran Associates, Inc.
- Gajane, P., Urvoy, T., and Kaufmann, E. (2016). Corrupt Bandits. In *EWRL 13*.
- Goldberg, K. Y., Roeder, T., Gupta, D., and Perkins, C. (2001). Eigentaste: A constant time collaborative filtering algorithm. *Inf. Retr.*, 4(2):133–151.
- Li, C., Fisher, E., Thomas, R., Pittard, S., Hertzberg, V., and Choi, J. D. (2020). Competence-level prediction and resume & job description matching using context-aware transformer models. In *EMNLP (1)*, pages 8456–8466. Association for Computational Linguistics.
- Li, L., Lu, Y., and Zhou, D. (2017). Provably optimal algorithms for generalized linear contextual bandits. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, pages 2071–2080. PMLR.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Provenzano, A. R., Trifirò, D., Datteo, A., Giada, L., Jean, N., Riciputi, A., Pera, G. L., Spadaccino, M., Massaron, L., and Nordio, C. (2020). Machine learning approach for credit scoring.
- Riquelme, C., Tucker, G., and Snoek, J. (2018). Deep bayesian bandits showdown: An empirical comparison of bayesian deep networks for thompson sampling. In *ICLR (Poster)*. OpenReview.net.
- Yang, L., Yang, J., and Ren, S. (2020). Multi-feedback bandit learning with probabilistic contexts. In *IJCAI*, pages 3087–3093. ijcai.org.
- Yun, S., Nam, J. H., Mo, S., and Shin, J. (2017). Contextual multi-armed bandits under feature uncertainty. *CoRR*, abs/1703.01347.

Supplementary Material: Top K Ranking for Multi-Armed Bandit with Noisy Evaluations

A LOWER BOUND

In this appendix, we prove why no algorithm can compute the top K arms at every time step. More, precisely we prove the following Lemma.

Lemma 3. *Let consider the linear case in Asm. 4 and parameters $K_t = K_{\max}$, $K = 1$, $\alpha_j = 1$ for all $j = 1, \dots, J$ and a noise distributed as $\mathcal{N}(0, \sigma_j)$ with $\sigma_j = \sigma_0$ for all $j = 1, \dots, J$. The learner \mathfrak{A} receives as input the evaluations $f_{i,t,j} = r_{i,t} + \epsilon_{i,t,j}$ and we denote by $\mathfrak{A}(\{f_{i,t,j}\})$ the arm returned by \mathfrak{A} . For all arms $i \leq K_t$ the reward $r_{i,t}$ is sampled from a Bernoulli distribution $\text{Ber}(1/2)$. At every step, let's define $I_t^* = \{i \leq K_t \mid r_{i,t} = 1\}$ the set of optimal arms. Then any learner \mathfrak{A} has a fixed non-zero probability to return the wrong ranking at each step, i.e.,*

$$\forall t, \forall \mathfrak{A}, \exists \delta > 0, \mathbb{P}_{\text{Ber}, \{\epsilon_{i,t,j}\}} \left[\mathfrak{A}(\{f_{i,t,j}\}) \notin I_t^* \right] \geq \delta. \quad (17)$$

Proof. We reason by contradiction and assume that there exists a deterministic algorithm \mathfrak{A} and a time step t such that

$$\forall \delta > 0, \quad \mathbb{P}_{\text{Ber}, \{\epsilon_{i,t,j}\}} \left[\mathfrak{A}(\{f_{i,t,j}\}) \notin I_t^* \right] \leq \delta. \quad (18)$$

Now given the rewards $(r_{i,t})_{i \leq K_t}$ the distribution of the evaluations $(f_{i,t,j})_j$ is distributed as $\mathcal{N}((r_{i,t})_{j \leq J}, \sigma_0 I_J)$ for each i . Now, let's consider the set of sub optimal arms, $I_t^- := \{i \leq K_t \mid r_{i,t} = 0\}$ because the rewards are independent then with probability at least $\frac{1}{4}$ the set I_t^+ and I_t^- are both non-empty. For an index $i_0 \leq K_t$, let $X^{0,1} \in \mathbb{R}^{K_t \times J}$ and $X^{1,0} \in \mathbb{R}^{K_t \times J}$ be two random variables distributed sampled from the same distributions as the $(f_{i,t,j})_{i,j} \mid I_t^* = \{i_0\}, I_t^- = \{i_{-0}\}$ and $(f_{i,t,j})_{i,j} \mid I_t^- = \{i_0\}, I_t^* = \{i_{-0}\}$ respectively. We then have that:

$$\mathbb{P}(X^{0,1} = X^{1,0}) > 0 \quad (19)$$

This is a direct consequence of the independence of the noise and rewards. Indeed,

$$\mathbb{P}(X^{0,1} = X^{1,0}) = \mathbb{P}(\forall (i, j) \in \{1, \dots, K_t\} \times \{1, \dots, J\} X_{i,j}^{0,1} = X_{i,j}^{1,0}) \quad (20)$$

$$= \prod_{i,j} \mathbb{P}(X_{i,j}^{0,1} = X_{i,j}^{1,0}) \quad (21)$$

But for a given couple $(i, j) \in \{1, \dots, K_t\} \times \{1, \dots, J\}$, $X_{i,j}^{0,1} \sim \mathcal{N}(\mathbb{1}_{\{i=i_0\}}, \sigma_0^2)$, $X_{i,j}^{1,0} \sim \mathcal{N}(\mathbb{1}_{\{i \neq i_0\}}, \sigma_0^2)$ and are independent. Therefore, the probability of those random variables being equal is:

$$\mathbb{P}(X_{i,j}^{0,1} = X_{i,j}^{1,0}) = \frac{1}{2\pi\sigma_0^2} \int_{-\infty}^{+\infty} \exp\left(-\frac{\left[(x - \mathbb{1}_{\{i=i_0\}})^2 + (x - \mathbb{1}_{\{i \neq i_0\}})^2\right]}{2\sigma_0^2}\right) dx > 0 \quad (22)$$

Let's note $\gamma > 0$ such that $\gamma \leq \mathbb{P}(X_{i,j}^{0,1} = X_{i,j}^{1,0})$ for all i then we have that: $\mathbb{P}(X^{0,1} = X^{1,0}) \geq \gamma^{K_t J} > 0$. This implies that there exists a δ_0 such that $\mathbb{P}_{\text{Ber}, \{\epsilon_{i,t,j}\}} \left[\mathfrak{A}(X^{0,1}) \in I_t^*, X^{0,1} = X^{1,0}, I_t^* = \{i_0\}, I_t^- = i_{-0} \right] \geq \delta_0 > 0$, thanks to Eq. (18) but then on this event we have that:

$$\delta_0 \leq \mathbb{P}_{\text{Ber}, \{\epsilon_{i,t,j}\}} \left[\mathfrak{A}(X^{0,1}) \in I_t^*, X^{0,1} = X^{1,0}, I_t^* = \{i_0\}, I_t^- = i_{-0} \right] = \mathbb{P}_{\text{Ber}, \{\epsilon_{i,t,j}\}} \left[\mathfrak{A}(X^{1,0}) = i_0, I_t^* = \{i_0\}, I_t^- = i_{-0} \right] \quad (23)$$

$$\leq \mathbb{P}_{\text{Ber}, \{\epsilon_{i,t,j}\}} \left[\mathfrak{A}(X^{1,0}) = i_0 \right] \quad (24)$$

$$\leq \mathbb{P}_{\text{Ber}, \{\epsilon_{i,t,j}\}} \left[\mathfrak{A}(X^{1,0}) \notin I_t^* \right] \leq \frac{\delta_0}{2} \quad (25)$$

thanks to Eq. 18. This is not possible therefore by contradiction we have the result. \square

Lem. 3 shows that there exists a instance of the problem studied in this paper where it is not possible to retrieve systematically the K actions maximiziing the reward $(r_{i,t})_{i \leq K_t}$. Therefore defining the regret with respect to the true top K actions.

B NOISE CORRELATION ISSUE

The standard result on GLM are based theorems with the same structure as the following proposition (Filippi et al., 2010):

Proposition 1. *Let $(\mathcal{F}_k)_{k \geq 0}$ be a filtration, $(m_k)_{k \geq 0}$ be an \mathbb{R}^d -valued stochastic process adapted to $(\mathcal{F}_k)_k$, $(\eta_k)_k$ be a real-valued martingale difference process adapted to $(\mathcal{F}_k)_k$. Assume that η_k is conditionnally sub-Gaussian in the sense that there exists some $R > 0$ such that for any $\gamma \geq 0$, $k \geq 1$,*

$$\mathbb{E}(\exp(\gamma \eta_k) \mid \mathcal{F}_{k-1}) \leq \exp\left(\frac{\gamma^2 R^2}{2}\right) \text{ a.s} \quad (26)$$

Consider the martingale $\xi_t = \sum_{k=1}^t m_{k-1} \eta_k$ and the process $M_t = \sum_{k=1}^t m_{k-1} m_{k-1}^\top$. Assume that with probability one the smallest of M_d is lower bounded by some positive constant λ_0 and that $\|m_k\|_2 \leq c_m$ holds a.s. for any $k \geq 0$. The following hold true: *Let*

$$\kappa = \sqrt{3 + 2 \log(1 + 2c_m^2/\lambda_0)}$$

For any $x \in \mathbb{R}^d$, $0 < \delta \leq e^{-1}$, $t \geq \max(d, 2)$, with probability at least $1 - \delta$,

$$|\langle x, \xi_t \rangle| \leq \kappa R \sqrt{2 \log(t) \log(1/\delta)} \|x\|_{M_t} \quad (27)$$

Further, for any $0 < \delta < \min\{1, d/e\}$, $t \geq \max\{d, 2\}$, with probability at least $1 - \delta$,

$$\|\xi_t\|_{M_t^{-1}} \leq \kappa R \sqrt{2d \log(t) \log(d/\delta)} \quad (28)$$

In our setting, we can not use the type of results presented above because estimators like ridge regression or maximum likelihood estimation rely on the assumption that the noise associated to the data is zero-mean. However, in our setting the learner⁸ takes a decision \mathcal{A}_t based on the noisy evaluations $(\phi_{i,t})_i$ and this introduces a statistical dependence that biases the distribution of the noise affecting the evaluations. For simplicity, assume $K = 1$ then the action a_t is a function of the noise in the observations $(\phi_{i,t})$ and, in general, we have that $\mathbb{E}(\epsilon_{a_t,t,j}) \neq 0$ for all evaluators $j \leq J$ (where $(\epsilon_{i,t,j})$ is the noise in the evaluation $f_{i,t,j}$) although for any fixed non random action $a \leq K_t$, $\mathbb{E}(\epsilon_{a,t,j}) = 0$.

More formally, let consider the GLM setting of Sec. 4 assuming $K = 1$ and a MLE estimator using all the data $\mathcal{H}_t := \{(f_{a_t,t,j})_j, r_{a_t,t}\}$ to estimate the parameter α . Let also assume that the action a_t is chosen as $a_t = \arg \max_i \langle w_t, \phi_{i,t} \rangle$ where for any t , w_t is a vector adapted to the σ -algebra $\sigma(\mathcal{H}_{t-1})$. For each evaluator j we compute the MLE, $\hat{\alpha}_j$ as the solution to

$$\sum_{l=1}^{t-1} (f_{a_t,t,j} - g(\hat{\alpha}_j r_{i_t,t})) r_{i_t,t} = 0 \quad (29)$$

with a_t is the arm selected by the algorithm and $r_{a_t,t}$ is the associated reward observed at the end of the round. In order to evaluate the accuracy of this estimator $\hat{\alpha}_j$, following the proof of Filippi et al. (2010) or (Li et al., 2017), we eventually need to control the term

$$\frac{\left| \sum_{l=1}^t r_{a_t,t} \epsilon_{a_t,t,j} \right|}{\sqrt{\sum_{l=1}^t r_{a_t,t}^2}}, \quad (30)$$

⁸Notice that the following discussion holds also for any non-learning algorithm, e.g., an oracle algorithm.

where $\epsilon_{a_l, l, j}$ is the noise associated with the evaluation j for arm a_l at round l . One may be tempted to apply Prop. 1 with $m_{t-1} = r_{a_t, t}$ and $\eta_t = \epsilon_{a_t, t, j}$. The issue is that $a_t = \arg \max_i \langle w_{t-1}, g^{-1}(\phi_{i, t}) \rangle$ where w_t is a measurable function for of the past (in addition, note that this action selection process is used in most optimistic algorithms for GLM). As a result,

$$\mathbb{E}(\epsilon_{a_t, t, j} \mid \mathcal{H}_{t-1}, (r_{i, t})_{i \leq K_t}) \neq 0, \quad (31)$$

Equation 31 can be further simplified when g is the identity function, $J = 1$, $\alpha_j = 0$, $\sigma_j = \sigma$ and $w_t \neq 0$ then Equation 31 reduces to

$$\mathbb{E}(\epsilon_{a_t, t, j} \mid \mathcal{H}_{t-1}, (r_{i, t})_{i \leq K_t}) = \mathbb{E}(\epsilon_{i, t} \mathbb{1}_{\{i = \arg \max w_t \epsilon_{i, t}\}} \mid \mathcal{H}_{t-1}) \approx \sigma, \quad (32)$$

thus showing that $\epsilon_{a_t, t}$ is no longer a zero-mean variable. A way to address this issue would be to take a union bound over all the arms chosen over time $a_1, \dots, a_t, \dots, a_T$. Unfortunately, this would lead to take a union bound over K^T elements, which would eventually lead to an additional \sqrt{T} factor in Prop 1 and ultimately a linear regret bound. We further confirm this effect in several empirical validations, where algorithms relying on data generated based on the evaluations introduce a bias in the estimation of the parameters α .

C GENERALIZED LINEAR MODEL

In this section, we present how to defined the oracle and the analysis of the regret in the case of a GLM model.

C.1 Oracle in the GLM Model (Proof of Lemma 1)

We consider an oracle \mathfrak{D} with access to the link function g , the parameters $(\alpha_j)_{j \leq J}$ of the evaluation function, and the parameters $(\sigma_j)_{j \leq J}$ of the noise distribution. Given the vector $\phi_{i, t} \in \mathbb{R}^J$ obtained by aggregating all the evaluations $(f_{i, t, \cdot})_j$ for each arm $i \leq K_t$, we recall that an oracle defines a set of weights $(w_j)_{j \leq J}$ and predicts the reward of arm i as

$$r_{i, t}^{\mathfrak{D}} = \langle w, g^{-1}(\phi_{i, t}) \rangle \quad (33)$$

We aim at minimizing the gap between reward of the true top- K arms i_1^*, \dots, i_K^* and the reward of the *estimated top- K arms* $i_1^{\mathfrak{D}}, \dots, i_K^{\mathfrak{D}}$ according to the estimated rewards $r_{i, t}^{\mathfrak{D}}$

$$(a) = \sum_{l=1}^K r_{i_l^*, t} - \sum_{l=1}^K r_{i_l^{\mathfrak{D}}, t}. \quad (34)$$

By leveraging the definition of estimated top- K arms, we can rewrite the previous expression as

$$\begin{aligned} (a) &= \sum_{l=1}^K r_{i_l^*, t} - \sum_{l=1}^K \widehat{r}_{i_l^*, t}^{\mathfrak{D}} + \sum_{l=1}^K \widehat{r}_{i_l^*, t}^{\mathfrak{D}} - \sum_{l=1}^K \widehat{r}_{i_l^{\mathfrak{D}}, t}^{\mathfrak{D}} + \sum_{l=1}^K \widehat{r}_{i_l^{\mathfrak{D}}, t}^{\mathfrak{D}} - \sum_{l=1}^K r_{i_l^{\mathfrak{D}}, t} \\ &\leq \sum_{l=1}^K r_{i_l^*, t} - \sum_{l=1}^K \widehat{r}_{i_l^*, t}^{\mathfrak{D}} + \sum_{l=1}^K \widehat{r}_{i_l^{\mathfrak{D}}, t}^{\mathfrak{D}} - \sum_{l=1}^K r_{i_l^{\mathfrak{D}}, t} \\ &= \sum_{l=1}^K (r_{i_l^*, t} - \widehat{r}_{i_l^*, t}^{\mathfrak{D}}) + \sum_{l=1}^K (\widehat{r}_{i_l^{\mathfrak{D}}, t}^{\mathfrak{D}} - r_{i_l^{\mathfrak{D}}, t}) \\ &\leq 2 \max_{i_1, \dots, i_K} \sum_{l=1}^K \underbrace{\left| \widehat{r}_{i_l, t}^{\mathfrak{D}} - r_{i_l, t} \right|}_{(b)}. \end{aligned}$$

Using the GLM model of Asm. 3, we have that for any $i = 1, \dots, K_t$, the estimated reward can be written as

$$\widehat{r}_{i, t}^{\mathfrak{D}} = \sum_{j=1}^J w_j g^{-1}(g(\alpha_j \cdot r_{i, t}) + \epsilon_{i, t, j}). \quad (35)$$

We can then measure the deviation between true reward and estimated reward as

$$(b) = \sum_{j=1}^J w_j \left[g^{-1}(g(\alpha_j \cdot r_{i,t}) + \epsilon_{i,t,j}) - g^{-1}(g(\alpha_j r_{i,t})) \right] + \left(\sum_{j=1}^J w_j \alpha_j - 1 \right) r_{i,t} \quad (36)$$

$$\leq c_g^{-1} \sum_{j=1}^J |w_j| \cdot |\epsilon_{i,t,j}| + \left(\sum_{j=1}^J w_j \alpha_j - 1 \right) r_{i,t}, \quad (37)$$

where in the first equality we introduce the term $\sum_{j=1}^K w_j \alpha_j = \sum_{j=1}^K w_j g^{-1}(g(\alpha_j r_{i,t}))$ and in the second step we leverage the fact that g^{-1} is c_g^{-1} -Lipschitz. The last expression above contains two random realizations, i.e., $\epsilon_{i,t,j}$ and $r_{i,t}$, thus preventing from computing a fixed set of weights to minimize the performance gap (a). Furthermore, while we assume the oracle to have prior knowledge about the parameters, we would like to avoid leveraging any knowledge on the true rewards ($r_{i,t}$). In order to remove any dependency on the rewards, we impose a constrain on the weights, such that $\sum_{j=1}^J w_j \alpha_j = 1$, thus removing the last term, and thus the rewards, in the previous expression. We can then focus on minimizing a high-probability upper bound of (a)

$$(a) \leq 2c_g^{-1} \max_{i_1, \dots, i_K} \sum_{l=1}^K \sum_{j=1}^J |w_j| \cdot |\epsilon_{i,t,j}| \quad (38)$$

where we leverage the fact that the noise $\epsilon_{i,t,j}$ are sub-Gaussian then $|\epsilon_{i,t,j}| - \mathbb{E}(|\epsilon_{i,t,j}|)$ is also sub-Gaussian with the same parameter and by Jansen inequality its mean is $\mathbb{E}(|\epsilon_{i,t,j}|) \leq \sqrt{\mathbb{E}(\epsilon_{i,t,j}^2)} = \sigma_j$. Therefore for any time t and any set $U \subset \{1, \dots, K_t\}$ of size K , we have by Chernoff inequality, for any $x > 0$

$$\mathbb{P} \left(\left| \sum_{u \in U} \sum_{j=1}^J |w_j \epsilon_{u,t,j}| - \mathbb{E}(|w_j \epsilon_{u,t,j}|) \right| \geq x \right) \leq 2 \exp \left(- \frac{x^2}{2K \sum_{j=1}^J (w_j \sigma_j)^2} \right) \quad (39)$$

Therefore we have taking a union over all the set of size K ,

$$\mathbb{P} \left(\max_{i_1, \dots, i_K} \sum_{l=1}^K \sum_{j=1}^J |w_j \epsilon_{i,t,j}| - \mathbb{E}(|w_j \epsilon_{i,t,j}|) \geq x \right) \leq \sum_{U, |U|=K} \mathbb{P} \left(\sum_{u \in U} \sum_{j=1}^J |w_j \epsilon_{u,t,j}| - \mathbb{E}(|w_j \epsilon_{u,t,j}|) \geq x \right) \quad (40)$$

$$\leq \sum_{U, |U|=K} 2 \exp \left(- \frac{x^2}{2 \sum_{j=1}^J (w_j \sigma_j)^2} \right) \quad (41)$$

$$= 2 \binom{K_t}{K} \exp \left(- \frac{x^2}{2K^2 \sum_{j=1}^J (w_j \sigma_j)^2} \right) \quad (42)$$

Therefore for any $\delta \in (0, 1)$ we have that with probability at least $1 - \delta$:

$$\max_{i_1, \dots, i_K} \sum_{l=1}^K \sum_{j=1}^J |w_j| \cdot |\epsilon_{i,t,j}| - \mathbb{E}(|w_j| \cdot |\epsilon_{i,t,j}|) \leq K \sqrt{2 \sum_{j=1}^J (w_j \sigma_j)^2 \ln \left(\frac{2 \binom{K_t}{K}}{\delta} \right)} \leq K \sqrt{2K \sum_{j=1}^J (w_j \sigma_j)^2 \ln \left(\frac{2eK_{\max}}{K\delta} \right)} \quad (43)$$

using the standard inequality $\binom{K_t}{K} \leq \left(\frac{eK_t}{K} \right)^K$ and $K_t \leq K_{\max}$. Therefore,

$$(a) \leq \frac{2K}{c_g} \sqrt{K \sum_{j=1}^J w_j^2 \sigma_j^2 \ln \left(\frac{K_t e}{\delta} \right)} + \frac{K}{c_g} \sum_{j=1}^J |w_j| \sigma_j \leq \frac{2K}{c_g} \sqrt{K \sum_{j=1}^J w_j^2 \sigma_j^2 \ln \left(\frac{K_{\max} e}{\delta} \right)} + \frac{K}{c_g} \sqrt{J \sum_{j=1}^J (w_j \sigma_j)^2}, \quad (44)$$

Finally, this leads to the optimization problem in Eq. 12. By plugging the optimal solution back into the optimization problem, we also obtain the stated upper bound on the suboptimality gap.

C.2 Regret of ϵ -greedy Algorithm for GLM Model (Proof of Thm. 1)

We now move on proving the regret upper bound of Thm. 1. The first step is to bound the norm of the noisy evaluations at every step. That is the object of the following lemma.

Lemma 4. *Let $\phi_{i,t} = (f_{i,t,1}, \dots, f_{i,t,j}, \dots, f_{i,t,J}) \in \mathbb{R}^J$ the vector summarizing all the evaluations observed at round t . Then for any $\delta \in (0, 1)$ with probability $1 - \delta$ for any $t \leq T$, for any set \mathcal{A}_t of K arms chosen in $\{1, \dots, K_{\max}\}$ possibly adaptive to the evaluations, it holds that*

$$\left\| \sum_{i \in \mathcal{A}_t} \phi_{i,t} \right\|_2 \leq \Phi := 2K \|\sigma\|_\infty \left(2\sqrt{J} + \sqrt{K \ln \left(\frac{2eK_{\max}T}{K\delta} \right)} \right) + 2KJ \|g\|_\infty$$

where $\|\sigma\|_\infty = \max_{j \leq J} \sigma_j$ and $\|g\|_\infty = \max_{j \leq J, x \in [0, C]} |g(\alpha_j x)|$

Proof. For a time $t \leq T$ and set of size K $\mathcal{A}_t \subset \{1, \dots, K_t\}$ and an evaluator $j \leq J$ the sum of features can be decomposed as,

$$\left(\sum_{i \in \mathcal{A}_t} \phi_{i,t,j} \right)^2 = \left(\sum_{i \in \mathcal{A}_t} g(\alpha_j r_{i,t}) + \sum_{i \in \mathcal{A}_t} \epsilon_{i,t,j} \right)^2 \leq 2 \left(\sum_{i \in \mathcal{A}_t} g(\alpha_j r_{i,t}) \right)^2 + 2 \left(\sum_{i \in \mathcal{A}_t} \epsilon_{i,t,j} \right)^2 \quad (45)$$

$$\leq 2K^2 \|g\|_\infty^2 + 2 \left(\sum_{i \in \mathcal{A}_t} \epsilon_{i,t,j} \right)^2 \quad (46)$$

Therefore, we just need to bound $\left\| \sum_{i \in \mathcal{A}_t} \epsilon_{i,t} \right\|_2$ to finish the proof. But thanks to the subGaussian assumption on the noise ϵ , we have that:

$$\left\| \sum_{i \in \mathcal{A}_t} \epsilon_{i,t} \right\|_2 \leq 4K \sqrt{J \max_{j \leq J} \sigma_j^2} + 2K \max_{j \leq J} \sigma_j \sqrt{\log \left(\frac{1}{\delta} \right)} \quad (47)$$

because the vector $\sum_{i \in \mathcal{A}_t} \epsilon_{i,t}$ is $K \max_j \sigma_j$ sub-Gaussian. The last step is to take a union bound over steps $t \leq T$ and subset of size K . \square

The first step to analyze the regret of Alg. 1 is to decompose the regret according to the steps where the algorithm selected a totally random set of arms or when it played greedy. Noting, as in Alg. 1, Z_t the Bernoulli random variable used by the algorithm to distinguish between exploratory and exploitative steps, the regret is

$$R_T = \underbrace{\sum_{t=1}^T Z_t \langle w^+, \sum_{i \in \mathcal{A}_t^+} \phi_{t,i} - \sum_{i \in \mathcal{A}_t} \phi_{i,t} \rangle}_{:=R_{T,1}} + \underbrace{\sum_{t=1}^T (1 - Z_t) \langle w^+, \sum_{i \in \mathcal{A}_t^+} \phi_{i,t} - \sum_{i \in \mathcal{A}_t} \phi_{i,t} \rangle}_{:=R_{T,2}} \quad (48)$$

In the following, we bound each term $R_{T,1}$ and $R_{T,2}$.

Bounding $R_{T,1}$. The term we analyze is the regret due to the random steps in the algorithm. Centering the variable Z_t we get that $R_{T,1}$ scales as:

$$\begin{aligned} R_{T,1} &= \sum_{t=1}^T (Z_t - \epsilon) \langle w^+, \sum_{i \in \mathcal{A}_t^+} \phi_{t,i} - \sum_{i \in \mathcal{A}_t} \phi_{i,t} \rangle + \epsilon \sum_{t=1}^T \langle w^+, \sum_{i \in \mathcal{A}_t^+} \phi_{t,i} - \sum_{i \in \mathcal{A}_t} \phi_{i,t} \rangle \\ &\leq \sum_{t=1}^T (Z_t - \epsilon) \langle w^+, \sum_{i \in \mathcal{A}_t^+} \phi_{t,i} - \sum_{i \in \mathcal{A}_t} \phi_{i,t} \rangle + \epsilon T \|w^+\| \left(\left\| \sum_{i \in \mathcal{A}_t^+} \phi_{i,t} \right\|_2 + \left\| \sum_{i \in \mathcal{A}_t} \phi_{i,t} \right\|_2 \right) \\ &\leq \sum_{t=1}^T (Z_t - \epsilon) \langle w^+, \sum_{i \in \mathcal{A}_t^+} \phi_{t,i} - \sum_{i \in \mathcal{A}_t} \phi_{i,t} \rangle + 2\Phi \epsilon T \|w^+\| \end{aligned} \quad (49)$$

thanks to Lemma 4. Now, the first term can be bounded thanks to a standard Azuma-Hoeffding inequality as for every time $t \leq T$, Z_t is independent from the evaluations $(\phi_{i,t})_i$ and the set of arms \mathcal{A}_t hence considering the filtration $(\mathcal{F}_t)_t$ being the history up until time t , we have,

$$\forall t \leq T, \left| \sum_{t=1}^T (Z_t - \epsilon) \langle w^+, \sum_{i \in \mathcal{A}_t^+} \phi_{t,i} - \sum_{i \in \mathcal{A}_t} \phi_{i,t} \rangle \right| \leq 4\Phi \|w^+\| \sqrt{2T \ln \left(\frac{4T}{\delta} \right)} \text{ w.p at least } 1 - \delta/2 \quad (50)$$

Putting everything together we have that for any $\delta \in (0, 1)$,

$$R_{T,1} \leq 4\Phi \|w^+\| \left(\sqrt{2T \ln \left(\frac{4T}{\delta} \right)} + \epsilon T \right) \quad (51)$$

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

Bounding $R_{T,2}$. Next, we bound the regret coming due to the estimation error in the exploitative steps, that is to say when $Z_t = 0$. We begin by using the greedy behavior of the algorithm to enough to study the error $\|w_t - w^+\|$ to bound $R_{T,2}$,

$$R_{T,2} \leq \sum_{t=1}^T (1 - Z_t) \left[\langle w^+ - w_t, \sum_{i \in \mathcal{A}_t^+} \phi_{i,t} \rangle + \langle w_t, \sum_{i \in \mathcal{A}_t^+} \phi_{i,t} - \sum_{i \in \mathcal{A}_t} \phi_{i,t} \rangle + \langle w_t - w^+, \sum_{i \in \mathcal{A}_t} \phi_{i,t} \rangle \right] \quad (52)$$

But because \mathcal{A}_t is the maximizer of the estimated reward with respect to w_t we have that

$$(1 - Z_t) \left\langle w_t, \sum_{i \in \mathcal{A}_t^+} \phi_{i,t} - \sum_{i \in \mathcal{A}_t} \phi_{i,t} \right\rangle \leq 0. \quad (53)$$

Furthermore, using Lemma 4 we have that with probability at least $1 - \delta/4$

$$\max \left\{ \left\langle w^+ - w_t, \sum_{i \in \mathcal{A}_t^+} \phi_{i,t} \right\rangle, \left\langle w_t - w^+, \sum_{i \in \mathcal{A}_t} \phi_{i,t} \right\rangle \right\} \leq \Phi \|w^+ - w_t\|_2. \quad (54)$$

As a result, the remaining step to obtain a bound on the regret is to build a concentration inequality for the weights w_t w.r.t. w^+ and properly tune the exploration factor ϵ .

We start by studying the concentration of the MLE estimator $\hat{\alpha}_t \in \mathbb{R}^J$ to the true parameter $\alpha \in \mathbb{R}^J$. Thanks to the random choice of \mathcal{A}_t in the explorative steps, all the data in \mathcal{H}_t are i.i.d. and we can directly apply the standard concentration inequality for MLE in GLM (Li et al., 2017).

Lemma 5. *For any $\delta \in (0, 1)$ with probability at least $1 - \delta/8$ for all $j \leq J$ and $t \leq T$:*

$$|\hat{\alpha}_{t,j} - \alpha_j| \leq \beta_{t,j} := \frac{\sqrt{\lambda} \|\alpha\| + \sigma_j \sqrt{\frac{1}{2} \log \left(1 + \frac{2T}{\lambda} \right) + \log \left(\frac{8JT}{\delta} \right)}}{c_g \sqrt{\sum_{l, Z_l=1} \sum_{i \in \mathcal{A}_t} r_{l,i}}}. \quad (55)$$

In other words,

$$\|\hat{\alpha}_t - \alpha_r\|_\infty \leq \beta_t := \max_j \beta_{t,j} = \frac{\sqrt{\lambda} \|\alpha\| + \|\sigma\|_\infty \sqrt{\frac{1}{2} \log \left(1 + \frac{2T}{\lambda} \right) + \log \left(\frac{8JT}{\delta} \right)}}{c_g \sqrt{\sum_{l, Z_l=1} \sum_{i \in \mathcal{A}_t} r_{l,i}}} \quad (56)$$

We leverage Lem. 5 to concentrate $\|w^* - w_t\|$. But first, recall that w_t is given by:

$$w_t = \frac{\sigma^{-1} \cdot \hat{\alpha}_t}{\|\sigma^{-1} \cdot \hat{\alpha}_t\|_2^2} \text{ and } w^+ = \frac{\sigma^{-1} \cdot \alpha}{\|\sigma^{-1} \cdot \alpha\|_2^2}$$

where $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_J)$. Therefore, the error between w^+ and w_t can be written as:

$$\|w_t - w^+\|_2 = \left\| \frac{\sigma^{-2} \cdot (\hat{\alpha}_t - \alpha)}{\|\sigma^{-1} \cdot \hat{\alpha}_t\|^2} + \sigma^{-2} \cdot \alpha \left(\frac{1}{\|\sigma^{-1} \cdot \hat{\alpha}_t\|^2} - \frac{1}{\|\sigma^{-1} \cdot \alpha\|^2} \right) \right\| \quad (57)$$

$$\leq \left\| \sigma^{-2} \cdot \alpha \left(\frac{1}{\|\sigma^{-1} \cdot \hat{\alpha}_t\|^2} - \frac{1}{\|\sigma^{-1} \cdot \alpha\|^2} \right) \right\|_2 + \left\| \frac{\sigma^{-2} \cdot (\hat{\alpha}_t - \alpha)}{\|\sigma^{-1} \cdot \hat{\alpha}_t\|^2} \right\|_2 \quad (58)$$

Proposition 2. For any time $t \leq T$, assuming that $\|\hat{\alpha}_t - \alpha\|_\infty \leq \beta_t$ and $\|\alpha\|_\infty \geq 8\beta_t$ then:

$$\frac{\|\sigma^{-2} \cdot (\hat{\alpha}_t - \alpha)\|_2}{\langle \hat{\alpha}_t, \sigma^{-2} \cdot \hat{\alpha}_t \rangle} \leq \frac{4\beta_t \sqrt{\sum_{j=1}^J \frac{1}{\sigma_j^4}}}{\sum_{j=1}^J \frac{\alpha_j^2 \mathbf{1}_{\{|\alpha_j| \geq 8\beta_t\}}}{\sigma_j^2}} = 4\beta_t \left(\frac{\|\sigma^{-1}\|_4}{\|\sigma^{-1} \cdot \alpha \mathbf{1}_{\{|\alpha| \geq 8\beta_t\}}\|_2} \right)^2 \quad (59)$$

Proof. Using the definition of σ^{-2} . we have,

$$\frac{\|\sigma^{-2} \cdot (\hat{\alpha}_t - \alpha)\|_2}{\langle \hat{\alpha}_t, \sigma^{-2} \cdot \hat{\alpha}_t \rangle} = \frac{\sqrt{\sum_{j=1}^J \frac{(\hat{\alpha}_{t,j} - \alpha_j)^2}{\sigma_j^4}}}{\sum_{r=1}^J \frac{\hat{\alpha}_{t,r}^2}{\sigma_r^2}} \quad (60)$$

However, we have that for any $j \leq J$:

$$\begin{aligned} \hat{\alpha}_{t,j}^2 &= \alpha_j^2 + (\hat{\alpha}_{t,j} - \alpha_j)^2 + 2(\hat{\alpha}_{t,j} - \alpha_j)\alpha_j \geq \alpha_j^2 + 2(\hat{\alpha}_{t,j} - \alpha_j)\alpha_j \\ &\geq \alpha_j^2 - 2\beta_t|\alpha_j| \\ &\geq \frac{\alpha_j^2}{2} \text{ when } |\alpha_j| \geq 8\beta_t \end{aligned}$$

using that $\|\hat{\alpha}_t - \alpha\|_\infty \leq \beta_t$. In addition, using again that $\|\hat{\alpha}_t - \alpha\|_\infty \leq \beta_t$,

$$\sum_{j=1}^J \frac{(\hat{\alpha}_{t,j} - \alpha_j)^2}{\sigma_j^4} \leq \sum_{j=1}^J \frac{\beta_t^2}{\sigma_j^4} \quad (61)$$

□

Using the same reasoning as in Prop. 2, we can show the following which bound the second term in Equation 57.

Proposition 3. For any time $t \leq T$, assuming that $\|\hat{\alpha}_t - \alpha\|_\infty \leq \beta_t$ and $\|\alpha\|_\infty \geq 8\beta_t$ then:

$$\left\| \left(\frac{\|\sigma^{-1} \cdot \alpha\|^2 - \|\sigma^{-1} \cdot \hat{\alpha}_t\|^2}{\|\sigma^{-1} \cdot \hat{\alpha}_t\|^2} \right) \frac{\sigma^{-2} \cdot \alpha}{\|\sigma^{-1} \cdot \alpha\|^2} \right\| \leq \frac{4\beta_t (\|\sigma^{-2} \cdot \alpha\|_2 + \beta_t \|\sigma^{-2}\|) \|\sigma^{-2} \cdot \alpha\|}{\|\sigma^{-1} \cdot \alpha \mathbf{1}_{\{|\alpha| \geq 8\beta_t\}}\|^2 \|\sigma^{-1} \cdot \alpha\|^2} \quad (62)$$

Therefore, using Prop. 2 and Prop. 3 on the event that the confidence intervals in Lem. 5 holds then with probability at least $1 - \delta/4$:

$$\|w_t - w^+\|_2 \leq \frac{4\beta_t (\|\sigma^{-2} \cdot \alpha\|_2 + \beta_t \|\sigma^{-2}\|) \|\sigma^{-2} \cdot \alpha\|}{\|\sigma^{-1} \cdot \alpha \mathbf{1}_{\{|\alpha| \geq 8\beta_t\}}\|^2 \|\sigma^{-1} \cdot \alpha\|^2} + 4\beta_t \left(\frac{\|\sigma^{-1}\|_4}{\|\sigma^{-1} \cdot \alpha \mathbf{1}_{\{|\alpha| \geq 8\beta_t\}}\|_2} \right)^2 \quad (63)$$

In order to complete the proof of the regret upper bound, we need to control how fast the confidence intervals width β_t decreases with t .

Lemma 6. For any $\delta \in (0, 1)$ and $\lambda > 0$ we have with probability at least $1 - \delta/4$,

$$\forall t \in \left\{ \frac{2C^4 \ln(\frac{8T}{\delta})}{\epsilon^2 K^2 \eta_\nu^4}, \dots, T \right\}, \beta_t \leq \frac{\|\sigma\|_\infty \sqrt{\ln(1 + \frac{2T}{\lambda}) + \ln(\frac{8TJ}{\delta})} + \sqrt{\lambda} \|\alpha\|}{c_g \sqrt{\epsilon t K \eta_\nu^2 - C^2 \sqrt{2t \ln(\frac{8T}{\delta})}}} \quad (64)$$

where for all $j \leq J$, $\eta_\nu^2 = \mathbb{E}_{r \sim \nu}(r^2)$

Proof. First, thanks to a standard Hoeffding inequality we have that with probability at least $1 - \delta/4$,

$$\forall t \leq T, \left| \sum_{l \leq t, Z_l=1} \sum_{i \in A_t} r_{l,i}^2 - \mathbb{E} \left(\sum_{l \leq t, Z_l=1} \sum_{i \in A_t} r_{l,i}^2 \right) \right| \leq C^2 \sqrt{2t \ln \left(\frac{8T}{\delta} \right)}$$

And $\mathbb{E} \left(\sum_{l \leq t, Z_l=1} \sum_{i \in A_t} r_{l,i}^2 \right) = \epsilon t K \eta_\nu^2$ where $\eta_\nu^2 := \mathbb{E}_{r \sim \nu}(r^2)$. Therefore, using the definition of β_t

$$\beta_t \leq \frac{\|\sigma\|_\infty \sqrt{\ln \left(1 + \frac{2T}{\lambda} \right) + \ln \left(\frac{8TJ}{\delta} \right) + \sqrt{\lambda} \|\alpha\|}}{c_g \sqrt{\epsilon t K \eta_\nu^2 - C^2 \sqrt{2t \ln \left(\frac{8T}{\delta} \right)}}} \quad (65)$$

□

As a consequence, for $t \geq \tau_0 := \frac{8 \ln \left(\frac{8T}{\delta} \right) C^4}{(\epsilon K \eta_\nu^2)^2}$, $\beta_t \leq \mathcal{O} \left(\frac{1}{\sqrt{\epsilon t}} \right)$. Therefore, using Lem. 6 with probability at least $1 - \delta/4$:

$$\begin{aligned} \sum_{t=\tau_0+1}^T \beta_t &\leq \sum_{t=\tau_0+1}^T \frac{\sqrt{2} \|\sigma\|_\infty \sqrt{\ln \left(1 + \frac{2T}{\lambda} \right) + \ln \left(\frac{8TJ}{\delta} \right) + \sqrt{\lambda} \|\alpha\|}}{c_g \sqrt{\epsilon t K \eta_\nu^2}} \\ &\leq \frac{3\sqrt{2T} \|\sigma\|_\infty}{c_g \sqrt{\epsilon K \eta_\nu^2}} \left(\sqrt{\ln \left(1 + \frac{2T}{\lambda} \right) + \ln \left(\frac{8TJ}{\delta} \right) + \sqrt{\lambda} \|\alpha\|} \right) \end{aligned} \quad (66)$$

We can finally bound $R_{T,1}$ with probability at least $1 - \delta/2$ by:

$$\begin{aligned} R_{T,1} &\leq \tau + \frac{24\Phi\sqrt{2T}\|\sigma\|_\infty}{c_g\sqrt{\epsilon K\eta_\nu^2}} \left(\sqrt{\ln \left(1 + \frac{2T}{\lambda} \right) + \ln \left(\frac{8TJ}{\delta} \right) + \sqrt{\lambda}\|\alpha\|} \right) \left(\frac{(\|\sigma^{-2} \cdot \alpha\|_2 + \|\sigma^{-2}\|) \|\sigma^{-2} \cdot \alpha\|}{\|\sigma^{-1} \cdot \alpha \mathbf{1}_{\{|\alpha| \geq 8\beta_t\}}\|^2 \|\sigma^{-1} \cdot \alpha\|^2} \right. \\ &\quad \left. + \left(\frac{\|\sigma^{-1}\|_4}{\|\sigma^{-1} \cdot \alpha \mathbf{1}_{\{|\alpha| \geq 8\beta_t\}}\|_2} \right)^2 \right) \end{aligned} \quad (67)$$

where $\tau := \max \left\{ \frac{128\|\sigma\|_\infty (\ln(1 + \frac{2T}{\lambda}) + \ln(\frac{8TJ}{\delta})) + \lambda\|\alpha\|^2}{\min_j |\alpha_j|^2 \epsilon K \eta_\nu^2}, \frac{8 \ln \left(\frac{8T}{\delta} \right) C^4}{(\epsilon K \eta_\nu^2)^2} \right\}$

Finally, the regret is bounded with probability at least $1 - \delta$ by:

$$\begin{aligned} R_T &\leq \frac{4\Phi}{\|\sigma^{-1} \cdot \alpha\|} \left(\sqrt{2T \ln \left(\frac{4T}{\delta} \right) + \epsilon T} \right) + \max \left\{ \frac{128\|\sigma\|_\infty (\ln(1 + \frac{2T}{\lambda}) + \ln(\frac{8TJ}{\delta})) + \lambda\|\alpha\|^2}{\min_{j, \alpha_j \neq 0} |\alpha_j|^2 \epsilon K \eta_\nu^2}, \frac{8 \ln \left(\frac{8T}{\delta} \right) C^4}{(\epsilon K \eta_\nu^2)^2} \right\} \\ &\quad + \frac{24\Phi\sqrt{2T}\|\sigma\|_\infty}{c_g\sqrt{\epsilon K\eta_\nu^2}} \left(\sqrt{\ln \left(1 + \frac{2T}{\lambda} \right) + \ln \left(\frac{8TJ}{\delta} \right) + \sqrt{\lambda}\|\alpha\|} \right) \left(\frac{(\|\sigma^{-2} \cdot \alpha\|_2 + \|\sigma^{-2}\|) \|\sigma^{-2} \cdot \alpha\|}{\|\sigma^{-1} \cdot \alpha\|^2 \|\sigma^{-1} \cdot \alpha\|^2} + \left(\frac{\|\sigma^{-1}\|_4}{\|\sigma^{-1} \cdot \alpha\|_2} \right)^2 \right) \end{aligned} \quad (68)$$

Remark. Compared to the bound reported in Thm. 1, we corrected an error in the analysis which removed the dependency on the norm $\|g\|_\infty$ and replaced it with the upper bound on the distribution of the reward C . In addition, here we report a bound depends on $\min_{j \leq J, \alpha_j \neq 0} |\alpha_j|^2$ however because of the condition $\mathbf{1}_{\{|\alpha_j| \geq 8\beta_t\}}$ it is possible to present a regret upper bound that scales inversly with $\max_j |\alpha_j|$ but at the cost of a much bigger problem-dependent quantity S . Fortunately, the difference between the bound report in Thm 1 and the one in Equation 68 has no impact on the discussion around the dependency on J . Simplifying the bound in Equation 68, we get that the regret is bounded with high probability by,

$$R_T \leq \tilde{\mathcal{O}} \left(T^{2/3} \left(\frac{(1 + \sqrt{J^{-1}} \|\alpha\|) \Phi S \|\sigma\|_\infty}{c_g \sqrt{K} \eta_\nu} + \max \left\{ \frac{\|\sigma\|_\infty}{\min_{j, \alpha_j \neq 0} |\alpha_j|^2 K \eta_\nu^2}, \frac{C^4}{(K \eta_\nu^2)^2} \right\} \right) \right) \quad (69)$$

when choosing $\epsilon = \frac{1}{T^{1/3}}$, $\lambda = J^{-1}$ and setting $S := \frac{(\|\sigma^{-2} \cdot \alpha\|_2 + \|\sigma^{-2}\|) \|\sigma^{-2} \cdot \alpha\|}{\|\sigma^{-1} \cdot \alpha\|^2 \|\sigma^{-1} \cdot \alpha\|^2} + \left(\frac{\|\sigma^{-1}\|_4}{\|\sigma^{-1} \cdot \alpha\|_2} \right)^2$.

D LINEAR MODEL

In this section, we show how the definition of the oracle changes with the linear setting and the proof of Thm. 2, that is to the regret of Alg. 2

D.1 Oracle in the Linear Model (Proof of Lem. 2)

Just as in Section C.1, we consider an oracle that has access to the parameters $(\alpha_j)_{j \leq J}$ but also knows the noise $(\sigma_j)_{j \leq J}$ for each evaluation functions. Given a vector of evaluations $\phi_{i,t}$ for some $t \leq T$ and $i \leq K_t$, an oracle is defined through a set of weights $(w_j)_{j \leq J}$ and predicts the reward, $r_{i,t}^\Delta = \langle w, \phi_{i,t} \rangle$. Following the reasoning in Section C.1, the difference between the true top- K arms, i_1^+, \dots, i_K^+ , and the top K according to the oracle, $i_1^\Delta, \dots, i_K^\Delta$ is bounded by

$$\sum_{l=1}^K r_{i_l^+, t} - r_{i_l^\Delta, t} \leq 2 \max_{i_1, \dots, i_K} \sum_{l=1}^K |\widehat{r}_{i_l, t}^\Delta - r_{i_l, t}| \quad (70)$$

Again following the same reasoning as in App. C.1, the difference between the estimated reward and the actual reward for any arm $i \leq K_t$ is bounded by

$$|r_{i,t} - \widehat{r}_{i,t}^\Delta| \leq \left| r_{i,t} \left(\sum_{j=1}^J w_j \alpha_j - 1 \right) \right| + \left| \sum_{j=1}^J w_j \epsilon_{i,t,j} \right| \leq C \left| \left(\sum_{j=1}^J w_j \alpha_j - 1 \right) \right| + \left| \sum_{j=1}^J w_j \epsilon_{i,t,j} \right| \quad (71)$$

The main difference compared to the proof of Lem. 1 is that thanks to the linear structure controlling the noise can be done directly without using any Lipschitz property. Therefore, in order to remove the dependency on the reward, we focus on weights w such that $\langle w, \alpha \rangle = 1$. Hence using Chernoff inequality the error is bounded for any $\delta \in (0, 1)$ by

$$|r_{i,t} - r_{i,t}^\Delta| \leq 2K \sqrt{K \sum_{j=1}^J w_j^2 \sigma_j^2 \ln \left(\frac{K_{\max} e}{\delta} \right)} \quad (72)$$

with probability at least $1 - \delta$. This concludes the proof of Lem. 1

D.2 Regret of ESAG (Proof of Thm. 2)

The first step to analyze the regret of Alg. 2 is to compute an upper bound on the deviation of the estimator $\widehat{\alpha}_t$.

Lemma 7. *For any $\delta \in (0, 1)$ and $t \leq T$, the error between $\widehat{\alpha}_t$ and $\mathbb{E}_{r \sim \nu}(r)\alpha$ is bounded with probability at least $1 - \delta/4$ by,*

$$\|\widehat{\alpha}_t - \mathbb{E}_{r \sim \nu}(r)\alpha\| \leq \beta_t^L := \frac{C \sqrt{2 \ln \left(\frac{16JT}{\delta} \right)}}{\sqrt{\sum_{l=1}^{t-1} K_l}} + \frac{4\|\sigma\|_\infty \sqrt{J}}{\sqrt{\sum_{l=1}^{t-1} K_l}} + \frac{2\|\sigma\|_\infty \sqrt{\ln \left(\frac{8TJ}{\delta} \right)}}{\sqrt{\sum_{l=1}^{t-1} K_l}} \quad (73)$$

Proof. For any $t \leq T$, $\widehat{\alpha}_t = \frac{\sum_{l=1}^{t-1} \sum_{i=1}^{K_l} \phi_{i,l}}{\sum_{l=1}^{t-1} K_l}$ but using the structure of the evaluations $\phi_{i,t}$ for every $j \leq J$,

$$\widehat{\alpha}_{t,j} = \frac{1}{\sum_{l=1}^{t-1} K_l} \left(\sum_{l=1}^{t-1} \sum_{i=1}^{K_l} \alpha_j r_{i,l} + \sum_{l=1}^{t-1} \sum_{i=1}^{K_l} \epsilon_{i,t,j} \right) \quad (74)$$

But the first term converges toward $\alpha_j \mathbb{E}_{r \sim \nu}(r)$ because at every time step the reward are independently sampled from the distribution ν (see Assumption 1). More precisely, with probability at least $1 - \delta/8$, for any $t \leq T$ and $j \leq J$

$$\left| \frac{1}{\sum_{l=1}^{t-1} K_l} \sum_{l=1}^{t-1} \sum_{i=1}^{K_l} r_{i,l} - \mathbb{E}_{r \sim \nu}(r) \right| \leq \frac{C \sqrt{2 \ln \left(\frac{16JT}{\delta} \right)}}{\sqrt{\sum_{l=1}^{t-1} K_l}} \quad (75)$$

In addition, the second term in Eq. (74) can be bounded with probability at least $1 - \delta/8$ by

$$\begin{aligned} \left\| \frac{1}{\sum_{l=1}^{t-1} K_l} \sum_{l=1}^{t-1} \sum_{i=1}^{K_l} \epsilon_{i,l} \right\| &\leq \frac{4\sqrt{J \max_j \sigma_j^2 \sum_{l=1}^{t-1} K_l} + 2\sqrt{\max_j \sigma_j^2 \sum_{l=1}^{t-1} K_l \ln\left(\frac{8TJ}{\delta}\right)}}{\sum_{l=1}^{t-1} K_l} \\ &= \frac{4\|\sigma\|_\infty \sqrt{J}}{\sqrt{\sum_{l=1}^{t-1} K_l}} + \frac{2\|\sigma\|_\infty \sqrt{\ln\left(\frac{8TJ}{\delta}\right)}}{\sqrt{\sum_{l=1}^{t-1} K_l}} \end{aligned} \quad (76)$$

Finally, to finish the proof, we simply need to note that

$$\|\widehat{\alpha}_t - \mathbb{E}_{r \sim \nu}(r)\alpha\| \leq \left\| \frac{\alpha}{\sum_{l=1}^{t-1} K_l} \sum_{l=1}^{t-1} \sum_{i=1}^{K_l} r_{i,l} - \alpha \mathbb{E}_{r \sim \nu}(r) \right\| + \left\| \frac{1}{\sum_{l=1}^{t-1} K_l} \sum_{l=1}^{t-1} \sum_{i=1}^{K_l} \epsilon_{i,l} \right\| \quad (77)$$

$$\leq \|\alpha\| \left| \frac{1}{\sum_{l=1}^{t-1} K_l} \sum_{l=1}^{t-1} \sum_{i=1}^{K_l} r_{i,l} - \mathbb{E}_{r \sim \nu}(r) \right| + \frac{4\|\sigma\|_\infty \sqrt{J}}{\sqrt{\sum_{l=1}^{t-1} K_l}} + \frac{2\|\sigma\|_\infty \sqrt{\ln\left(\frac{8TJ}{\delta}\right)}}{\sqrt{\sum_{l=1}^{t-1} K_l}} \quad (78)$$

□

Once ESAG has computed the weights $\widehat{\alpha}_t$ it then computes a set of weights thanks to this estimator and based on the shape of the optimal weights of the oracle. That is to say the weights w used by the algorithm are:

$$w_t = \frac{\sigma^{-2} \cdot \widehat{\alpha}_t}{\|\sigma^{-1} \cdot \widehat{\alpha}_t\|_2^2} \quad (79)$$

The optimal weights for the parameter $\mathbb{E}_{r \sim \nu}(r)\alpha$ is $\frac{w^+}{\mathbb{E}_{r \sim \nu}(r)}$. Therefore, in the following we bound the error between the weights w_t and the rescaled optimal weights $\frac{w^+}{\mathbb{E}_{r \sim \nu}(r)}$. Following the same reasoning as in Prop. 2 and Prop. 3, we can show the following proposition

Proposition 4. *For any $t \leq T$ let's assume $\mathbb{E}_{r \sim \nu}(r)\|\alpha\|_\infty \geq 8\beta_t^L$ with β_t^L defined in Eq. (74). Then for any $\delta \in (0, 1)$ with probability at least $1 - \delta/4$*

$$\left\| w_t - \frac{w^+}{\mathbb{E}_{r \sim \nu}(r)} \right\|_2 \leq \frac{4\beta_t^L (\mathbb{E}_{r \sim \nu}(r)\|\sigma^{-2} \cdot \alpha\|_2 + \beta_t^L \|\sigma^{-2}\|) \|\sigma^{-2} \cdot \alpha\|}{\mathbb{E}_{r \sim \nu}(r)^3 \|\sigma^{-1} \cdot \alpha \mathbf{1}_{\{\mathbb{E}_{r \sim \nu}(r)|\alpha| \geq 8\beta_t^L}\|_2^2 \|\sigma^{-1} \cdot \alpha\|^2} + \frac{4\beta_t^L}{\mathbb{E}_{r \sim \nu}(r)^2} \left(\frac{\|\sigma^{-1}\|_4}{\|\sigma^{-1} \cdot \alpha \mathbf{1}_{\{\mathbb{E}_{r \sim \nu}(r)|\alpha| \geq 8\beta_t^L}\|_2} \right)^2 \quad (80)$$

Thanks to Lemma 7 and Prop. 4, we can now analyze the regret of ESAG. Recall the regret is defined as follows,

$$R_T = \sum_{t=1}^T \sum_{i \in \mathcal{A}_t^+} \langle w^+, \phi_{i,t} \rangle - \sum_{i \in \mathcal{A}_t} \langle w^+, \phi_{i,t} \rangle \quad (81)$$

where $\mathcal{A}_t^+ := \arg \max_i^K \langle w^+, \phi_{i,t} \rangle$ and \mathcal{A}_t is the set of jobs selected by the learner at time t .

However, the set \mathcal{A}_t^+ is only defined as the arg max of some values, it is invariant by multiplying the weights w^+ by some constant independent of the arms. In particular, we have that for any time $t \leq T$

$$\mathcal{A}_t^+ = \arg \max_i^K \langle w^+, \phi_{i,t} \rangle = \arg \max_i^K \left\langle \frac{w^+}{\mathbb{E}_{r \sim \nu}(r)}, \phi_{i,t} \right\rangle \quad (82)$$

Hence the regret can be rewritten as,

$$\begin{aligned} R_T &= \sum_{t=1}^T \mathbb{E}_{r \sim \nu}(r) \left(\sum_{i \in \mathcal{A}_t^+} \left\langle \frac{w^+}{\mathbb{E}_{r \sim \nu}(r)}, \phi_{i,t} \right\rangle - \sum_{i \in \mathcal{A}_t} \left\langle \frac{w^+}{\mathbb{E}_{r \sim \nu}(r)}, \phi_{i,t} \right\rangle \right) \\ &= \sum_{t=1}^T \mathbb{E}_{r \sim \nu}(r) \left(\sum_{i \in \mathcal{A}_t^+} \left\langle \frac{w^+}{\mathbb{E}_{r \sim \nu}(r)} - w_t, \phi_{i,t} \right\rangle + \sum_{i \in \mathcal{A}_t^+} \langle w_t, \phi_{i,t} \rangle - \sum_{i \in \mathcal{A}_t} \langle w_t, \phi_{i,t} \rangle + \sum_{i \in \mathcal{A}_t} \left\langle w_t - \frac{w^+}{\mathbb{E}_{r \sim \nu}(r)}, \phi_{i,t} \right\rangle \right) \end{aligned}$$

But for any $t \leq T$, because of the greedy arm selection process of ESAG:

$$\sum_{i \in \mathcal{A}_t^+} \langle w_t, \phi_{i,t} \rangle - \sum_{i \in \mathcal{A}_t} \langle w_t, \phi_{i,t} \rangle \leq 0 \quad (83)$$

But using Lem. 4 with probability at least $1 - \delta/4$

$$\max \left\{ \left\langle \frac{w^+}{\mathbb{E}_{r \sim \nu}(r)} - w_t, \sum_{i \in \mathcal{A}_t^+} \phi_{i,t} \right\rangle, \left\langle w_t - \frac{w^+}{\mathbb{E}_{r \sim \nu}(r)}, \sum_{i \in \mathcal{A}_t} \phi_{i,t} \right\rangle \right\} \leq \Phi \left\| \frac{w^+}{\mathbb{E}_{r \sim \nu}(r)} - w_t \right\|_2. \quad (84)$$

Thanks to Lem.7 and that for all $t \leq T$, $K_t \geq K$ therefore for $t \geq t_0 := 1 + 4K^{-1} (C^2 \ln(\frac{16JT}{\delta}) + 2\|\sigma\|_\infty (4J + \ln(\frac{8JT}{\delta}))) \max \left\{ \frac{64}{\mathbb{E}_{r \sim \nu}(r)^2 \min_{j, \alpha_j \neq 0} \alpha_j^2}, 1 \right\}$,

$$\left\| \frac{w^+}{\mathbb{E}_{r \sim \nu}(r)} - w_t \right\| \leq \frac{4\beta_t^L (\mathbb{E}_{r \sim \nu}(r) \|\sigma^{-2} \cdot \alpha\|_2 + \|\sigma^{-2}\|) \|\sigma^{-2} \cdot \alpha\|}{\mathbb{E}_{r \sim \nu}(r)^3 \|\sigma^{-1} \cdot \alpha\|^2 \|\sigma^{-1} \cdot \alpha\|^2} + \frac{4\beta_t^L}{\mathbb{E}_{r \sim \nu}(r)^2} \left(\frac{\|\sigma^{-1}\|_4}{\|\sigma^{-1} \cdot \alpha\|_2} \right)^2 \quad (85)$$

But using the definition of β_t^L :

$$\begin{aligned} \sum_{t=t_0+1}^T \beta_t^L &\leq \left(C \sqrt{2 \ln \left(\frac{16JT}{\delta} \right)} + 4\|\sigma\|_\infty \sqrt{J} + 2\|\sigma\|_\infty \sqrt{\ln \left(\frac{8TJ}{\delta} \right)} \right) \sum_{t=t_0+1}^T \frac{1}{\sqrt{\sum_{l=1}^{t-1} K_l}} \\ &\leq \left(C \sqrt{2 \ln \left(\frac{16JT}{\delta} \right)} + 4\|\sigma\|_\infty \sqrt{J} + 2\|\sigma\|_\infty \sqrt{\ln \left(\frac{8TJ}{\delta} \right)} \right) \sqrt{\frac{T \sum_{t=1}^{T-1} \frac{1}{t}}{\bar{K}_T}} \\ &\leq \left(C \sqrt{2 \ln \left(\frac{16JT}{\delta} \right)} + 4\|\sigma\|_\infty \sqrt{J} + 2\|\sigma\|_\infty \sqrt{\ln \left(\frac{8TJ}{\delta} \right)} \right) \sqrt{\frac{T \ln(T)}{\bar{K}_T}} \end{aligned} \quad (86)$$

where $\bar{K}_T := \frac{T}{\sum_{t=t_0}^{T-1} \frac{t-1}{\sum_{l=1}^{t-1} K_l}}$ is the average number of arms presented over the T steps, $K_T \geq K$. Therefore the regret is bounded with probability at least $1 - \delta$

$$R_T \leq \frac{2}{\|\sigma^{-1} \cdot \alpha\|} \left(1 + \frac{2L^2}{K} \max \left\{ \frac{64}{\mathbb{E}_{r \sim \nu}(r)^2 \min_{j, \alpha_j \neq 0} \alpha_j^2}, 1 \right\} \right) + \frac{16LS\Phi}{\mathbb{E}_{r \sim \nu}(r)} \max \left\{ 1, \frac{1}{\mathbb{E}_{r \sim \nu}(r)} \right\} \sqrt{\frac{T \ln(T)}{\bar{K}_T}} \quad (87)$$

where $S = (\|\alpha \cdot \sigma^{-2}\|_2 + \|\sigma^{-2}\|_2) \frac{\|\sigma^{-2} \cdot \alpha\|_2}{\|\sigma^{-1} \cdot \alpha\|_2^2} + \left(\frac{\|\sigma^{-1}\|_4}{\|\sigma^{-1} \cdot \alpha\|_2} \right)^2$ and $L = C \sqrt{2 \ln \left(\frac{16JT}{\delta} \right)} + 8\|\sigma\|_\infty \sqrt{J} + 4\|\sigma\|_\infty \sqrt{\ln \left(\frac{8TJ}{\delta} \right)}$

E REGRET ANALYSIS

In Section 2, the regret has been defined as the difference between the top- K arms according to the oracle and the K arms selected by a learner, see Equation 10. However one may argue that a more interesting notion of regret is to compare the actions of the oracle to the action of the learner with respect to the true reward of each arm. That is to say, to define the notion of *absolute* regret,

$$R_T^{\text{abs}} = \sum_{t=1}^T \sum_{i \in \mathcal{A}_t^+} r_{i,t} - \sum_{i \in \mathcal{A}_t} r_{i,t} \quad (88)$$

The two notions of regret, that is to say the *relative* and *absolute* regret are related,

$$R_T - R_T^{\text{abs}} = \sum_{t=1}^T \sum_{i \in \mathcal{A}_t^+} \langle w^+, g^{-1}(\phi_{i,t}) - g^{-1}(g(\alpha \cdot r_{i,t})) \rangle - \sum_{i \in \mathcal{A}_t} \langle w^+, g^{-1}(\phi_{i,t}) - g^{-1}(g(\alpha \cdot r_{i,t})) \rangle \quad (89)$$

However controlling the deviation here is not possible through standard concentration inequalities. Indeed the set of arms, \mathcal{A}_t^+ is computed by taking the argmax over the noise therefore $(\epsilon_{i,t})_{i \in \mathcal{A}_t^+}$ are not centered anymore, see App. B. Nonetheless, one can still provide some guarantees on the regret \tilde{R}_T^{abs} .

Theorem 3. For every $t \leq T$ let $(\phi_{(i,t)})_{i \leq K_t}$ be the ordered set of arms with respect to the oracle $\widehat{\mathfrak{D}}$, that is to say $\langle w^*, g^{-1}(\phi_{t,(1)}) \rangle \geq \langle w^*, g^{-1}(\phi_{t,(2)}) \rangle \geq \dots \geq \langle w^*, g^{-1}(\phi_{t,(K_t)}) \rangle$ let's now define the gap between the top- K according the oracle and the rest of the arms, $\Delta_{t,K} := \langle w^+, g^{-1}(\phi_{(K),t}) - g^{-1}(\phi_{(K+1),t}) \rangle$. For any weights $w \in \mathbb{R}^J$ such that:

$$\forall i \leq K_t, \quad |\langle w^+ - w, g^{-1}(\phi_{i,t}) \rangle| \leq \frac{\Delta_{t,K}}{4} \quad (90)$$

Then the top- K arms, $\mathcal{A} := \arg \max_i^K \langle w, \phi_{i,t} \rangle$ is exactly the ranking of the oracle, that is to say $\mathcal{A} = \mathcal{A}_t^+$.

Proof. For an arm $i \in \mathcal{A}_t^+$,

$$\begin{aligned} \langle w_t, g^{-1}(\phi_{i,t}) \rangle &\geq \langle w^*, g^{-1}(\phi_{i,t}) \rangle - \frac{\Delta_{t,K}}{4} \geq \langle w^*, g^{-1}(\phi_{(K),t}) \rangle - \frac{\Delta_{t,K}}{4} \\ &\geq \langle w^*, g^{-1}(\phi_{(K+1),t}) \rangle + \frac{3\Delta_{t,K}}{4} \\ &\geq \langle w_t, g^{-1}(\phi_{(K+1),t}) \rangle + \frac{\Delta_{t,K}}{2} > \langle w_t, g^{-1}(\phi_{(K+1),t}) \rangle \end{aligned}$$

Therefore $\mathcal{A}_t^+ \subset \mathcal{A}$. On the other hand for any $i \notin \mathcal{A}_t^+$, we have that:

$$\begin{aligned} \langle w_t, g^{-1}(\phi_{i,t}) \rangle &\leq \langle w^*, g^{-1}(\phi_{(K+1),t}) \rangle + \frac{\Delta_{t,K}}{4} \leq \langle w^*, g^{-1}(\phi_{(K),t}) \rangle - \frac{3\Delta_{t,K}}{4} \\ &\leq \langle w_t, g^{-1}(\phi_{(K),t}) \rangle - \frac{\Delta_{t,K}}{2} < \langle w_t, g^{-1}(\phi_{(K),t}) \rangle \end{aligned}$$

Therefore, the set $\mathcal{A}_t^+ \subset \mathcal{A}$ but because both are of size K , we have that $\mathcal{A} = \mathcal{A}_t^+$. \square

The main implication of Thm. 3 is that we can now provide a bound on the absolute regret of GLM- ε -GREEDY. Indeed, thanks to the bound of Lem. 5, for $t \geq$

$$\max \left\{ \left(\frac{2C^2 \sqrt{2 \ln(\frac{8T}{\delta})}}{\varepsilon K \eta_v^2} \right)^2, \frac{32 \left(\sqrt{\lambda} \|\alpha\| + \|\sigma\|_\infty \left(\sqrt{\ln(1 + \frac{2T}{\lambda})} + \ln(\frac{8JT}{\delta}) \right) \right)^2}{\varepsilon K \eta_v^2 \min_{t \leq T} \Delta_{t,K}^2} \right\},$$

the error between the estimated weights w_t and the optimal oracle weights w^+ is bounded by $\Delta_{t,K}/4$, therefore the absolute regret of Alg. 1 is bounded by

$$R_T^{\text{abs}} \leq C \max \left\{ \left(\frac{2C^2 \sqrt{2 \ln(\frac{8T}{\delta})}}{\varepsilon K \eta_v^2} \right)^2, \frac{32 \left(\sqrt{\lambda} \|\alpha\| + \|\sigma\|_\infty \left(\sqrt{\ln(1 + \frac{2T}{\lambda})} + \ln(\frac{8JT}{\delta}) \right) \right)^2}{\varepsilon K \eta_v^2 \min_{t \leq T} \Delta_{t,K}^2} \right\} \quad (91)$$

$$+ \frac{4\Phi}{\|\alpha \cdot \sigma^{-1}\|} \left(\sqrt{2T \ln\left(\frac{4T}{\delta}\right)} + \varepsilon T \right) \quad (92)$$

Choosing $\varepsilon = \min \left\{ T^{-1/3}, \min_{t \leq T} \Delta_{t,K}^2, \frac{1}{\sqrt{T} \min_{t \leq T} \Delta_{t,K}} \right\}$ yields an absolute regret of $R_T^{\text{abs}} = \widetilde{O} \left(T^{2/3}, \frac{\sqrt{T}}{\min_{t \leq T} \Delta_{t,K}}, \frac{1}{\min_{t \leq T} \Delta_{t,K}^4} \right)$. Therefore, when $\min_{t \leq T} \Delta_{t,K} \geq T^{-1/6}$ the absolute regret and relative regret of Alg. 1 scales similarly with a relative and absolute regret of order $\widetilde{O}(T^{2/3})$.

F ADDITIONAL EXPERIMENTS

In this section, we present in details the algorithms used in Sec. 6, the experimental protocol and additional experimental results.

F.1 Baseline Algorithms

In Sec. 6, we compare our algorithms, GLM- ε -GREEDY (Alg. 1) and ESAG (Alg. 2), to different baselines.

Algorithm 3: GLM- ε -GREEDY-ALL SAMPLES algorithm

Input: Noise parameters $\{\sigma_j\}_{j \leq J}$, confidence level δ
Parameters: exploration level ε ; number of arms to pull K ; regularization λ

 Set $\mathcal{H}_0 = \emptyset$, $\hat{\alpha} = 0$ and $w_0 = 0$
for $t = 1, \dots, T$ **do**

 Sample $Z_t \sim \text{Ber}(\varepsilon)$

 Observe evaluations for each arm $(\phi_{i,t})_{i \leq K_t}$
if $Z_t = 1$ **then**

 Pull arms in \mathcal{A}_t obtained by sampling K arms uniformly in $\{1, \dots, K_t\}$

 Observe rewards $r_{i,t}$ for all $i \in \mathcal{A}_t$
else

 Select $\mathcal{A}_t = \arg \max_i^K \langle w_t, g^{-1}(\phi_{i,t}) \rangle$

 Add sample to dataset $\mathcal{H}_t = \mathcal{H}_{t-1} \cup (\cup_{i \in \mathcal{A}_t} \{(\phi_{i,t}, r_{i,t})\})$

 Update estimators $\hat{\alpha}_{j,t}$ by solving

$$\sum_{\phi, r \in \mathcal{H}_t} r(g(\hat{\alpha}_{t,j} \cdot r) - \phi_j) - \lambda \hat{\alpha}_{t,j} = 0 \quad (93)$$

 Update weights $w_{t,j} = \hat{\alpha}_{t,j} / (\sigma_j^2) \|\hat{\alpha}_t \cdot \sigma^{-1}\|^2$
Algorithm 4: GLM-EVALBASEDUCB algorithm

Input: Noise parameters $\{\sigma_j\}_{j \leq J}$, confidence level δ

 Set $\mathcal{H}_0 = \emptyset$, $\hat{\alpha} = 0$ and $w_0 = 0$
for $t = 1, \dots, T$ **do**

 Observe evaluations for each arm $(\phi_{i,t})_{i \leq K_t}$

 Compute $\hat{\alpha}_t = \frac{\sum_{l=1}^{t-1} \sum_{i \in \mathcal{A}_l} g^{-1}(\phi_{i,l})}{\sum_{l=1}^{t-1} \sum_{i \in \mathcal{A}_l} r_{i,l}}$ and

$$\beta_t^E = \frac{2\sqrt{2J \ln(2/\delta)}}{3} + \sqrt{(t-1)K \sum_{j=1}^J \sigma_j^2 \ln(2/\delta)} \quad (94)$$

if $\|\hat{\alpha}_t\|_2 \leq \beta_t^E$ **then**

 Compute weights $w_t = \frac{\sigma^{-2} \hat{\alpha}_t}{\|\sigma^{-1} \hat{\alpha}_t\|^2}$
else

 Compute $\lambda_t > 0$ solving $\sum_{j=1, \sigma_j > 0}^J \left(\frac{\hat{\alpha}_{t,j}}{\lambda_t \sigma_j^2 + (\beta_t^E)^2} \right)^2 = \frac{1}{\beta_t^2}$

 Compute $\tilde{\alpha}_t = \left(I_J - \beta_t^2 (\lambda_t I_J + (\beta_t^E)^2 \sigma^{-2})^{-1} \sigma^{-2} \right) \hat{\alpha}_t$ and $w_t = \frac{\sigma^{-2} \tilde{\alpha}_t}{\|\sigma^{-1} \tilde{\alpha}_t\|_2}$

 Select $\mathcal{A}_t = \arg \max_i^K \langle w_t, g^{-1}(\phi_{i,t}) \rangle$ and observe rewards $r_{i,t}$ for all $i \in \mathcal{A}_t$

GLM- ε -greedy-ALL (Alg. 3). This algorithm is the s Alg. 1 and illustrate the problem with the noise correlation of App. B. This algorithm thus updates its own MLE estimator at every time step instead of step with a random exploration.

GLM-EvalBasedUCB (Alg. 4). This algorithm does not rely on an ε -greedy type of exploration but on an optimistic approach. Similar to other algorithms, it first computes an MLE estimator of α using all samples collected over time. Then it build a ‘‘confidence’’ set around the estimated $\hat{\alpha}$ and it computes weights w_t as those resulting from a worst-case choice of a parameter $\tilde{\alpha}$ in terms of error. Notice that the confidence sets are not theoretically justified because they are designed by ignoring the correlation between decisions and evaluations. As such, this algorithm has not theoretical guarantee and it should be considered as a heuristic. We provide further details in Sect. F.2.

RAND (Alg. 5). This baseline algorithm is simply selecting one random evaluators every step and computing a ranking of the arms based on this random evaluators.

GLM-Greedy. This algorithm is a variant of GLM- ε -GREEDY-ALL with $\varepsilon = 0$.

Algorithm 5: RAND algorithm

Parameters: number of arms to pull K
for $t = 1, \dots, T$ **do**

- Randomly select $j \leq J$
- Observe evaluations for each arm $(\phi_{i,t})_{i \leq K_t}$
- Select $\mathcal{A}_t = \arg \max_i^K g^{-1}(\phi_{i,t,j})$

F.2 Derivation of GLM-EvalBasedUCB

Here, we briefly present how we derive the GLM-EVALBASEDUCB algorithm. This algorithm takes a *optimistic approach* w.r.t. the set of plausible α values and compute a set of weights w_t as the set of oracle weights but for a different value of α . The first step is to compute an estimator $\tilde{\alpha}_t$ such that:

$$\tilde{\alpha}_t = \arg \max_{z \in \mathbb{R}^J: \|\hat{\alpha}_t - z\| \leq \beta_t^E} \min_{w \in \mathbb{R}^J, \langle w, z \rangle = 1} \sum_{j=1}^J w_j^2 \sigma_j^2 \quad (95)$$

where β_t^E is the width of the confidence intervals defined in Eq. (73) Solving the minimization problem in the previous equation implies that $\tilde{\alpha}_t$ is solution to the following equivalent optimization problem,

$$\min_{\|z - \hat{\alpha}_t\|_2 \leq \beta_t^E} \|\sigma^{-1} \cdot z\|_2^2 = \min_{\|u\| \leq 1} \|\sigma^{-1} \cdot (\hat{\alpha}_t + \beta_t^E u)\|_2^2 \quad (96)$$

with $\hat{\alpha}_t = \frac{\sum_{l=1}^{t-1} \sum_{i \in \mathcal{A}_l} g^{-1}(\phi_{i,l})}{\sum_{l=1}^{t-1} \sum_{i \in \mathcal{A}_l} r_{i,l}}$ the average features observed until time t . Eq. (96) is a convex problem and using KKT conditions, we distinguish two different situations:

1. If $\|\hat{\alpha}_t\| \leq \beta_t^E$ then $\min_{\|z - \hat{\alpha}_t\|_2 \leq \beta_t^E} \|\sigma^{-1} \cdot z\|_2^2 = 0$ thus the problem is ill-defined and the weights are also not defined (A necessary condition for this condition to be fulfilled is that $\|\alpha\| \leq \beta_t^E (1 + \sqrt{1 + (\beta_t^E)^2})$)
2. If $\|\hat{\alpha}_t\| > \beta_t^E$ then the solution is such that $u = -\beta_t^E (\lambda_t I_J + (\beta_t^E)^2 \text{Diag}(\sigma^{-2}))^{-1} \sigma^{-2} \cdot \hat{\alpha}_t$ where $\lambda_t > 0$ is solution to⁹:

$$\sum_{j=1, \sigma_j > 0}^J \left(\frac{\hat{\alpha}_{t,j}}{\lambda_t \sigma_j^2 + (\beta_t^E)^2} \right)^2 = \frac{1}{(\beta_t^E)^2} \quad (97)$$

Overall, $\tilde{\alpha}_t$ is defined as:

$$\tilde{\alpha}_t = \left(I_J - (\beta_t^E)^2 (\lambda_t I_J + (\beta_t^E)^2 \text{Diag}(\sigma^{-2}))^{-1} \sigma^{-2} \right) \hat{\alpha}_t, \quad (98)$$

that is to say the weights used by the algorithm is $w_t = \frac{\Sigma^{-2} \tilde{\alpha}_t}{\|\Sigma^{-1} \tilde{\alpha}_t\|_2^2}$

F.3 Additional Synthetic Experiments

We report here the regret and estimation bias for the in the logistic and linear setting for the three different regimes of $\frac{\alpha_0}{\sigma_0} \in \{0.1, 1, 10\}$ as reported in Section 6. For both the linear and logistic case we choose $K = 10$ and $K_t = 60$. For GLM- ε -GREEDY and GLM- ε -GREEDY-ALL SAMPLES, we use different values of $\varepsilon \in \{0.1T^{-1/3}, 0.1T^{-1/2}\}$. The rewards are sampled from a centered unit variance normal distribution truncated between $[0, 20]$.

Logistic and Linear Experiments. As highlighted in Section 6, for a small ratio $\frac{\alpha_0}{\sigma_0}$, the algorithms GLM- ε -GREEDY-ALL SAMPLES or GLM-EVALBASEDUCB that use all samples suffer from a linear regret compared to our algorithm Alg. 1 for $\varepsilon = \mathcal{O}(T^{-1/3})$. In addition, we observe that as the ratio $\frac{\alpha_0}{\sigma_0}$ increases the correlation effect highlighted in App. B becomes less and less relevant and GLM-EVALBASEDUCB performs particularly well.

⁹A solution λ_t always exist when there is at least one evaluator j such that $\sigma_j > 0$ because the function $f : \lambda \mapsto \sum_{j=1, \sigma_j > 0}^J \left(\frac{\hat{\alpha}_{t,j}}{\lambda \sigma_j^2 + (\beta_t^E)^2} \right)^2$ is strictly non-increasing over \mathbb{R}^+ and for $\lambda = 0$, we have $f(0) = \|\hat{\alpha}_t\|^2 / (\beta_t^E)^4 > (\beta_t^E)^{-2}$ and $\lim_{\lambda \rightarrow +\infty} f(\lambda) = 0$ so by continuity there exists a unique solution $\lambda_t > 0$.




Figure 3: $\frac{\alpha_0}{\sigma_0} = 0.1$




Figure 4: $\frac{\alpha_0}{\sigma_0} = 1$




Figure 5: $\frac{\alpha_0}{\sigma_0} = 10$

Figure 6: Regret (*Left*) and zoomed in regret (*Right*, removing RAND and EXP4.P algorithms) wrt to the oracle \mathcal{D} in Section 4 for different values of $\frac{\alpha_0}{\sigma_0}$




Figure 7: $\frac{\alpha_0}{\sigma_0} = 0.1$




Figure 8: $\frac{\alpha_0}{\sigma_0} = 1$






Figure 9: $\frac{\alpha_0}{\sigma_0} = 10$

Figure 10: Estimation error for the α parameters for different values of $\frac{\alpha_0}{\sigma_0}$


 Figure 11: Regret w.r.t. to the oracle \mathfrak{D} in Section 5 for different values of $\frac{\alpha_0}{\sigma_0}$

 Figure 12: Estimation Bias in the linear case. For ESAG, we compute the error $\|\tilde{\alpha}_t - \mathbb{E}_{r \sim \nu}(r)\alpha\|$

F.4 Second Content Review Prioritization Dataset.

In order to validate our observations of Section 6 on the content review prioritization dataset, we consider a second small dataset (D_2) of content similar to the one used in Section 6. Table ?? sums up the cumulative reward of each algorithms at $T = 2000$ steps selecting $K = 10$ out of $K_t = 200$ arms. Similarly to Table 2 linear algorithms performs best compared to the logistic algorithms with the same three algorithms getting the best performance.

F.5 Jester Experiments.

We consider the Jester dataset (Goldberg et al., 2001), which consists of joke ratings in a continuous range from -10 to 10 for a total of 100 jokes and 73421 users. We consider the same set of users and jokes as in (Riquelme

 Table 2: Cumulative badness on D_2 for $T = 2000$ steps

Alg.	Badness
RAND	60025.8
GLM- ϵ -GREEDY	105965.8
GLM- ϵ -GREEDY-ALL	105499.1
GLM-EVALBASEDUCB	105545.4
GLM-LINUCB	60347.4
GLM-ESAG	106237.5
GLM-GREEDY	105264.3
EXP4.P	88979.4
EVALBASEDUCB	138230.1
LINUCB	144009.4
ESAG	141172.1
GREEDY	106195.4

Table 3: Statistics about the evaluators. We report the R2 score of the evaluators w.r.t. the true ratings. The column $\hat{\sigma}$ is the maximum estimate standard deviation.

algo	R2	$\hat{\sigma}$
rf5-s18	0.864	0.115
dt-s50	0.448	0.174
linear	0.348	0.167
adaboost	0.440	0.048
mlp	0.917	0.110
rf	0.936	0.064
adaboost20-s60	0.309	0.086

Table 4: Average reward after $T = 6000$ steps in the Jester experiment averaged over 50 runs.

Algorithm	Average Reward $\frac{1}{T} \sum_{i=1}^T r_i$
EVALBASEDUCB	0.87
ESAG	0.86
LINUCB	0.85
GLM-EVALBASEDUCB	0.69
GLM-GREEDY	0.69
GLM-ESAGesaglogistic	0.69
GLM- ε -GREEDY	0.69
RAND	0.57
GLM-LINUCB	0.51

et al., 2018). For a subset of 40 jokes and 19181 users rating all these 40 jokes, we build evaluators as follows. We fit 7 *contextual* models to predict the ratings from features –obtained as concatenation of user and joke features– extracted via a low-rank factorization of the full matrix (of dimension 36). Ratings are normalized and transformed through the logarithmic function. We trained the following models:¹⁰

- rf5-s18: a random forest with 5 trees trained over 5 randomly selected features;
- dt-s50: a decision tree with max depth 10 trained over 50 randomly selected features;
- linear: a linear model with intercept;
- adaboost: an implementation of AdaBoost.R2 with 20 trees;
- mlp: a neural network with two hidden layers ([512, 128]) and ReLu activation;
- adaboost20-s60: AdaBoost.R2 with 20 trees trained over 60 randomly selected features.

We use the predictions of each model as inputs for our algorithms. As rewards, we use the real ratings and we add to them zero-mean Gaussian noise (standard deviation is 0.5). The resulting problem is thus misspecified since the evaluators are not perfect, as shown in Tab. 3.

We estimate the parameters of the algorithms by computing statistics about the relationship between true ratings and predicted ones. For ε -greedy, we use the value $T^{1/3}$. From Tab. 4, we can see that linear algorithms outperform the logistic algorithms. EVALBASEDUCB and ESAG behave similarly and perform better than LINUCB.

¹⁰We used the implementations provided by scikit-learn (Pedregosa et al., 2011).

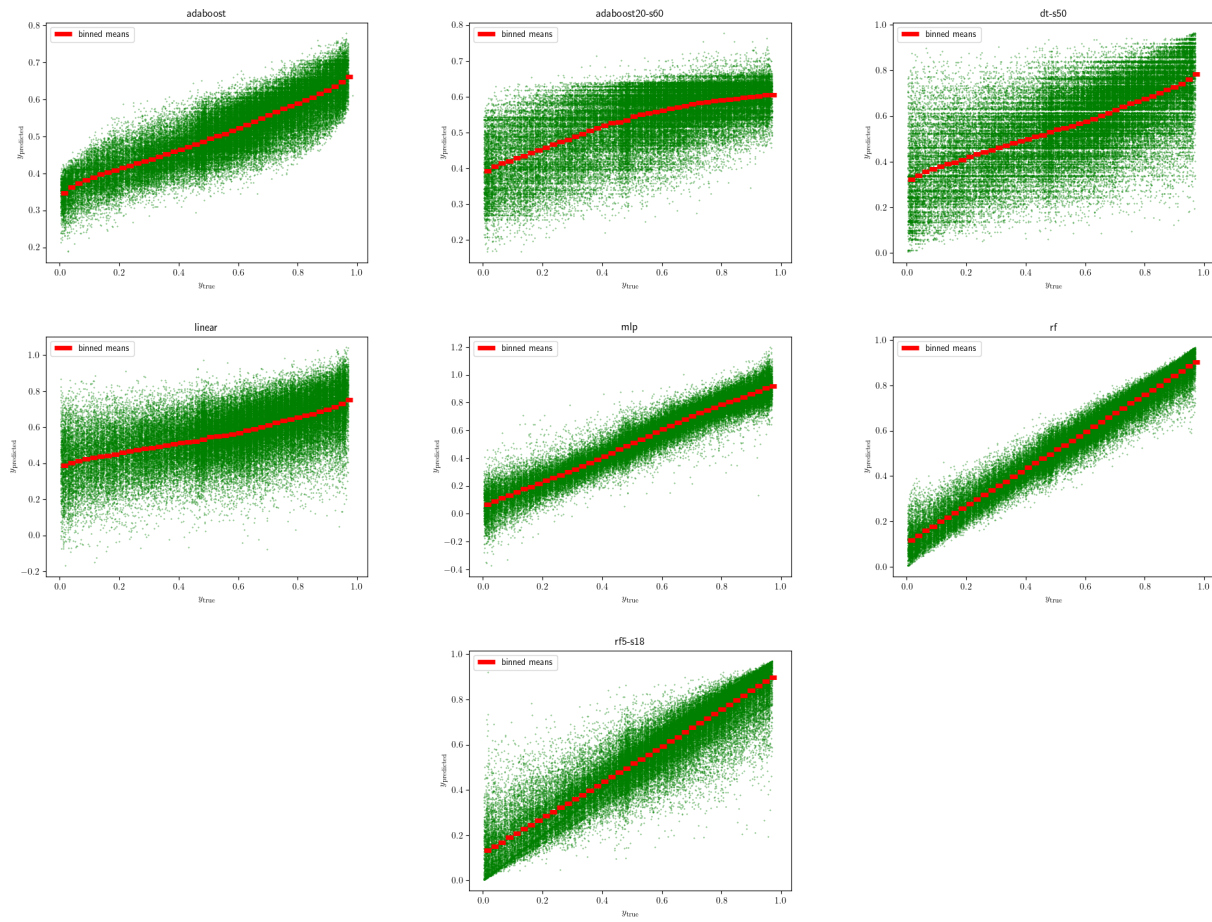


Figure 13: We report the predictions of the evaluators as a function of the true ratings (for clarity we reported only 10% of the samples). We also report the average value over a discretization of the ratings into 40 bins.