
Online Learning in Contextual Second-Price Pay-Per-Click Auctions

Mengxiao Zhang
University of Southern California

Haipeng Luo
University of Southern California

Abstract

We study online learning in contextual pay-per-click auctions where at each of the T rounds, the learner receives some context along with a set of ads and needs to make an estimate on their click-through rate (CTR) in order to run a second-price pay-per-click auction. The learner’s goal is to minimize her regret, defined as the gap between her total revenue and that of an oracle strategy that always makes perfect CTR predictions. We first show that \sqrt{T} -regret is obtainable via a computationally inefficient algorithm and that it is unavoidable since our algorithm is no easier than the classical multi-armed bandit problem. A by-product of our results is a \sqrt{T} -regret bound for the simpler non-contextual setting, improving upon a recent work of [Feng et al. \(2023\)](#) by removing the inverse CTR dependency that could be arbitrarily large. Then, borrowing ideas from recent advances on efficient contextual bandit algorithms, we develop two practically efficient contextual auction algorithms: the first one uses the exponential weight scheme with optimistic square errors and maintains the same \sqrt{T} -regret bound, while the second one reduces the problem to online regression via a simple epsilon-greedy strategy, albeit with a worse regret bound. Finally, we conduct experiments on a synthetic dataset to showcase the effectiveness and superior performance of our algorithms.

1 Introduction

The rapid growth of internet-based advertising has led to increasing reliance on online auctions to efficiently

allocate advertisement slots. Pay-per-click (PPC) auctions, in particular, have become a prevalent mechanism in the digital advertising landscape, where advertisers are charged based on the number of clicks on their ads. In these auctions, a platform’s primary goal is to deliver the relevant experience while maximizing value to the advertiser and the publisher.

Existing literature on online PPC auctions mostly focuses on the non-contextual setup ([Devanur and Kakade, 2009](#); [Buccapatnam et al., 2014](#); [Babaioff et al., 2015](#); [Feng et al., 2023](#)), where the same set of ads repeatedly participates in an auction, each with a click-through rate (CTR) fixed over time. In reality, however, the set of participating ads and their CTR vary in each auction based on the user query, user preferences/history, ad relevance, and other contextual information. To tackle such practical scenarios, in this work, we consider online *contextual* PPC auctions with unknown CTRs and study how the auction platform can leverage the contextual information to make a revenue close to that of an oracle strategy that runs a second-price auction with perfect knowledge of the CTRs.

More concretely, we formulate this problem as an online learning problem over T rounds. At each round, the auction platform (learner) first observes some context and a set of participating ads, and then makes an estimate for the CTR of each ad without seeing their bid. Afterwards, a second-price PPC auction, a truthful and widely used mechanism ([Aggarwal et al., 2006](#)), is run: each ad is assigned a score equal to the product of its bid and its estimated CTR, and the ad with the highest score wins the auction with the payment-per-click equal to the *critical price* (that is, the lowest price that still guarantees a win). The learner’s goal is to decide the estimated CTRs in a way so that the total revenue is close to what one would have received if the CTR estimations were always perfect — we call the gap between them the *regret* of the learner. To make sublinear regret possible, we make a standard realizability assumption that the learner is given access to a CTR predictor class that contains a perfect and unknown predictor, but we do not make any assumption

on how the contexts and the bids are generated — they can even be maliciously decided by an adversary.

To our knowledge, our work is the first to consider online learning for such contextual PPC auctions. However, similar to its non-contextual version, the problem has deep connections with the heavily studied (contextual) multi-armed bandit problem (Lai et al., 1985; Auer et al., 2002). In particular, because we only observe feedback on the winner, balancing exploration and exploitation, the infamous dilemma originated from multi-armed bandits, is also the key challenge of our problem. What makes our problem even more difficult is that we cannot explore/exploit whichever ad/arm we want but instead have to do so implicitly via setting the estimated CTRs, which themselves by definition also affect the reward of each ad/arm. Despite these difficulties, by extending ideas from contextual bandits and making careful adjustment tailored to our problem structure, we obtain a series of positive results both theoretically and empirically. Specifically, our contributions are:

1. As the first step, in Section 3, we provide a characterization of the optimal regret of our problem via a computationally inefficient algorithm and a simple lower-bound argument showing that our problem is no easier than multi-armed bandits. Our algorithm is based on the well-known exponential weight strategy, with an inverse propensity score (IPS) weighted loss estimator that is similar to the Exp4 algorithm (Auer et al., 2002) for contextual bandits. For a finite predictor class \mathcal{F} , our algorithm achieves $\mathcal{O}(\sqrt{NT} \log |\mathcal{F}|)$ near-optimal regret with N being the number of ads. Notably, our result immediately implies $\tilde{\mathcal{O}}(N\sqrt{T})$ regret for the non-contextual setting, improving upon the $\tilde{\mathcal{O}}(\sum_{i=1}^N \frac{1}{\rho_i} \sqrt{T})$ regret of a recent work by Feng et al. (2023) where $\rho_i \in [0, 1]$ is the CTR of ad i (whose inverse could be arbitrarily large).
2. To address the computational inefficiency, we then develop two practically efficient algorithms, taking inspiration from recent developments in designing efficient contextual bandit algorithms. Our first approach (Section 4) replaces the IPS estimator with an *optimistic square error estimator* that is efficiently computable and shares similar ideas with the Feel-Good Thompson Sampling algorithm of Zhang (2022). The resulting algorithm not only still enjoys a \sqrt{T} -regret bound, but also admits an efficient (approximate) implementation by applying stochastic gradient Langevin dynamics (SGLD) (Welling and Teh, 2011).
3. Our second approach (Section 5) follows another trend of recent studies that reduce contextual ban-

dit to an easier regression problem where efficient algorithms already exist (Foster and Rakhlin, 2020; Foster and Krishnamurthy, 2021; Foster et al., 2021). We adopt the general framework of (Foster et al., 2021) in attempt to find the optimal reduction from our contextual auction problem to online regression. We provide a simple solution based on an epsilon-greedy strategy that is efficiently implementable and achieves $\mathcal{O}(T^{\frac{2}{3}}(N\mathbf{Reg}_{\text{Sq}})^{\frac{1}{3}})$ regret with \mathbf{Reg}_{Sq} being the regret of the regression problem. While this method leads to a worse regret bound, we conjecture that it might be the best one can do using this reduction approach due to its distinct property: unlike the last two algorithms, this approach does not use the bid information from previous rounds to decide the next CTR estimates (which sometimes might be desirable in practice).

4. In Section 6, we also test our two efficient algorithms on a synthetic dataset, demonstrating their superior performance against several baseline algorithms.

Related works. One line of closely related work is the study on the non-contextual counterpart, such as (Devanur and Kakade, 2009; Babaioff et al., 2014, 2015; Feng et al., 2023), all of which consider designing a globally truthful no-regret mechanism so that bidders are incentivized to bid their true valuation throughout all T rounds. (Devanur and Kakade, 2009; Babaioff et al., 2014) achieves so via an explore-then-commit strategy with $\mathcal{O}(N^{\frac{1}{3}}T^{\frac{2}{3}})$ regret, which is shown to be optimal for globally truthful mechanisms. (Babaioff et al., 2015) further considers the setting where each bidder’s bid is stochastic and designs a randomized auction that enjoys $\mathcal{O}(\sqrt{T})$ regret and is globally truthful only in expectation. Further extensions include multi-slot mechanism design (Gatti et al., 2012), valuations unknown to the bidders (Kandasamy et al., 2023), and others.

A recent work by Feng et al. (2023) considers the myopic bidder setting with adversarial bids and designs a UCB-based algorithm which leads to a per-round truthful auction and achieves $\tilde{\mathcal{O}}(\sum_{i=1}^N \frac{1}{\rho_i} \sqrt{T})$ regret. They also consider the fixed valuation setting and design a globally truthful auction with $-\Omega(T)$ regret when there exists a time-independent constant gap between the winner and the runner up. More recently, (Xu et al., 2023) generalizes this non-contextual setting to the stochastic context setting and derives an ε -greedy-based algorithm achieving $\tilde{\mathcal{O}}(N^{\frac{4}{3}}T^{\frac{2}{3}} + \frac{1}{\alpha^2}T^{\frac{1}{3}}N^{\frac{2}{3}})$ regret under α -rational bidders. We note that in our setting, similar to (Feng et al., 2023), the learner is allowed to adjust the auction on the fly based on previous observation (but not the bids for the current round), which makes it truthful per round but not necessarily

globally. However, since we allow *adversarial* contexts that might not be manipulatable by the bidders, we do not find global truthfulness a meaningful requirement for our setting (see Section 2 for more discussion).

Another line of work on online learning in auction considers designing auto-bidding algorithms under different types of auction mechanisms (such as first-price auction (Wang et al., 2023; Han et al., 2020), second-price auction (Balseiro and Gur, 2019; Balseiro et al., 2023), and core auction (Gaitonde et al., 2022)) or different types of resource and return-on-investment constraints (Balseiro and Gur, 2019; Balseiro et al., 2021; Lucier et al., 2023). Since we allow adversarial bids, our results hold whether the platform is facing these auto-bidding algorithms or not.

As mentioned, our problem bares some similarity with the heavily-studied contextual bandit problem but is generally more difficult due to the fact that both the way to select an arm and its reward are determined through the estimated CTRs. (Auer et al., 2002) proposes the first contextual bandit algorithm Exp4 which achieves the optimal regret but is computationally inefficient. Given the vast number of real-life applications of contextual bandits, starting from (Langford and Zhang, 2007), many studies focus on developing practically efficient algorithm under reasonable computational assumptions, such as access to classification/regression oracles (Dudik et al., 2011; Agarwal et al., 2014; Foster et al., 2018; Foster and Rakhlin, 2020; Xu and Zeevi, 2020; Foster and Krishnamurthy, 2021; Simchi-Levi and Xu, 2022; Zhu and Mineiro, 2022), or ability to sample from a certain distribution using Markov chain Monte Carlo methods (Zhang, 2022). As mentioned, we follow and extend the ideas of these two approaches to design efficient contextual auction algorithms.

2 Notations and Problem Setup

General notations. For a positive integer m , we use $[m]$ to denote the set $\{1, 2, \dots, m\}$, Δ_m to denote the $(m-1)$ -dimensional simplex, and $\mathbb{R}_{\geq 0}^m$ to denote the set of m -dimensional vectors with non-negative entries. For a vector $v \in \mathbb{R}^m$, we use $\max_{i \in [m]} v_i$ and $\text{smax}_{i \in [m]} v_i$ to denote the largest and the second largest entry of v respectively, and $\arg\max_{i \in [m]} v_i$ and $\text{argsmax}_{i \in [m]} v_i$ to denote their index.¹ We also use $\mathbf{1}$ to denote the all-one vector and e_i to denote the i -th standard basis vector (both with an appropriate dimension depending on the context).

¹In this definition, we break ties by an arbitrary fixed deterministic rule. Note that when there is a tie, it is possible that the “second largest” entry in fact has the largest value.

Problem setup. The formal setup of the contextual second-price pay-per-click auctions we consider is as follows. An ad auction platform (called the learner) sequentially interacts with some bidders/advertisers for T rounds. At each round $t \in [T]$:

1. The learner first observes a context x_t from some context space \mathcal{X} and a set of N_t bidders participating in the current campaign with their ads. Here, the context x_t encodes any available information about the current campaign, such as the user query (in the case of a search engine) and features of the participating ads. We denote the maximum number of bidders by $N = \max_t N_t$.
2. What the learner needs to decide is an estimated CTR vector: $\tilde{\rho}_t \in [0, 1]^{N_t}$, with each entry $\tilde{\rho}_{t,i}$ being an estimation of the true and unknown CTR $\rho_{t,i} \in [0, 1]$ of ad i under context x_t .
3. Simultaneously, each bidder i decides their own bid $b_{t,i} \in [0, 1]$ (without knowing $\tilde{\rho}_t$ or ρ_t).
4. A second-price pay-per-click auction (Aggarwal et al., 2006) is then run: the winner of this campaign is $i_t = \arg\max_{i \in [N_t]} b_{t,i} \tilde{\rho}_{t,i}$, that is, the bidder with the highest estimated expected cost per impression; the payment per click of the winner is $d_t = \frac{b_{t,j_t} \tilde{\rho}_{t,j_t}}{\tilde{\rho}_{t,i_t}}$ where $j_t = \text{argsmax}_{i \in [N_t]} b_{t,i} \tilde{\rho}_{t,i}$ is the runner-up (note $d_t \leq b_{t,i_t}$ by definition, so the winner never pays more than their bid); the winner’s ad is then displayed, and is clicked with probability ρ_{t,i_t} .
5. The feedback of the learner includes all the bids $b_{t,1}, \dots, b_{t,N_t}$ and a binary variable $c_t \in \{0, 1\}$, which is 1 if the displayed ad is clicked and 0 otherwise (by definition, c_t is a Bernoulli random variable with mean ρ_{t,i_t}). The payment that the learner receives in the end is thus $c_t d_t$.

The goal of the learner is to minimize her regret, which measures how much more she would have received if she had perfect knowledge of the true CTRs and set $\tilde{\rho}_t = \rho_t$ all the time. More concretely, since in this imaginary situation with $\tilde{\rho}_t = \rho_t$, the expected payment received at round t is $\mathbb{E}[c_t d_t] = \rho_{t,i_t} \cdot \frac{b_{t,j_t} \rho_{t,j_t}}{\rho_{t,i_t}} = b_{t,j_t} \rho_{t,j_t} = \text{smax}_{i \in [N_t]} b_{t,i} \rho_{t,i}$, the regret is formally defined as

$$\text{Reg} \triangleq \sum_{t=1}^T \text{smax}_{i \in [N_t]} b_{t,i} \rho_{t,i} - \sum_{t=1}^T c_t d_t. \quad (1)$$

Achieving sublinear (in T) regret thus implies that the learner is on average performing almost as well as the ideal benchmark even though she does not know the true CTRs. This is a strong requirement that intuitively is possible only if there is some connection

between the context x_t and the true CTR vector ρ_t so that over time the learner can gradually improve her estimation $\tilde{\rho}_t$ based on prior observations. To this end, we make the following realizable assumption that is analogous to those usually made in the contextual bandit literature (see for example (Foster and Rakhlin, 2020; Foster and Krishnamurthy, 2021; Foster et al., 2021; Zhu and Mineiro, 2022; Simchi-Levi and Xu, 2022)).

Assumption 1 (Realizability). *A function class $\mathcal{F} = \{f : \mathcal{X} \times [N] \rightarrow [0, 1]\}$, given to the learner, contains a perfect (but unknown) CTR predictor f^* such that $f^*(x_t, i) = \rho_{t,i}$ for all t and i .*

Our goal is to derive algorithms with a regret bound that is sublinear in T and polynomial in N and some common complexity measure of \mathcal{F} , such as $\ln|\mathcal{F}|$ for the case of a finite class or the number of parameters for the case of a parametric class. In practice, such a function class \mathcal{F} can be any common machine learning model (for example, a linear class or a class of neural nets), as long as it is believed to be complex enough to predict the true CTRs reasonably well for the specific problem on hand.

We emphasize that we do not make any assumptions on how the contexts and the bids are generated. In particular, they can even be generated by an adaptive adversary who knows the learner’s algorithm and observes her decisions in previous rounds. This generality allows us to handle strategic bidders.

Instantaneous truthfulness versus global truthfulness. The reason to consider such a second-price auction is that it is well-known to be *instantaneous truthful*, that is, looking only at a particular round, each bidder is incentivized to use their true valuation as their bid (Aggarwal et al., 2006). Prior work such as (Devanur and Kakade, 2009; Feng et al., 2023) considers designing a (non-contextual) auction mechanism that is globally truthful so that even looking at all rounds together, each bidder is still incentivized to bid their true valuation. However, this is often achieved by an explore-then-commit strategy which fixes the auction mechanism (such as the estimated CTRs) after a certain period of pure exploration. Such strategies not only are impractical but also would not make sense at all in our contextual setting with potentially adversarial contexts. In fact, exactly because the contexts in our problem are partially decided by exogenous factors such as the user queries (that are not manipulable by the bidders), we do not find global truthfulness a reasonable requirement for our problem. We thus stick with only instantaneous truthfulness and allow the learner to change the auction on the fly based on the context.

Comparisons to contextual bandits. One key difference between our problem and the well-studied contextual bandit problem is that we *cannot* freely pick an “arm” (corresponding to the winner in our context), but have to do so via proposing a particular estimated CTR vector $\tilde{\rho}_t$, which itself affects the “reward” of an arm according to the definition of the winner’s payment, making the problem more complicated. On the other hand, what is common in both problems is that a good balance between exploration and exploitation is clearly necessary due to the limited feedback on the selected arm/winner only. Because of these similarities and differences, in what follows we will show that some ideas from the contextual bandit literature are readily applicable to our problem, while others require more adjustments.

3 Achieving $\mathcal{O}(\sqrt{T})$ Regret Inefficiently

In this section, as the first step, we show how adopting the idea of a classical contextual bandit algorithm called Exp4 (Auer et al., 2002) leads to, for example, $\mathcal{O}(\sqrt{NT \log |\mathcal{F}|})$ regret for a finite predictor class \mathcal{F} . The algorithm is computationally inefficient, but it illustrates that \sqrt{T} -type regret is obtainable for this problem information-theoretically.

The idea of Exp4 is to maintain a distribution q_t over all predictors in \mathcal{F} , defined via a classical exponential weight scheme: $q_{t,f} \propto \exp(-\eta \sum_{s < t} \hat{\ell}_{s,f})$ where $\hat{\ell}_{s,f}$ is an estimator for some loss $\ell_{s,f} \in [0, 1]$ of predictor f at round s . With such a distribution, the algorithm simply samples a predictor f_t from q_t for round t and follows its suggestion.

In our problem, “following f_t ’s suggestion” means setting the estimated CTR $\tilde{\rho}_{t,i}$ directly as $f_t(x_t, i)$ for each $i \in [N_t]$. The loss $\ell_{t,f}$ of predictor f at round t is intuitively the negative expected payment if one follows f ’s suggestion. To ensure a range of $[0, 1]$, we shift it by 1, leading to $\ell_{t,f} = 1 - \rho_{t,i_{t,f}} \frac{\text{smax}_j b_{t,j} f(x_t, j)}{f(x_t, i_{t,f})}$ where $i_{t,f} = \text{argmax}_i b_{t,i} f(x_t, i)$. It remains to construct the loss estimator $\hat{\ell}_{t,f}$ based on the learner’s observations b_t and c_t . Even though our loss structure is quite different from contextual bandits, we find the standard inverse propensity score weighting still applicable, leading to a natural loss estimator defined in the (IPS) option of Eq. (2). It is clear that $\hat{\ell}_{t,f}$ is unbiased since $\mathbb{E}[\mathbb{1}\{i_t = i_{t,f}\}] = p_{t,i_{t,f}}$ and $\mathbb{E}[c_t] = \rho_{t,i_t}$ implies $\mathbb{E}[\hat{\ell}_{t,f}] = \ell_{t,f}$. See Algorithm 1 for the complete pseudocode. Following standard analysis of Exp4, we show the following regret bound when \mathcal{F} is finite (whose proof is deferred to Appendix A).

Theorem 3.1. *Algorithm 1 with learning rate $\eta = \sqrt{\frac{\log |\mathcal{F}|}{\sum_{t=1}^T N_t}}$ and (IPS) estimators guarantees $\mathbb{E}[\text{Reg}] =$*

Algorithm 1 Exponential Weights for Contextual AuctionInput: learning rate $\eta > 0$ and CTR predictor class \mathcal{F} **for** $t = 1, 2, \dots, T$ **do**Sample a function f_t from q_t , a distribution over \mathcal{F} defined via $q_{t,f} \propto \exp(-\eta \sum_{s < t} \widehat{\ell}_{s,f})$.Receive context x_t and the set of N_t bidders.Set the estimated CTR to be $\widetilde{\rho}_{t,i} = f_t(x_t, i)$ for all $i \in [N_t]$ and receive bids $b_t \in [0, 1]^{N_t}$.Select the bidder $i_t = \operatorname{argmax}_{i \in [N_t]} b_{t,i} \widetilde{\rho}_{t,i}$ as the winner, with payment per click $d_t = \frac{b_{t,j_t} \widetilde{\rho}_{t,j_t}}{\widetilde{\rho}_{t,i_t}}$ where $j_t = \operatorname{argsmax}_{i \in [N_t]} b_{t,i} \widetilde{\rho}_{t,i}$ is the runner up.Receive feedback $c_t = \mathbf{1}\{\text{ad } i_t \text{ is clicked}\}$ and payment $c_t d_t$.Define loss estimator $\widehat{\ell}_{t,f}$ for each $f \in \mathcal{F}$ as

$$\widehat{\ell}_{t,f} = \begin{cases} \frac{\mathbf{1}\{i_t = \operatorname{argmax}_{i \in [N_t]} b_{t,i} f(x_t, i)\}}{p_{t,i_t}} \left(1 - c_t \cdot \frac{\operatorname{smax}_{j \in [N_t]} b_{t,j} f(x_t, j)}{f(x_t, i_t)}\right), & \text{(IPS)} \\ \frac{1}{4\eta} (f(x_t, i_t) - c_t)^2 - \operatorname{smax}_{j \in [N_t]} b_{t,j} f(x_t, j), & \text{(OptSq)} \end{cases} \quad (2)$$

and $p_{t,i} = \Pr_{f_t \sim q_t}\{i_t = i\}$ is the probability of i being the winner.

$$\mathcal{O}\left(\sqrt{\sum_{t=1}^T N_t \log |\mathcal{F}|}\right) = \mathcal{O}(\sqrt{NT \log |\mathcal{F}|}).^2$$

Now, we argue that such \sqrt{T} dependence in the regret is unavoidable. To see this, we first discuss the implication of [Theorem 3.1](#) for a special non-contextual case where for each i , $\rho_{t,i} = \rho_i$ stays the same for all t (so the context x_t plays no role in predicting the CTRs), which is essentially the setting of ([Devanur and Kakade, 2009](#)). In this case, a class of constant predictors (that ignore the context input) $\mathcal{F} = \{(x, i) \rightarrow \theta_i : \theta \in [0, 1]^N\}$ trivially satisfies [Assumption 1](#). By applying [Algorithm 1](#) to a discretized version this class: $\widehat{\mathcal{F}} = \{(x, i) \rightarrow \theta_i : \theta \in \{0, \frac{1}{T}, \frac{2}{T}, \dots, 1\}^N\}$, which has a size of $\mathcal{O}(T^N)$ and can approximate \mathcal{F} up to error $\mathcal{O}(1/T)$, we immediately obtain the following corollary (see [Appendix A](#) for the proof).

Corollary 3.1. *In the non-contextual setting described above, [Algorithm 1](#) with predictor class $\widehat{\mathcal{F}}$, learning rate $\eta = \sqrt{\frac{\log T}{T}}$, and (IPS) estimators guarantees $\mathbb{E}[\text{Reg}] \leq \mathcal{O}(N\sqrt{T \log T})$.³*

Notably, ([Feng et al., 2023](#)) considers the same non-contextual setting and designs a UCB-based algorithm with $\mathcal{O}(\sum_{i=1}^N \frac{1}{\rho_i} \sqrt{T \log(NT)})$ regret. Our result thus strictly improves upon theirs by removing the dependence on $\frac{1}{\rho_i}$, which can be arbitrarily large as long as there exists an ad with a very low CTR.

Next, we show in the following theorem that the \sqrt{T}

²For simplicity, we set the learning rate in terms of the unknown quantity $\sum_{t=1}^T N_t$, but this can be easily resolved by applying a standard doubling trick. The same holds for other results in this work.

³We note that this is not a contradiction with the $\Omega(T^{2/3})$ lower bound of ([Devanur and Kakade, 2009](#)) since they insist on global truthfulness while we do not; see related discussions in [Section 2](#).

dependence is unavoidable even in this non-contextual case.

Theorem 3.2. *In the non-contextual setting described above with $T \geq N \geq 3$, for any algorithm \mathcal{A} , there exists a sequence of bids $\{b_t\}_{t=1}^T$ and CTRs $\{\rho_i\}_{i=1}^N$ such that the expected regret suffered by \mathcal{A} is at least $\Omega(\sqrt{NT})$.*

The formal proof of [Theorem 3.2](#) is deferred to [Appendix A](#). The general idea of the proof is to transform the auction problem in this non-contextual setting to a hard instance of the classical multi-armed bandit problem, whose minimax regret is well-known to be of order $\Theta(\sqrt{NT})$. To see this, consider the case where $N_t = N$ and $b_t = \mathbf{1}$ for all t , and $\rho_i = 1/2$ for all i except for two uniformly at random chosen ads, whose CTR are always $1/2 + \Theta(\sqrt{N/T})$. Obverse that: 1) the oracle strategy with $\widetilde{\rho}_{t,i} = \rho_i$ receives expected payment $1/2 + \Theta(\sqrt{N/T})$ per round; 2) no matter which ad (arm) is selected as the winner by the learner, since the payment per click can not exceed the bid 1 according to the auction design, the expected payment of the learner is at most the CTR of the chosen winner (arm); 3) the feedback to the learner at each round is only a Bernoulli random variable with mean being the CTR of the chosen winner (arm). These facts together show that the instance above is no easier than an N -armed bandit problem with reward means being the same as the CTR configuration above, which is well-known to incur $\Omega(\sqrt{NT})$ regret.

While [Algorithm 1](#) with (IPS) achieves optimal regret guarantee, as mentioned, the caveat of [Algorithm 1](#) with (IPS) loss estimators is its computational inefficiency. Indeed, the complicated form of the estimator, in particular the p_{t,i_t} term, makes it difficult to compute efficiently. In the next two sections, we address this issue via two different approaches.

4 An Efficient Approach via Optimistic Squared Errors

Our first approach to derive a practically efficient algorithm is to replace the (IPS) estimator $\widehat{\ell}_{t,f}$ in [Algorithm 1](#) with an optimistic squared error that not only is easy to compute itself, but also allows efficient gradient computation so that one can directly apply stochastic gradient Langevin dynamics (SGLD) ([Welling and Teh, 2011](#)) to approximately sample from q_t . The optimistic squared error is defined in (OptSq) of [Eq. \(2\)](#), whose design is largely inspired by the Feel-Good Thompson Sampling algorithm of a recent work ([Zhang, 2022](#)) for contextual bandits. It contains a natural squared error term $(f(x_t, i_t) - c_t)^2$ (scaled by $1/4\eta$), which measures how good f is in predicting the CTR of the selected winner i_t under context x_t . However, the squared error itself is insufficient in encouraging exploration; see ([Zhang, 2022](#)) for detailed discussions even for the easier contextual bandit problem. To ensure exploration, we propose to subtract from the squared error a “bonus” term $\text{smax}_j b_{t,j} f(x_t, j)$, which is the expected payment we would receive if f ’s predictions were exactly the true CTRs (and we followed these predictions in setting $\widetilde{\theta}_t$). This serves as a form of optimism: the larger the payment in f ’s predicted world, the smaller its optimistic squared error $\widehat{\ell}_{t,f}$, and consequently the larger weight f gets in the sampling distribution q_t to ensure that it is properly explored. The necessity of such optimism is not only for regret analysis, but also verified in our experiments (see [Section 6](#)).

Efficient (approximate) implementation. The advantage of the (OptSq) estimator over (IPS) is that it enables efficient sampling for a parametrized predictor class. Specifically, consider a parametrized and differentiable predictor class $\mathcal{F} = \{f_\theta : \theta \in \Theta\}$ for some d -dimensional parameter space $\Theta \subseteq \mathbb{R}^d$. Then the gradient of $\widehat{\ell}_{t,f_\theta}$ with respect to θ (at any differentiable point) is

$$\frac{\partial \widehat{\ell}_{t,f_\theta}}{\partial \theta} = \frac{1}{2\eta} (f_\theta(x_t, i_t) - c_t) \cdot \frac{\partial f_\theta(x_t, i_t)}{\partial \theta} - b_{t,k} \frac{\partial f_\theta(x_t, k)}{\partial \theta}$$

where $k = \text{argmax}_j b_{t,j} f_\theta(x_t, j)$. We can then apply SGLD to approximately sample from q_t as follows: start from a random initialization $\theta \in \Theta$, then repeatedly sample uniformly at random an $s \in [t-1]$ and update

$$\theta \leftarrow \theta - \alpha \eta \frac{\partial \widehat{\ell}_{s,f_\theta}}{\partial \theta} + \sqrt{\frac{2\alpha}{t}} \varepsilon \quad (3)$$

where ε is a fresh d -dimensional standard Gaussian noise and α is a step size. Even though the existing theory for SGLD does not necessarily tell us how many steps are needed to ensure an accurate enough sample

(which is already the case for contextual bandits ([Zhang, 2022](#))), this is clearly a practically efficient gradient-based method in line with standard machine learning practice.

Regret guarantees. Next, we show that the resulting algorithm also enjoys similar regret guarantees as the version with (IPS). We start by the following general theorem, whose analysis extends the idea of ([Zhang, 2022](#)) to our more complicated loss structure and the new bonus term that is tailored to second-price auctions and importantly involves possibly adversarial bids b_t (see [Appendix B](#)).

Theorem 4.1. *Algorithm 1 with learning rate $\eta \leq 1$ and (OptSq) estimators ensures $\mathbb{E}[\text{Reg}] \leq \frac{Z_T}{\eta} + \mathcal{O}\left(\eta \sum_{t=1}^T N_t\right)$, where $Z_T = -\mathbb{E}\left[\log \mathbb{E}_{f \sim q_1} \exp\left(-\eta \sum_{t=1}^T \left(\widehat{\ell}_{t,f} - \widehat{\ell}_{t,f^*}\right)\right)\right]$.*

Here, Z_T should be treated as some complexity measure of the predictor class \mathcal{F} . As a concrete example, when \mathcal{F} is finite, we have the following corollary that exactly matches [Theorem 3.1](#).

Corollary 4.1. *For a finite \mathcal{F} , we have $Z_T \leq \log |\mathcal{F}|$, and thus the regret bound in [Theorem 4.1](#) becomes $\mathbb{E}[\text{Reg}] = \mathcal{O}\left(\log |\mathcal{F}| + \sqrt{\sum_{t=1}^T N_t \log |\mathcal{F}|\right)}$ after picking the optimal η .*

Proof. Since q_1 is uniform, we have

$$\begin{aligned} Z_T &= -\mathbb{E}\left[\log \sum_{f \in \mathcal{F}} \frac{1}{|\mathcal{F}|} \exp\left(-\eta \sum_{t=1}^T \left(\widehat{\ell}_{t,f} - \widehat{\ell}_{t,f^*}\right)\right)\right] \\ &\leq -\mathbb{E}\left[\log \frac{1}{|\mathcal{F}|} \exp\left(-\eta \sum_{t=1}^T \left(\widehat{\ell}_{t,f^*} - \widehat{\ell}_{t,f^*}\right)\right)\right] \\ &= \log |\mathcal{F}|, \end{aligned}$$

finishing the proof. \square

As another example, we consider a Lipschitz class in the next result (proof deferred to [Appendix B](#)).

Corollary 4.2. *For some constants $\alpha, B \geq 0$ and a parametrized class*

$$\begin{aligned} \mathcal{F} &= \{f_\theta : \theta \in \Theta \subseteq [-B, B]^d, \\ &\quad f_\theta \text{ is } \alpha\text{-Lipschitz with respect to } \|\cdot\|_\infty \text{ in } \theta\}, \end{aligned}$$

we have $Z_T = \mathcal{O}(d \log(\alpha B T))$, and thus the regret bound in [Theorem 4.1](#) becomes

$$\mathbb{E}[\text{Reg}] = \mathcal{O}\left(\sqrt{d \log(\alpha B T) \sum_{t=1}^T N_t}\right)$$

after picking the optimal η .

5 An Efficient Approach via a Regression Oracle

The second approach we take to derive a practically efficient algorithm follows a different trend of recent studies that reduce contextual bandits to some online regression problem and only access the predictor class \mathcal{F} through an efficient regression oracle (Foster and Rakhlin, 2020; Foster and Krishnamurthy, 2021; Zhu and Mineiro, 2022). More concretely, we assume access to an online regression oracle AlgSq which follows the following protocol: at each round $t \in [T]$, the oracle selects randomly a function $f_t \in \mathcal{F}$, then it receives a tuple $(x_t, i_t, c_t) \in \mathcal{X} \times [N] \times [0, 1]$, potentially generated by an adaptive adversary, and suffers a squared error $(f_t(x_t, i_t) - c_t)^2$.⁴ We assume that the oracle ensures some regret bound against the best predictor in \mathcal{F} :

Assumption 2 (Bounded Squared Error Regression Regret). *The regression oracle AlgSq guarantees that for any (potentially adaptively chosen) sequence $\{(x_t, i_t, c_t)\}_{t \in [T]}$, the following is bounded by \mathbf{Reg}_{Sq} :*

$$\mathbb{E} \left[\sum_{t=1}^T (f_t(x_t, i_t) - c_t)^2 - \inf_{f \in \mathcal{F}} \sum_{t=1}^T (f(x_t, i_t) - c_t)^2 \right].$$

There are many examples of such a regression oracle. For instance, when \mathcal{F} is a d -dimensional linear class, Online Newton Step (ONS) (Hazan et al., 2007) achieves $\mathbf{Reg}_{\text{Sq}} = \mathcal{O}(d \log T)$. We refer the reader to (Foster and Rakhlin, 2020) for other examples, and point out that in practice (and in our experiments), such an oracle is often implemented by a simple gradient-based method. The important point is that the regression problem does not require a balance between exploration and exploitation, and is thus generally easier than the original problem and is better studied with known efficient algorithms. This also means that to reduce the original problem to its regression counterpart, the key is to figure out an appropriate exploration strategy.

Foster et al. (2021) provide a general framework to find the optimal exploration strategy using a concept called *Decision-Estimation Coefficient* (DEC). For our problem, it boils down to understanding the following quantity

$$\text{dec}_\gamma(\hat{\rho}) = \min_Q \max_{\rho \in [0,1]^N, b \in [0,1]^N} \mathbb{E}_{\hat{\rho} \in Q} \left[\underbrace{\left(\text{smax}_{i \in [N]} b_i \rho_i - \rho_{i^*} \frac{\text{smax}_{j \in [N]} b_j \tilde{\rho}_j}{\tilde{\rho}_{i^*}} \right)}_{\text{per-round regret}} - \gamma \underbrace{(\rho_{i^*} - \hat{\rho}_{i^*})^2}_{\text{squared error}} \right] \quad (4)$$

⁴For simplicity, we consider a *proper* oracle here, but one can also use an *improper* oracle that makes a prediction after seeing (x_t, i_t) , not necessarily following some $f_t \in \mathcal{F}$, as in (Foster and Rakhlin, 2020).

where the min is over all possible distributions Q over the set $[0, 1]^N$, $i^* = \text{argmax}_i b_i \tilde{\rho}_i$ is the winner according to $\tilde{\rho}$, $\gamma > 0$ is some coefficient, and $\hat{\rho} \in [0, 1]^N$ is a given CTR prediction (provided by the regression oracle). In words, $\text{dec}_\gamma(\hat{\rho})$ measures how small the gap can be made between the per-round regret of the learner using a randomized strategy Q and the squared prediction error of the oracle in the worst case. Understanding this quantity provides both an algorithm and a corresponding regret bound, as shown below.

Proposition 5.1. *Suppose that at each round t , 1) the oracle AlgSq outputs a CTR predictor $f_t \in \mathcal{F}$; 2) the learner then randomly sets $\tilde{\rho}_t$ according to the distribution that realizes the \min_Q in the definition of $\text{dec}_\gamma(\hat{\rho}_t)$ where $\hat{\rho}_{t,i} = f_t(x_t, i)$; 3) finally the oracle is fed with the tuple (x_t, i_t, c_t) . Then the regret of the learner satisfies $\mathbb{E}[\text{Reg}] \leq \gamma \mathbf{Reg}_{\text{Sq}} + \mathbb{E} \left[\sum_{t=1}^T \text{dec}_\gamma(\hat{\rho}_t) \right]$.*

For example, if $\text{dec}_\gamma(\hat{\rho}_t)$ turns out to be $\mathcal{O}(\frac{N}{\gamma})$ for all t , then picking the optimal γ leads to $\mathbb{E}[\text{Reg}] = \mathcal{O}(\sqrt{NT \mathbf{Reg}_{\text{Sq}}})$ (which, for linear class, implies $\mathbb{E}[\text{Reg}] = \mathcal{O}(\sqrt{dNT \log T})$ if ONS is used as the oracle). For contextual bandits, the corresponding DEC is indeed shown to be $\mathcal{O}(\frac{N}{\gamma})$ (Foster and Rakhlin, 2020).

An upper bound and a simple algorithm. Unfortunately, we are unable to show that our DEC is of order $1/\gamma$ due to the very complicated second-price structure. Nevertheless, we can show a worse bound $\text{dec}_\gamma(\hat{\rho}) = \mathcal{O}(\sqrt{N/\gamma})$. This is achieved by a simple ε -greedy strategy (with $\varepsilon = \sqrt{N/\gamma}$) — let Q concentrate on $N + 1$ CTR predictions: the greedy one $\hat{\rho}$ with probability $1 - \varepsilon$ (exploitation), and the one-hot CTR predictions e_1, \dots, e_N , each with probability ε/N (exploration, since when $\tilde{\rho} = e_i$, ad i will always be selected as the winner, and we observe signal on its CTR).⁵ Put together, this leads to Algorithm 2 and the following result, whose analysis requires a careful treatment to the second-price structure despite the simplicity of the algorithm (see Appendix C).

Theorem 5.1. *For any $\hat{\rho}_t \in [0, 1]^N$, we have $\text{dec}_\gamma(\hat{\rho}) = \mathcal{O}(\sqrt{N/\gamma})$, evidenced by the ε -greedy strategy described above. Consequently, Algorithm 2 with $\varepsilon = T^{-\frac{1}{3}}(N \mathbf{Reg}_{\text{Sq}})^{\frac{1}{3}}$ guarantees $\mathbb{E}[\text{Reg}] = \mathcal{O}(T^{\frac{2}{3}}(N \mathbf{Reg}_{\text{Sq}})^{\frac{1}{3}})$.*

Discussion and conjecture. Despite the worse regret bound, one nice property of Algorithm 2 is that it

⁵Note that in this case the payment is always 0, which, practically speaking, might not be desirable. However, since in reality there is always a minimum allowed bid $\sigma > 0$, we can in fact replace e_i by $\frac{\sigma}{2} \cdot \mathbf{1} + (1 - \frac{\sigma}{2})e_i$ instead, which still ensures exploration of ad i while leading to at least $\sigma^2/2$ payment per click.

Algorithm 2 Epsilon Greedy with an Online Regression Oracle

Input: exploration parameter $\varepsilon > 0$ and an online regression oracle AlgSq satisfying [Assumption 2](#)

for $t = 1, 2, \dots, T$ **do**

 Obtain a CTR predictor f_t from the oracle AlgSq.

 Receive context x_t and the set of N_t bidders.

 With ε probability, pick $i \in [N_t]$ uniformly at random and set $\tilde{\rho}_t = e_i \in \mathbb{R}^{N_t}$; with the remaining $1 - \varepsilon$ probability, set $\tilde{\rho}_{t,i} = f_t(x_t, i)$ for all $i \in [N_t]$

 Receive bids $b_t \in [0, 1]^{N_t}$.

 Select the bidder $i_t = \operatorname{argmax}_{i \in [N_t]} b_{t,i} \tilde{\rho}_{t,i}$ as the winner, with payment per click $d_t = \frac{b_{t,j_t} \tilde{\rho}_{t,j_t}}{\tilde{\rho}_{t,i_t}}$ where $j_t = \operatorname{argmax}_{i \in [N_t]} b_{t,i} \tilde{\rho}_{t,i}$ is the runner-up.

 Receive feedback $c_t = \mathbf{1}\{\text{ad } i_t \text{ is clicked}\}$ and payment $c_t d_t$.

 Feed the tuple (x_t, i_t, c_t) to the oracle AlgSq.

in fact never uses the bid information in deciding $\tilde{\rho}_t$, since the oracle never receives the bid information as feedback. This is in sharp contrast with [Algorithm 1](#) where both (IPS) and (OptSq) are defined in terms of the bids. This property is useful when the second-price auction is actually run by a third party who collects the bids from the advertisers without revealing them to the platform that makes the CTR predictions.

In fact, this property is inherent in this DEC approach and not just because of our choice of a potentially suboptimal ε -greedy strategy. Indeed, even for the algorithm described in [Proposition 5.1](#) that solves the \min_Q part of [Eq. \(4\)](#) exactly, it still by definition does not use the bid information but only the predictions from the oracle (which, again, is independent of any bids from prior rounds). We conjecture that because of this aspect, the optimal bound of $\operatorname{dec}_\gamma(\hat{\rho})$ might indeed be $\Theta(\sqrt{N/\gamma})$, achieved by the simple ε -greedy strategy. We leave this question for future investigation.

6 Experiments

We implement the two efficient algorithms we propose, namely [Algorithm 1](#) with (OptSq) loss estimators and [Algorithm 2](#), and demonstrate their superior performance on a synthetic dataset compared to three simple baselines: 1) a strategy that always make random CTR predictions, 2) a strategy that always make the same CTR prediction for all ads, which is equivalent to running a second-price auction using only the bids, and 3) a variant of [Algorithm 1](#) that uses the (OptSq) estimator but removes its optimistic part $\operatorname{smax}_j f(x_t, j) b_{t,j}$ (denoted by (Sq)).

For [Algorithm 1](#), η is chosen from $\{\frac{1}{16}, \frac{1}{8}, \frac{1}{4}\}$, and SGLD [Eq. \(3\)](#) is run for 32 steps per round with a step size α chosen from $\{0.0005, 0.001, 0.005, 0.01, 0.05\}$. For [Algorithm 2](#), the parameter ε is chosen from $\{T^{-\frac{1}{3}}, 2T^{-\frac{1}{3}}, 4T^{-\frac{1}{3}}\}$, and the regression oracle is simply implemented by online gradient descent ([Zinkevich, 2003](#)) with a learning rate from $\{0.001, 0.005\}$.

We use one-hidden-layer neural nets for the CTR predictors. Specifically, at each time t , our context x_t is a matrix in $\mathbb{R}^{d \times (1+N_t)}$, where the first column $x_{t,0}$ represents some common information shared by all advertisers (such as the user query), and the $(i+1)$ -th column $x_{t,i}$ represents the feature of ad i . Our predictor class is then $\mathcal{F} = \{f_\theta : f_\theta(x, i) = \operatorname{Sigmoid}(\theta_0^\top x_0 + \theta^\top x_i), \theta_0, \theta \in [-1, 1]^d\}$ where for any $u \in \mathbb{R}$, $\operatorname{Sigmoid}(u) = \frac{1}{1+e^{-u}}$.

Synthetic dataset construction. We generate a synthetic dataset by choosing $d = 128$ and $T = 10^4$ and uniformly at random sampling N_t from 5 to 10, x_t from $[-1, 1]^{d \times (1+N_t)}$, and b_t from $[0.1, 1]$. To generate the underlying CTR predictor f^* , we first generate some fake CTRs uniformly at random from $[0.2, 1]$ for each t and then use the full dataset to fit a model from \mathcal{F} as the final f^* . To avoid the trivial fixed CTR strategy already performing very well, we also make a final adjustment and set the bid of the ad with the lowest CTR to be 1 (imagine a little-known new brand trying to bid high to get more impressions).

Results. We repeat our experiment with 4 different random seeds and plot in [Figure 1](#) the average and the standard deviation (as shaded areas) of the cumulative regret over these 4 trials under the best hyperparameters. From [Figure 1](#), we see that 1) as our theory indicates, [Algorithm 1](#) (OptSq) indeed suffers lower regret than [Algorithm 2](#); 2) both of our algorithms beat the three baselines with a large margin; 3) [Algorithm 1](#) (Sq) is only comparable to the random CTR strategy, demonstrating that the optimistic part of (OptSq) is indeed both theoretically and empirically critical.

Acknowledgements

MZ and HL are supported by NSF Award IIS-1943607.




Figure 1: Performance of different algorithms on synthetic data

References

- Alekh Agarwal, Daniel Hsu, Satyen Kale, John Langford, Lihong Li, and Robert Schapire. Taming the monster: A fast and simple algorithm for contextual bandits. In *International Conference on Machine Learning*, pages 1638–1646. PMLR, 2014.
- Gagan Aggarwal, Ashish Goel, and Rajeev Motwani. Truthful auctions for pricing search keywords. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, pages 1–7, 2006.
- Peter Auer, Nicolo Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multiarmed bandit problem. *SIAM journal on computing*, 32(1): 48–77, 2002.
- Moshe Babaioff, Yogeshwer Sharma, and Aleksandrs Slivkins. Characterizing truthful multi-armed bandit mechanisms. *SIAM Journal on Computing*, 43(1): 194, 2014.
- Moshe Babaioff, Robert D Kleinberg, and Aleksandrs Slivkins. Truthful mechanisms with implicit payment computation. *Journal of the ACM (JACM)*, 62(2): 1–37, 2015.
- Santiago Balseiro, Yuan Deng, Jieming Mao, Vahab Mirrokni, and Song Zuo. Robust auction design in the auto-bidding world. *Advances in Neural Information Processing Systems*, 34:17777–17788, 2021.
- Santiago R Balseiro and Yonatan Gur. Learning in repeated auctions with budgets: Regret minimization and equilibrium. *Management Science*, 65(9):3952–3968, 2019.
- Santiago R Balseiro, Haihao Lu, and Vahab Mirrokni. The best of many worlds: Dual mirror descent for online allocation problems. *Operations Research*, 71(1):101–119, 2023.
- Swapna Bucapatnam, Atilla Eryilmaz, and Ness B Shroff. Stochastic bandits with side observations on networks. In *The 2014 ACM international conference on Measurement and modeling of computer systems*, pages 289–300, 2014.
- Nikhil R Devanur and Sham M Kakade. The price of truthfulness for pay-per-click auctions. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 99–106, 2009.
- Miroslav Dudik, Daniel Hsu, Satyen Kale, Nikos Karampatziakis, John Langford, Lev Reyzin, and Tong Zhang. Efficient optimal learning for contextual bandits. In *Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence*, pages 169–178, 2011.
- Zhe Feng, Christopher Liaw, and Zixin Zhou. Improved online learning algorithms for ctr prediction in ad auctions. In *Proceedings of the 40th International Conference on Machine Learning (ICML)*, 2023.
- Dylan Foster and Alexander Rakhlin. Beyond ucb: Optimal and efficient contextual bandits with regression oracles. In *International Conference on Machine Learning*, pages 3199–3210. PMLR, 2020.
- Dylan Foster, Alekh Agarwal, Miroslav Dudik, Haipeng Luo, and Robert Schapire. Practical contextual bandits with regression oracles. In *International Conference on Machine Learning*, pages 1539–1548. PMLR, 2018.
- Dylan J Foster and Akshay Krishnamurthy. Efficient first-order contextual bandits: Prediction, allocation, and triangular discrimination. *Advances in Neural Information Processing Systems*, 34:18907–18919, 2021.
- Dylan J Foster, Sham M Kakade, Jian Qian, and Alexander Rakhlin. The statistical complexity of interactive decision making. *arXiv preprint arXiv:2112.13487*, 2021.
- Jason Gaitonde, Yingkai Li, Bar Light, Brendan Lucier, and Aleksandrs Slivkins. Budget pacing in repeated auctions: Regret and efficiency without convergence. *arXiv e-prints*, pages arXiv–2205, 2022.
- Nicola Gatti, Alessandro Lazaric, and Francesco Trovò. A truthful learning mechanism for contextual multi-slot sponsored search auctions with externalities. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 605–622, 2012.
- YanJun Han, Zhengyuan Zhou, Aaron Flores, Erik Ordentlich, and Tsachy Weissman. Learning to bid optimally and efficiently in adversarial first-price auctions. *arXiv preprint arXiv:2007.04568*, 2020.
- Elad Hazan, Amit Agarwal, and Satyen Kale. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 69(2-3):169–192, 2007.

- Kirthevasan Kandasamy, Joseph E Gonzalez, Michael I Jordan, and Ion Stoica. Vcg mechanism design with unknown agent values under stochastic bandit feedback. *Journal of Machine Learning Research*, 24(53): 1–45, 2023.
- Tze Leung Lai, Herbert Robbins, et al. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- John Langford and Tong Zhang. The epoch-greedy algorithm for multi-armed bandits with side information. *Advances in neural information processing systems*, 20, 2007.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Brendan Lucier, Sarath Pattathil, Aleksandrs Slivkins, and Mengxiao Zhang. Autobidders with budget and roi constraints: Efficiency, regret, and pacing dynamics. *arXiv preprint arXiv:2301.13306*, 2023.
- David Simchi-Levi and Yunzong Xu. Bypassing the monster: A faster and simpler optimal algorithm for contextual bandits under realizability. *Mathematics of Operations Research*, 47(3):1904–1931, 2022.
- Qian Wang, Zongjun Yang, Xiaotie Deng, and Yuqing Kong. Learning to bid in repeated first-price auctions with budgets. *arXiv preprint arXiv:2304.13477*, 2023.
- Chen-Yu Wei, Haipeng Luo, and Alekh Agarwal. Taking a hint: How to leverage loss predictors in contextual bandits? In *Conference on Learning Theory*, pages 3583–3634. PMLR, 2020.
- Max Welling and Yee W Teh. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 681–688, 2011.
- Yinglun Xu, Bhuvish Kumar, and Jacob Abernethy. On the robustness of epoch-greedy in multi-agent contextual bandit mechanisms. *arXiv preprint arXiv:2307.07675*, 2023.
- Yunbei Xu and Assaf Zeevi. Upper counterfactual confidence bounds: a new optimism principle for contextual bandits. *arXiv preprint arXiv:2007.07876*, 2020.
- Tong Zhang. Feel-good thompson sampling for contextual bandits and reinforcement learning. *SIAM Journal on Mathematics of Data Science*, 4(2):834–857, 2022.
- Yinglun Zhu and Paul Mineiro. Contextual bandits with smooth regret: Efficient learning in continuous action spaces. In *International Conference on Machine Learning*, pages 27574–27590. PMLR, 2022.
- Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th international conference on machine learning*, pages 928–936, 2003.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Yes]
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Yes]
 - (b) Complete proofs of all theoretical results. [Yes]
 - (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Yes]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. [Not Applicable]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
 - (d) Information about consent from data providers/curators. [Not Applicable]

- (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
- (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

A Omitted Proofs in Section 3

Proof of Theorem 3.1. According to standard analysis of exponential weight (see for example (Wei et al., 2020, Lemma 14)), we have for any $f_0 \in \mathcal{F}$:

$$\langle q_t - e_{f_0}, \widehat{\ell}_t \rangle \leq \frac{1}{\eta} D(e_{f_0}, q_t) - \frac{1}{\eta} D(e_{f_0}, q_{t+1}) + \frac{\eta}{2} \sum_{f \in \mathcal{F}} q_{t,f} \widehat{\ell}_{t,f}^2,$$

where $D(u, v) = \sum_{f \in \mathcal{F}} u_f \log \frac{u_f}{v_f}$ is the KL divergence between u and v . Taking summation over $t \in [T]$, we have with $i_{t,f} = \arg \max_i b_{t,i} f(x_t, i)$:

$$\begin{aligned} & \sum_{t=1}^T \langle q_t - e_{f_0}, \widehat{\ell}_t \rangle \\ & \leq \frac{D(e_{f_0}, q_1)}{\eta} + \frac{\eta}{2} \sum_{t=1}^T \sum_{f \in \mathcal{F}} q_{t,f} \widehat{\ell}_{t,f}^2 \\ & = \frac{D(e_{f_0}, q_1)}{\eta} + \frac{\eta}{2} \sum_{t=1}^T \sum_{i \in [N_t]} \sum_{f \in \mathcal{F}, i=i_{t,f}} q_{t,f} \widehat{\ell}_{t,f}^2 \\ & = \frac{D(e_{f_0}, q_1)}{\eta} + \frac{\eta}{2} \sum_{t=1}^T \sum_{i \in [N_t]} \sum_{f \in \mathcal{F}, i=i_{t,f}} \frac{q_{t,f}}{p_{t,i}^2} \left(1 - c_t \cdot \frac{\text{smax}_j b_{t,j} f(x_t, j)}{f(x_t, i_{t,f})} \right)^2 \cdot \mathbb{1}\{i_t = i\} \\ & \leq \frac{D(e_{f_0}, q_1)}{\eta} + \frac{\eta}{2} \sum_{t=1}^T \sum_{i \in [N_t]} \sum_{f \in \mathcal{F}, i=i_{t,f}} \frac{q_{t,f} \mathbb{1}\{i_t = i\}}{p_{t,i}^2} \\ & = \frac{D(e_{f_0}, q_1)}{\eta} + \frac{\eta}{2} \sum_{t=1}^T \sum_{i \in [N_t]} \frac{\mathbb{1}\{i_t = i\}}{p_{t,i}} \quad (p_{t,i} = \sum_{f \in \mathcal{F}: i=i_{t,f}} q_{t,f}) \end{aligned}$$

where the inequality is because $\text{smax}_j b_{t,j} f(x_t, j) \leq b_{t,i_{t,f}} f(x_t, i_{t,f})$ and thus $1 - c_t \cdot \frac{\text{smax}_j b_{t,j} f(x_t, j)}{f(x_t, i_{t,f})} \in [0, 1]$. Taking expectation over both sides and noticing

$$\begin{aligned} \mathbb{E} \left[\frac{\mathbb{1}\{i_t = i\}}{p_{t,i}} \right] &= \mathbb{E} \left[\frac{p_{t,i}}{p_{t,i}} \right] = 1, \\ \mathbb{E} \left[\widehat{\ell}_{t,f} \right] &= \mathbb{E} \left[\frac{\mathbb{1}\{i_t = i_{t,f}\}}{p_{t,i_t}} \left(1 - \rho_{t,i_{t,f}} \cdot \frac{\text{smax}_j b_{t,j} f(x_t, j)}{f(x_t, i_{t,f})} \right) \right] \\ &= 1 - \rho_{t,i_{t,f}} \cdot \frac{\text{smax}_j b_{t,j} f(x_t, j)}{f(x_t, i_{t,f})} \\ &= 1 - f^*(x_t, i_{t,f}) \cdot \frac{\text{smax}_j b_{t,j} f(x_t, j)}{f(x_t, i_{t,f})}, \end{aligned}$$

we obtain

$$\begin{aligned} & \mathbb{E} \left[\sum_{t=1}^T f^*(x_t, i_{t,f_0}) \cdot \frac{\text{smax}_j b_{t,j} f_0(x_t, j)}{f_0(x_t, i_{t,f_0})} - \sum_{t=1}^T f^*(x_t, i_{t,f_t}) \cdot \frac{\text{smax}_j b_{t,j} f_t(x_t, j)}{f_t(x_t, i_{t,f_t})} \right] \\ &= \mathbb{E} \left[\sum_{t=1}^T f^*(x_t, i_{t,f_0}) \cdot \frac{\text{smax}_j b_{t,j} f_0(x_t, j)}{f_0(x_t, i_{t,f_0})} - \sum_{t=1}^T c_t d_t \right] \leq \frac{D(e_{f_0}, q_1)}{\eta} + \frac{\eta}{2} \sum_{t=1}^T N_t. \end{aligned} \quad (5)$$

Picking $f_0 = f^*$, which is in the class \mathcal{F} by [Assumption 1](#), the left-hand side above becomes exactly the expected regret $\mathbb{E}[\text{Reg}]$. Finally, since q_1 is a uniform distribution over \mathcal{F} , we know that $D(e_{f^*}, q_1) = \log |\mathcal{F}|$, and picking $\eta = \sqrt{\frac{\log |\mathcal{F}|}{\sum_{t=1}^T N_t}}$ then leads to the claimed regret bound. \square

Proof of Corollary 3.1. For any $u \in [0, 1]$, define $\lfloor u \rfloor_T = \frac{1}{T} \lfloor Tu \rfloor$. Since we consider the non-contextual setting, we have $\rho = \rho_t$ for all $t \in [T]$ and the policy $f \in \widehat{\mathcal{F}}$ can be represented as a vector in $[0, 1]^N$. To show that [Algorithm 1](#)

guarantees $\mathcal{O}(N\sqrt{T\log N})$ regret, let $\hat{f}^* = ([\rho_1]_T, \dots, [\rho_N]_T) \in \hat{\mathcal{F}}$. Using Eq. (5) and picking $f_0 = \hat{f}^*$, we know that

$$\mathbb{E} \left[\sum_{t=1}^T \rho_{i_t, \hat{f}^*} \cdot \frac{b_{t, j_{t, \hat{f}^*}} \hat{f}_{j_{t, \hat{f}^*}}^*}{\hat{f}_{i_t, \hat{f}^*}^*} - \sum_{t=1}^T c_t d_t \right] \leq \mathcal{O} \left(\frac{N \log T}{\eta} \right) + \frac{\eta NT}{2},$$

where $i_{t, \rho} = \operatorname{argmax}_i b_{t, i} \rho_{t, i}$ and $j_{t, \rho} = \operatorname{argsmax}_i b_{t, i} \rho_{t, i}$. It remains to argue that the first term above is close to $\sum_{t=1}^T \operatorname{smax}_i b_{t, i} \rho_i$, the revenue of the oracle strategy. To this end, for each time t we consider two cases depending on whether \hat{f}^* and ρ selects the same winner. If $i_{t, \hat{f}^*} = i_{t, \rho}$, we have

$$\begin{aligned} \rho_{i_t, \hat{f}^*} \cdot \frac{b_{t, j_{t, \hat{f}^*}} \hat{f}_{j_{t, \hat{f}^*}}^*}{\hat{f}_{i_t, \hat{f}^*}^*} &\geq b_{t, j_{t, \hat{f}^*}} \hat{f}_{j_{t, \hat{f}^*}}^* && \text{(since } \hat{f}_i^* = [\rho_i]_T \leq \rho_i) \\ &\geq b_{t, j_{t, \rho}} \hat{f}_{j_{t, \rho}}^* && \text{(since } b_{t, j_{t, \hat{f}^*}} \hat{f}_{j_{t, \hat{f}^*}}^* = \operatorname{smax}_i b_{t, i} \hat{f}_i^* \text{ and } j_{t, \rho} \neq i_{t, \rho} = i_{t, \hat{f}^*}) \\ &\geq b_{t, j_{t, \rho}} \rho_{j_{t, \rho}} - \frac{1}{T} && \text{(since } \hat{f}_i^* = [\rho_i]_T \geq \rho_i - \frac{1}{T}) \\ &= \operatorname{smax}_i b_{t, i} \rho_i - \frac{1}{T}. \end{aligned}$$

Otherwise, we know that $b_{t, j_{t, \hat{f}^*}} \hat{f}_{j_{t, \hat{f}^*}}^* \geq b_{t, i_{t, \rho}} \hat{f}_{i_{t, \rho}}^* \geq b_{t, i_{t, \rho}} \rho_{i_{t, \rho}} - \frac{1}{T} \geq \operatorname{smax}_i b_{t, i} \rho_i - \frac{1}{T}$. Combining both cases, we know that

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^T \operatorname{smax}_i b_{t, i} \rho_i - \sum_{t=1}^T c_t d_t \right] &\leq \mathbb{E} \left[\sum_{t=1}^T \rho_{i_t, \hat{f}^*} \cdot \frac{b_{t, j_{t, \hat{f}^*}} \hat{f}_{j_{t, \hat{f}^*}}^*}{\hat{f}_{i_t, \hat{f}^*}^*} - \sum_{t=1}^T c_t d_t \right] + 1 \\ &\leq \mathcal{O} \left(\frac{N \log T}{\eta} \right) + \frac{\eta NT}{2} + 1. \end{aligned}$$

Picking $\eta = \sqrt{\frac{\log T}{T}}$ achieves the claimed bound. \square

Proof of Theorem 3.2. The proof follows the idea of the lower bound construction in multi-armed bandits. Consider an instance of the auction problem with $N_t = N$ bidders and CTR $\{\rho_i\}_{i=1}^N$. The bid vector b_t at each round t is $\mathbf{1}$. Denote $\rho^{(i, j)}$ ($i \neq j$) to be the environment where $\rho_k^{(i, j)} = \frac{1}{2}$ for all $k \in [N] \setminus \{i, j\}$ and $\rho_i^{(i, j)} = \rho_j^{(i, j)} = \frac{1}{2} + \varepsilon$ for some $\varepsilon > 0$ to be specified later. We also denote $\rho^{(0)}$ to be the environment where $\rho_i^{(0)} = \frac{1}{2}$ for all $i \in [N]$.

Consider the environment \mathcal{E} where ρ is uniformly drawn from the $\frac{N(N-1)}{2}$ environments $\{\rho^{(i, j)}\}_{i \neq j, i, j \in [n]}$. For notational convenience, we denote $\mathbb{E}_{i, j}[\cdot]$ (\mathbb{E}_0) to be $\mathbb{E}_{\rho^{(i, j)}}[\cdot]$ ($\mathbb{E}_{\rho^{(0)}}$). Let n_i be the number of rounds ad i is selected as the winner (via picking an estimated CTR). Since the payment per click can not exceed the bid 1 according to the auction design, the expected revenue of the learner is at most the CTR of the winning ad. In addition, the benchmark revenue in environment $\rho^{(i, j)}$ is $(\frac{1}{2} + \varepsilon)T$ by picking the estimated CTR as the true CTR. Therefore, the expected regret with respect to environment \mathcal{E} is lower bounded as follows:

$$\begin{aligned} \mathbb{E}_{\mathcal{E}}[\text{Reg}] &\geq \frac{2}{N(N-1)} \sum_{i \neq j} \mathbb{E}_{i, j} \left[\left(\frac{1}{2} + \varepsilon \right) T - \frac{1}{2} \sum_{k \neq i, j} n_k - \left(\frac{1}{2} + \varepsilon \right) (n_i + n_j) \right] \\ &= \frac{2}{N(N-1)} \sum_{i \neq j, i, j \in [N]} \mathbb{E}_{i, j} [\varepsilon(T - n_i - n_j)], \end{aligned} \quad (6)$$

where the last equality uses the fact that $\sum_{i=1}^N n_i = T$. According to Exercise 15.2.(a) of (Lattimore and Szepesvári, 2020), we have

$$\mathbb{E}_{i, j} [n_i + n_j] \leq \mathbb{E}_0 [n_i + n_j] + T \sqrt{\frac{1}{4} \varepsilon^2 \mathbb{E}_0 [n_i + n_j]}.$$

Taking summation over all $i \neq j, i, j \in [N]$, we obtain that

$$\begin{aligned}
 \sum_{i \neq j, i, j \in [N]} \mathbb{E}_{i,j}[n_i + n_j] &\leq \sum_{i \neq j, i, j \in [N]} \mathbb{E}_0[n_i + n_j] + \frac{T\varepsilon}{2} \sum_{i \neq j, i, j \in [N]} \sqrt{\mathbb{E}_0[n_i + n_j]} \\
 &\leq (N-1) \sum_{i=1}^N \mathbb{E}_0[n_i] + \frac{T\varepsilon}{2} (N-1) \sum_{i=1}^N \sqrt{\mathbb{E}_0[n_i]} \\
 &\leq (N-1)T + \frac{T\varepsilon(N-1)}{2} \sqrt{NT}, \tag{7}
 \end{aligned}$$

where the second inequality uses $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$ and the third inequality is due to Cauchy-Schwarz inequality.

Applying Eq. (7) to Eq. (6), we obtain that

$$\begin{aligned}
 \mathbb{E}_{\mathcal{E}}[\text{Reg}] &\geq \varepsilon T - \frac{2\varepsilon}{N(N-1)} \cdot \left[(N-1)T + \frac{T\varepsilon(N-1)}{2} \sqrt{NT} \right] \\
 &= \varepsilon T - \frac{2\varepsilon T}{N} - \varepsilon^2 N^{-\frac{1}{2}} T^{\frac{3}{2}} \\
 &\geq \frac{1}{3} \varepsilon T - \varepsilon^2 N^{-\frac{1}{2}} T^{\frac{3}{2}}. \tag{N \geq 3}
 \end{aligned}$$

Picking $\varepsilon = \frac{1}{4} \sqrt{\frac{N}{T}}$ leads to $\mathbb{E}_{\mathcal{E}}[\text{Reg}] \geq \frac{1}{48} \sqrt{NT}$. Therefore, there exists one environment among $\{\rho^{(i,j)}\}_{i \neq j, i, j \in [N]}$ such that $\mathbb{E}_{i,j}[\text{Reg}] \geq \frac{1}{48} \sqrt{NT}$, which finishes the proof. \square

B Omitted Proofs in Section 4

Proof of Theorem 4.1. Let $i_t^* = \text{argmax}_{i \in [N_t]} b_{t,i} \rho_{t,i}$, $j_t^* = \text{argmax}_{j \in [N_t]} b_{t,i} \rho_{t,i}$, $i_{t,f} = \text{argmax}_{i \in [N_t]} b_{t,i} f(x_t, i)$, and $j_{t,f} = \text{argmax}_{i \in [N_t]} b_{t,i} f(x_t, i)$. Also recall the notation $\tilde{i}_t = \text{argmax}_{i \in [N_t]} b_{t,i} \tilde{\rho}_{t,i}$ and $\tilde{j}_t = \text{argmax}_{j \in [N_t]} b_{t,i} \tilde{\rho}_{t,i}$. Then, we decompose the regret as follows:

$$\begin{aligned}
 \mathbb{E}[\text{Reg}] &= \mathbb{E} \left[\sum_{t=1}^T \text{smax}_{i \in [N_t]} b_{t,i} \rho_{t,i} - \sum_{t=1}^T c_t d_t \right] \\
 &= \mathbb{E} \left[\sum_{t=1}^T \rho_{t,j_t^*} b_{t,j_t^*} - \sum_{t=1}^T \frac{\tilde{\rho}_{t,j_t} b_{t,j_t}}{\tilde{\rho}_{t,i_t}} \rho_{t,i_t} \right] \\
 &= \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_{i_t, j_t} \left[\frac{\tilde{\rho}_{t,j_t} b_{t,j_t}}{\tilde{\rho}_{t,i_t}} (\tilde{\rho}_{t,i_t} - \rho_{t,i_t}) \right] \right] - \mathbb{E} \left[\sum_{t=1}^T (\tilde{\rho}_{t,j_t} b_{t,j_t} - \rho_{t,j_t^*} b_{t,j_t^*}) \right] \\
 &= \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_{f \sim q_t} \left[\frac{f(x_t, j_{t,f}) b_{t,j_{t,f}}}{f(x_t, i_{t,f})} (f(x_t, i_{t,f}) - f^*(x_t, i_{t,f})) \right] \right] \\
 &\quad - \mathbb{E} \left[\sum_{t=1}^T \underbrace{(f_t(x_t, j_t) b_{t,j_t} - f^*(x_t, j_t^*) b_{t,j_t^*})}_{\triangleq \text{FG}_t} \right] \\
 &\leq \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_{f \sim q_t} [|f(x_t, i_{t,f}) - f^*(x_t, i_{t,f})|] \right] - \mathbb{E} \left[\sum_{t=1}^T \text{FG}_t \right]
 \end{aligned}$$

where the inequality is because $\frac{f(x_t, j_{t,f}) b_{t,j_{t,f}}}{f(x_t, i_{t,f})} \leq b_{t,i_{t,f}} \leq 1$. Note that the ‘‘Feel-Good’’ term FG_t is the difference between what f_t believes its payment is in its predicted world and that of the perfect predictor f^* . We now analyze the other term for each time t , which can be written as:

$$\sum_{i \in [N_t]: p_{t,i} \neq 0} \mathbb{E}_{f \sim q_t} [\mathbb{1}\{i_{t,f} = i\} |f(x_t, i) - f^*(x_t, i)|].$$

where $p_{t,i}$ is the probability of i_t being i . Now consider a fixed i . We have for any $\mu > 0$,

$$\begin{aligned} & \mathbb{E}_{f \sim q_t} [\mathbb{1}\{i_{t,f} = i\} |f(x_t, i) - f^*(x_t, i)|] \\ & \leq \mathbb{E}_{f \sim q_t} \left[\frac{\mathbb{1}\{i_{t,f} = i\}}{4\mu p_{t,i}} + \mu p_{t,i} (f(x_t, i) - f^*(x_t, i))^2 \right] \quad (\text{AM-GM inequality}) \\ & = \frac{1}{4\mu} + \mu p_{t,i} \mathbb{E}_{f \sim q_t} \left[(f(x_t, i) - f^*(x_t, i))^2 \right], \quad (p_{t,i} = \mathbb{E}_{f \sim q_t} [\mathbb{1}\{i_{t,f} = i\}]) \end{aligned}$$

which, after taking the summation over i , becomes

$$\frac{N_t}{4\mu} + \mu \mathbb{E}_{i_t \sim p_t} \mathbb{E}_{f \sim q_t} \left[(f(x_t, i_t) - f^*(x_t, i_t))^2 \right]$$

(note the important decoupling effect here: f is not f_t). With notation $\text{LS}_t \triangleq (f(x_t, i_t) - f^*(x_t, i_t))^2$ (“Least Squares”), we have thus shown for any $\mu > 0$:

$$\mathbb{E}[\text{Reg}] \leq \frac{1}{4\mu} \sum_{t=1}^T N_t + \mu \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_{i_t \sim p_t} \mathbb{E}_{f \sim q_t} [\text{LS}_t] \right] - \mathbb{E} \left[\sum_{t=1}^T \mathbb{E}_{f_t \sim q_t} [\text{FG}_t] \right]. \quad (8)$$

Next, using the fact that c_t is a Bernoulli random variable with mean $f^*(x_t, i_t)$ and following a similar analysis of Lemma 4 in (Zhang, 2022), we can show that

$$\frac{1}{16\eta} \mathbb{E}_{f \sim q_t} [\text{LS}_t] - \mathbb{E}_{f_t \sim q_t} [\text{FG}_t] \leq -\frac{1}{\eta} \log \mathbb{E}_{c_t | x_t, i_t} \mathbb{E}_{f \sim q_t} \left[\exp \left(-\eta \left(\widehat{\ell}_{t,f} - \widehat{\ell}_{t,f^*} \right) \right) \right] + 4\eta, \quad (9)$$

with $\widehat{\ell}_{t,f}$ defined as the (OptSq) option shown in Algorithm 1. For completeness, we include the proof of Eq. (9) in Lemma B.1. Comparing Eq. (8) and Eq. (9) naturally suggests picking $\mu = \frac{1}{16\eta}$, so that

$$\mathbb{E}[\text{Reg}] \leq 4\eta \sum_{t=1}^T N_t + 4T\eta - \frac{1}{\eta} \sum_{t=1}^T \log \mathbb{E}_{c_t | x_t, i_t} \mathbb{E}_{f \sim q_t} \left[\exp \left(-\eta \left(\widehat{\ell}_{t,f} - \widehat{\ell}_{t,f^*} \right) \right) \right]. \quad (10)$$

Note that the analysis so far holds for any q_t . To complete the proof, we now use the specific form of q_t defined in Algorithm 1: $q_{t,f} \propto \exp \left(-\eta \sum_{s=1}^{t-1} \widehat{\ell}_{s,f} \right)$. Using standard analysis of the multiplicative weight update, we show in Lemma B.2 the following:

$$-\mathbb{E} \left[\mathbb{E}_{i_t} \log \mathbb{E}_{c_t | x_t, i_t} \mathbb{E}_{f \sim q_t} \left[\exp \left(-\eta \left(\widehat{\ell}_{t,f} - \widehat{\ell}_{t,f^*} \right) \right) \right] \right] \leq Z_t - Z_{t-1}, \quad (11)$$

where $Z_t = -\mathbb{E} \left[\log \mathbb{E}_{f \sim q_t} \exp \left(-\eta \sum_{\tau=1}^t \left(\widehat{\ell}_{\tau,f} - \widehat{\ell}_{\tau,f^*} \right) \right) \right]$. Combining this fact with Eq. (10), we arrive at

$$\begin{aligned} \mathbb{E}[\text{Reg}] & \leq 4\eta \sum_{t=1}^T N_t + 4T\eta + \frac{1}{\eta} \sum_{t=1}^T (Z_t - Z_{t-1}) \\ & \leq 4\eta \sum_{t=1}^T N_t + 4T\eta + \frac{Z_T}{\eta}, \quad (\text{telescoping and } Z_0 = 0) \end{aligned}$$

which finishes the proof. \square

Lemma B.1. *Suppose that $\eta \leq 1$. For any distribution q_t over \mathcal{F} and $b_t \in [0, 1]^N$, we have*

$$\frac{1}{16\eta} \mathbb{E}_{f \sim q_t} [\text{LS}_t] - \mathbb{E}_{f_t \sim q_t} [\text{FG}_t] \leq -\frac{1}{\eta} \log \mathbb{E}_{c_t | x_t, i_t} \mathbb{E}_{f \sim q_t} \exp \left(-\eta \left(\widehat{\ell}_{t,f} - \widehat{\ell}_{t,f^*} \right) \right) + 4\eta,$$

where LS_t and FG_t are defined in the proof of Theorem 4.1, and $\widehat{\ell}_{t,f}$ is defined in the (OptSq) option of Eq. (2).

Proof. Since $c_t \in \{0, 1\}$ is a Bernoulli random variable with mean $f^*(x_t, i_t)$, c_t satisfies the following sub-Gaussian random variable property: for any ρ ,

$$\mathbb{E}_{c_t|x_t, i_t} [\exp(\rho(c_t - f^*(x_t, i_t)))] \leq \exp\left(\frac{\rho^2}{8}\right).$$

Let $\varepsilon_t = c_t - f^*(x_t, i_t)$. For any f , setting $\rho = -\frac{1}{2}(f^*(x_t, i_t) - f(x_t, i_t))$, we thus have

$$\mathbb{E}_{c_t|x_t, i_t} \exp\left(-\frac{1}{2}\varepsilon_t(f^*(x_t, i_t) - f(x_t, i_t))\right) \leq \exp\left(\frac{1}{32}\text{LS}_t\right). \quad (12)$$

On the other hand, consider the following equalities:

$$\begin{aligned} & -\eta(\widehat{\ell}_{t,f} - \widehat{\ell}_{t,f^*}) \\ &= -\frac{1}{4}(f(x_t, i_t) - c_t)^2 + \frac{1}{4}(f^*(x_t, i_t) - c_t)^2 + \eta \cdot \text{smax}_j b_{t,j} f(x_t, j) - \eta \cdot \text{smax}_j b_{t,j} f^*(x_t, j) \\ &= -\frac{1}{4}(\varepsilon_t + f^*(x_t, i_t) - f(x_t, i_t))^2 + \frac{1}{4}\varepsilon_t^2 + \eta \cdot \text{smax}_j b_{t,j} f(x_t, j) - \eta \cdot \text{smax}_j b_{t,j} f^*(x_t, j) \\ &= -\frac{1}{2}\varepsilon_t(f^*(x_t, i_t) - f(x_t, i_t)) - \frac{1}{4}\text{LS}_t + \eta\text{FG}_t(f), \end{aligned}$$

where we define $\text{FG}_t(f) = \text{smax}_j(f(x_t, j)b_{t,j}) - \text{smax}_j(f^*(x_t, j)b_{t,j})$ (so $\text{FG}_t = \text{FG}_t(f_t)$). Combining the above with Eq. (12), we obtain

$$\mathbb{E}_{c_t|x_t, i_t} \left[\exp\left(-\eta(\widehat{\ell}_{t,f} - \widehat{\ell}_{t,f^*})\right) \right] \leq \exp\left(-\frac{7}{32}\text{LS}_t + \eta\text{FG}_t(f)\right).$$

This further shows:

$$\begin{aligned} & \log \mathbb{E}_{f \sim q_t} \left[\mathbb{E}_{c_t|x_t, i_t} \left[\exp\left(-\eta(\widehat{\ell}_{t,f} - \widehat{\ell}_{t,f^*})\right) \right] \right] \\ & \leq \log \mathbb{E}_{f \sim q_t} \left[\exp\left(-\frac{7}{32}\text{LS}_t + \eta\text{FG}_t(f)\right) \right] \\ & \leq \frac{1}{2} \log \mathbb{E}_{f \sim q_t} \left[\exp\left(-\frac{7}{16}\text{LS}_t\right) \right] + \frac{1}{2} \log \mathbb{E}_{f \sim q_t} [\exp(2\eta\text{FG}_t(f))], \end{aligned} \quad (13)$$

where the last inequality is due to Cauchy-Schwarz inequality. For the first term, using the facts $0 \leq \text{LS}_t \leq 1$ and $e^x \leq 1 + x + \frac{x^2}{2}$ for $x \leq 0$, we know

$$\mathbb{E}_{f \sim q_t} \left[\exp\left(-\frac{7}{16}\text{LS}_t\right) \right] \leq 1 - \frac{7}{16}\mathbb{E}_{f \sim q_t}[\text{LS}_t] + \frac{49}{256}\mathbb{E}_{f \sim q_t}[\text{LS}_t^2] \leq 1 - \frac{1}{8}\mathbb{E}_{f \sim q_t}[\text{LS}_t],$$

where the last inequality is because $\text{LS}_t^2 \leq \text{LS}_t$. Further using $\log(1+x) \leq x$ gives

$$\frac{1}{2} \log \mathbb{E}_{f \sim q_t} \left[\exp\left(-\frac{7}{16}\text{LS}_t\right) \right] \leq \frac{1}{2} \log \left(1 - \frac{1}{8}\mathbb{E}_{f \sim q_t}[\text{LS}_t]\right) \leq -\frac{1}{16}\mathbb{E}_{f \sim q_t}[\text{LS}_t]. \quad (14)$$

Moreover, since $\eta \leq 1$ and $|\text{FG}_t(f)| \leq 1$, using $e^x \leq 1 + x + 2x^2$ for $x \leq 2$, we have

$$\begin{aligned} \frac{1}{2} \log \mathbb{E}_{f \sim q_t} [\exp(2\eta\text{FG}_t(f))] & \leq \frac{1}{2} \log (1 + 2\eta\mathbb{E}_{f \sim q_t}[\text{FG}_t(f)] + 2(2\eta)^2) \\ & \leq \eta\mathbb{E}_{f \sim q_t}[\text{FG}_t(f)] + 4\eta^2 && (\log(1+x) \leq x) \\ & = \eta\mathbb{E}_{f_t \sim q_t}[\text{FG}_t] + 4\eta^2. && (f_t \text{ is drawn from } q_t) \end{aligned}$$

Plugging the last bound and Eq. (14) into Eq. (13) and rearranging finishes the proof. \square

Lemma B.2. *Algorithm 1 guarantees that for each $t \in [T]$,*

$$-\mathbb{E} \left[\mathbb{E}_{i_t \sim p_t} \log \mathbb{E}_{c_t|x_t, i_t} \mathbb{E}_{f \sim q_t} \left[\exp(-\eta(\widehat{\ell}_{t,f} - \widehat{\ell}_{t,f^*})) \right] \right] \leq Z_t - Z_{t-1},$$

where $Z_t = -\mathbb{E} \left[\log \mathbb{E}_{f \sim q_t} \left[\exp\left(-\eta \sum_{\tau=1}^t (\widehat{\ell}_{\tau,f} - \widehat{\ell}_{\tau,f^*})\right) \right] \right]$.

Proof. Let $W_{t,f} \triangleq \exp\left(-\eta \sum_{\tau=1}^t (\widehat{\ell}_{t,f} - \widehat{\ell}_{t,f^*})\right)$. According to [Algorithm 1](#), we know that

$$q_{t,f} = \frac{\exp\left(-\eta \sum_{\tau=1}^{t-1} \widehat{\ell}_{\tau,f}\right)}{\int_{f' \in \mathcal{F}} \exp\left(-\eta \sum_{\tau=1}^{t-1} \widehat{\ell}_{\tau,f'}\right) df'} = \frac{W_{t-1,f}}{\int_{f' \in \mathcal{F}} W_{t-1,f'} df'}.$$

Then, according to the definition of Z_t , we have

$$\begin{aligned} & Z_{t-1} - Z_t \\ &= \mathbb{E} \left[\log \frac{\int_{f \in \mathcal{F}} W_{t,f} df}{\int_{f \in \mathcal{F}} W_{t-1,f} df} \right] \\ &= \mathbb{E} \left[\log \frac{\int_{f \in \mathcal{F}} W_{t-1,f} \exp(-\eta(\widehat{\ell}_{t,f} - \widehat{\ell}_{t,f^*})) df}{\int_{f \in \mathcal{F}} W_{t-1,f} df} \right] \\ &= \mathbb{E} \left[\log \mathbb{E}_{f \sim q_t} \left[\exp(-\eta(\widehat{\ell}_{t,f} - \widehat{\ell}_{t,f^*})) \right] \right] \\ &\leq \mathbb{E} \left[\mathbb{E}_{i_t \sim p_t} \log \mathbb{E}_{c_t | x_t, i_t} \mathbb{E}_{f \sim q_t} \left[\exp(-\eta(\widehat{\ell}_{t,f} - \widehat{\ell}_{t,f^*})) \right] \right], \end{aligned}$$

where the last inequality is due to Jensen's inequality. Rearranging the terms finishes the proof. \square

Proof of [Corollary 4.2](#). To show that $Z_T = \mathcal{O}(d \log(\alpha BT))$, we consider a small cube around the true parameter θ^* : $\Omega_T = \{\theta : \|\theta - \theta^*\|_\infty \leq \frac{1}{\alpha T}\}$. Since \mathcal{F} is α -Lipschitz with respect to $\|\cdot\|_\infty$, we know that $\|f_\theta - f_{\theta^*}\|_\infty \leq \frac{1}{T}$. Therefore, for any $\theta \in \Omega_T$,

$$\begin{aligned} & -\eta(\widehat{\ell}_{t,f_\theta} - \widehat{\ell}_{t,f_{\theta^*}}) \\ &= -\frac{1}{4}(f_\theta(x_t, i_t) - c_t)^2 + \frac{1}{4}(f_{\theta^*}(x_t, i_t) - c_t)^2 \\ &\quad + \eta \cdot \text{smax}_j b_{t,j} f_\theta(x_t, j) - \eta \cdot \text{smax}_j b_{t,j} f_{\theta^*}(x_t, j) \\ &\geq -\frac{1}{2}|f_\theta(x_t, i_t) - f_{\theta^*}(x_t, i_t)| + \eta \cdot \text{smax}_j b_{t,j} f_\theta(x_t, j) - \eta \cdot \text{smax}_j b_{t,j} f_{\theta^*}(x_t, j). \end{aligned} \tag{15}$$

For the second and the third term, if $\text{argmax}_i b_{t,i} f_\theta(x_t, i) = \text{argmax}_i b_{t,i} f_{\theta^*}(x_t, i)$, let $j_{\theta^*} = \text{argmax}_j b_{t,j} f_{\theta^*}(x_t, j)$ and we know that

$$\begin{aligned} & \text{smax}_j b_{t,j} f_\theta(x_t, j) - \text{smax}_j b_{t,j} f_{\theta^*}(x_t, j) \\ &\geq b_{t,j_{\theta^*}} f_\theta(x_t, j_{\theta^*}) - \text{smax}_j b_{t,j} f_{\theta^*}(x_t, j) \\ &\geq -b_{t,j_{\theta^*}} |f_\theta(x_t, j_{\theta^*}) - f_{\theta^*}(x_t, j_{\theta^*})|. \end{aligned}$$

Otherwise, let $i_{\theta^*} = \text{argmax}_i b_{t,i} f_{\theta^*}(x_t, i)$ and we know that

$$\begin{aligned} & \text{smax}_j b_{t,j} f_\theta(x_t, j) - \text{smax}_j b_{t,j} f_{\theta^*}(x_t, j) \\ &\geq \text{smax}_j b_{t,j} f_\theta(x_t, j) - b_{t,i_{\theta^*}} f_{\theta^*}(x_t, i_{\theta^*}) \\ &\geq b_{t,i_{\theta^*}} f_\theta(x_t, i_{\theta^*}) - b_{t,i_{\theta^*}} f_{\theta^*}(x_t, i_{\theta^*}) \\ &\geq -b_{t,i_{\theta^*}} |f_\theta(x_t, i_{\theta^*}) - f_{\theta^*}(x_t, i_{\theta^*})|. \end{aligned}$$

Combining the two cases above and plugging them into [Eq. \(15\)](#), we obtain

$$-\eta(\widehat{\ell}_{t,f_\theta} - \widehat{\ell}_{t,f_{\theta^*}}) \geq -\frac{1}{2T} - \frac{\eta}{T}.$$

This means that

$$Z_T = -\mathbb{E} \left[\log \mathbb{E}_{f \sim q_1} \exp \left(-\eta \sum_{t=1}^T (\widehat{\ell}_{t,f_\theta} - \widehat{\ell}_{t,f_{\theta^*}}) \right) \right]$$

$$\begin{aligned}
 &\leq -\mathbb{E} \left[\log \frac{1}{(\alpha BT)^d} \inf_{\theta \in \Omega_T} \exp \left(-\eta \sum_{t=1}^T (\hat{\ell}_{t, f_\theta} - \hat{\ell}_{t, f_{\theta^*}}) \right) \right] \\
 &\leq d \log(\alpha BT) + \frac{1}{2} + \eta = \mathcal{O}(d \log(\alpha BT)),
 \end{aligned}$$

which finishes the proof. \square

C Omitted Proofs in Section 5

Proof of Proposition 5.1. By the definition of dec_γ , we have

$$\mathbb{E}[\text{Reg}] \leq \mathbb{E} \left[\gamma \sum_{t=1}^T (\rho_{t, i_t} - \hat{\rho}_{t, i_t})^2 + \sum_{t=1}^T \text{dec}_\gamma(\hat{\rho}_t) \right].$$

It remains to notice that the first term is bounded by \mathbf{Reg}_{Sq} under [Assumption 2](#):

$$\begin{aligned}
 \mathbf{Reg}_{\text{Sq}} &\geq \mathbb{E} \left[\sum_{t=1}^T (f_t(x_t, i_t) - c_t)^2 - \sum_{t=1}^T (f^*(x_t, i_t) - c_t)^2 \right] \\
 &= \mathbb{E} \left[\sum_{t=1}^T (f_t(x_t, i_t) - f^*(x_t, i_t)) (f_t(x_t, i_t) + f^*(x_t, i_t) - 2c_t) \right] \\
 &= \mathbb{E} \left[\sum_{t=1}^T (f_t(x_t, i_t) - f^*(x_t, i_t))^2 \right] \quad (\text{the conditional expectation of } c_t \text{ is } f^*(x_t, i_t)) \\
 &= \mathbb{E} \left[\sum_{t=1}^T (\rho_{t, i_t} - \hat{\rho}_{t, i_t})^2 \right].
 \end{aligned}$$

This finishes the proof. \square

Proof of Theorem 5.1. Consider the ε -greedy strategy described above [Theorem 5.1](#) and any $\rho, b \in [0, 1]^N$. We have

$$\begin{aligned}
 &\mathbb{E}_{\tilde{\rho} \in Q} \left[\left(\text{smax}_{i \in [N]} b_i \rho_i - \rho_{i^*} \frac{\text{smax}_{j \in [N]} b_j \tilde{\rho}_j}{\tilde{\rho}_{i^*}} \right) - \gamma (\rho_{i^*} - \hat{\rho}_{i^*})^2 \right] \\
 &= \left(\text{smax}_{i \in [N]} b_i \rho_i - \mathbb{E}_{\tilde{\rho} \in Q} \left[\rho_{i^*} \frac{\text{smax}_{j \in [N]} b_j \tilde{\rho}_j}{\tilde{\rho}_{i^*}} \right] \right) - \gamma \sum_{i=1}^N p_i (\rho_i - \hat{\rho}_i)^2 \\
 &\leq \left(\text{smax}_{i \in [N]} b_i \rho_i - (1 - \varepsilon) \rho_{\hat{i}^*} \frac{\text{smax}_{j \in [N]} b_j \hat{\rho}_j}{\hat{\rho}_{\hat{i}^*}} \right) - \gamma \sum_{i=1}^N p_i (\rho_i - \hat{\rho}_i)^2 \\
 &\leq \left(\text{smax}_{i \in [N]} b_i \rho_i - \rho_{\hat{i}^*} \frac{\text{smax}_{j \in [N]} b_j \hat{\rho}_j}{\hat{\rho}_{\hat{i}^*}} \right) - \gamma \sum_{i=1}^N p_i (\rho_i - \hat{\rho}_i)^2 + \varepsilon
 \end{aligned} \tag{16}$$

where in the first step we introduce the notation p_i which is the probability of i^* being i , in the second step we ignore the revenue obtained from exploration and use \hat{i}^* to denote $\text{argmax}_i b_i \hat{\rho}_i$, and in the last step we use the fact $\frac{\text{smax}_{j \in [N]} b_j \hat{\rho}_j}{\hat{\rho}_{\hat{i}^*}} \leq 1$. We now make the following two observations for any $i \in [N]$, both due to AM-GM inequality and the fact $p_i \geq \varepsilon/N$:

$$|b_i \rho_i - b_i \hat{\rho}_i| - \frac{\gamma}{2} p_i (\rho_i - \hat{\rho}_i)^2 \leq \frac{b_i^2}{2\gamma p_i} \leq \frac{N}{2\gamma \varepsilon}, \tag{17}$$

$$\begin{aligned}
 \frac{\text{smax}_{j \in [N]} b_j \hat{\rho}_j}{\hat{\rho}_i} |\rho_{\hat{i}^*} - \hat{\rho}_{\hat{i}^*}| - \frac{\gamma}{2} p_{\hat{i}^*} (\rho_{\hat{i}^*} - \hat{\rho}_{\hat{i}^*})^2 &\leq \frac{1}{2\gamma p_{\hat{i}^*}} \left(\frac{\text{smax}_{j \in [N]} b_j \hat{\rho}_j}{\hat{\rho}_{\hat{i}^*}} \right)^2 \\
 &\leq \frac{b_{\hat{i}^*}^2}{2\gamma p_{\hat{i}^*}} \leq \frac{N}{2\gamma \varepsilon}.
 \end{aligned} \tag{18}$$

We are now ready to bound [Eq. \(16\)](#) by considering two cases: $\widehat{i}^* \neq j^*$ or $\widehat{i}^* = j^*$ where $j^* = \operatorname{argmax}_i b_i \rho_i$. In the first case ($\widehat{i}^* \neq j^*$), we use [Eq. \(17\)](#) with $i = j^*$ and [Eq. \(18\)](#) to continue to bound [Eq. \(16\)](#) as

$$b_{j^*} \widehat{\rho}_{j^*} - \operatorname{smax}_{j \in [N]} b_j \widehat{\rho}_j + \frac{N}{\gamma \varepsilon} + \varepsilon \leq \frac{N}{\gamma \varepsilon} + \varepsilon,$$

where the last step is because j^* is not \widehat{i}^* and thus $b_{j^*} \widehat{\rho}_{j^*}$ is not $\max_j b_j \widehat{\rho}_j$ and must be at most the second max. In the second case ($\widehat{i}^* = j^*$), we know $\widehat{i}^* \neq \operatorname{argmax}_i b_i \rho_i \triangleq k$, and thus we first bound $\operatorname{smax}_{i \in [N]} b_i \rho_i$ by $b_k \rho_k$ and then apply [Eq. \(17\)](#) with $i = k$ and [Eq. \(18\)](#) to obtain the following upper bound:

$$b_k \widehat{\rho}_k - \operatorname{smax}_{j \in [N]} b_j \widehat{\rho}_j + \frac{N}{\gamma \varepsilon} + \varepsilon \leq \frac{N}{\gamma \varepsilon} + \varepsilon,$$

where the last step is again because k is not \widehat{i}^* . Picking $\varepsilon = \sqrt{N/\gamma}$, we have thus proven $\operatorname{dec}_\gamma(\widehat{\rho}) \leq 2\sqrt{N/\gamma}$.

The second statement of theorem is then a direct consequence. Indeed, following the analysis of [Proposition 5.1](#), we know that [Algorithm 2](#) ensures

$$\mathbb{E}[\operatorname{Reg}] \leq \gamma \mathbf{Reg}_{\text{Sq}} + 2 \sum_{t=1}^T \sqrt{\frac{N}{\gamma}}$$

where $\gamma = N/\varepsilon^2$. Plugging the value of ε finishes the proof. □