

Uplifting Bandits

Yu-Guan Hsieh*

Shiva Prasad Kasiviswanathan[†]

Branislav Kveton[‡]

Abstract

We introduce a multi-armed bandit model where the reward is a sum of multiple random variables, and each action only alters the distributions of some of them. After each action, the agent observes the realizations of all the variables. This model is motivated by marketing campaigns and recommender systems, where the variables represent outcomes on individual customers, such as clicks. We propose UCB-style algorithms that estimate the *uplifts* of the actions over a baseline. We study multiple variants of the problem, including when the baseline and affected variables are unknown, and prove sublinear regret bounds for all of these. We also provide lower bounds that justify the necessity of our modeling assumptions. Experiments on synthetic and real-world datasets show the benefit of methods that estimate the uplifts over policies that do not use this structure.

*Université Grenoble Alpes, yu-guan.hsieh@univ-grenoble-alpes.fr. Work done during internship at Amazon.

[†]Amazon, kasivisw@gmail.com.

[‡]Amazon, bkveton@amazon.com.

Contents

1	Introduction	3
2	Problem Description	4
3	Case of Known Baseline and Known Affected Variables	6
4	Lower Bounds	8
5	Case of Unknown Baseline	9
6	Case of Unknown Affected Variables	11
6.1	Known Upper Bound on Number of Affected Variables	11
6.1.1	Known Baseline Payoffs	11
6.1.2	Unknown Baseline Payoffs	11
6.1.3	Regret Bounds	12
6.2	Known Lower Bound on Individual Uplift	13
6.2.1	Known Baseline Payoffs	13
6.2.2	Unknown Baseline Payoffs	15
7	Numerical Experiments	16
7.1	Ablation Study on Gaussian Uplifting Bandit	17
8	Contextual Extensions	18
9	Related Work	20
10	Concluding Remarks	22
A	Proofs for Upper Bound Results	25
A.1	Concentration Bounds	25
A.2	Unknown Baseline: Proof of Theorem 1	26
A.3	Unknown Affected Variables: Proof of Theorem 2	28
A.4	Unknown Affected Variables: Proofs of Theorems 3 and 4	32
B	Lower Bounds	35
B.1	A General Information-theoretic Lower Bound	35
B.2	Lower Bounds for Uplifting Bandits– Proof	36
C	Contextual Combinatorial Uplifting Bandits	38
C.1	Model and Algorithm Template	38
C.2	Linear Expected Payoffs	38
C.3	Experiments for Targeted Campaign with Criteo Uplifting Dataset	42
D	Additional Experimental Details	44
D.1	Choice of Algorithm Parameters	45
D.2	More Details on Bernoulli Uplifting Bandit	46

1 Introduction

Multi-armed bandit (MAB) is an important framework for sequential decision making under uncertainty [Lai and Robbins, 1985, Bubeck and Cesa-Bianchi, 2012, Lattimore and Szepesvári, 2020]. In this problem, a learner repeatedly takes action and receives their rewards, while the outcomes of the other actions are unobserved. The goal of the learner is to maximize their cumulative reward over time by balancing exploration (select actions with uncertain reward estimates) and exploitation (select actions with high reward estimates). MAB has applications in online advertisement, recommender systems, portfolio management, and dynamic channel selection in wireless networks.

One prominent question in the MAB literature is how the dependencies between the actions can be exploited to improve statistical efficiency. Popular examples working toward this goal include *linear bandits* [Dani et al., 2008] and *combinatorial bandits* [Cesa-Bianchi and Lugosi, 2012]. In this work, we study a different structured bandit problem with the following three features: (i) the reward is the sum of the payoffs of a fixed set of variables; (ii) these payoffs are observed; and (iii) each action only affects a small subset of the variables. This structure arises in many applications, as discussed below.

(a) Online Marketing. An ecommerce platform can opt among several marketing strategies (actions) to encourage customers to make more purchases on their website. As different customers can be sensitive to a different marketing strategy, regarding each of them as a variable, it is natural to expect that each action would likely affect only a subset of the variables. The payoff associated to a customer can be for example the revenue generated by that customer; then the reward is just the total revenue received by the platform.

(b) Product Discount. Consider a company that uses discount strategies to increase their sales. It is common to design discount strategies that only apply to a small subset of products. In this case, we can view the sales of each product as a variable and each discount strategy as an action, and assume that each action only has a significant impact on the sales of those products discounted by this action.

(c) Movie Recommendation. Consider a bandit model for movie recommendation where actions correspond to different recommendation algorithms, and the variables are all the movies in the catalog. For each user, define a set of binary payoffs that indicate whether the user watches a movie in the catalog, and the reward is the number of movies from the catalog that this user watches. Since a recommendation (action) for this user contains (promotes) only a subset of movies, it is reasonable to assume that only their associated payoffs are affected by the action.

Our Contributions. To begin with, we formalize an *uplifting bandit* model that captures the aforementioned structure. The term *uplift* is borrowed from the field of *uplift modeling* [Radcliffe, 2007, Gutierrez and Gérardy, 2017] and indicates that the actions’ effects are incremental over a certain baseline. We consider a stochastic model, that is, the payoffs of the variables follow an unknown distribution that depends on the chosen action, and study this problem under various assumptions on the learners’ prior knowledge about (1) the baseline payoff of each variable and (2) the set of affected variables of each action. Our first result (Section 3) shows that when both (1) and (2) are known, a simple modification of the upper confidence bound (UCB) algorithm [Auer, 2002] (Algorithm UPUCB (bl)) for estimating the uplifts has a much lower regret than its standard implementation. Roughly speaking, when m is the number of variables and L is the number of variables affected by each action, we get a $O(L^2)$ regret bound instead of $O(m^2)$. This results in major regret reduction for $L \ll m$, and is a distinguishing feature of all our results.

Going one step further, we design algorithms that have minimax optimal regret bounds without assuming that either (1) or (2) is known in Sections 5 and 6. When (1) the baseline payoffs are unknown and (2) the

Algorithm	UCB	UpUCB (bl)	UpUCB	UpUCB-nAff	UpUCB-iLift
Affected variables known	No	Yes	Yes	No	No
Baseline payoffs known	No	Yes	No	No	No
Regret Bound	$\frac{Km^2}{\Delta}$	$\frac{KL^2}{\Delta}$		$\frac{KL^2}{\Delta}$	$\frac{K \text{ clip}(\Delta/\Delta_{\text{up}}, L, m)^2}{\Delta}$

Table 1: Summary of our regret bounds for uplifting bandits. Constant and logarithmic factors are ignored throughout. For simplicity, we assume here all actions affect exactly L of the m variables and all the suboptimal actions have the same suboptimality gap Δ . K and Δ_{up} are respectively the number of actions and a lower bound on individual uplift (Δ_{up} is formally introduced in Section 6.2). The operator clip restricts the value of its first variable to the range defined by its second and third variable, i.e., $\text{clip}(x, \alpha, \beta) = \max(\alpha, \min(\beta, x))$.

affected variables are known, we compute differences of UCBs to estimate the uplifts (Algorithm UpUCB). In contrast to standard UCB methods, these differences are *not* optimistic; this is because the feedback of the actions also reveals important information about the baseline. When (2) the affected variables are also unknown, we identify them on the fly to maintain suitable estimates of the uplifts. We study two approaches, which differ in what they know about the affected variables. Algorithm UpUCB-nAff knows an upper bound on number of affected variables, whereas Algorithm UpUCB-iLift knows a lower bound on individual uplift. Our regret bounds are summarized in Table 1. These results are further completed with lower bounds that justify the need for our modeling assumptions (Section 4).

Organization. We introduce our uplifting bandit model along with its various variations in Section 2. Over Sections 3, 5, 6, we provide regret bounds for these variations. Lower bounds and numerical experiments are presented in Section 4 and Section 7. To demonstrate the generality of our setup and how our algorithmic ideas extend beyond vanilla multi-armed bandits, we further discuss contextual extensions of our model in Section 8. Comparison to related work is deferred to Section 9.

2 Problem Description

We start by formally introducing our uplifting bandit model. We illustrate it in Figure 1 and summarize our notations in Table 2. Contextual extensions of this model are discussed in Section 8.

A (K, m) -uplifting bandit is a stochastic bandit with K actions and m underlying variables. Each action $a \in \mathcal{A} := \{1, \dots, K\}$ is associated with a distribution \mathcal{P}^a on \mathbb{R}^m . At each round t , the learner chooses an action $a_t \in \mathcal{A}$ and receives reward $r_t = \sum_{i \in \mathcal{V}} y_t(i)$ where $\mathcal{V} := \{1, \dots, m\}$ the set of all variables and $y_t = (y_t(i))_{i \in \mathcal{V}} \sim \mathcal{P}^{a_t}$ is the payoff vector.¹ Our model is distinguished by the two assumptions that we describe below.

(I) Limited Number of Affected Variables. Let $\mathcal{V}^a \subseteq \mathcal{V}$ be the subset of variables affected by action a and \mathcal{P}^0 be the baseline distribution that the variables’ payoffs follow when no action is taken. By definition \mathcal{P}^a and \mathcal{P}^0 have the same marginal distribution on $\overline{\mathcal{V}}^a := \mathcal{V} \setminus \mathcal{V}^a$, the variables unaffected by action a .² While the above

¹The terms *reward* and *payoff* distinguish r_t and $(y_t(i))_{i \in \mathcal{V}}$.

²If the actions in fact have small impact on the variables in $\overline{\mathcal{V}}^a$, our model is misspecified and incurs additional linear regret whose size is proportional to the impact of the actions on $\overline{\mathcal{V}}^a$. An interesting question is how we can design algorithms that self-adapt to the degree of misspecification.

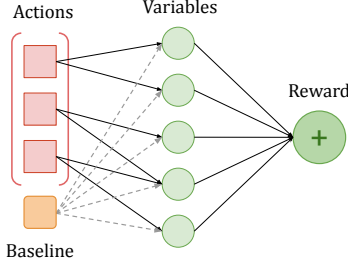
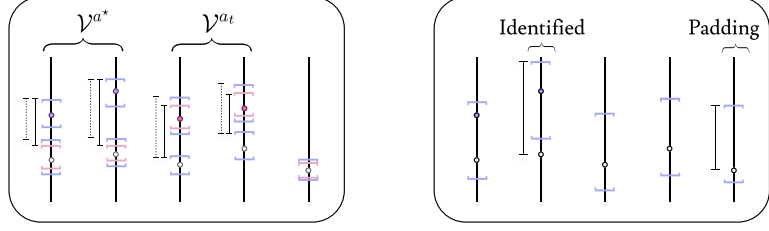


Figure 1: Illustration of the uplifting bandit model. This example has $K = 3$ arms, $m = 5$ variables and each action affects $L^a = 2$ variables. Dash lines indicate the variables' payoffs follow the baseline distribution \mathcal{P}^0 by default.



(a) UPUCB: After suboptimal action a_t is taken, the confidence intervals of $(\mu^a(i))_{i \in \mathcal{V}^{a_t}}$ and $(\mu^0(i))_{i \in \mathcal{V}^{a^*} \setminus \mathcal{V}^{a_t}}$ shrink (from purple to pink). Hence the uplifting index of a_t decreases while that of a^* increases (from dash to solid).

(b) UPUCB-nAff (bl): To compute the uplifting index of an action, we identify a set of variables whose associated confidence intervals do not contain the baseline payoff and pad it with the variables with the largest UCBs on uplifts.

Figure 2: Explanation about UPUCB and UPUCB-nAff (bl) using the model illustrated in Figure 1.

condition is always satisfied with $\mathcal{V}^a = \mathcal{V}$, we are interested in the case of $L^a := \text{card}(\mathcal{V}^a) \ll m$, meaning only a few variables are affected by a . We define L as a uniform bound on L^a , so that $\max_{a \in \mathcal{A}} L^a \leq L$. For convenience of notation, we write $\mathcal{A}_0 = \mathcal{A} \cup \{0\}$ and $[T] := \{1, \dots, T\}$.

(II) Observability of Individual Payoff. In addition to the reward r_t , we assume that the learner observes all the payoffs $(y_t(i))_{i \in \mathcal{V}}$ after an action is taken in round t .

Uplift and Noise. Let $y^a = (y^a(i))_{i \in \mathcal{V}}$ be a random variable with distribution \mathcal{P}^a . We define $\mu^a(i) = \mathbb{E}[y^a(i)]$ and $\xi^a(i) = y^a(i) - \mu^a(i)$ respectively as the expected value of and the noise in $y^a(i)$. We use μ^a (resp. $\mu^0(i)$) to denote the vector of $(\mu^a(i))_{i \in \mathcal{V}}$ (resp. $(\mu^0(i))_{i \in \mathcal{V}}$), and refer to μ^0 as the baseline payoffs. The *individual* uplift associated to a pair $(a, i) \in \mathcal{A} \times \mathcal{V}$ is defined as $\mu_{\text{up}}^a(i) = \mu^a(i) - \mu^0(i)$. An individual uplift can be *positive* or *negative*. We obtain the (total) uplift of an action by summing its individual uplifts over all the variables affected by that action

$$r_{\text{up}}^a = \sum_{i \in \mathcal{V}^a} \mu_{\text{up}}^a(i) = \sum_{i \in \mathcal{V}^a} (\mu^a(i) - \mu^0(i)). \quad (1)$$

Let $r^a = \sum_{i \in \mathcal{V}} \mu^a(i)$ be the expected reward of an action or of pure observation. We also have $r_{\text{up}}^a = r^a - r^0$ since $\mu^a(i) = \mu^0(i)$ as long as $i \notin \mathcal{V}^a$.

A real-value random variable X is said to be σ -sub-Gaussian if for all $\gamma \in \mathbb{R}$, it holds $\mathbb{E}[\exp(\gamma X)] \leq \exp(\sigma^2 \gamma^2 / 2)$. Throughout the paper, we assume that $\xi^a(i)$ is 1-sub-Gaussian for all $a \in \mathcal{A}_0$ and $i \in \mathcal{V}$. Note that we do not assume that $(y^a(i))_{i \in \mathcal{V}}$ are independent, i.e., the elements in the noise vector $(\xi^a(i))_{i \in \mathcal{V}}$ may be correlated, for the following two reasons:

- The independence assumption is not always realistic. In our first example, it excludes any potential correlation between two customers' purchases.
- While a learner can exploit their knowledge on the noise covariance matrix to reduce the regret, as for example shown in [Degenne and Perchet, 2016], incorporating such knowledge complicates the algorithm design and the accompanying analysis. We however believe this is an interesting future direction to pursue.

Regret. The learner's performance is characterized by their regret. To define it, we denote by $a^* \in \arg \max_{a \in \mathcal{A}} r^a$ an action with the largest expected reward and $r^* = r^{a^*} = \max_{a \in \mathcal{A}} r^a$ the largest expected reward. The regret of the learner after T rounds is then given by

$$\text{Reg}_T = r^* T - \sum_{t=1}^T r^{a_t} = \sum_{a \in \mathcal{A}} \sum_{t=1}^T \mathbb{1}\{a_t = a\} \Delta^a, \quad (2)$$

where $\Delta^a = r^* - r^a$ is the suboptimality gap of action a . In words, we compare the learner's cumulative (expected) reward against the best we can achieved by taking an optimal action at each round. For posterity, we write $\Delta = \min_{a \in \mathcal{A}, \Delta^a > 0} \Delta^a$ for the minimum non-zero suboptimality gap and refer to it as the suboptimality gap of the problem.

UCB for Uplifting Bandits. The UCB algorithm [Auer, 2002], at each round, constructs an *upper confidence bound* (UCB) on the expected reward of each action and chooses the action with the highest UCB. When applied to our model, we get a regret of $\mathcal{O}(Km^2 \log T / \Delta)$, as the noise in the reward is m -sub-Gaussian (recall that we do not assume independence of the payoffs). However, this approach completely ignores the structure of our problem. As we show in Section 3, we can achieve much smaller regret by focusing on the uplifts.

Problem Variations. In the following sections, we study several variants of the above basic problem that differ in the prior knowledge about $(\mathcal{P}^a)_{a \in \mathcal{A}_0}$ that the learner possesses.

(a) Knowledge of Baseline Payoffs. We consider two scenarios based on whether the learner knows the baseline payoffs $\mu^0 := (\mu^0(i))_{i \in \mathcal{V}}$ or not.

(b) Knowledge of Affected Variables. We again consider two scenarios based on whether the learner knows the affected variables associated with each action $(\mathcal{V}^a)_{a \in \mathcal{A}}$ or not.

3 Case of Known Baseline and Known Affected Variables

We start with the simplest setting where both the affected variables and the baseline payoffs are known. To address this problem, we make the crucial observation that for any two actions, the difference in their rewards equals to that of their uplifts. As an important consequence, uplift maximization has the same optimal action as total reward maximization, and we can replace rewards by uplifts in the definition of the regret (2). Formally, we write $r_{\text{up}}^* = r_{\text{up}}^{a^*} = \max_{a \in \mathcal{A}} r_{\text{up}}^a$ for the largest uplift; then

$$\text{Reg}_T = r_{\text{up}}^* T - \sum_{t=1}^T r_{\text{up}}^{a_t}.$$

By making this transformation, we gain in statistical efficiency because $r_{\text{up}}^a = \mathbb{E}[\sum_{i \in \mathcal{V}^a} (y^a(i) - \mu^0(i))]$ can now be estimated much more efficiently under our notion of sparsity. Since both μ^0 and $(\mathcal{V}^a)_{a \in \mathcal{A}}$ are known, we can directly construct a UCB on r_{up}^a . For this, we define for all rounds $t \in [T]$, actions $a \in \mathcal{A}$, and variables $i \in \mathcal{V}$ the following quantities.

$$N_t^a = \sum_{s=1}^t \mathbb{1}\{a_s = a\}, \quad c_t^a = \sqrt{\frac{2 \log(1/\delta')}{N_t^a}}, \quad \hat{\mu}_t^a(i) = \sum_{s=1}^t \frac{y_s(i) \mathbb{1}\{a_s = a\}}{\max(1, N_t^a)}, \quad (3)$$

Algorithm 1 UPUCB (bl)

- 1: **Input:** Error probability δ' , Baseline payoffs μ^0 , Sets of affected variables $\{\mathcal{V}^a : a \in \mathcal{A}\}$
 - 2: **Initialization:** Take each action once
 - 3: **for** $t = K + 1, \dots, T$ **do**
 - 4: Compute the statistics following (3)
 - 5: For $a \in \mathcal{A}$, compute uplifting index $\tau_t^a \leftarrow \sum_{i \in \mathcal{V}^a} (\hat{\mu}_{t-1}^a(i) + c_{t-1}^a - \mu^0(i))$
 - 6: Select action $a_t \in \arg \max_{a \in \mathcal{A}} \tau_t^a$
 - 7: **end for**
-

where $\delta' > 0$ is a tunable parameter. In words, N_t^a , $\hat{\mu}_t^a(i)$, and c_t^a represent respectively the number of times that action a is taken, the empirical estimate of $\mu^a(i)$, and the associated radius of confidence interval calculated at the end of round t . The UCB for action $a \in \mathcal{A}$ and variable $i \in \mathcal{V}^a$ at round t is $U_t^a(i) = \hat{\mu}_{t-1}^a(i) + c_{t-1}^a$. We further define $\tau_t^a = \sum_{i \in \mathcal{V}^a} (U_t^a(i) - \mu^0(i))$ as the *uplifting index*, and refer to UPUCB (bl) as the algorithm that takes an action with the largest uplifting index τ_t^a at each round t (Algorithm 1). Here, the suffix (bl) indicates that the method operates with the knowledge of the baseline payoffs.

Since UPUCB (bl) is nothing but a standard UCB with transformed rewards $r'_t = \sum_{i \in \mathcal{V}_{a_t}} (y_t(i) - \mu^0(i))$, and $\mathbb{E}[r'_t] = r_{\text{up}}^{a_t}$ defines exactly the same regret as the original reward, a standard analysis for UCB yields the following result.

Proposition 1. *Let $\delta' = \delta/(2KT)$. Then the regret of UPUCB (bl) (Algorithm 1), with probability at least $1 - \delta$, satisfies:*

$$\text{Reg}_T \leq \sum_{a \in \mathcal{A}: \Delta^a > 0} \left(\frac{8(L^a)^2 \log(2KT/\delta)}{\Delta^a} + \Delta^a \right). \quad (4)$$

As expected, the regret bound does not depend on m and scales with L^2 . This is because the transformed reward of action a is only L^a -sub-Gaussian. The improvement is significant when $L \ll m$. However, this also comes at a price, as both the baseline payoffs and affected variables need to be known. We address these shortcomings in Sections 5 and 6.

Estimating Baseline Payoffs from Observational Data. In practice, the baseline payoffs μ^0 can be estimated from observational data, which gives confidence intervals where $\mu^0(i)$ lie with high probability. The uplifting indices can then be constructed by subtracting the lower confidence bounds of $\mu^0(i)$ from $U_t^a(i)$. If the number of samples is n , the widths of the confidence intervals are in $\mathcal{O}(1/\sqrt{n})$. Identification of a^* is possible only when the sum of these widths over L variables is at most Δ , which requires $n = \Omega(L^2/\Delta^2)$. Otherwise, these errors persist at each iteration and n must be in the order of T to ensure $\mathcal{O}(\sqrt{T})$ regret.

About Expected and Gap-Free Bounds. In Proposition 1, we state a high-probability instance-dependent regret bound, and we will continue to do so for all the regret upper bounds that we present for the non-contextual variants. This type of result can be directly transformed to bound on $\mathbb{E}[\text{Reg}_T]$ by taking $\delta = 1/T$. Following the routine of separating Δ^a into two groups depending on their scale, most of our proofs can also be easily modified to obtain a gap-free regret bound, which is typically in the order of $\mathcal{O}(L\sqrt{KT \log(T)})$. We will not state these results to avoid unnecessary repetitions.

4 Lower Bounds

In this section, we discuss the necessity of our modeling assumptions for obtaining the improved regret bounds of [Proposition 1](#). We prove these results in [Appendix B](#), where we also introduce a general information-theoretic lower bounds for bandit problems with similar design.

Intuitively, the regret can be improved both because the noise in the effect of an action is small, and because the observation of this effect does not heavily deteriorate with noise. These two points correspond respectively to assumptions [\(I\)](#) and [\(II\)](#). Moreover, the knowledge on $(\mathcal{V}^a)_{a \in \mathcal{A}}$ allows the learner to distinguish between problems with different structures. Without such distinction, there is no chance that the learner can leverage the underlying structure. Therefore, the aforementioned three points are crucial for obtaining [\(4\)](#). Below, we establish this formally for algorithms that are *consistent* [[Lai and Robbins, 1985](#)] over the class of 1-sub-Gaussian uplifting bandits, which means the induced regret of the algorithm on any uplifting bandit with 1-sub-Gaussian noise satisfies $\text{Reg}_t = o(t^p)$ for all $p > 0$.

For conciseness, in the following we say that an uplifting bandit has parameters $(K, m, (\Delta^a, L^a)_{a \in \mathcal{A}})$ if the number of actions, the number of variables, the suboptimality gaps, and the number of affected variables of this instance are respectively $K, m, (\Delta^a)_{a \in \mathcal{A}}$, and $(L^a)_{a \in \mathcal{A}}$. We focus on instance-dependent lower-bounds and assume that the learner has full knowledge of the baseline distribution. (Of course, the problem only becomes more challenging if the learner does not know the baseline distribution.)

Proposition 2. *Let π be a consistent algorithm over the class of 1-sub-Gaussian uplifting bandits that at most uses knowledge of \mathcal{P}^0 , $(\mathcal{V}^a)_{a \in \mathcal{A}}$, and the fact that the noise is 1-sub-Gaussian. Let $K, m > 0$ and sequence $(L^a, \Delta^a)_{1 \leq a \leq K} \in ([m] \times \mathbb{R}_+)^K$ satisfy $\Delta^1 = 0$. Assume either of the following holds.*

- (a) $L^a = m$ for all $a \in \mathcal{A}$, so that in the bandits considered below all actions affect all variables.
- (b) Only the reward is observed.
- (c) The algorithm π does not make use of any prior knowledge about the arms' expected payoffs $(\mu^a)_{a \in \mathcal{A}}$ (in particular, the knowledge of $(\mathcal{V}^a)_{i \in \mathcal{V}}$ is not used by π).

Then, there exists a 1-sub-Gaussian uplifting bandit with parameters $(K, m, (\Delta^a, L^a)_{a \in \mathcal{A}})$ such that the regret induced by π on it satisfies:

$$\liminf_{T \rightarrow +\infty} \frac{\mathbb{E}[\text{Reg}_T]}{\log T} \geq \sum_{a \in \mathcal{A}: \Delta^a > 0} \frac{2m^2}{\Delta^a}.$$

As discussed above, [Proposition 2](#) shows that there would be a higher regret in absence of any of the following: (a) Limited number of affected variables, (b) Side information about variables beyond reward, and (c) Some prior knowledge on how the variables are affected. We discuss below these three points in detail.

(a) This point justifies the necessity of assumption [\(I\)](#). It can be further refined to account for non-uniform numbers of affected variables.

Proposition 3. *Let π be a consistent algorithm over the class of 1-sub-Gaussian uplifting bandits that at most uses knowledge of \mathcal{P}^0 , $(\mathcal{V}^a)_{a \in \mathcal{A}}$, and the fact that the noise is 1-sub-Gaussian. Then, for any $K, m > 0$ and sequence $(L^a, \Delta^a)_{1 \leq a \leq K} \in ([m] \times \mathbb{R}_+)^K$ with $\Delta^1 = 0$, there exists a 1-sub-Gaussian uplifting bandit $(K, m, (\Delta^a, L^a)_{a \in \mathcal{A}})$ such that the regret induced by π on it satisfies*

$$\liminf_{T \rightarrow +\infty} \frac{\mathbb{E}[\text{Reg}_T]}{\log T} \geq \sum_{a \in \mathcal{A}: \Delta^a > 0} \frac{2(L^a)^2}{\Delta^a}. \quad (5)$$

[Proposition 3](#) essentially posits that the regret must scale with the magnitudes of the noises. The stated lower bound (5) matches an expected version of (4) up to a constant factor, suggesting the minimax optimality of [Proposition 4](#). One way to argue this lower bound holds is by reducing a carefully designed uplifting bandits with parameters $(K, m, (\Delta^a, L^a)_{a \in \mathcal{A}})$ to a K -arm bandit with the corresponding noise scales. We will take a different approach in [Appendix B](#) that allows us prove all the lower bounds following the same schema.

(b) This point partially justifies assumption [\(II\)](#). Note that it is however not always necessary to observe all the payoffs. For example, observing the aggregated payoff $\sum_{i \in \mathcal{V}^{a_t}} y_t(i)$ is sufficient for computing the uplifting indices of UPUCB (bl). Having observation of the entire payoff vector $(y_t(i))_{i \in \mathcal{V}}$ is particularly relevant for the case of unknown affected variables since it helps identify the affected variables relatively easily, as we will see in [Section 6](#).

(c) This point justifies the learner using knowledge about the affected variables. To get some more intuition on the meaning of this condition, let us consider two bandits with exactly the same baseline and noise distributions and very close expected payoffs of the arms. Then, with finite observations, they provide very similar feedback which prevents the learner from distinguishing between the two problems. However, the two bandits may have completely different structures, and in particular, one of them may have all variables affected by all actions since slight perturbation to the baseline would mean that a variable is affected. Therefore, the lower bound from [\(a\)](#) applies to this instance. Now, provided that the learner is doomed to treat the two instances equally, the lower bound applies to the other instance as well. Of course, one can argue that it is unnatural that the regrets of two similar problems differ significantly even when the learner is provided some prior knowledge. We believe this is the drawback of the asymptotic analysis that we look at here, since these differences may indeed be small within a reasonable number of runs. This is exactly related to the problem of misspecified model that we mentioned in [Footnote 2](#).

With [Proposition 2](#), what remains unclear at this point is whether similar improvement of the regret is still possible when the baseline is unknown or when the learner only has access to more restricted knowledge than $(\mathcal{V}^a)_{a \in \mathcal{A}}$. We give affirmative answers to these two questions in the next two sections.

Remark 1. Our lower bounds are derived with respect to the worst correlation structure of the noise. It is also possible to obtain lower bounds that depend on the covariance of the noise. For example, if the variables $(y^a(i))_{i \in \mathcal{V}}$ are independent, the lower bound in (5) only scales linearly with L^a .

5 Case of Unknown Baseline

In this section, we consider the situation where the sets of affected variables $(\mathcal{V}^a)_{a \in \mathcal{A}}$ are known by the learner, but not the baseline payoffs. Missing proofs are in [Appendix A](#).

Since the actual uplift at each round is never observed, the uplifts of the actions can hardly be estimated directly without knowing the baseline payoffs. Instead, we follow a two-model approach. Leveraging the fact that \mathcal{P}^a and \mathcal{P}^0 have the same marginal distribution on $\overline{\mathcal{V}}^a$, we can estimate the baseline payoffs by aggregating information from the feedback of different actions. This leads to

$$N_t^0(i) = \sum_{s=1}^t \mathbb{1}\{i \notin \mathcal{V}^{a_s}\}, \quad c_t^0(i) = \sqrt{\frac{2 \log(1/\delta')}{N_t^0(i)}}, \quad \hat{\mu}_t^0(i) = \frac{\sum_{s=1}^t y_s(i) \mathbb{1}\{i \notin \mathcal{V}^{a_s}\}}{\max(1, N_t^0(i))}. \quad (6)$$

Compared to (3), we notice that both N_t^0 and c_t^0 are functions of i . This is because for each taken action a , we only increase the counters $N_t^0(i)$ for those $i \notin \mathcal{V}^a$, which causes a non-uniform increase of $(N_t^0(i))_{i \in \mathcal{V}}$.

Algorithm 2 UPUCB

- 1: **Input:** Error probability δ' , the sets of variables each action affects $\{\mathcal{V}^a : a \in \mathcal{A}\}$
 - 2: **Initialization:** Take each action once
 - 3: **for** $t = K + 1, \dots, T$ **do**
 - 4: Compute the UCB indices following (7)
 - 5: For $a \in \mathcal{A}$, set $\tau_t^a \leftarrow \sum_{i \in \mathcal{V}^a} (U_t^a(i) - U_t^0(i))$
 - 6: Select action $a_t \in \arg \max_{a \in \mathcal{A}} \tau_t^a$
 - 7: **end for**
-

Then, by looking at all the rounds that variable i is not influenced by the chosen action, we manage to compute $\hat{\mu}_t^0(i)$, an estimate of $\mu^0(i)$. An alternative here is to consider a specific ‘action 0’ and estimate the baseline payoffs only using feedback from this action. However, this necessarily requires action 0 to be taken sufficiently often and would thus result in a very high regret. As we will see, our algorithm does not make use of this specific action and achieves regret comparable to the one presented in [Proposition 1](#). To proceed, we define the following UCB indices for all the pairs $(a, i) \in \mathcal{A}_0 \times \mathcal{V}$

$$U_t^a(i) = \begin{cases} 0 & \text{if } a = 0 \text{ and } i \in \bigcap_{a \in \mathcal{A}} \mathcal{V}^a, \\ \hat{\mu}_{t-1}^a(i) + c_{t-1}^a & \text{if } a \in \mathcal{A} \text{ and } i \in \mathcal{V}^a, \\ \hat{\mu}_{t-1}^0(i) + c_{t-1}^0(i) & \text{otherwise.} \end{cases} \quad (7)$$

The second and the third lines of (7) contain the usual definition of UCBs using the empirical estimates and the radii of the confidence intervals defined in (3) and (6). In the special case that a variable is affected by all the actions (recall that we do not use ‘action 0’), it is impossible to estimate $U^0(i)$ but it is enough to compare $U^a(i)$ directly against $U^{a'}(i)$ for any two actions $a, a' \in \mathcal{A}$, so we just set $U_t^0(i)$ to 0 in this case.

We outline the proposed method, UPUCB, in [Algorithm 2](#). The uplifting indices are given by $\tau_t^a = \sum_{i \in \mathcal{V}^a} (U_t^a(i) - U_t^0(i))$. It may be counter-intuitive to compare the differences between two UCBs. Indeed, τ_t^a is no longer an optimistic estimate of the uplifting effect r_{up}^a , but it captures the essential trade-off between learning action a ’s payoffs and learning the baseline μ^0 . To provide some intuition, we give an informal justification of UPUCB in [Figure 2a](#): If a suboptimal action a is taken in round t , the estimates of all $U_t^a(i)$ move closer to the actual mean from above. As a result, τ_t^a decreases, since all $U_t^a(i)$ for $i \in \mathcal{V}^a$ do. Thus action a is less likely to be taken next. Moreover, $\tau_t^{a^*}$ increases, since $U_t^0(i)$ decrease for any i affected by a^* but not a . Thus a^* is more likely to be taken next. The effectiveness of UPUCB is confirmed by the following theorem.

Theorem 1. *Let $\delta' = \delta/(4KLT)$. Then the regret of UPUCB ([Algorithm 2](#)) with probability at least $1 - \delta$, satisfies:*

$$\text{Reg}_T \leq \sum_{a \in \mathcal{A}: \Delta^a > 0} \left(\frac{8(L^a + L^{a^*})^2 \log(4KLT/\delta)}{\Delta^a} + \Delta^a \right). \quad (8)$$

As in [Proposition 1](#), the regret is $\mathcal{O}(KL^2 \log T / \Delta)$. In fact, all decisions in UPUCB are made on uncertain estimates of at most L variables; thus the statistical efficiency scales with L and not m . Moreover, for any action a , the width of the confidence intervals for baseline estimates $c_t^0(i)$ are lower than those of a whenever $i \in \overline{\mathcal{V}^a}$. The factors $c_t^0(i)$ thus play a secondary role in the analysis of the regret. A detailed comparison with [Proposition 1](#) reveals that the difference is in the order of $K(L^{a^*})^2 \log T / \Delta$; this is only significant when the optimal actions affect many more variables than the suboptimal ones. Hence, the price of not knowing the baseline payoffs is generally quite small.

6 Case of Unknown Affected Variables

Now we study the more challenging setting where the affected variables $(\mathcal{V}^a)_{a \in \mathcal{A}}$ are unknown to the learner. [Proposition 2](#) states that improvement is impossible if the learner does not possess any prior knowledge that allow them to exploit the structure. To circumvent this negative result, we study two weak assumptions motivated by practice: the learner has either access to (i) an upper bound on the number of affected variables or (ii) a lower bound on individual uplift.

6.1 Known Upper Bound on Number of Affected Variables

Here we assume that we have a known upper bound on number of affected variables L (i.e., $L \geq \max_{a \in \mathcal{A}} L^a$). We design algorithms with $\mathcal{O}(KL^2 \log T / \Delta)$ regret bounds that takes L as input. We consider the cases of known and unknown baseline payoffs.

6.1.1 Known Baseline Payoffs

To illustrate our ideas, we start by assuming that the baseline payoffs are known. We propose an optimistic algorithm that maintains a UCB on the total uplift with an overestimate in the order of L . Let N_t^a , c_t^a , and $\hat{\mu}_t^a(i)$ be defined as in (3). The UCB, uplifting indices, and the confidence intervals for each (action, variable) pair $(a, i) \in \mathcal{A} \times \mathcal{V}$ are

$$\begin{aligned} U_t^a(i) &= \hat{\mu}_{t-1}^a(i) + c_{t-1}^a, \quad \rho_t^a(i) = U_t^a(i) - \mu^0(i), \\ \mathcal{C}_t^a(i) &= [\hat{\mu}_{t-1}^a(i) - c_{t-1}^a, \hat{\mu}_{t-1}^a(i) + c_{t-1}^a]. \end{aligned} \quad (9)$$

In the sequel, we refer to $\rho_t^a(i)$ as the *individual uplifting index* of the pair (a, i) . It serves as an estimate of the individual uplift $\mu_{\text{up}}^a(i)$. Our algorithm, UPUCB-nAff(bl), leverages two important procedures to compute an optimistic estimate of the uplift r_{up}^a : identification of affected variables, and padding with variables with the highest individual uplifting indices.

To begin, in line 6, UPUCB-nAff(bl) constructs the set of *identified variables*

$$\widehat{\mathcal{V}}_t^a = \{i \in \mathcal{V} : \mu^0(i) \notin \mathcal{C}_t^a(i)\}.$$

This set is contained in \mathcal{V}^a with high probability. In fact, by concentration of measure, with high probability $\mu^a(i) \in \mathcal{C}_t^a(i)$, in which case $\mu^0(i) \notin \mathcal{C}_t^a(i)$ indicates $i \in \mathcal{V}^a$. However, $\widehat{\mathcal{V}}_t^a$ is not guaranteed to capture all the affected variables, so we also need to provide an upper bound for $\sum_{i \in \mathcal{V}^a \setminus \widehat{\mathcal{V}}_t^a} \mu_{\text{up}}^a(i)$, the uplift contributed by the unidentified affected variables. Since the individual uplifting index $\rho_t^a(i)$ is in fact a UCB on the individual uplift $\mu_{\text{up}}^a(i)$ here and $\text{card}(\mathcal{V}^a) \leq L$, we can simply choose the $L - \text{card}(\widehat{\mathcal{V}}_t^a)$ variables in $\mathcal{V} \setminus \widehat{\mathcal{V}}_t^a$ with the largest individual uplifting index $\rho_t^a(i)$, as done in line 8 of [Algorithm 3](#). We refer to this set as \mathcal{L}_t^a . We then get a proper UCB on the uplift of action a by computing $\tau_t^a = \sum_{i \in \widehat{\mathcal{V}}_t^a \cup \mathcal{L}_t^a} \rho_t^a(i)$ in line 9. This process is summarized in [Algorithm 3](#) and illustrated in [Figure 2b](#).

6.1.2 Unknown Baseline Payoffs

Now we focus on the most challenging setting, where also the baseline payoffs are unknown. In this case, neither the sets of identified variables nor the uplifting indices of UPUCB-nAff(bl) can be defined. We also cannot estimate the baseline payoffs using (6) since the sets of affected variables are unknown. To overcome these challenges, we note that for any two actions $a, a' \in \mathcal{A}$, μ^a and $\mu^{a'}$ only differ on $\mathcal{V}^a \cup \mathcal{V}^{a'}$,

Algorithm 3 UPUCB-nAff (bl)

```
1: Input: Error probability  $\delta'$ , Baseline payoffs  $\mu^0$ , Upper bound  $L$  on the number of affected variables
2: Initialization: Take each action once
3: for  $t = K + 1, \dots, T$  do
4:   for  $a \in \mathcal{A}$  do
5:     Compute UCBs, uplifting indices, and confidence intervals following (9)
6:     Set  $\widehat{\mathcal{V}}_t^a \leftarrow \{i \in \mathcal{V} : \mu^0(i) \notin \mathcal{C}_t^a(i)\}$ 
7:     Set  $L_t^a \leftarrow \max(0, L - \text{card}(\widehat{\mathcal{V}}_t^a))$ 
8:     Set  $\mathcal{L}_t^a \leftarrow \arg \max_{\substack{\mathcal{L} \subseteq \mathcal{V} \setminus \widehat{\mathcal{V}}_t^a \\ \text{card}(\mathcal{L}) = L_t^a}} \sum_{i \in \mathcal{L}} \rho_t^a(i)$ 
9:     Compute uplifting index  $\tau_t^a \leftarrow \sum_{i \in \widehat{\mathcal{V}}_t^a \cup \mathcal{L}_t^a} \rho_t^a(i)$ 
10:   end for
11:   Select action  $a_t \in \arg \max_{a \in \mathcal{A}} \tau_t^a$ 
12: end for
```

and $\text{card}(\mathcal{V}^a \cup \mathcal{V}^{a'}) \leq 2L$. Therefore, μ^a and $\mu^{a'}$ differ in at most $2L$ variables, and we recover a similar problem structure by taking the payoffs of any action as the baseline.

Combining this idea with the elements that we have introduced previously, we obtain UPUCB-nAff (Algorithm 4). In each round, UPUCB-nAff starts by picking a most frequently taken action b_t (Line 4) whose payoffs are treated as the baseline during this round. Then, in Line 6, UPUCB-nAff chooses variables that are guaranteed to be either in \mathcal{V}^a or \mathcal{V}^{b_t} . This generalizes the identification step of UPUCB-nAff (bl). The individual uplifting indices $\rho_t^a(i) = U_t^a(i) - U_t^{b_t}(i)$ are computed in Line 7. The definition as differences of UCBs is inspired by similar construction in UPUCB. Lines 8 and 9 constitute the padding step, during which variables with the highest uplifting indices are selected, and finally in Line 10 we combine the above to get the uplifting index of the action.

The similarity between UPUCB-nAff and UPUCB-nAff (bl) can be shown by supposing that one action has been taken frequently. Then the baseline payoffs are precisely estimated and do not change much between consecutive rounds.

6.1.3 Regret Bounds

Both UPUCB-nAff (bl) and UPUCB-nAff choose $\mathcal{O}(L)$ variables for estimating the uplift of an action, and the decisions are made based on these estimates. Therefore, the statistical efficiencies of these algorithms only scale with L and not m . This in turn translates into an improvement of the regret, as demonstrated by the theorem below.

Theorem 2. *Let $\delta' = \delta/(2KmT)$. Then the regret of UPUCB-nAff (bl) (Algorithm 3) (resp. UPUCB-nAff, Algorithm 4), with probability at least $1 - \delta$, satisfies:*

$$\text{Reg}_T \leq \sum_{a \in \mathcal{A} : \Delta^a > 0} \left(\frac{\alpha L^2 \log(2KmT/\delta)}{\Delta^a} + \Delta^a \right), \quad (10)$$

where $\alpha = 32$ (resp. 512) in the above inequality.

The proof of Theorem 2 (Appendix A.3) is notable for two reasons. First, tracking of the identified variables guarantees that the uplifting index τ_t^a does not overestimate the uplift r_{up}^a much. Take UPUCB-nAff (bl) as an example. An alternative to constructing a UCB on r_{up}^a is to choose the L variables with

Algorithm 4 UPUCB-nAff

```
1: Input: Error probability  $\delta'$ , Upper bound  $L$  on the number of variables that each action affects
2: Initialization: Take each action once
3: for  $t = K + 1, \dots, T$  do
4:   Choose  $b_t \in \arg \max_{a \in \mathcal{A}} N_{t-1}^a$  and compute UCBs and confidence intervals using (9)
5:   for  $a \in \mathcal{A}$  do
6:     Set  $\widehat{\mathcal{V}}_t^a \leftarrow \{i \in \mathcal{V} : \mathcal{C}_t^a(i) \cap \mathcal{C}_t^{b_t}(i) = \emptyset\}$ 
7:     For  $i \in \mathcal{V}$ , compute  $\rho_t^a(i) \leftarrow U_t^a(i) - U_t^{b_t}(i)$ 
8:     Set  $L_t^a \leftarrow \max(0, 2L - \text{card}(\widehat{\mathcal{V}}_t^a))$ 
9:     Set  $\mathcal{L}_t^a \leftarrow \arg \max_{\substack{\mathcal{L} \subseteq \mathcal{V} \setminus \widehat{\mathcal{V}}_t^a, \\ \text{card}(\mathcal{L}) \leq L_t^a}} \sum_{i \in \mathcal{L}} \rho_t^a(i)$ 
10:    Compute uplifting index  $\tau_t^a \leftarrow \sum_{i \in \widehat{\mathcal{V}}_t^a \cup \mathcal{L}_t^a} \rho_t^a(i)$ 
11:  end for
12:  Select action  $a_t \in \arg \max_{a \in \mathcal{A}} \tau_t^a$ 
13: end for
```

the highest individual uplifting indices $\rho_t^a(i)$. However, this would result in a severe overestimate when a negative individual uplift is present. Second, to prove (10) for UPUCB-nAff, we use that the widths of confidence intervals of the chosen b_t are always smaller than those of the taken action. This is ensured by taking b_t as the most frequent action (Line 4 in Algorithm 4).

6.2 Known Lower Bound on Individual Uplift

In this second part of this section, we assume a lower bound on individual uplift is known. This means the learner has access to $\Delta_{\text{up}} > 0$ satisfying that for all $a \in \mathcal{A}$ and $i \in \mathcal{V}^a$, $|\mu^a(i) - \mu^0(i)| \geq \Delta_{\text{up}}$. This gives us an indicator on how many times we need to take each action in order to identify all the affected variables. Similar assumption was made by Lu et al. [2021] to identify which node in the causal graph affects the reward variable.

Again in this scenario, we derive regret bounds for both when the baseline payoffs are known or not to the learner. Our algorithms combine UCB, successive elimination [Even-Dar et al., 2006], and the idea that the affected variables can be identified after an action is taken sufficiently many times. Our regret bounds feature the ratio $\Delta^a/\Delta_{\text{up}}$, suggesting that this knowledge is helpful when Δ_{up} is relatively large.

Hereinafter, for any $x, \alpha, \beta \in \mathbb{R}$ with $\alpha \leq \beta$, we use $\text{clip}(x, \alpha, \beta)$ to denote the clipping function that restricts x to the interval $[\alpha, \beta]$, i.e., $\text{clip}(x, \alpha, \beta) = \max(\alpha, \min(\beta, x))$. Missing proofs are collected in Appendix A.4.

6.2.1 Known Baseline Payoffs

As we have seen in Section 6.1.1, if the baseline payoffs μ^0 are known, we can directly check whether $\mu^0(i)$ is in the confidence interval. Moreover, with the knowledge $|\mu^a(i) - \mu^0(i)| \geq \Delta_{\text{up}}$ for all $i \in \mathcal{V}^a$, we deduce that \mathcal{V}^a can be identified correctly with probability $1 - 2m\delta'$ after action a is taken $N_0 = \lceil (8/\Delta_{\text{up}}^2) \log(1/\delta') \rceil$ times. In fact, let $y_1^a, \dots, y_{N_0}^a$ be N_0 i.i.d. realizations of $y^a \sim \mathcal{P}^a$ and $\tilde{\mu}_{N_0}^a(i) = (\sum_{t=1}^{N_0} y_t^a(i))/N_0$ be the

Algorithm 5 UPUCB-iLift (bl)

1: **Input:** Error probability δ' , Baseline payoffs μ^0 , Lower bound on individual uplift Δ_{up}
2: **Initialization:** Take each action once; Set $N_0 \leftarrow \lceil (8/\Delta_{\text{up}}^2) \log(1/\delta') \rceil$ and $\widehat{\mathcal{V}}^a \leftarrow \mathcal{V}$
3: **for** $t = K + 1, \dots, T$ **do**
4: Compute the UCB indices following (9)
5: For $a \in \mathcal{A}$, compute uplifting index $\tau_t^a \leftarrow \sum_{i \in \widehat{\mathcal{V}}^a} (U_t^a(i) - \mu^0(i))$
6: Select action $a_t \in \arg \max_{a \in \mathcal{A}} \tau_t^a$
7: **if** $N_t^{a_t} \geq N_0$ **then**
8: Set $\widehat{\mathcal{V}}^{a_t} \leftarrow \{i \in \mathcal{V} : |\hat{\mu}_t^{a_t}(i) - \mu^0(i)| > \Delta_{\text{up}}/2\}$
9: **end if**
10: **end for**

empirical mean of variable i . Applying the sub-Gaussian concentration inequality (17, Appendix A.1) gives

$$\mathbb{P}\left(|\tilde{\mu}_{N_0}^a(i) - \mu^a(i)| \geq \frac{\Delta_{\text{up}}}{2}\right) \leq 2 \exp\left(-\frac{N_0 \Delta_{\text{up}}^2}{8}\right) \leq 2\delta'. \quad (11)$$

With a union bound we conclude that with probability at least $1 - 2m\delta'$ it holds for all i that $|\tilde{\mu}_{N_0}^a(i) - \mu^a(i)| < \Delta_{\text{up}}/2$. We distinguish between the following two situations

1. $i \in \mathcal{V}^a$, then $|\tilde{\mu}_{N_0}^a(i) - \mu^0(i)| \geq |\mu^a(i) - \mu^0(i)| - |\tilde{\mu}_{N_0}^a(i) - \mu^a(i)| > \Delta_{\text{up}} - \frac{\Delta_{\text{up}}}{2} = \frac{\Delta_{\text{up}}}{2}$.
2. $i \notin \mathcal{V}^a$, then $|\tilde{\mu}_{N_0}^a(i) - \mu^0(i)| = |\tilde{\mu}_{N_0}^a(i) - \mu^a(i)| < \frac{\Delta_{\text{up}}}{2}$.

This implies that with the choice $\widehat{\mathcal{V}}^a = \{i \in \mathcal{V} : |\tilde{\mu}_{N_0}^a(i) - \mu^0(i)| > \Delta_{\text{up}}/2\}$, we have effectively $\mathbb{P}(\widehat{\mathcal{V}}^a = \mathcal{V}^a) \geq 1 - 2m\delta'$. One natural algorithm is to thus first pull each arm N_0 times and run UPUCB (bl) (Algorithm 1) in the remaining rounds by regarding $\widehat{\mathcal{V}}^a$ as the variables that action a affects. A direct computation shows that the regret of this algorithm is in $\mathcal{O}(\sum_{a \in \mathcal{A}: \Delta^a > 0} (\Delta^a/\Delta_{\text{up}} + L^a)^2 \log T/\Delta^a)$ with high probability.

However, N_0 can be arbitrarily large when Δ_{up} gets closer to 0, and taking an action much fewer times is probably sufficient to deduce that action is not optimal (with high probability). We thus propose to implement the above idea in a more flexible manner: We run UPUCB (bl) with the uplifting indices $\tau_t^a = \sum_{i \in \widehat{\mathcal{V}}_t^a} (U_t^a(i) - \mu^0(i))$, where

$$\widehat{\mathcal{V}}_t^a = \begin{cases} \mathcal{V} & \text{if } N_t^a < N_0, \\ \{i \in \mathcal{V} : |\hat{\mu}_t^{a_t}(i) - \mu^0(i)| > \Delta_{\text{up}}/2\} & \text{otherwise.} \end{cases}$$

We refer to this algorithm as UPUCB-iLift (bl), with iLift standing for individual uplift, and summarize it in Algorithm 5. UPUCB-iLift (bl) has the following regret guarantee.

Theorem 3. *Let $\delta' = \delta/(2KmT)$. Then the regret of UPUCB-iLift (bl) (Algorithm 5), with probability at least $1 - \delta$, satisfies:*

$$\text{Reg}_T \leq \sum_{a \in \mathcal{A}: \Delta^a > 0} \left(\frac{8 \text{clip}(\Delta^a/\Delta_{\text{up}}, L^a, m)^2 \log(2KmT/\delta)}{\Delta^a} + \Delta^a \right). \quad (12)$$

Algorithm 6 UpUCB-iLift

```
1: Input: Error probability  $\delta'$ , Lower bound on effect gap  $\Delta_{\text{up}}$ 
2: Initialization: Set  $t \leftarrow 0$ ,  $\mathcal{A}_1 \leftarrow \mathcal{A}$ , and  $N_0 \leftarrow \lceil (32/\Delta_{\text{up}}^2) \log(1/\delta') \rceil$ 
3: Phase I: Successive elimination
4: for  $\ell = 1, \dots, N_0$  do
5:   if  $T - t \geq \text{card}(\mathcal{A}_\ell)$  then
6:     Take each action  $a \in \mathcal{A}_\ell$  once
7:     Set  $t \leftarrow t + \text{card}(\mathcal{A}_\ell)$ 
8:     Compute the empirical estimates  $\tilde{r}_\ell^a$  for  $a \in \mathcal{A}_\ell$  and the radius of confidence interval  $c_\ell$ 
9:     Update active set  $\mathcal{A}_{\ell+1} \leftarrow \{a \in \mathcal{A}_\ell : \tilde{r}_\ell^a + 2c_\ell \geq \max_{a'} \tilde{r}_\ell^{a'}\}$ 
10:  else
11:    Take  $T - t$  actions from  $\mathcal{A}_\ell$  and terminate the algorithm
12:  end if
13: end for
14: Phase II: UpUCB
15: Construct the sets  $\widehat{\mathcal{S}}(i)$  as in (13) and  $\widehat{\mathcal{V}}^a \leftarrow \{i \in \mathcal{V} : a \notin \widehat{\mathcal{S}}(i)\}$ 
16: for  $t = t + 1, \dots, T$  do
17:   Compute the UCB indices following (9),(14)
18:   For  $a \in \mathcal{A}_{N_0+1}$ , compute uplifting index  $\tau_t^a \leftarrow \sum_{i \in \widehat{\mathcal{V}}^a} (U_t^a(i) - U_t^0(i))$ 
19:   Select action  $a_t \in \arg \max_{a \in \mathcal{A}_\ell} \tau_t^a$ 
20: end for
```

[Theorem 3](#) shows that [Algorithm 5](#) provides a smooth transition between standard UCB and the naive strategy that tries to identify the affected variables of all the actions. When $\Delta^a/\Delta_{\text{up}} \geq m$, the loss of taking action a is so large that UCB prevents it from being taken further even if $N_t^a \leq N_0$. Otherwise, the affected variables get identified after an action is taken sufficient number of times and the algorithm benefits from this knowledge to improve the estimate of the uplift. It is also quite straightforward to combine [Algorithm 3](#) with [Algorithm 5](#) when both L and Δ_{up} are known; this results in a regret bound that replaces $\text{clip}(\Delta^a/\Delta_{\text{up}}, L^a, m)$ by $\text{clip}(\Delta^a/\Delta_{\text{up}}, L^a, L)$ in (12).

6.2.2 Unknown Baseline Payoffs

Several additional challenges emerge when we want to adapt the aforementioned strategies to the case where the baseline payoffs are unknown. First, without additional assumption, it is impossible to tell whether two arbitrarily close estimates indicate the same effect of two actions on a variable. Then, without a proper baseline to compare with, we can never exclude the possibility that all the variables are affected by all the actions. To circumvent this issue, we make the following more stringent assumption.

Assumption 1. The effect of any two actions on a variable differ by at least Δ_{up} as long as one of them affects this variable, i.e., for all $a \in \mathcal{A}$, $i \in \mathcal{V}^a$, and $a' \in \mathcal{A} \setminus \{a\}$, it holds $|\mu^a(i) - \mu^{a'}(i)| \geq \Delta_{\text{up}}$.

In other words, Δ_{up} is not only a lower bound on individual uplift, but also a lower bound on (individual) effect gap. Second, with a UCB-style algorithm, the optimal actions may not be taken sufficiently many times, in which case we can only conservatively assume that they affect all the variables. This would incur additional regret in the analysis when μ^0 is unknown. Therefore, we would like to ensure that all the (potentially

optimal) actions are taken the same number of times at the point that the affected variables get identified. This leads to a two-stage method as described in [Algorithm 6](#).

In the first stage, we perform successive elimination until the uneliminated actions are taken sufficiently many times. Here \tilde{r}_ℓ^a is the empirical estimate of the reward associated to action a and $c_\ell = m\sqrt{2\log(1/\delta')/\ell}$ is the associated radius of confidence interval. Subsequently, for each variable $i \in \mathcal{V}$ we construct the set

$$\widehat{\mathcal{S}}(i) = \{a \in \mathcal{A} : \exists a' \in \mathcal{A} \setminus \{a\}, C_{t_0+1}^a(i) \cap C_{t_0+1}^{a'}(i) \neq \emptyset\}, \quad (13)$$

where t_0 is the number of iterations in the elimination phase and the confidence intervals are defined as in (9). If [Assumption 1](#) is verified, then with high probability $a \in \widehat{\mathcal{S}}(i)$ only if $i \notin \mathcal{V}^a$. Therefore, the observed results from these arms can be used to estimate the baseline payoff of i . On the other hand, the identified affected variables for action $a \in \mathcal{A}$ are then $\widehat{\mathcal{V}}^a = \{i \in \mathcal{V} : a \notin \widehat{\mathcal{S}}(i)\}$.

The second stage of the algorithm reuses the idea from UPUCB. We define the UCB indices following (9) for $a \in \mathcal{A}$ and for the baseline estimate we choose among $\widehat{\mathcal{S}}(i)$ the action that is taken the most frequently whenever this set is non-empty. Otherwise we arbitrarily set it to 0. That is,

$$U_t^0(i) = \begin{cases} 0 & \text{if } \widehat{\mathcal{S}}(i) = \emptyset, \\ U_t^{a_t^0(i)}(i) \text{ where } a_t^0(i) = \arg \max_{a \in \widehat{\mathcal{S}}(i)} N_{t-1}^a & \text{otherwise.} \end{cases} \quad (14)$$

The uplifting indices that guide the decision of the algorithm are defined by $\tau_t^a = \sum_{i \in \widehat{\mathcal{V}}_t^a} (U_t^a(i) - U_t^0(i))$.

As shown below, the regret of [Algorithm 6](#) takes roughly the same form as the one in [Theorem 3](#), but as in [Theorem 1](#), the number of variables that an optimal action affects also plays a crucial role in the regret bound.

Theorem 4. *Let [Assumption 1](#) hold and $\delta' = \delta/(2KmT)$. Then the regret of UPUCB-iLift ([Algorithm 6](#)), with probability at least $1 - \delta$, satisfies:*

$$\text{Reg}_T \leq \sum_{a \in \mathcal{A} : \Delta^a > 0} \left(\frac{8 \text{clip}(2\Delta^a/\Delta_{\text{up}}, L^a + L^{a^*}, 2m)^2 \log(2KmT/\delta)}{\Delta^a} + \Delta^a \right). \quad (15)$$

Remark 2. A natural candidate for $U_t^0(i)$ is the one that uses aggregated information as in (6). The problem of this choice is that how we aggregate the observations collected in the first phase are only determined at the end of the the phase, while this choice itself depends on these observations. Mathematically, this prevents us from defining a suitable martingale difference sequence for which we establish concentration bound. An easy fix is to only aggregate observations from the second phase, while the use of the observations from the first phase should be limited to those coming from just one arm.

7 Numerical Experiments

In this section, we present numerical experiments to demonstrate the benefit of estimating uplifts in our model. We compare UPUCB (bl), UPUCB, UPUCB-nAff (bl), and UPUCB-nAff ([Algorithms 1–4](#)) against UCB and Thompson sampling with Gaussian prior and Gaussian noise that only use the observed rewards $(r_t)_{t \in [T]}$. To ensure a fair comparison, we tune all the considered algorithms and report the results for the parameters that yield the best averaged performance. The experimental setup details are provided in [Appendix D](#). The experiments for the contextual combinatorial uplifting model introduced in [Section 8](#) are presented in [Appendix C.3](#).

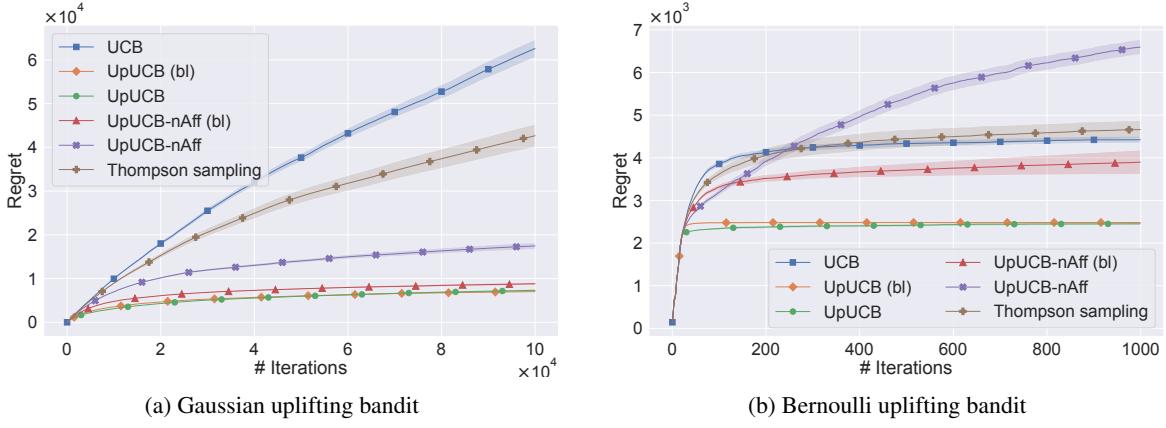


Figure 3 Experimental results on a synthetic and real-world dataset. All the curves are averaged over 100 runs and the shaded areas represent the standard errors.

Gaussian Uplifting Bandit. We first study our algorithms in a Gaussian uplifting bandit, i.e., the noise vectors $(\xi^a)_{a \in \mathcal{A}}$ are Gaussian, with $K = 10$ actions, $m = 100$ variables, and $L^a \equiv 10$ meaning that each action affects 10 variables. The expected payoffs are contained in $[0, 1]$, and the covariance matrix of the noise is taken the same for all the actions. The suboptimality gap of the created problem is around 0.2, and the variance of the total noise $\sum_{i \in \mathcal{V}} \xi^a(i)$ is around 80.

Bernoulli Uplifting Bandit with Criteo Uplift. We use the Criteo Uplift Prediction Dataset [Diemert et al., 2018] with ‘visit’ as the outcome variable to build a Bernoulli uplifting bandit, where the payoff of each variable has a Bernoulli distribution. This dataset is designed for uplift modeling, and has outcomes for both treated and untreated individuals. Thus it is suitable for our simulations. To build the model, we sample 10^5 examples from the dataset, and use K-means to partition these samples into 20 clusters of various sizes. The 10^5 examples are taken as our variables. We consider 20 actions that correspond to treating individuals of each cluster. Assuming that each cluster contains a single type of users that react in the same pattern, we sample the payoffs $(y_t(i))_{i \in \mathcal{V}}$ from independent Bernoulli distributions with means that equal the average visit rates of the treated or untreated individuals of the clusters, depending on which action is chosen. Here, $L = 12654$ and Δ is around 30.

Results. Figure 3 confirms that we can effectively achieve much smaller regret by restricting our attention to the uplift. Moreover, when the sets of affected are known, the loss of not knowing the baseline payoffs seems to be minimal. Not knowing the affected variables has a more severe effect, and in particular causes much larger regret in the second experiment. In fact, the design of UPUCB-nAff and UPUCB-nAff (bl) heavily rely on the additive structure of the uplifting index, and can thus hardly benefit from the payoff independence which allows the other four algorithms to achieve smaller regret in this case.

7.1 Ablation Study on Gaussian Uplifting Bandit

Below we further present ablation studies on the Gaussian uplifting bandit model.

Baseline UCB versus Baseline LCB. In Algorithm 2, we suggest subtracting the UCBs of the baseline payoffs from the UCBs of the actions. This could be counter-intuitive because to get an optimistic estimate of the uplift, we should instead subtract the LCBs of the baseline payoffs. However, the latter strategy prevents

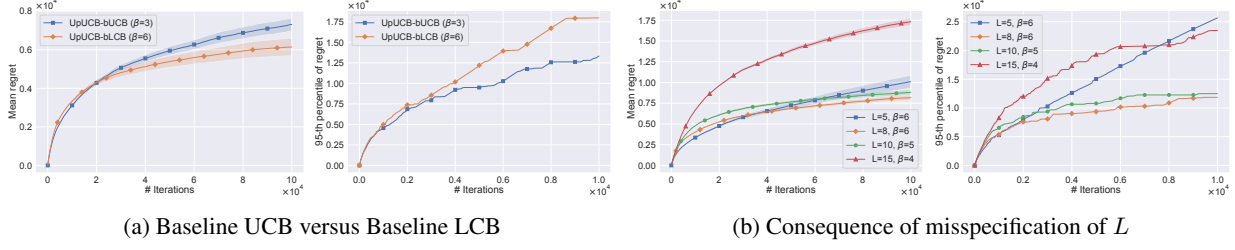


Figure 4: Ablation study on Gaussian uplifting bandit to investigate two aspects of our algorithms. Besides mean regret with standard error, we also plot the 95-th percentile of the regrets as an indicator for the robustness of a method.

the learner from learning more about the baselines that are badly estimated, and can thus lead to failure more frequently. In the following, we refer to the two strategies respectively as UPUCB (bl)-bUCB and UPUCB (bl)-bLCB.

In Figure 4a (Left), we illustrate the mean regret of the two designs when their parameters are optimally tuned. It turns out there the two strategies perform similarly here. We would just like to highlight that to achieve the optimal performance, UPUCB (bl)-bLCB needs to use a larger exploration parameter. Therefore, in terms of average performance, it seems that the effect of using LCBs of the baseline is close to a shrinkage of the exploration parameter.

Next, we further inspect the robustness of the algorithms by plotting the 95th percentile of the regrets (which roughly corresponds to the 5th largest regret among those recorded in the 100 runs). The same parameters give the optimal values for this metric. However, as shown in Figure 4a (Right), it turns out that UPUCB (bl)-bLCB achieves a much higher 95th percentile of regret, which indicates that it indeed fails more frequently.

Misspecification of L . We also investigate the effect of misspecification of L on Algorithm 3. For the bandit instance that we employ here, $\max_{a \in \mathcal{A}} L^a = 10$ and thus according to Proposition 4, Algorithm 3 guarantees sublinear regret as long as it is provided with $L \geq 10$. On the flip side, the algorithm may incur linear regret if L is underestimated.

To verify this, we run the algorithm with $L \in \{5, 8, 10, 15\}$ and plot the averaged regret and the 95-th percentile of regret for optimally tuned exploration parameters in Figure 4b. We effectively observe sublinear regret when L is overspecified, but the regret becomes much larger because the algorithm is more conservative and explores more. On the other hand, under severe underspecification of L , the algorithm is doomed to fail whatever the exploration parameter used. Finally, it seems that slight underspecification of L can be compensated by choosing a larger exploration parameter. In fact, in the figure we see that for $L = 8$ choosing $\beta = 6$ gives the smallest regret among the presented curves.

8 Contextual Extensions

In this section, we briefly discuss potential contextual extensions of our model. As in contextual bandits, context is a side information that helps the learner to make a more informed decision, which results in a higher reward. To incorporate context, one possibility is to associate each variable with a feature vector $x_t(i) \in \mathbb{R}^d$. The subscript t indicates that the context can change from one round to another. We also associate each action with a function f^a so that the expected payoff of action a acting on a variable with feature $x_t(i)$ is $f^a(x_t(i))$. The expected reward of choosing a at round t is then $r^a(x_t) = \sum_{i \in \mathcal{V}} f^a(x_t(i))$. The optimal action of round

Algorithm 7 C2UPUCB

```
1: Input: Budget constraint  $L$ , Parameters for computing the UCBs indices
2: for  $t = 1, \dots, T$  do
3:   for  $i \in \mathcal{V}$  do
4:     Compute UCB indices  $(U_t^z(i))_{z \in \{0,1\}}$ 
5:     Compute uplifting index  $\rho_t(i)$   $\{\rho_t(i) = U_t^1(i) - U_t^0(i)\}$ 
6:   end for
7:   Set  $\mathcal{V}_t \leftarrow \arg \max_{S \subset \mathcal{V}, \text{card}(S) \leq L} \sum_{i \in S} \rho_t(i)$ 
8:   Select action  $a_t = (\mathbb{1}\{i \in \mathcal{V}_t\})_{i \in \mathcal{V}}$   $\{\text{this defines the set of treated customers}\}$ 
9: end for
```

t is $a_t^* = \arg \max_{a \in \mathcal{A}} r^a(x_t)$ and the regret of a learner that takes the actions $(a_t)_{t \in [T]}$ is given by

$$\text{Reg}_T = \sum_{t=1}^T \sum_{i \in \mathcal{V}} (f^{a_t^*}(x_t(i)) - f^{a_t}(x_t(i))).$$

As in [Section 2](#), upon taking an action, the learner observes the noisy payoffs of all the variables. The key structure in our model is there exists a baseline payoff vector μ^0 such that for any given action a , the equality $\mu^a(i) = \mu^0(i)$ holds for most $i \in \mathcal{V}$. Given context, this can be translated into the existence of a baseline function f^0 such that for any given a and t , the equality $f^a(x_t(i)) = f^0(x_t(i))$ holds for most $i \in \mathcal{V}$. The uplift of action a is defined as

$$r_{\text{up}}^a(x_t) = \sum_{i \in \mathcal{V}} f^a(x_t(i)) - f^0(x_t(i)).$$

We provide two concrete examples of the contextual extension below.

Contextual Uplifting Bandit with Linear Payoffs. In this model, each action is associated with an unknown parameter θ^a and the expected payoff is the scalar product of θ^a and the feature of the variable, i.e., $f^a(x_t(i)) = \langle \theta^a, x_t(i) \rangle$. We also assume the aforementioned equality to hold for the baseline function f^0 and we use

$$\mathcal{V}_t^a = \{i \in \mathcal{V} : \langle \theta^a, x_t(i) \rangle \neq \langle \theta^0, x_t(i) \rangle\}$$

for the variables affected by action a at round t . One sufficient condition for \mathcal{V}_t^a to be small is sparsity in both the parameter difference $\theta_{\text{up}}^a = \theta^a - \theta^0$ and the context vector $x_t(i)$. In fact, $\langle \theta^a, x_t(i) \rangle = \langle \theta^0, x_t(i) \rangle$ as long as the supports of θ_{up}^a and $x_t(i)$ are disjoint. In the sequel, we assume $\text{card}(\mathcal{V}_t^a)$ is uniformly bounded by L .

Clearly, our algorithms can be directly applied as long as we can construct a UCB on $\langle \theta^a, x_t(i) \rangle$. This can for example be done using the construction of Linear UCB [[Abbasi-Yadkori et al., 2011](#)]. In this way, the decision of the learner is again based on the uncertain estimates of at most $\mathcal{O}(L)$ variables, and we thus expect the improvements that we have shown in our theorems would persist. As an example, when both θ^0 and \mathcal{V}_t^a are known, UPUCB (bl) adapted to this situation constructs UCB for $\sum_{i \in \mathcal{V}_t^a} \langle \theta_{\text{up}}^a, x_t(i) \rangle$, and it is straightforward to show that the regret of such algorithm can be as small as $\mathcal{O}(Ld\sqrt{KT})$.³ In contrast, if the learner works directly with the total reward, the regret is in $\mathcal{O}(md\sqrt{KT})$.

Contextual Combinatorial Uplifting Bandit. To demonstrate the flexibility of our framework, we also consider a different model that tackles the problem of targeted campaign. Here, we refer to the variables

³As in the previous sections, this bounds applies to the case where the noises are not necessarily independent.

as individuals (customers) that we may want to treat (serve) with some treatment (campaign). To improve the overall return on investment, it is often beneficial to treat only a subset of the individuals. The bandit problem thus consists of finding the right set of individuals to treat at each round. Concretely, if the learner selects at most L individuals to treat in a round, the action set can be written as $\mathcal{A} = \{a \in \{0, 1\}^m : \|a\|_0 \leq L\}$, where $\|a\|_0$ is the number of non-zero coordinates of a . In this form, the i -th coordinate of a is 1 if and only if individual i receives the treatment. Similar to before, we observe the payoffs of all the individuals after treatments are applied.

Here, we make a practical *no interference* assumption, that the payoff of each individual only depends on whether they are treated or not (i.e., independent of others). We can thus define functions g^z for $z \in \{0, 1\}$ so that $f^a(x_t(i)) = g^{a(i)}(x_t(i))$.⁴ Considering the action of not treating any individual as the baseline, the uplift of action a in round t is given by

$$r_{\text{up}}^a(x_t) = \sum_{i \in \mathcal{V}^a} g^{a(i)}(x_t(i)) - g^0(x_t(i)),$$

where \mathcal{V}^a is the set of individuals treated by action a . For UCB-type methods, we should thus provide upper confidence bounds $U_t^z(i)$ for $g^z(x_t(i))$. Depending on whether the baseline function is known or not, we may subtract $g^0(x_t(i))$ or $U_t^0(i)$ from $U_t^1(i)$ to compute the uplifting indices for each individual. The algorithm then chooses the L individuals with the highest uplifting indices. We refer to the resulting algorithm as C2UPUCB, where C2 is the abbreviation of “contextual combinatorial”. We provide its pseudo code in [Algorithm 7](#) (where we suppose the baseline is unknown).

In [Appendix C](#), we further present a generalization of the algorithm to the case of multiple treatments, a theoretical analysis thereof for linear payoffs, and accompanying numerical experiments.

9 Related Work

In this section, we link our work to other existing bandit models. To begin, as suggested by the name, the concept of *uplift* plays a crucial role in our approach. From a high-level perspective, the goals of both uplifting modeling and MAB is to help selecting the optimal action. The former achieves it by modelling the incremental effect of an action on an individual’s behavior. Despite this apparent connection between the two concepts, few papers explicitly link them together. We believe that this is because the resulting problem can often be solved by classic bandit algorithms with a redefined reward. This approach was taken in [\[Sawant et al., 2018, Berrevoets et al., 2019\]](#).

Instead, our paper focuses on bandit problems in which estimating the uplifts improves the statistical efficiency of the algorithms, and this is possible thanks to the ‘sparsity’ of the actions’ effects. Prior to our work, sparsity assumptions in bandits primarily concerned the sparsity of the parameter vectors in linear bandits, which could be of great interest in high-dimensional settings [\[Abbasi-Yadkori et al., 2012, Bastani and Bayati, 2020, Hao et al., 2021\]](#). A notable exception is [Kwon et al. \[2017\]](#) who studied a variant of the MAB where the sparsity is reflected by the fact that the number of arms with positive reward is small. Our work is orthogonal to all of these in that we look at a different form of sparsity. As we saw in [Section 8](#), while sparsity in parameter can be a cause of sparsity in action effect, the improvement of regret is established with a different mechanism. In the upcoming paragraphs, we further exhibit how our model is related to causal and combinatorial bandits.

⁴Strictly speaking, f^a is a function of $(x_t(i), i)$ here.

Causal Bandits. The term causal bandits captures the general idea that in a bandit model the generation of reward may be governed by some causal mechanism while the actions correspond to interventions on the underlying causal graph. This idea was first formalized in [Lattimore et al., 2016] and subsequently explored in a series of work [Lee and Bareinboim, 2018, Lu et al., 2020, Nair et al., 2021] that cover multiples variants of the problem, including both the minimization of simple and cumulative regrets. Our problem fits into the causal bandit framework by regarding Figure 1 as a simple causal graph where the variables consist of all the directed causes of the reward and the actions correspond to atomic interventions on the action nodes. The observability of these direct causes are also key assumptions in many of the aforementioned works. Moreover, in the unknown affected case our causal graph is in fact partially unknown. This is an intriguing problem in causal bandits as argued by [Lu et al., 2021]. While we have shown improvement of the algorithms for this specific model, a interesting question is whether similar improvement can still be shown in more general causal bandit models under the assumption that each actions only affects a limited number of variables of the causal graph.

Combinatorial Bandits. In combinatorial bandits [Cesa-Bianchi and Lugosi, 2012], the learner selects at each round a subset of ground items depending on which the reward is computed. In its simplest form, the reward is just the sum of the payoffs of the selected items. Of particular relevance to us is the semi-bandit setting [Chen et al., 2013, Kveton et al., 2015], which assumes that the payoffs of the selected items are also observed. This structure of the reward and observability of individual payoffs both suggest a strong similarity between our model and semi-bandits. However, this comes with an important conceptual difference: the set of items, or variables, for which we observe the outcomes are selected in semi-bandits, while for us they are fixed and inherent to the reward generation mechanism. As an example, in example (c), combinatorial bandit models optimize the number of movies that are watched among the recommended ones, while we also consider non-recommended movies. In fact, as argued in [Wang et al., 2020], a recommendation is only effective if the user actually watches the movie and would not watch it without the recommendation

On the other hand, mathematically, it is possible to fit our model into the stochastic combinatorial semi-bandit framework when the sets of affected variables are known. For this, consider a semi-bandit problem with $m + \sum_{a \in \mathcal{A}} L^a$ items that correspond to the random variables $\{y^0(i) : i \in \mathcal{V}\} \cup \bigcup_{a \in \mathcal{A}} \{y^a(i) : i \in \mathcal{V}^a\}$. At each round, the learner selects m items that depends on the taken action and observes the payoffs of the selected items. More precisely, when action a is taken, the selected items are $\{y^0(i) : i \in \overline{\mathcal{V}^a}\} \cup \{y^a(i) : i \in \mathcal{V}^a\}$.

In terms of algorithm, let us consider the UPUCB algorithm described in Algorithm 2, which is based on selecting the action with the largest uplift index $\tau_t^a = \sum_{i \in \mathcal{V}^a} (U_t^a(i) - U_t^0(i))$. We notice that the quantity $U_t^a = \sum_{i \in \mathcal{V}} U_t^a(i) = \tau_t^a + \sum_{i \in \mathcal{V}} U_t^0(i)$ is indeed an upper confidence bound on the reward of action a and UPUCB is equivalent to choosing the action with the largest U_t^a in each round. This observation indeed reduces UPUCB to CombUCB [Chen et al., 2013, Kveton et al., 2015] on the transformed model. Nonetheless, the UPUCB formulation has several advantages:

1. It captures the fact that we really care about is the uplifting effect and the variables affected by each action.
2. The uplifting interpretation facilitates algorithm design in more involved situations as we have shown in Section 6.
3. While the general analysis of CombUCB yields a regret in $\mathcal{O}((m^2 + mKL) \log T / \Delta)$, we have shown in Theorem 1 that the regret of UPUCB is in $\mathcal{O}(L^2 \log T / \Delta)$. Our analysis presented in Appendix A.2 also greatly benefits from the UPUCB formulation of the algorithm.

10 Concluding Remarks

This paper addresses multi-armed bandit problems whose rewards can be written as a sum of some observable variables. We showed that when each of the action only affects a limited number of these variables, much smaller regret can be achieved, and we developed algorithms with such guarantee under different variations of knowledge the learner possesses. While the focus of this paper is on UCBs for stochastic bandits, we believe understanding how similar problem structure can be exploited for other algorithms such as Thompson sampling, information directed sampling, or even methods for adversarial bandits is also important.

Table 2: Main notations used in the paper.

	Notation	Explanation
Bandits	K	Number of actions
	a	An action
	a^*	An optimal action, $a^* \in \arg \max_{a \in \mathcal{A}} r^a$
	\mathcal{A}	The set of actions, $\mathcal{A} = \{1, \dots, K\}$
	t	Round index
	T	Number of rounds / Time horizon
	$[T]$	The set of all the rounds, $[T] = \{1, \dots, T\}$
	r_t	Reward received at round t
	r^a	Expected reward of action a
	r^*	Largest expected reward, $r^* = r^{a^*} = \max_{a \in \mathcal{A}} r^a$
Uplifting Bandits	Δ^a	Suboptimality gap of action a , $\Delta^a = r^* - r^a$
	Δ	Minimum non-zero suboptimality gap, i.e., $\Delta = \min_{a \in \mathcal{A}, \Delta^a > 0} \Delta^a$
	m	Number of underlying variables
	i	A variable
	\mathcal{V}	The set of variables, $\mathcal{V} = \{1, \dots, m\}$
	L^a	Number of variables affected by action a
	L	Upper bound on number of affected variables, i.e., $L \geq \max_{a \in \mathcal{A}} L^a$
	\mathcal{V}^a	The set of variables affected by action a
	0	A convenient notation to represent all quantities related to the baseline
	\mathcal{A}_0	A notation to include both the actions and the baseline
	r^0	Baseline reward
	r_{up}^a	Expected uplift of action a , $r_{\text{up}}^a = r^a - r^0$
	r_{up}^*	Largest expected uplift, $r_{\text{up}}^* = r_{\text{up}}^{a^*} = \max_{a \in \mathcal{A}} r_{\text{up}}^a$
	y_t	Payoffs observed at round t , $y_t = (y_t(i))_{i \in \mathcal{V}}$
	μ^a	Expected payoffs of action a , $\mu^a = (\mu^a(i))_{i \in \mathcal{V}}$
	μ^0	Baseline payoffs, $\mu^0 = (\mu^0(i))_{i \in \mathcal{V}}$
	μ_{up}^a	Expected individual uplifts of action a , $\mu_{\text{up}}^a = (\mu_{\text{up}}^a(i))_{i \in \mathcal{V}}$
	\mathcal{P}^a	The distribution of the payoffs associated to action a
	\mathcal{P}^0	Baseline distribution of the payoffs
	Δ_{up}	Lower bound on individual uplift, i.e., $\forall a \in \mathcal{A}, i \in \mathcal{V}^a, \mu^a(i) - \mu^0(i) \geq \Delta_{\text{up}}$
UCB	N_t^a	Number of times action a is taken in the first t rounds
	$\hat{\mu}_t^a$	Empirical estimates of expected payoffs
	c_t^a	Widths of confidence intervals
	\mathcal{C}_t^a	Confidence intervals
	\mathcal{U}_t^a	Upper confidence bounds
	τ_t^a	Uplifting index of action a
	ρ_t^a	Individual uplifting indices
	δ^j	Error probability provided as input for UCB-type algorithms

References

- Yasin Abbasi-Yadkori, Dávid Pál, and Csaba Szepesvári. Improved algorithms for linear stochastic bandits. *Advances in Neural Information Processing Systems*, 24:2312–2320, 2011.
- Yasin Abbasi-Yadkori, David Pal, and Csaba Szepesvari. Online-to-confidence-set conversions and application to sparse stochastic bandits. In *Artificial Intelligence and Statistics*, pages 1–9. PMLR, 2012.
- Peter Auer. Using confidence bounds for exploitation-exploration trade-offs. *Journal of Machine Learning Research*, 3(Nov):397–422, 2002.
- Hamsa Bastani and Mohsen Bayati. Online decision making with high-dimensional covariates. *Operations Research*, 68(1):276–294, 2020.
- Jeroen Berrevoets, Sam Verboven, and Wouter Verbeke. Optimising individual-treatment-effect using bandits. *arXiv preprint arXiv:1910.07265*, 2019.
- Sébastien Bubeck and Nicolo Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Machine Learning*, 5(1):1–122, 2012.
- Apostolos N Burnetas and Michael N Katehakis. Optimal adaptive policies for sequential allocation problems. *Advances in Applied Mathematics*, 17(2):122–142, 1996.
- Nicolo Cesa-Bianchi and Gábor Lugosi. Combinatorial bandits. *Journal of Computer and System Sciences*, 78(5):1404–1422, 2012.
- Wei Chen, Yajun Wang, and Yang Yuan. Combinatorial multi-armed bandit: General framework and applications. In *International Conference on Machine Learning*, pages 151–159. PMLR, 2013.
- Kai Lai Chung and Kailai Zhong. *A course in probability theory*. Academic press, 2001.
- Varsha Dani, Thomas P Hayes, and Sham M Kakade. Stochastic linear optimization under bandit feedback. In *Conference on Learning Theory*, pages 355–366, 2008.
- Rémy Degenne and Vianney Perchet. Combinatorial semi-bandit with known covariance. In *Advances in Neural Information Processing Systems*, pages 2972–2980, 2016.
- Eustache Diemert, Artem Betlei, Christophe Renaudin, and Massih-Reza Amini. A large scale benchmark for uplift modeling. In *International Conference on Knowledge Discovery and Data Mining*. ACM, 2018.
- Joseph L. Doob. *Stochastic processes*. John Wiley & Sons, 1953.
- Eyal Even-Dar, Shie Mannor, Yishay Mansour, and Sridhar Mahadevan. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of machine learning research*, 7(6), 2006.
- Sarah Filippi, Olivier Cappé, Aurélien Garivier, and Csaba Szepesvári. Parametric bandits: The generalized linear case. In *Advances in Neural Information Processing Systems*, volume 23, pages 586–594, 2010.
- Pierre Gutierrez and Jean-Yves Gérardy. Causal inference and uplift modelling: A review of the literature. In *International Conference on Predictive Applications and APIs*, pages 1–13. PMLR, 2017.

- Botao Hao, Tor Lattimore, and Wei Deng. Information directed sampling for sparse linear bandits. In *Advances in Neural Information Processing Systems*, 2021.
- Branislav Kveton, Zheng Wen, Azin Ashkan, and Csaba Szepesvári. Tight regret bounds for stochastic combinatorial semi-bandits. In *Artificial Intelligence and Statistics*, pages 535–543. PMLR, 2015.
- Joon Kwon, Vianney Perchet, and Claire Vernade. Sparse stochastic bandits. In *Conference on Learning Theory*, volume 65, 2017.
- Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- Finnian Lattimore, Tor Lattimore, and Mark D Reid. Causal bandits: Learning good interventions via causal inference. In *Advances in neural information processing systems*, 2016.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.
- Sanghack Lee and Elias Bareinboim. Structural causal bandits: where to intervene? *Advances in Neural Information Processing Systems 31*, 31, 2018.
- Yangyi Lu, Amirhossein Meisami, Ambuj Tewari, and William Yan. Regret analysis of bandit problems with causal background knowledge. In *Conference on Uncertainty in Artificial Intelligence*, pages 141–150. PMLR, 2020.
- Yangyi Lu, Amirhossein Meisami, and Ambuj Tewari. Causal bandits with unknown graph structure. In *Advances in Neural Information Processing Systems*, 2021.
- Vineet Nair, Vishakha Patil, and Gaurav Sinha. Budgeted and non-budgeted causal bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 2017–2025. PMLR, 2021.
- Lijing Qin, Shouyuan Chen, and Xiaoyan Zhu. Contextual combinatorial bandit and its application on diversified online recommendation. In *Proceedings of the 2014 SIAM International Conference on Data Mining*, pages 461–469. SIAM, 2014.
- Nicholas Radcliffe. Using control groups to target on predicted lift: Building and assessing uplift model. *Direct Marketing Analytics Journal*, pages 14–21, 2007.
- Neela Sawant, Chitti Babu Namballa, Narayanan Sadagopan, and Houssam Nassif. Contextual multi-armed bandits for causal marketing. In *International Conference on Machine Learning*, 2018.
- István Szita and Csaba Szepesvári. Agnostic KWIK learning and efficient approximate reinforcement learning. In *Conference on Learning Theory*, pages 739–772. JMLR Workshop and Conference Proceedings, 2011.
- Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.
- Yixin Wang, Dawen Liang, Laurent Charlin, and David M Blei. Causal inference for recommender systems. In *Fourteenth ACM Conference on Recommender Systems*, pages 426–431, 2020.

A Proofs for Upper Bound Results

In this appendix, we provide proofs of the regret bounds from [Sections 5](#) and [6](#). Our proofs have the following structure.

1. Leveraging [Lemmas 2](#) and [3](#) presented in [Appendix A.1](#), we define a high-probability event \mathcal{E} on which all the expected payoffs belong to their confidence intervals.
2. For UCB-like methods, we use the fact that the uplifting index of the taken action is larger than that of any optimal action by design to show that on event \mathcal{E} a suboptimal action can only be taken when its associated width of confidence interval is sufficiently large.
3. The previous point allows us to bound the number of times that a suboptimal action is taken on \mathcal{E} .
4. We conclude with the following regret decomposition (which was already stated in [\(2\)](#)).

$$\text{Reg}_T = \sum_{a \in \mathcal{A}} N_t^a \Delta^a. \quad (16)$$

A.1 Concentration Bounds

To prove the regret upper bounds, we will make use of concentration bounds that apply to quantities defined in [\(3\)](#) and [\(6\)](#). Underlying all these results is the following fundamental concentration inequality for σ -sub-Gaussian random variables [[Vershynin, 2018](#)]

$$\mathbb{P}(|X| \geq \gamma) \leq 2 \exp(-\gamma^2/(2\sigma^2)) \text{ for all } \gamma > 0. \quad (17)$$

However, the estimates in [\(3\)](#) and [\(6\)](#) aggregate observed values in a way that depend on the decisions that have been made. Therefore, concentration bounds for the sum of independent variables, which can be directly derived from [\(17\)](#), are not sufficient for our purposes. The standard trick to circumvent this issue in bandit literature is to combine Doob's optional skipping ([\[Doob, 1953, Chap. 3\]](#), [\[Chung and Zhong, 2001, Chap. 9\]](#)) with concentration inequality for martingale difference sequence. The following lemma adapted from [\[Szita and Szepesvári, 2011, Lem. A.1\]](#) is proved exactly using this technique.

Lemma 1. *Let $(\mathcal{F}_t)_{t \in \mathbb{N}}$ be a filtration and let $(\epsilon_t, X_t)_{t \in \mathbb{N}}$ be a sequence of $\{0, 1\} \times \mathbb{R}$ -valued random variables such that ϵ_t is \mathcal{F}_{t-1} -measurable, and X_t is \mathcal{F}_t -measurable and conditionally σ -sub-Gaussian, i.e., $\mathbb{E}[\exp(\gamma X_t) | \mathcal{F}_{t-1}] \leq \exp(\gamma^2 \sigma^2 / 2)$ for all $\gamma \in \mathbb{R}$. Let $N > 0$ and let $\tau = \min\{t \geq 1 : \sum_{s=1}^t \epsilon_s = N\}$, where we take $\tau = +\infty$ when $\sum_{s=1}^{+\infty} \epsilon_s < N$. Then, for any $\delta \in (0, 1)$,*

$$\mathbb{P}\left(\left|\sum_{t=1}^{\tau} \epsilon_t X_t\right| \geq \sigma \sqrt{2N \log(2/\delta)}\right) \leq \delta. \quad (18)$$

Proof. The proof of [\[Szita and Szepesvári, 2011, Lem. A.1\]](#) essentially applies. We just need to replace the use of Hoeffding-Azuma inequality by the use of concentration inequality for conditionally sub-Gaussian martingale difference sequences, which is itself derived from [\(17\)](#) and the fact that the sum of conditionally sub-Gaussian martingale differences is sub-Gaussian. \square

In the following, $(\mathcal{F}_t)_t$ will represent the natural filtration associated to $(y_t, a_{t+1})_t$ such that a_t is \mathcal{F}_{t-1} -measurable, and y_t is \mathcal{F}_t -measurable. Applying [Lemma 1](#) to properly defined $(\epsilon_t, X_t)_t$ along with the use of a union bound gives the following concentration results that are key to our analyses.

Lemma 2. Let $a \in \mathcal{A}$, $i \in \mathcal{V}$, and let $\mu_t^a(i)$ be defined as in (3). Then for any $\delta \in (0, 1)$, it holds

$$\mathbb{P}\left(\exists t \in T, |\hat{\mu}_t^a(i) - \mu^a(i)| \geq \sqrt{\frac{2 \log(2/\delta)}{N_t^a}}\right) \leq T\delta.$$

Proof. For $N \in [T]$, we define the random variables

$$\tau_N = \min \left\{ t \geq 1 : \sum_{s=1}^t \mathbb{1}\{a_s = a\} = N \right\} \quad \text{and} \quad S_N^a(i) = \sum_{s=1}^{\tau_N} y_s(i) \mathbb{1}\{a_s = a\}.$$

We then define the failure events

$$\begin{aligned} \mathcal{E}(i) &= \left\{ \exists t \in [T], |\hat{\mu}_t^a(i) - \mu^a(i)| \geq \sqrt{\frac{2 \log(2/\delta)}{N_t^a}} \right\}, \\ \tilde{\mathcal{E}}(i) &= \left\{ \exists N \in [T], \left| S_N^a(i) - \mu^a(i) \left(\sum_{s=1}^{\tau_N} \mathbb{1}\{a_s = a\} \right) \right| \geq \sqrt{2N \log(2/\delta)} \right\}. \end{aligned}$$

It holds $\mathcal{E}(i) \subseteq \tilde{\mathcal{E}}(i)$. In fact, if $\mathcal{E}(i)$ happens then $|\hat{\mu}_t^a(i) - \mu^a(i)| \geq \sqrt{2 \log(2/\delta) / N_t^a(i)}$ for some t and this is only possible when $N_t^a(i) \geq 1$. With the definition $N_t^a(i) = \sum_{s=1}^t \mathbb{1}\{a_s = a\}$, we further see that $N_t^a(i) \leq t \leq T$ and $\tau_{N_t^a(i)} \leq t < +\infty$. The latter means that we indeed have $\sum_{s=1}^{\tau_{N_t^a(i)}} \mathbb{1}\{a_s = a\} = N_t^a(i)$. Thus multiplying $|\hat{\mu}_t^a(i) - \mu^a(i)| \geq \sqrt{2 \log(2/\delta) / N_t^a(i)}$ by $N = N_t^a(i) \in [T]$ gives exactly

$$\left| S_N^a(i) - \mu^a(i) \left(\sum_{s=1}^{\tau_N} \mathbb{1}\{a_s = a\} \right) \right| \geq \sqrt{2N \log(2/\delta)},$$

which shows that $\tilde{\mathcal{E}}(i)$ happens as well. It is thus sufficient to prove that $\mathbb{P}(\tilde{\mathcal{E}}(i)) \leq T\delta$.

For $N \in [T]$, we apply [Lemma 1](#) to $\epsilon_t \leftarrow \mathbb{1}\{a_t = a\}$ and $X_t \leftarrow \xi_t^a(i) = y_t^a(i) - \mu^{a_t}(i)$. Since $\mu^{a_t}(i) \mathbb{1}\{a_t = a\} = \mu^a(i) \mathbb{1}\{a_t = a\}$, we recover the exact inequality that appears in the definition of $\tilde{\mathcal{E}}(i)$. We conclude with a union bound taking N ranging from 1 to T . \square

Lemma 3. Let $i \in \mathcal{V}$, and let $\mu_t^0(i)$ be defined as in (6). Then for any $\delta \in (0, 1)$, it holds

$$\mathbb{P}\left(\exists t \in T, |\hat{\mu}_t^0(i) - \mu^0(i)| \geq \sqrt{\frac{2 \log(2/\delta)}{N_t^0(i)}}\right) \leq T\delta.$$

Proof. [Lemma 3](#) is proved in the same way as [Lemma 2](#). In particular, we apply [Lemma 1](#) to $\epsilon_t \leftarrow \mathbb{1}\{i \notin \mathcal{V}^{a_t}\}$ and $X_t \leftarrow \xi_t^0(i)$. \square

A.2 Unknown Baseline: Proof of [Theorem 1](#)

Theorem 1. Let $\delta' = \delta / (4KLT)$. Then the regret of UPUCB ([Algorithm 2](#)) with probability at least $1 - \delta$, satisfies:

$$\text{Reg}_T \leq \sum_{a \in \mathcal{A}: \Delta^a > 0} \left(\frac{8(L^a + L^{a^*})^2 \log(4KLT/\delta)}{\Delta^a} + \Delta^a \right). \quad (8)$$

Proof. We first notice that while there are m variables, not all of them are involved in estimating the uplifts. Therefore, we only need to focus on the set of relevant variables $\mathcal{V}' = \bigcup_{a \in \mathcal{A}} \mathcal{V}^a$. Clearly, $\text{card}(\mathcal{V}') \leq KL$. With this in mind, we define the events

$$\begin{aligned}\mathcal{E}_1 &= \{\forall t \in [T], \forall a \in \mathcal{A}, \forall i \in \mathcal{V}^a, |\hat{\mu}_t^a(i) - \mu^a(i)| \leq c_t^a\}, \\ \mathcal{E}_2 &= \{\forall t \in [T], \forall i \in \mathcal{V}', |\hat{\mu}_t^0(i) - \mu^0(i)| \leq c_t^0(i)\}.\end{aligned}$$

Lemma 2 and Lemma 3 in Appendix A.1 along with union bounds over a and i guarantee

$$\mathbb{P}(\mathcal{E}_1 \cup \mathcal{E}_2) \geq 1 - \left(\sum_{a \in \mathcal{A}} L^a + \text{card}(\mathcal{V}') \right) T(2\delta') \geq 1 - 4KLT\delta' = 1 - \delta.$$

It is thus sufficient to show that (8) holds on $\mathcal{E}_1 \cup \mathcal{E}_2$.

In the remainder of the proof, we consider an arbitrary realization of $\mathcal{E}_1 \cup \mathcal{E}_2$ and prove that (8) indeed holds for this realization. From (16) it is clear that we just need to bound $N_T^a = \sum_{t=1}^T \mathbb{1}\{a_t = a\}$ from above for all suboptimal action $a \in \mathcal{A}$ with $\Delta^a > 0$. This is done by providing a lower bound on c_t^a for any t such that $a_t = a$. To proceed, suppose that a suboptimal action a is taken for the last time at round $t \in \{K+1, \dots, T\}$.⁵ This implies $\tau_t^a \geq \tau_t^{a^*}$ and $N_T^a = N_{t-1}^a + 1$. By the definition of the uplifting indices, we get

$$\sum_{i \in \mathcal{V}^a} (U_t^a(i) - U_t^0(i)) \geq \sum_{i \in \mathcal{V}^{a^*}} (U_t^{a^*}(i) - U_t^0(i)).$$

Rearranging and dropping $\sum_{i \in \mathcal{V}^a \cap \mathcal{V}^{a^*}} U_t^0(i)$ from both sides, we get

$$\sum_{i \in \mathcal{V}^a} U_t^a(i) + \sum_{i \in \mathcal{V}^{a^*} \setminus \mathcal{V}^a} U_t^0(i) \geq \sum_{i \in \mathcal{V}^{a^*}} U_t^{a^*}(i) + \sum_{i \in \mathcal{V}^a \setminus \mathcal{V}^{a^*}} U_t^0(i). \quad (19)$$

This is similar to the type of inequality that defines standard UCB. In fact, the left and the right hand sides of (19) are respectively upper confidence bounds on $\sum_{i \in \mathcal{V}^a \cup \mathcal{V}^{a^*}} \mu^a(i)$ and upper confidence bounds on $\sum_{i \in \mathcal{V}^a \cup \mathcal{V}^{a^*}} \mu^{a^*}(i)$. We next bound these two quantities respectively from above and from below. Since both \mathcal{E}_1 and \mathcal{E}_2 hold, we have the following inequalities for the UCB indices

$$\begin{aligned}\forall s \in [T], \forall a \in \mathcal{A}, \forall i \in \mathcal{V}^a, \mu^a(i) + 2c_{s-1}^a &\geq U_s^a(i) \geq \mu^a(i), \\ \forall s \in [T], \forall i \in \mathcal{V}' \setminus \bigcap_{a \in \mathcal{A}} \mathcal{V}^a, \mu^0(i) + 2c_{s-1}^0(i) &\geq U_s^0(i) \geq \mu^0(i).\end{aligned}$$

Combining the above, we obtain

$$\sum_{i \in \mathcal{V}^a} (\mu^a(i) + 2c_{t-1}^a) + \sum_{i \in \mathcal{V}^{a^*} \setminus \mathcal{V}^a} (\mu^0(i) + 2c_{t-1}^0(i)) \geq \sum_{i \in \mathcal{V}^{a^*}} \mu^{a^*}(i) + \sum_{i \in \mathcal{V}^a \setminus \mathcal{V}^{a^*}} \mu^0(i).$$

Adding $\sum_{i \in \mathcal{V}^a \cap \mathcal{V}^{a^*}} \mu^0(i)$ to both sides and rearranging leads to

$$\sum_{i \in \mathcal{V}^{a^*}} \mu^{a^*}(i) - \sum_{i \in \mathcal{V}^a} \mu^a(i) + \sum_{i \in \mathcal{V}^a} \mu^0(i) - \sum_{i \in \mathcal{V}^{a^*}} \mu^0(i) \leq \sum_{i \in \mathcal{V}^a} 2c_{t-1}^a + \sum_{i \in \mathcal{V}^{a^*} \setminus \mathcal{V}^a} 2c_{t-1}^0(i).$$

⁵If the action is not taken anymore after the first T runs. We have $N_T^a = 1$ so the upper bound that we show later trivially holds.

From this we can derive an inequality between the suboptimal gap Δ^a and the widths of the confidence intervals as follows

$$\begin{aligned}
\Delta^a &= r_{\text{up}}^{a^*} - r_{\text{up}}^a \\
&= \sum_{i \in \mathcal{V}^{a^*}} (\mu^{a^*}(i) - \mu^0(i)) - \sum_{i \in \mathcal{V}^a} (\mu^a(i) - \mu^0(i)) \\
&= \sum_{i \in \mathcal{V}^{a^*}} \mu^{a^*}(i) - \sum_{i \in \mathcal{V}^a} \mu^a(i) + \sum_{i \in \mathcal{V}^a} \mu^0(i) - \sum_{i \in \mathcal{V}^{a^*}} \mu^0(i) \\
&\leq 2L^a c_{t-1}^a + \sum_{i \in \mathcal{V}^{a^*} \setminus \mathcal{V}^a} 2c_{t-1}^0(i).
\end{aligned} \tag{20}$$

We next argue that for all $i \in \mathcal{V}^{a^*} \setminus \mathcal{V}^a$ we have $N_{t-1}^0(i) \geq N_{t-1}^a$ and hence $c_{t-1}^0(i) \leq c_{t-1}^a$. In fact, if $i \notin \mathcal{V}^a$, then whenever action a is taken, the count of $N_s^0(i)$ also increases by 1. Formally,

$$N_{t-1}^0(i) = \sum_{s=1}^{t-1} \mathbb{1}\{i \notin \mathcal{V}^{a_s}\} \geq \sum_{s=1}^{t-1} \mathbb{1}\{a_s = a\} \mathbb{1}\{i \notin \mathcal{V}^{a_s}\} = \sum_{s=1}^{t-1} \mathbb{1}\{a_s = a\} = N_{t-1}^a.$$

Therefore, we further deduce $\Delta^a \leq 2(L^a + L^{a^*})c_{t-1}^a$. Equivalently,

$$N_{t-1}^a \leq \frac{8(L^a + L^{a^*})^2 \log(1/\delta')}{(\Delta^a)^2}.$$

We conclude that action a is taken at most $\lfloor 8(L^a + L^{a^*})^2 \log(1/\delta')/(\Delta^a)^2 \rfloor + 1$ times during the execution of the algorithm. Plugging this into (16) gives exactly (8). \square

A.3 Unknown Affected Variables: Proof of Theorem 2

Below we prove Theorem 2 separately for UPUCB-nAff(bl) and UPUCB-nAff.

Proposition 4 (Regret of UPUCB-nAff(bl)). *Let $\delta' = \delta/(2KmT)$. If the learner knows and an upper bound L on maximum number of variables an action affects and the baseline payoffs, then UPUCB-nAff(bl) (Algorithm 3), with probability at least $1 - \delta$, has regret bounded by:*

$$\text{Reg}_T \leq \sum_{a \in \mathcal{A}: \Delta^a > 0} \left(\frac{32L^2 \log(2KmT/\delta)}{\Delta^a} + \Delta^a \right). \tag{21}$$

Proof. Let us consider the event

$$\mathcal{E} = \{\forall t \in [T], \forall a \in \mathcal{A}, \forall i \in \mathcal{V}, |\hat{\mu}_t^a(i) - \mu^a(i)| < c_t^a\}, \tag{22}$$

By Lemma 3 and a union bound we get immediately $\mathbb{P}(\mathcal{E}) \geq 1 - KmT(2\delta') = 1 - \delta$. In the following, we assume that \mathcal{E} occurs and prove (21).

First, by the definition of $\widehat{\mathcal{V}}_t^a$ we have clearly $\widehat{\mathcal{V}}_t^a \subseteq \mathcal{V}^a$. In fact, if $i \notin \mathcal{V}^a$ then $\mu^0(i) = \mu^a(i) \in \mathcal{C}_t^a(i)$. The inclusion $\widehat{\mathcal{V}}_t^a \subseteq \mathcal{V}^a$ along with the condition $\text{card}(\mathcal{V}^a) \leq L$ imply $\text{card}(\mathcal{V}^a \setminus \widehat{\mathcal{V}}_t^a) \leq L - \text{card}(\widehat{\mathcal{V}}_t^a)$; in particular, $L - \text{card}(\widehat{\mathcal{V}}_t^a) \geq 0$ and thus $L_t^a = L - \text{card}(\widehat{\mathcal{V}}_t^a)$.

Next, for any $i \in \mathcal{V} \setminus \widehat{\mathcal{V}}_t^a$, it holds $\mu^0(i) \in \mathcal{C}_t^a(i)$. This implies $\mu^0(i) \leq U_t^a(i)$, and accordingly $\rho_t^a(i) \geq 0$. We can then write

$$\sum_{i \in \mathcal{V}^a \setminus \widehat{\mathcal{V}}_t^a} \rho_t^a(i) \leq \max_{\substack{\mathcal{L} \subseteq \mathcal{V} \setminus \widehat{\mathcal{V}}_t^a \\ \text{card}(\mathcal{L}) = \text{card}(\mathcal{V} \setminus \widehat{\mathcal{V}}_t^a)}} \sum_{i \in \mathcal{L}} \rho_t^a(i) \leq \max_{\substack{\mathcal{L} \subseteq \mathcal{V} \setminus \widehat{\mathcal{V}}_t^a \\ \text{card}(\mathcal{L}) = L - \text{card}(\widehat{\mathcal{V}}_t^a)}} \sum_{i \in \mathcal{L}} \rho_t^a(i) = \sum_{i \in \mathcal{L}_t^a} \rho_t^a(i).$$

Since $U_t^a(i) \geq \mu^a(i)$ on \mathcal{E} , we have $\rho_t^a(i) \geq \mu_{\text{up}}^a(i)$; subsequently

$$r_{\text{up}}^a = \sum_{i \in \mathcal{V}^a} \mu_{\text{up}}^a(i) = \sum_{i \in \widehat{\mathcal{V}}_t^a} \mu_{\text{up}}^a(i) + \sum_{i \in \mathcal{V}^a \setminus \widehat{\mathcal{V}}_t^a} \mu_{\text{up}}^a(i) \leq \sum_{i \in \widehat{\mathcal{V}}_t^a} \rho_t^a(i) + \sum_{i \in \mathcal{L}_t^a} \rho_t^a(i) = \tau_t^a.$$

This shows that τ_t^a is effectively an upper bound of r_{up}^a . With this in mind, we are ready to bound the number of times that a suboptimal action is taken.

Let $a \in \mathcal{A}$ with $\Delta^a > 0$ and assume that it is taken at round $t \in \{K+1, \dots, T\}$, which means that $\tau_t^a \geq \tau_t^{a^*}$. We have shown that $\tau_t^{a^*} \geq r_{\text{up}}^{a^*}$. It remains to provide an upper bound for τ_t^a that involves r_{up}^a . For this, we again use $\widehat{\mathcal{V}}_t^a \subseteq \mathcal{V}^a$ and the fact that $\rho_t^a(i) \geq 0$ for all $i \in \mathcal{V} \setminus \widehat{\mathcal{V}}_t^a$ to decompose

$$\tau_t^a = \sum_{i \in \widehat{\mathcal{V}}_t^a \cup \mathcal{L}_t^a} \rho_t^a(i) \leq \sum_{i \in \mathcal{V}^a \cup \mathcal{L}_t^a} \rho_t^a(i) = \sum_{i \in \mathcal{V}^a} \rho_t^a(i) + \sum_{i \in \mathcal{L}_t^a \setminus \mathcal{V}^a} \rho_t^a(i). \quad (23)$$

The summation of the individual uplifting indices over $i \in \mathcal{V}^a$ is naturally related to r_{up}^a . Since $\mu^a(i) \geq \hat{\mu}_{t-1}^a(i) - c_{t-1}^a$, for all $i \in \mathcal{V}$ it holds

$$\rho_t^a(i) = U_t^a(i) - \mu^a(i) \leq \mu^a(i) - \mu^0(i) + 2c_{t-1}^a = \mu_{\text{up}}^a(i) + 2c_{t-1}^a.$$

In consequence,

$$\sum_{i \in \mathcal{V}^a} \rho_t^a(i) \leq \sum_{i \in \mathcal{V}^a} (\mu_{\text{up}}^a(i) + 2c_{t-1}^a) \leq r_{\text{up}}^a + 2Lc_{t-1}^a. \quad (24)$$

What we need to show next is that the definition of $\widehat{\mathcal{V}}_t^a$ guarantees $\rho_t^a(i)$ to be small whenever $i \notin \widehat{\mathcal{V}}_t^a$. In fact, if $i \in \mathcal{V} \setminus \widehat{\mathcal{V}}_t^a$, we have $\mu^0(i) \in \mathcal{C}_t^a(i)$, and in particular $\mu^0(i) \geq \hat{\mu}_{t-1}^a(i) - c_{t-1}^a$, implying that

$$\rho_t^a(i) = U_t^a(i) - \mu^0(i) \leq (\hat{\mu}_{t-1}^a(i) + c_{t-1}^a) - (\hat{\mu}_{t-1}^a(i) - c_{t-1}^a) = 2c_{t-1}^a.$$

By definition, $\mathcal{L}_t^a \subseteq \mathcal{V} \setminus \widehat{\mathcal{V}}_t^a$, and hence

$$\sum_{i \in \mathcal{L}_t^a \setminus \mathcal{V}^a} \rho_t^a(i) \leq \sum_{i \in \mathcal{L}_t^a \setminus \mathcal{V}^a} 2c_{t-1}^a \leq 2Lc_{t-1}^a. \quad (25)$$

Putting (23), (24), and (25) together, we get

$$\tau_t^a \leq \sum_{i \in \mathcal{V}^a} \rho_t^a(i) + \sum_{i \in \mathcal{L}_t^a \setminus \mathcal{V}^a} \rho_t^a(i) \leq r_{\text{up}}^a + 4Lc_{t-1}^a.$$

With $\tau_t^a \geq \tau_t^{a^*} \geq r_{\text{up}}^{a^*}$ we deduce

$$r_{\text{up}}^{a^*} \leq r_{\text{up}}^a + 4Lc_{t-1}^a.$$

Therefore, $\Delta^a = r_{\text{up}}^{a^*} - r_{\text{up}}^a \leq 4Lc_{t-1}^a$, and equivalently

$$N_{t-1}^a \leq \frac{32L^2 \log(1/\delta')}{(\Delta^a)^2}.$$

Following the argument in the proof of [Theorem 1](#), this translates into an upper bound on N_T^a and we conclude by invoking the decomposition of (16). \square

We notice that it is quite straightforward to adapt UPUCB-nAff(bl) and its analysis to the setting where the learner knows an action-dependent upper bound $\bar{L}^a \geq L^a$. The only change in (21) is that L is replaced by \bar{L}^a . Thus we recover (4) up to a multiplicative constant when $\bar{L}^a = \Theta(L^a)$. To derive the regret bound for UPUCB-nAff, we combine idea from the proofs of [Theorem 1](#) and [Proposition 4](#).

Proposition 5 (Regret of UPUCB-nAff). *Let $\delta' = \delta/(2KmT)$. Then the regret of UPUCB-nAff (Algorithm 4), with probability at least $1 - \delta$, satisfies:*

$$\text{Reg}_T \leq \sum_{a \in \mathcal{A}: \Delta^a > 0} \left(\frac{512L^2 \log(2KmT/\delta)}{\Delta^a} + \Delta^a \right), \quad (26)$$

Proof. We again place ourselves in the event \mathcal{E} defined in (22), and prove (26) under this condition. Recall that $b_t \in \arg \max_{a \in \mathcal{A}} N_{t-1}^a$ is an action that is taken most frequently during the first $t - 1$ rounds and the associated estimates serve as baselines at round t . We define $\tilde{\mathcal{V}}_t^a = \mathcal{V}^a \cup \mathcal{V}^{b_t}$ as the set of variables that are affected by either action a or b_t . Clearly, $\text{card}(\tilde{\mathcal{V}}_t^a) \leq 2L$.

We first show that $\widehat{\mathcal{V}}_t^a \subseteq \tilde{\mathcal{V}}_t^a$ whenever \mathcal{E} occurs. Let $i \in \mathcal{V} \setminus \tilde{\mathcal{V}}_t^a$. Under \mathcal{E} we have both $\mu^0(i) = \mu^a(i) \in \mathcal{C}_t^a(i)$ and $\mu^0(i) = \mu^{b_t}(i) \in \mathcal{C}_t^{b_t}(i)$. Therefore, $\mathcal{C}_t^a(i) \cap \mathcal{C}_t^{b_t}(i) \neq \emptyset$; subsequently $i \notin \widehat{\mathcal{V}}_t^a$. The contrapositive of the above gives exactly $\widehat{\mathcal{V}}_t^a \subseteq \tilde{\mathcal{V}}_t^a$.

We shall now bound N_T^a for any suboptimal action a , from which the proof concludes by applying (16). Let $a \in \mathcal{A}$ satisfying $\Delta^a > 0$. If it is taken at round $t \in \{K + 1, \dots, T\}$, we have $\tau_t^a \geq \tau_t^{a^*}$. This can be written as

$$\sum_{i \in \widehat{\mathcal{V}}_t^a} \rho_t^a(i) + \sum_{i \in \mathcal{L}_t^a} \rho_t^a(i) \geq \sum_{i \in \tilde{\mathcal{V}}_t^{a^*}} \rho_t^{a^*}(i) + \sum_{i \in \mathcal{L}_t^{a^*}} \rho_t^{a^*}(i). \quad (27)$$

Decomposition of Suboptimality Gap using the New Baseline. Similar to before, we will use (27) to relate the widths of the confidence intervals to the suboptimality gap Δ^a . For this, we notice that the suboptimality gap of a can be rewritten as

$$\begin{aligned} \Delta^a &= \sum_{i \in \mathcal{V}} \mu^{a^*}(i) - \sum_{i \in \mathcal{V}} \mu^a(i) \\ &= \sum_{i \in \mathcal{V}} (\mu^{a^*}(i) - \mu^{b_t}(i)) - \sum_{i \in \mathcal{V}} (\mu^a(i) - \mu^{b_t}(i)) \\ &= \sum_{i \in \tilde{\mathcal{V}}_t^{a^*}} (\mu^{a^*}(i) - \mu^{b_t}(i)) - \sum_{i \in \tilde{\mathcal{V}}_t^a} (\mu^a(i) - \mu^{b_t}(i)) \end{aligned} \quad (28)$$

The last equality is true because $\mu^{a^*}(i) = \mu^0(i)$ outside $\tilde{\mathcal{V}}_t^{a^*}$, and the same holds when we replace a^* by a . Therefore, our goal is to lower bound the right-hand side of (27) by an expression that contains $\sum_{i \in \tilde{\mathcal{V}}_t^{a^*}} (\mu^{a^*}(i) - \mu^{b_t}(i))$ and upper bound the left-hand side of (27) by an expression that contains $\sum_{i \in \tilde{\mathcal{V}}_t^a} (\mu^a(i) - \mu^{b_t}(i))$.

Lower Bound on RHS of (27). Since $\sum_{i \in \mathcal{L}_t^{a^*}} \rho_t^{a^*}(i)$ is the maximum padding, we can lower bound the right-hand side of the inequality by $\sum_{i \in \tilde{\mathcal{V}}_t^{a^*}} \rho_t^{a^*}(i)$. Formally, with $\widehat{\mathcal{V}}_t^{a^*} \subseteq \tilde{\mathcal{V}}_t^{a^*}$ we have $\text{card}(\tilde{\mathcal{V}}_t^{a^*} \setminus \widehat{\mathcal{V}}_t^{a^*}) \leq 2L - \text{card}(\widehat{\mathcal{V}}_t^{a^*})$ and $L_t^{a^*} = 2L - \text{card}(\widehat{\mathcal{V}}_t^{a^*})$. Thus, by definition of $\mathcal{L}_t^{a^*}$,

$$\sum_{i \in \mathcal{L}_t^{a^*}} \rho_t^{a^*}(i) = \max_{\substack{\mathcal{L} \subseteq \mathcal{V}^a \setminus \widehat{\mathcal{V}}_t^a \\ \text{card}(\mathcal{L}) \leq 2L - \text{card}(\widehat{\mathcal{V}}_t^{a^*})}} \sum_{i \in \mathcal{L}} \rho_t^{a^*}(i) \geq \sum_{i \in \tilde{\mathcal{V}}_t^{a^*} \setminus \widehat{\mathcal{V}}_t^{a^*}} \rho_t^{a^*}(i).$$

To further provide a lower bound on $\sum_{i \in \tilde{\mathcal{V}}_t^{a^*}} \rho_t^{a^*}(i)$, we recall that $\rho_t^{a^*}(i) = U_t^{a^*}(i) - U_t^{b_t}(i)$, so with $U_t^{a^*}(i) \geq \mu^{a^*}(i)$, $U_t^{b_t}(i) \leq \mu^{b_t}(i) + 2c_{t-1}^{b_t}$, and $\text{card}(\tilde{\mathcal{V}}_t^{a^*}) \leq 2L$, we get

$$\begin{aligned} \sum_{i \in \tilde{\mathcal{V}}_t^{a^*}} \rho_t^{a^*}(i) + \sum_{i \in \mathcal{L}_t^{a^*}} \rho_t^{a^*}(i) &\geq \sum_{i \in \tilde{\mathcal{V}}_t^{a^*}} \rho_t^{a^*}(i) \\ &\geq \sum_{i \in \tilde{\mathcal{V}}_t^{a^*}} (\mu^{a^*}(i) - \mu^{b_t}(i) - 2c_{t-1}^{b_t}) \\ &\geq \sum_{i \in \tilde{\mathcal{V}}_t^{a^*}} (\mu^{a^*}(i) - \mu^{b_t}(i)) - 4Lc_{t-1}^{b_t}. \end{aligned} \quad (29)$$

Upper Bound on LHS of (27). To make the term $\sum_{i \in \tilde{\mathcal{V}}_t^a} (\mu^a(i) - \mu^{b_t}(i))$ appear, we first relate it to the sum of the uplifting indices using $U_t^a(i) \leq \mu^a(i) + 2c_{t-1}^a$.

$$\sum_{i \in \tilde{\mathcal{V}}_t^a} \rho_t^a(i) \leq \sum_{i \in \tilde{\mathcal{V}}_t^a} (\mu^a(i) + 2c_{t-1}^a - \mu^{b_t}(i)) \leq \sum_{i \in \tilde{\mathcal{V}}_t^a} (\mu^a(i) - \mu^{b_t}(i)) + 4Lc_{t-1}^a. \quad (30)$$

The last inequality uses $\text{card}(\tilde{\mathcal{V}}_t^a) \leq 2L$. Next, to relate $\sum_{i \in \tilde{\mathcal{V}}_t^a} \rho_t^a(i)$ to the left-hand side of (27), we want to show (i) $\sum_{i \in \mathcal{L}_t^a} \rho_t^a(i)$ is small, and (ii) We can add additional terms from $\tilde{\mathcal{V}}_t^a \setminus \widehat{\mathcal{V}}_t^a$ to the left-hand side of (27). These two points correspond respectively to inequalities (25) and (23) in the proof of Proposition 4. Again, this can be done by looking at the definition of $\widehat{\mathcal{V}}_t^a$. In fact, if $i \in \mathcal{V} \setminus \widehat{\mathcal{V}}_t^a$, then by definition $\mathcal{C}_t^a(i) \cap \mathcal{C}_t^{b_t}(i) \neq \emptyset$, implying that

$$\begin{aligned} \hat{\mu}_{t-1}^{b_t}(i) + c_{t-1}^{b_t} &\geq \hat{\mu}_{t-1}^a(i) - c_{t-1}^a, \\ \hat{\mu}_{t-1}^a(i) + c_{t-1}^a &\geq \hat{\mu}_{t-1}^{b_t}(i) - c_{t-1}^{b_t}. \end{aligned}$$

The above two inequalities respectively translate into

$$\begin{aligned} U_t^a(i) - U_t^{b_t}(i) &\leq 2c_{t-1}^a, \\ U_t^a(i) - U_t^{b_t}(i) &\geq -2c_{t-1}^{b_t}, \end{aligned}$$

In other words, for any $i \in \mathcal{V} \setminus \widehat{\mathcal{V}}_t^a$ we have $-2c_{t-1}^{b_t} \leq \rho_t^a(i) \leq 2c_{t-1}^a$. Subsequently,

$$\sum_{i \in \mathcal{L}_t^a} \rho_t^a(i) \leq \sum_{i \in \mathcal{L}_t^a} 2c_{t-1}^a, \quad (31)$$

$$\sum_{i \in \tilde{\mathcal{V}}_t^a \setminus \widehat{\mathcal{V}}_t^a} -2c_{t-1}^{b_t} \leq \sum_{i \in \tilde{\mathcal{V}}_t^a \setminus \widehat{\mathcal{V}}_t^a} \rho_t^a(i). \quad (32)$$

Therefore

$$\begin{aligned} \sum_{i \in \tilde{\mathcal{V}}_t^a} \rho_t^a(i) + \sum_{i \in \mathcal{L}_t^a} \rho_t^a(i) &\leq \sum_{i \in \tilde{\mathcal{V}}_t^a} \rho_t^a(i) + \sum_{i \in \mathcal{L}_t^a} 2c_{t-1}^a + \sum_{i \in \tilde{\mathcal{V}}_t^a \setminus \widehat{\mathcal{V}}_t^a} (\rho_t^a(i) + 2c_{t-1}^{b_t}) \\ &= \sum_{i \in \tilde{\mathcal{V}}_t^a} \rho_t^a(i) + \sum_{i \in \mathcal{L}_t^a} 2c_{t-1}^a + \sum_{i \in \tilde{\mathcal{V}}_t^a \setminus \widehat{\mathcal{V}}_t^a} 2c_{t-1}^{b_t} \\ &\leq \sum_{i \in \tilde{\mathcal{V}}_t^a} \rho_t^a(i) + 4Lc_{t-1}^a + 4Lc_{t-1}^{b_t}. \end{aligned} \quad (33)$$

The first inequality makes use of (31) and (32), and in particular the last term on the right-hand side is non-negative thanks to (32). The equality in the second line is true because $\widehat{\mathcal{V}}_t^a \subseteq \tilde{\mathcal{V}}_t^a$. Finally, for the last inequality we use the fact that $\text{card}(\mathcal{L}_t^a) \leq 2L$ and $\text{card}(\tilde{\mathcal{V}}_t^a) \leq 2L$. Along with (33) we get an upper bound on $\sum_{i \in \tilde{\mathcal{V}}_t^a} \rho_t^a(i) + \sum_{i \in \mathcal{L}_t^a} \rho_t^a(i)$ that involves $\sum_{i \in \tilde{\mathcal{V}}_t^a} (\mu^a(i) - \mu^{b_t}(i))$.

Conclusion. Combining (27), (29), (30), and (33), we obtain

$$\sum_{i \in \tilde{\mathcal{V}}_t^a} (\mu^a(i) - \mu^{b_t}(i)) + 8Lc_{t-1}^a + 4Lc_{t-1}^{b_t} \geq \sum_{i \in \tilde{\mathcal{V}}_t^{a^*}} (\mu^{a^*}(i) - \mu^{b_t}(i)) - 4Lc_{t-1}^{b_t}.$$

By the choice of b_t we have $c_{t-1}^a \geq c_{t-1}^{b_t}$. With (28), it follows that $16Lc_{t-1}^a \geq \Delta^a$ and thus

$$N_{t-1}^a \leq \frac{512L^2 \log(1/\delta')}{(\Delta^a)^2}.$$

Repeating the argument in the proof of [Theorem 1](#) (i.e., choose t to be the last time that action a is taken if it is taken more than once and apply decomposition (16)) gives the desired result. \square

Theorem 2. Let $\delta' = \delta/(2KmT)$. Then the regret of UPUCB-nAff (bl) ([Algorithm 3](#)) (resp. UPUCB-nAff, [Algorithm 4](#)), with probability at least $1 - \delta$, satisfies:

$$\text{Reg}_T \leq \sum_{a \in \mathcal{A}: \Delta^a > 0} \left(\frac{\alpha L^2 \log(2KmT/\delta)}{\Delta^a} + \Delta^a \right), \quad (10)$$

where $\alpha = 32$ (resp. 512) in the above inequality.

Proof. The proof follows by combining [Propositions 4](#) and [5](#). \square

A.4 Unknown Affected Variables: Proofs of [Theorems 3](#) and [4](#)

Theorem 3. Let $\delta' = \delta/(2KmT)$. Then the regret of UPUCB-iLift (bl) ([Algorithm 5](#)), with probability at least $1 - \delta$, satisfies:

$$\text{Reg}_T \leq \sum_{a \in \mathcal{A}: \Delta^a > 0} \left(\frac{8 \text{clip}(\Delta^a/\Delta_{\text{up}}, L^a, m)^2 \log(2KmT/\delta)}{\Delta^a} + \Delta^a \right). \quad (12)$$

Proof. As before, we prove (12) under the condition that \mathcal{E} defined in (22) occurs. Recall that we have shown in the proof of [Proposition 4](#) that $\mathbb{P}(\mathcal{E}) \geq 1 - \delta$ with our choice of δ' , so this effectively leads to a high-probability regret bound.

In the following we assume that \mathcal{E} occurs and show that for any suboptimal action $a \in \mathcal{A}$, it holds

$$N_T^a \leq \max \left(\min \left(\frac{8 \log(1/\delta')}{\Delta_{\text{up}}^2}, \frac{8m^2 \log(1/\delta')}{(\Delta^a)^2} \right), \frac{8(L^a)^2 \log(1/\delta')}{(\Delta^a)^2} \right) + 1. \quad (34)$$

The inequality (12) then follows directly from (16) and a rearrangement of the terms.

We first claim that for any $a \in \mathcal{A}$ if $N_t^a \geq N_0$ then $\widehat{\mathcal{V}}_t^a = \mathcal{V}^a$. In fact, by the definition of N_0 , $N_t^a \geq N_0$ implies that $c_t^a < \Delta_{\text{up}}/2$. Therefore, as argued in the text, $i \in \widehat{\mathcal{V}}_t^a$ if and only if $i \in \mathcal{V}^a$. Since $\rho_t^a(i) \geq \mu_{\text{up}}^a(i)$ for any a and i , this also proves that $\rho_t^a \geq r_{\text{up}}^a$ always holds (we have either $\widehat{\mathcal{V}}_t^a = \mathcal{V}$ or $\widehat{\mathcal{V}}_t^a = \mathcal{V}^a$).

Let $t \in \{K+1, \dots, T\}$ be the last time that a suboptimal action a is taken so that $N_T^a = N_{t-1}^a + 1$. This indicates $\tau_t^a \geq \tau_t^{a^*}$, and hence $\tau_t^a \geq r_{\text{up}}^{a^*}$. We distinguish between the following two cases:

1. $N_{t-1}^a < N_0$: Then $\widehat{\mathcal{V}}_t^a = \mathcal{V}$ so $\tau_t^a = \sum_{i \in \mathcal{V}} (U_t^a(i) - \mu^0) \leq r_{\text{up}}^a + 2mc_{t-1}^a$, which in turn leads to $c_{t-1}^a \geq \Delta^a/(2m)$, or equivalently $N_{t-1}^a \leq 8m^2 \log(1/\delta')/(\Delta^a)^2$.
2. $N_{t-1}^a \geq N_0$: Then $\widehat{\mathcal{V}}_t^a = \mathcal{V}^a$ so $\tau_t^a = \sum_{i \in \mathcal{V}^a} (U_t^a(i) - \mu^0) \leq r_{\text{up}}^a + 2L^a c_{t-1}^a$, which in turn leads to $c_{t-1}^a \geq \Delta^a/(2L^a)$, or equivalently $N_{t-1}^a \leq 8(L^a)^2 \log(1/\delta')/(\Delta^a)^2$.

Since we also have $N_{t-1}^a < N_0$ in the first case, the above gives

$$N_{t-1}^a \leq \max \left(\min \left(N_0 - 1, \frac{8m^2 \log(1/\delta')}{(\Delta^a)^2} \right), \frac{8(L^a)^2 \log(1/\delta')}{(\Delta^a)^2} \right).$$

Plugging in the definition of N_0 we get immediately (34). \square

Theorem 4. Let *Assumption 1* hold and $\delta' = \delta/(2KmT)$. Then the regret of UPUCB-iLift (*Algorithm 6*), with probability at least $1 - \delta$, satisfies:

$$\text{Reg}_T \leq \sum_{a \in \mathcal{A}: \Delta^a > 0} \left(\frac{8 \text{clip}(2\Delta^a/\Delta_{\text{up}}, L^a + L^{a^*}, 2m)^2 \log(2KmT/\delta)}{\Delta^a} + \Delta^a \right). \quad (15)$$

Proof. We shall again consider the event \mathcal{E} defined in (22), which holds with probability $1 - \delta$ as argued in the proof of *Proposition 4*. Leveraging the decomposition in (16) it is sufficient to prove that on \mathcal{E} we have

$$N_T^a \leq \max \left(\min \left(\frac{32 \log(1/\delta')}{\Delta_{\text{up}}^2}, \frac{32m^2 \log(1/\delta')}{(\Delta^a)^2} \right), \frac{8(L^a + L^{a^*})^2 \log(1/\delta')}{(\Delta^a)^2} \right) + 1. \quad (35)$$

We assume \mathcal{E} happens in the rest of the proof.

Let us first focus on the elimination phase, notice that event \mathcal{E} implies $|\tilde{r}_\ell^a - r^a| < c_\ell$ for every $a \in \mathcal{A}_\ell$. In this case, the optimal actions never get eliminated. In fact, if $a^* \in \mathcal{A}_\ell$, then for all $a \in \mathcal{A}_\ell$ it holds that

$$\tilde{r}_\ell^{a^*} + 2c_\ell \geq r^{a^*} + c_\ell \geq r^a + c_\ell \geq \tilde{r}_\ell^a,$$

As a consequence, we also have $a^* \in \mathcal{A}_{\ell+1}$ and we conclude by the principle of induction. With this we further deduce that all action a with $\Delta^a \geq 4c_\ell$ gets eliminated after it is taken at most ℓ times. To see this, notice that for such action we have

$$\tilde{r}_\ell^a + 2c_\ell < r^a + 3c_\ell = r^{a^*} - \Delta^a + 3c_\ell < \tilde{r}_\ell^{a^*} - \Delta^a + 4c_\ell \leq \tilde{r}_\ell^{a^*}.$$

Therefore, for all t in phase I of the algorithm, it is true that

$$N_t^a \leq \min \left(\left\lceil \frac{32 \log(1/\delta')}{\Delta_{\text{up}}^2} \right\rceil, \left\lceil \frac{32m^2 \log(1/\delta')}{(\Delta^a)^2} \right\rceil \right). \quad (36)$$

This suggests that inequality (35) holds if the algorithm terminates in the elimination phase.

Otherwise, the algorithm constructs $\mathcal{S}(i)$'s, $\widehat{\mathcal{V}}^a$'s, and switches to the UPUCB phase. Let us denote by $\mathcal{S}(i) = \{a \in \mathcal{A} : i \notin \mathcal{V}^a\}$ the set of actions that do not affect variable i and let $\mathcal{A}_{\text{res}} = \mathcal{A}_{N_0+1}$ be the actions that remain after phase I. We argue that the sets $\widehat{\mathcal{S}}(i)$ are constructed such that when \mathcal{E} occurs,

$$(i) \quad \widehat{\mathcal{S}}(i) \subseteq \mathcal{S}(i).$$

$$(ii) \quad \text{If there exist } a, a' \in \mathcal{A}_{\text{res}} \cap \mathcal{S}(i) \text{ such that } a \neq a' \text{ then we also have } a, a' \in \widehat{\mathcal{S}}(i).$$

The second point is proved by observing that under \mathcal{E} we must have $\mu^0(i) \in \mathcal{C}_{t_0+1}^a(i) \cap \mathcal{C}_{t_0+1}^{a'}(i)$ if $a, a' \in \mathcal{A}_{\text{res}} \cap \mathcal{S}(i)$. To show the first point, notice that after the elimination phase, each of the action in \mathcal{A}_{res} is taken exactly $\lceil (32/\Delta_{\text{up}}^2) \log(1/\delta') \rceil$ times and thus $c_{t_0}^a \leq \Delta_{\text{up}}/4$. By *Assumption 1*, if $a \notin \mathcal{S}(i)$, then for all $a' \in \mathcal{A}_{\text{res}} \setminus \{a\}$ we have $|\mu^a(i) - \mu^{a'}(i)| \geq \Delta_{\text{up}}$, which along with the definition of \mathcal{E} implies $\mathcal{C}_{t_0+1}^a(i) \cap \mathcal{C}_{t_0+1}^{a'}(i) = \emptyset$, and therefore $a \notin \widehat{\mathcal{S}}(i)$.

The rest of the proof follows closely that of [Theorem 1](#). Suppose that a suboptimal action a is taken at round t during phase II of the algorithm. This means $\tau_t^a \geq \tau_t^{a^*}$ (recall that $a^* \in \mathcal{A}_{\text{res}}$ since \mathcal{E} happens), i.e.,

$$\sum_{i \in \widehat{\mathcal{V}}^a} (U_t^a(i) - U_t^0(i)) \geq \sum_{i \in \widehat{\mathcal{V}}^{a^*}} (U_t^{a^*}(i) - U_t^0(i)).$$

Rearranging the terms we get

$$\sum_{i \in \widehat{\mathcal{V}}^a} U_t^a(i) + \sum_{i \in \widehat{\mathcal{V}}^{a^*} \setminus \widehat{\mathcal{V}}^a} U_t^0(i) \geq \sum_{i \in \widehat{\mathcal{V}}^{a^*}} U_t^{a^*}(i) + \sum_{i \in \widehat{\mathcal{V}}^a \setminus \widehat{\mathcal{V}}^{a^*}} U_t^0(i).$$

If $i \in \widehat{\mathcal{V}}^{a^*} \setminus \widehat{\mathcal{V}}^a$ then $a \in \widehat{\mathcal{S}}(i)$. In particular, $\widehat{\mathcal{S}}(i) \neq \emptyset$ so $a' = a_t^0(i) = \arg \max_{\tilde{a} \in \widehat{\mathcal{S}}(i)} N_{t-1}^{\tilde{a}}$ is well-defined and $U_t^0(i) = U_t^{a'}(i)$. With $\widehat{\mathcal{S}}(i) \subseteq \mathcal{S}(i)$ we know that $\mu^{a'}(i) = \mu^0(i)$ and accordingly $|\mu_{t-1}^{a'}(i) - \mu^0(i)| \leq c_{t-1}^{a'}$ by the fact that \mathcal{E} happens. This shows $U_t^0(i) \geq \mu^0(i)$. Moreover, $N_t^{a'} \geq N_t^a$ by the choice of a' ; hence, $U_t^0(i) \leq \mu^0(i) + 2c_{t-1}^{a'} \leq \mu^0(i) + 2c_{t-1}^a$. The same argument applies to $i \in \widehat{\mathcal{V}}^a \setminus \widehat{\mathcal{V}}^{a^*}$, and the UCB indices of a and a^* can be bounded from above and below as in previous proofs. In summary, we obtain

$$\sum_{i \in \widehat{\mathcal{V}}^a} (\mu^a(i) + 2c_{t-1}^a) + \sum_{i \in \widehat{\mathcal{V}}^{a^*} \setminus \widehat{\mathcal{V}}^a} (\mu^0(i) + 2c_{t-1}^a) \geq \sum_{i \in \widehat{\mathcal{V}}^{a^*}} \mu^{a^*}(i) + \sum_{i \in \widehat{\mathcal{V}}^a \setminus \widehat{\mathcal{V}}^{a^*}} \mu^0(i).$$

Equivalently,

$$2 \text{card}(\widehat{\mathcal{V}}^a \cup \widehat{\mathcal{V}}^{a^*}) c_{t-1}^a \geq \sum_{i \in \widehat{\mathcal{V}}^a \cup \widehat{\mathcal{V}}^{a^*}} \mu^{a^*}(i) - \sum_{i \in \widehat{\mathcal{V}}^a \cup \widehat{\mathcal{V}}^{a^*}} \mu^a(i). \quad (37)$$

To conclude, we claim that

$$\widehat{\mathcal{V}}^a \cup \widehat{\mathcal{V}}^{a^*} = \mathcal{V}^a \cup \mathcal{V}^{a^*}. \quad (38)$$

Let us first show $\mathcal{V}^a \subseteq \widehat{\mathcal{V}}^a \subseteq \mathcal{V}^a \cup \mathcal{V}^{a^*}$. The first inclusion is a consequence of $\widehat{\mathcal{S}}(i) \subseteq \mathcal{S}(i)$:

$$i \in \mathcal{V}^a \implies a \notin \mathcal{S}(i) \implies a \notin \widehat{\mathcal{S}}(i) \implies i \in \widehat{\mathcal{V}}^a.$$

As for the second inclusion we prove it's contrapositive:

$$i \notin \mathcal{V}^a \cup \mathcal{V}^{a^*} \implies a, a^* \in \mathcal{A}_{\text{res}} \cap \mathcal{S}(i) \implies a, a^* \in \widehat{\mathcal{S}}(i) \implies i \notin \widehat{\mathcal{V}}^a$$

Notice that the second implication holds because a is suboptimal and in particular $a \neq a^*$. The inclusion $\mathcal{V}^{a^*} \subseteq \widehat{\mathcal{V}}^{a^*} \subseteq \mathcal{V}^a \cup \mathcal{V}^{a^*}$ can be proved in a similar way, and combining these results gives (38). Therefore, the right-hand side of (37) is exactly Δ^a while the left-hand side of (37) is bounded from above by $2(L^a + L^{a^*})c_{t-1}^a$. In consequence,

$$N_{t-1}^a \leq \frac{8(L^a + L^{a^*})^2 \log(1/\delta')}{(\Delta^a)^2}.$$

Therefore, any suboptimal action that gets taken in the UPUCB phase satisfies

$$N_T^a \leq \frac{8(L^a + L^{a^*})^2 \log(1/\delta')}{(\Delta^a)^2} + 1.$$

On the other hand, if an suboptimal action is not taken in the UPUCB phase then (36) applies, so in either case (35) is verified, which concludes the proof. \square

B Lower Bounds

To prove our lower bounds, we first establish in [Appendix B.1](#) a general lemma for deriving instance-dependent lower-bounds for bandit problems with underlying variables, side observations, and prior knowledge on the distribution of the variables. Subsequently, in [Appendix B.2](#) we regard uplifting bandit as a special case of this model and apply the lemma to prove [Propositions 2](#) and [3](#).

B.1 A General Information-theoretic Lower Bound

Let $y = (y(1), \dots, y(m))$ be a random vector of m variables that entirely determines the reward so that we can write $r = \varphi(y)$ for some deterministic reward function $\varphi: \mathbb{R}^m \rightarrow \mathbb{R}$. Let \mathcal{P}^a be the distribution of action a on the underlying variables and \mathcal{P}^0 be the baseline distribution on these variables. At each round after taking an action, the learner observes a vector in \mathbb{R}^d that is itself a function of the variables, written as $o(y)$ where $o: \mathbb{R}^m \rightarrow \mathbb{R}^d$ is the observation function. The decision of which action to take can then only be based on the interaction history $(a_1, o(y_1), \dots, a_{t-1}, o(y_{t-1}))$ and some prior knowledge of the learner about the distributions of the variables $\Upsilon(\mathcal{P}^0, \dots, \mathcal{P}^K)$.⁶

A class of action distributions is defined as $\mathcal{B} = \mathcal{M}^0 \times \dots \times \mathcal{M}^K$ where each \mathcal{M}^a is a set of distributions on \mathbb{R}^m . We say that \mathcal{B} is indistinguishable under Υ if for all $\mathcal{P}_* = (\mathcal{P}^a)_{a \in \mathcal{A}_0}$, $\mathcal{Q}_* = (\mathcal{Q}^a)_{a \in \mathcal{A}_0} \in \mathcal{B}$ we have $\Upsilon(\mathcal{P}^0, \dots, \mathcal{P}^K) = \Upsilon(\mathcal{Q}^0, \dots, \mathcal{Q}^K)$. A bandit problem with underlying variables, side observations, and prior knowledge on the variable distribution is thus defined by the quadruple $(\mathcal{B}, \varphi, o, \Upsilon)$. A policy π is said to be *consistent* over (\mathcal{B}, φ) if for all $p > 0$ and all instance of (\mathcal{B}, φ) , the regret of π on the instance satisfies $\text{Reg}_t = o(t^p)$. A policy π is compatible with observation o and prior knowledge Υ if it can be implemented by a learner that observes $o(y)$ and has prior knowledge $\Upsilon(\mathcal{P}^0, \dots, \mathcal{P}^K)$.

By abuse of notation, for any distributions \mathcal{P} on \mathbb{R}^m we will write $\varphi(\mathcal{P}) = \mathbb{E}_{y \sim \mathcal{P}}(\varphi(y))$ as the expected reward when the variables follow distribution \mathcal{P} whenever this quantity is defined. We also write \mathcal{P}_o as the pushforward of \mathcal{P} along o . Let \mathcal{M} be a set of distributions such that $\varphi(\mathcal{P})$ is well-defined for any $\mathcal{P} \in \mathcal{M}$. Let $r^* \in \mathbb{R}$ and $\mathcal{P} \in \mathcal{M}$ such that $\varphi(\mathcal{P}) < r^*$. The following quantity is crucial to the analysis

$$d_{\inf}(\mathcal{P}, \mathcal{M}, r^*, \varphi, o) = \inf_{\mathcal{Q} \in \mathcal{M}} \{D_{\text{KL}}(\mathcal{P}_o, \mathcal{Q}_o) : \varphi(\mathcal{Q}) > r^*\},$$

where $D_{\text{KL}}(\mathcal{P}_o, \mathcal{Q}_o)$ denotes the KL divergence of \mathcal{P}_o from \mathcal{Q}_o . Intuitively, it quantifies how difficult it is to learn that the expected reward $\varphi(\mathcal{P})$ is smaller than r^* (smaller the value of d_{\inf} the more difficult it is). The next lemma is a straightforward adaptation of [\[Burnetas and Katehakis, 1996, Th. 1\]](#), [\[Lattimore and Szepesvári, 2020, Th. 16.2\]](#) to our model. It provides a lower bound on the asymptotic regret

Lemma 4. *Let $\mathcal{B} = \mathcal{M}^0 \times \dots \times \mathcal{M}^K$ be a class of action distributions and φ be the reward function. Let π be a consistent algorithm over (\mathcal{B}, φ) that is compatible with prior knowledge function Υ and observation function o . Then, if \mathcal{B} is indistinguishable under Υ , for all $\mathcal{P}_* \in \mathcal{B}$, it holds that*

$$\liminf_{T \rightarrow +\infty} \frac{\mathbb{E}[\text{Reg}_T]}{\log T} \geq \sum_{a \in \mathcal{A}: \Delta^a > 0} \frac{\Delta^a}{d_{\inf}(\mathcal{P}^a, \mathcal{M}^a, r^*, \varphi, o)}. \quad (39)$$

Proof. The proof follows closely the one presented in [\[Lattimore and Szepesvári, 2020\]](#). The key observation is that the same proof still carries out if we have a random vector and our reward and observation depends on this random vector. In particular, let $\mathbb{Q}(\pi, \mathcal{P}_*)$ and $\mathbb{Q}(\pi, \mathcal{Q}_*)$ be the probability measure on the actions and the observations induced by the interaction of the policy with the bandit instances (\mathcal{P}_*, φ) and (\mathcal{Q}_*, φ) for

⁶For our purpose, we can regard this as a partition of the set of all bandit instances that the learner may encounter.

$\mathcal{P}_*, \mathcal{Q}_* \in \mathcal{B}$. The assumption that π is compatible with o and Υ and that \mathcal{B} is indistinguishable under Υ allows us to write $q_{\pi, \mathcal{P}_*}(a_1, o_1, \dots, a_T, o_T) = \prod_{t=1}^T \pi_t^{\mathcal{B}}(a_t | a_1, o_1, \dots, a_{t-1}, o_{t-1}) p_o^{a_t}(o_t)$, where o_1, \dots, o_T are the observations of the learner, q_{π, \mathcal{P}_*} and $p_o^{a_t}$ are respectively the probability density functions of $\mathbb{Q}(\pi, \mathcal{P}_*)$ and \mathcal{P}_o^a , and $\pi_t^{\mathcal{B}} = \pi_{t, \Upsilon}(\mathcal{P}_*) = \pi_{t, \Upsilon}(\mathcal{Q}_*)$ is a probability kernel that can be chosen using prior knowledge.⁷ We have a similar decomposition for q_{π, \mathcal{Q}_*} . Then, following the proof of [Lattimore and Szepesvári, 2020, Lem. 15.1] we get immediately

$$D_{\text{KL}}(\mathbb{Q}(\pi, \mathcal{P}_*), \mathbb{Q}(\pi, \mathcal{Q}_*)) = \sum_{a \in \mathcal{A}} \mathbb{E}_{\pi, \mathcal{P}}[N_T^a] D_{\text{KL}}(\mathcal{P}_o^a, \mathcal{Q}_o^a).$$

In the above, $\mathbb{E}_{\pi, \mathcal{P}}$ means that the expectation is taken with respect to $\mathbb{Q}(\pi, \mathcal{P}_*)$. The proof can be completed in the same way as done for [Lattimore and Szepesvári, 2020, Th. 16.2]. \square

B.2 Lower Bounds for Uplifting Bandits– Proof

The following proposition includes both Proposition 2 and Proposition 3.

Proposition 6. *Let π be a consistent algorithm over the class of 1-sub-Gaussian uplifting bandits that at most uses knowledge of \mathcal{P}^0 , $(\mathcal{V}^a)_{a \in \mathcal{A}}$, and the fact that the noise is 1-sub-Gaussian. Then, for any $K, m > 0$ and sequence $(L^a, \Delta^a)_{1 \leq a \leq K} \in ([m] \times \mathbb{R}_+)^K$ with $\Delta^1 = 0$, there exists a 1-sub-Gaussian uplifting bandit with parameters $(K, m, (\Delta^a, L^a)_{a \in \mathcal{A}})$ such that the regret induced by π on it satisfies*

$$\liminf_{T \rightarrow +\infty} \frac{\mathbb{E}[\text{Reg}_T]}{\log T} \geq \sum_{a \in \mathcal{A}: \Delta^a > 0} \frac{2(L^a)^2}{\Delta^a}. \quad (40)$$

Moreover, if π can be implemented under either of the following conditions

- (a) Only the reward is observed.
- (b) The learner does not have any prior knowledge about the arms' expected payoffs.

Then, there exists a 1-sub-Gaussian uplifting bandit with parameters $(K, m, (\Delta^a, L^a)_{a \in \mathcal{A}})$ where the regret of π is

$$\liminf_{T \rightarrow +\infty} \frac{\mathbb{E}[\text{Reg}_T]}{\log T} \geq \sum_{a \in \mathcal{A}: \Delta^a > 0} \frac{2m^2}{\Delta^a}. \quad (41)$$

Proof. Let $r^* = 1 + \max_{a \in \mathcal{A}} \Delta^a > 0$. Throughout the proofs, the lower bound will be shown for problem instances with $\mu^0(i) \equiv 0$ and the following mean values for $a \in \mathcal{A}$

$$\mu^a(i) = \begin{cases} \frac{r^* - \Delta^a}{L^a} & \text{if } 1 \leq i \leq L^a, \\ 0 & \text{if } L^a + 1 \leq i \leq m. \end{cases} \quad (42)$$

By construction, we have $\mu^a(i) \neq \mu^0(i)$ if and only if $1 \leq i \leq L^a$ and thus the number of mean-affected variables of a is exactly L^a . Clearly, 1 is an optimal action and the suboptimality gap of action a is Δ^a .

Next, we will need to define the classes of action distributions $\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3$ that we will use for proving (40), (a)-(41), and (b)-(41). To begin, we want these classes to be indistinguishable under a certain prior knowledge function, and since we allow the learner to know the baseline distribution, \mathcal{M}^0 must be a singleton. As for

⁷Many technical details are omitted here. We refer the interested readers to [Lattimore and Szepesvári, 2020] for a rigorous treatment of such proof.

\mathcal{M}^a , it will take the form $\mathcal{M}^a = \{\mathcal{N}(\mu, \Sigma^a) : \mu \in \mathcal{U}^a\}$ for some mean vector set $\mathcal{U}^a \subset \mathbb{R}^m$ and covariance matrix $\Sigma^a \in [0, 1]^{m \times m}$ to be specified. This implies that every noise variable $\xi^a(i)$ is 1-sub-Gaussian and thus π is consistent in each of these classes. Moreover, the classes will also be defined in such a way that when the mean values are taken as in (42), \mathcal{P}^0 and \mathcal{P}^a indeed have the same marginal distribution on all but the first L^a variables.

The Importance of Limited Number of Affected Variables. To prove (40), we let the unique element of \mathcal{M}_0 be the Dirac measure that assigns full weight to the point $(0, \dots, 0)$. For any $a \in \mathcal{A}$, we define $\mathcal{U}^a = (\mathbb{R} \setminus \{0\})^{L^a} \times \{0\}^{m-L^a}$ and block matrix Σ^a such that $\Sigma^a(i, j) = \mathbb{1}\{\max(i, j) \leq L^a\}$. In other words, the random variable y^a can be written in the following form

$$y^a(i) = \begin{cases} \mu^a(i) + \xi^a & \text{if } 1 \leq i \leq L^a, \\ 0 & \text{if } L^a + 1 \leq i \leq m, \end{cases}$$

where ξ^a is Gaussian with zero mean and unit variance. Notice that the same noise applies to all the affected variables. By construction, every $\mathcal{P}^a \in \mathcal{M}^a$ has exactly the first L^a variables affected, i.e., $\mathcal{V}^a = \{1, \dots, L^a\}$. Thus \mathcal{B}_1 defined in this way is indistinguishable under prior knowledge on baseline distribution \mathcal{P}^0 and the sets of affected variables $(\mathcal{V}^a)_{a \in \mathcal{A}}$. To conclude, we would like to apply Lemma 4. For this, we observe that for any $\mathcal{P}^a, \mathcal{Q}^a \in \mathcal{M}^a$ with means μ^a and ν^a , it holds⁸

$$D_{\text{KL}}(\mathcal{P}^a, \mathcal{Q}^a) = \begin{cases} \frac{\tilde{\Delta}^2}{2} & \text{if } \mu^a(i) - \nu^a(i) = \tilde{\Delta} \text{ for all } i \in \{1, \dots, L^a\}, \\ +\infty & \text{otherwise.} \end{cases} \quad (43)$$

With this we see immediately that for \mathcal{P}^a with mean values (42), we have $d_{\text{inf}}(\mathcal{P}^a, \mathcal{M}^a, r^*, (+), \text{Id}) = \Delta^a / (2(L^a)^2)$, where $(+)$ and Id are respectively the sum and the identity function that gives the reward and the observations of the learner. Plugging this into (39) gives exactly (40).

The importance of Side Observation. To construct \mathcal{B}_2 , we draw the same noise for all the variables. In terms of covariance matrix, this means all the entries of Σ is equal to 1. The baseline distribution is $\mathcal{P}^0 = \mathcal{N}(\mathbf{0}, \Sigma)$ and for every a , we set $\Sigma^a = \Sigma$ and choose again $\mathcal{U}^a = (\mathbb{R} \setminus \{0\})^{L^a} \times \{0\}^{m-L^a}$. This time, only the reward, that is, the sum of the variables, is observed, and the distribution of the reward at action a is $\mathcal{N}(r^* - \Delta^a, m^2)$. For any two real numbers γ, γ' , one has $D_{\text{KL}}(\mathcal{N}(\gamma, m^2), \mathcal{N}(\gamma', m^2)) = (\gamma - \gamma')^2 / (2m^2)$. Therefore $d_{\text{inf}}(\mathcal{P}^a, \mathcal{M}^a, r^*, (+), (+)) = \Delta^a / (2m^2)$ for \mathcal{P}^a with mean values (42). Notice that \mathcal{B}_2 is also indistinguishable under prior knowledge on the baseline distribution and on the sets of affected variables $(\mathcal{V}^a)_{a \in \mathcal{A}}$, and thus Lemma 4 can be applied, leading to (41).

The Importance of Prior Knowledge on How the Variables are Affected. The class \mathcal{B}_3 is defined almost in the same way as \mathcal{B}_2 except for the fact that we now set $\mathcal{U}^a = \mathbb{R}^m$ for every $a \in \mathcal{A}$. Our condition (b) postulates that \mathcal{B}_3 is indistinguishable under the prior knowledge of the learner. Similar to (43), for any $\mathcal{P}^a, \mathcal{Q}^a \in \mathcal{M}^a$ with means μ^a and ν^a , we now have

$$D_{\text{KL}}(\mathcal{P}^a, \mathcal{Q}^a) = \begin{cases} \frac{\tilde{\Delta}^2}{2} & \text{if } \mu^a(i) - \nu^a(i) = \tilde{\Delta} \text{ for all } i \in \{1, \dots, m\}, \\ +\infty & \text{otherwise.} \end{cases}$$

⁸This can for example be proved by applying the chain rule of KL divergence.

In particular, if μ^a is defined as in (42) and $\nu^a(i) = \mu^a(i) + (\Delta^a + \varepsilon)/m$ for all i for some $\varepsilon \in \mathbb{R}$, we have $D_{\text{KL}}(\mathcal{P}^a, \mathcal{Q}^a) = (\Delta^a + \varepsilon)^2/(2m^2)$. We deduce immediately $d_{\text{inf}}(\mathcal{P}^a, \mathcal{M}^a, r^*, (+), \text{Id}) = \Delta^a/(2m^2)$. Invoking Lemma 4 completes the proof. \square

C Contextual Combinatorial Uplifting Bandits

In this appendix, we discuss in more detail the problem of contextual combinatorial uplifting bandit that we introduced in Section 8. We extend the model and the algorithm to the case of multiple treatments in Appendix C.1. Subsequently, we specify the algorithm for the case of linear expected payoffs and present some associated analysis in Appendix C.2. Numerical experiments conducted with the Criteo Uplift Prediction Dataset are presented in Appendix C.3.

C.1 Model and Algorithm Template

Let us consider a pool of m individuals that may change from round to round, associated with (time-dependent) contexts $(x_t(i))_{i \in \mathcal{V}}$ with $x_t(i) \in \mathbb{R}^d$, and a set of treatments that can be applied to the individuals. For convenient, we also write $\mathcal{Z}_0 = \mathcal{Z} \cup \{0\}$, where, as before, 0 indicates that no treatment is applied to an individual. Following Section 8, we consider the action set $\mathcal{A} = \{a \in (\mathcal{Z}_0)^m : \|a\|_0 \leq L\}$, which encodes the constraint that at most L individuals can be treated in each run. Similarly, we associate each treatment with a function g^z . The reward and the uplift of choosing action a in round t are given then by

$$r^a(x_t) = \sum_{i \in \mathcal{V}} g^{a(i)}(x_t(i)), \quad r_{\text{up}}^a(x_t) = \sum_{i \in \mathcal{V}^a} g^{a(i)}(x_t(i)) - g^0(x_t(i)).$$

The problem thus consists in finding both the set of customers to treat and the treatments to apply to these customers at each round. To adapt C2UPUCB to this situation, we define a treatment-dependent uplifting index $\rho_t^z(i) = U_t^z(i) - U_t^0(i)$ where $U_t^z(i)$ and $U_t^0(i)$ are respectively upper confidence bounds on $g^{a(i)}(x_t(i))$ and $g^0(x_t(i))$. Again, if g^0 is known, $U_t^0(i)$ can be simply taken as $g^0(x_t(i))$. Next, at each time step t , we define the presumed optimal treatment for individual i and the associated uplifting index as

$$z_t(i) = \arg \max_{z \in \mathcal{Z}} \rho_t^z(i), \quad \rho_t(i) = \rho_t^{z_t(i)}(i) = \max_{z \in \mathcal{Z}} \rho_t^z(i). \quad (44)$$

The adapted C2UPUCB algorithm then chooses the L customers with the largest uplifting index and for each of these customers chooses their presumed optimal treatment. The resulting algorithm (referred to as C2UPUCB-M) is summarized in Algorithm 8.

C.2 Linear Expected Payoffs

To provide a concrete algorithm and its analysis thereof, let us assume that g^z is linear in the feature vector with some unknown coefficient vector θ^z , i.e., $g^z(x_t(i)) = \langle \theta^z, x_t(i) \rangle$. Then, we can construct upper confidence bounds for $g^z(x_t(i))$ following the ideas of linear UCB [Abbasi-Yadkori et al., 2011]. To begin, for each $z \in \mathcal{Z}_0$ and $t \in \{1, \dots, T\}$, we define

$$\mathcal{V}_t^z = \{i \in \mathcal{V} : a_t(i) = z\}$$

as the set of customers that receive treatment z in round t . Fix a regularization parameter $\lambda > 0$, we further define the following $d \times d$ matrix and d -dimensional vector

$$V_t^z = \lambda I + \sum_{s=1}^t \sum_{i \in \mathcal{V}_s^z} x_s(i) x_s(i)^\top, \quad b_t^z = \sum_{s=1}^t \sum_{i \in \mathcal{V}_s^z} y_s(i) x_s(i). \quad (45)$$

Algorithm 8 C2UPUCB-M

```
1: Input: Budget constraint  $L$ , Parameters for computing the UCBs indices, Treatment set  $\mathcal{Z}_0$ 
2: for  $t = 1, \dots, T$  do
3:   for  $i \in \mathcal{V}$  do
4:     Compute UCB indices  $(U_t^z(i))_{z \in \mathcal{Z}_0}$  for  $(g^z(x_t(i)))_{z \in \mathcal{Z}_0}$            {use (46) for linear payoffs}
5:     Compute uplifting indices  $(\rho_t^z(i))_{z \in \mathcal{Z}}$                              { $\rho_t^z(i) = U_t^z(i) - U_t^0(i)$ }
6:     Compute presumed optimal treatment  $z_t(i)$  and uplifting index  $\rho_t(i)$        {use (44)}
7:   end for
8:   Set  $\mathcal{V}_t \leftarrow \arg \max_{S \subset \mathcal{V}, \text{card}(S) \leq L} \sum_{i \in S} \rho_t(i)$            {this defines the set of treated customers}
9:   Select action  $a_t \leftarrow (z_t(i) \mathbb{1}_{\{i \in \mathcal{V}_t\}})_{i \in \mathcal{V}}$                    { $a_t(i) = z_t(i)$  if  $i \in \mathcal{V}_t$ ;  $a_t(i) = 0$  otherwise}
10: end for
```

The regularized least-square estimate of θ^z is then given by $\hat{\theta}_t^z = (V_t^z)^{-1}b_t^z$. The UCB index is of the form

$$U_t^z(i) = \langle \hat{\theta}_{t-1}^z, x_t(i) \rangle + \alpha_t \|x_t(i)\|_{(V_{t-1}^z)^{-1}}. \quad (46)$$

for some α_t that is properly chosen.

Regret Guarantee. In this part, we provide regret analysis for Algorithm 8 when *i*) the baseline payoffs are unknown, *ii*) the expected payoffs are linear, and *iii*) the UCB indices are computed following (46). For technical reasons, we will assume the noises $\xi_t(i)$ to be mutually independent (here $\xi_t(i) = y_t(i) - \langle \theta^z, x_t(i) \rangle$). We comment on this assumption in Remark 3.

The theorem below establishes a high-probability regret bound for C2UPUCB-M when the above conditions are satisfied.

Theorem 5. Assume that $\|x_t(i)\| \leq 1$ for all t and i and $(\xi_t(i))_{i \in \mathcal{V}}$ are mutually independent and 1-sub-Gaussian (conditioning on the past). Let $S > 0$ such that $\|\theta^z\| \leq S$ for all $z \in \mathcal{Z}_0$. Then if C2UPUCB-M (Algorithm 8) is run with the UCB indices defined in (46) under the parameter choices

$$\alpha_t = \sqrt{\lambda}S + \sqrt{2 \log \left(\frac{q+1}{\delta} \right) + d \log \left(1 + \frac{mt}{d\lambda} \right)},$$

and $\lambda \geq L$, it holds, with probability at least $1 - \delta$, that

$$\text{Reg}_T \leq 4 \sqrt{dL(q+1)T \log \left(1 + \frac{LT}{d\lambda} \right)} \left(\sqrt{\lambda}S + \sqrt{2 \log \left(\frac{q+1}{\delta} \right) + d \log \left(1 + \frac{mT}{d\lambda} \right)} \right). \quad (47)$$

Theorem 5 provides a regret bound in $\tilde{\mathcal{O}}(\sqrt{\max(d, L)dqLT})$. We observe that the dependence on m is only logarithmic. In fact, m plays no role in either the number of parameters to estimate (which is $(q+1)d$) or the scale of the noise (which is in $\mathcal{O}(\sqrt{L})$). Ideally, we would like to remove completely the dependence on m . However, we are not able to achieve this in our proof, and we leave whether this is possible or not as an open question. Finally, due the mutual independence of noises, the regret only scales with \sqrt{L} as long as $L \leq d$.

Remark 3. The $\log(m)$ factor appears because $\hat{\theta}_t^0$ is estimated using feedback from roughly mT samples (see Lemma 5 below). This is also the reason why we need mutual independence of the noises, since otherwise, this estimation may cause the regret to get \sqrt{m} larger. It can be easily verified that the dependence on m can

be completely removed when the baseline payoffs are known and used for computing the uplifting indices. In this case, we can also show if the noises are not mutually independent, multiplying the second term of α_t by \sqrt{L} still guarantees a regret in $\tilde{O}(dL\sqrt{qT})$. What is at stake here is [Lemma 5](#) that we present below, as the size of the confidence ellipsoid varies for different underlying assumptions.

Analysis. To prove [Theorem 5](#), we need to show that $U_t^a(i)$ as defined in (46) is indeed an upper confidence bound on $\langle \theta^z, x_t(i) \rangle$. This step is based on the following lemma which establishes the confidence ellipsoids of the estimated parameters. In fact, $U_t^a(i)$ defined in (46) can be regarded as the largest estimated payoff that one can get from choosing a parameter from the confidence ellipsoid.

Lemma 5. *Assume that $\|x_t(i)\| \leq 1$ for all t and i and $(\xi_t(i))_{i \in \mathcal{V}}$ are mutually independent and 1-sub-Gaussian (conditioning on the past). For all $z \in \mathcal{Z}_0$ and $\delta' \in (0, 1)$, with probability at least $1 - \delta'$ it holds for all t that*

$$\begin{aligned} \|\theta^z - \hat{\theta}_t^z\|_{V_t^z} &\leq \sqrt{\lambda} \|\theta^z\| + \sqrt{2 \log\left(\frac{1}{\delta'}\right) + \log\left(\frac{\det(V_t^z)}{\lambda^d}\right)} \\ &\leq \sqrt{\lambda} \|\theta^z\| + \sqrt{2 \log\left(\frac{1}{\delta'}\right) + d \log\left(1 + \frac{mt}{d\lambda}\right)}. \end{aligned}$$

Proof. This is an immediate corollary of [[Abbasi-Yadkori et al., 2011](#), Th. 2]. Since $(\xi_t(i))_{i \in \mathcal{V}}$ are mutually independent, we can apply this theorem by arranging $(\xi_t(i))_{i \in \mathcal{V}_t^z}$ in some arbitrary order $\xi_t(i_1), \dots, \xi_t(i_{\text{card}(\mathcal{V}_t^z)})$ so that $\xi_t(i_k)$ is 1-sub-Gaussian conditioning on $(\bigcup_{s < t} \{\xi_s(i) : i \in \mathcal{V}_s^z\}) \cup \{\xi_t(i_{k'})\}_{1 \leq k' \leq k}$.

For the second inequality, we notice that $\text{tr}(V_t^z) \leq d\lambda + \sum_{s=1}^t \sum_{i \in \mathcal{V}_s^z} \|x_t(i)\|^2 \leq d\lambda + mt$. It thus follows from the AM–GM inequality on the eigenvalues of V_t^z , $\det(V_t^z) \leq (\text{tr}(V_t^z)/d)^d$. \square

The lemma below will also be useful for bounding the regret.

Lemma 6. *Consider subsets of the variables $\mathcal{V}_1, \dots, \mathcal{V}_T \subset \mathcal{V}$ and define $V_t = \lambda I + \sum_{s=1}^t \sum_{i \in \mathcal{V}_s} x_s(i) x_s(i)^\top$. If $\lambda \geq \max_{t \in [T]} \text{card}(\mathcal{V}_t)$ and $\|x_t(i)\| \leq 1$ for all t and i , it holds that*

$$\sum_{t=1}^T \sum_{i \in \mathcal{V}_t} \|x_t(i)\|_{V_{t-1}^{-1}}^2 \leq 2 \log\left(\frac{\det(V_T)}{\lambda^d}\right) \leq 2d \log\left(1 + \frac{\sum_{t=1}^T \text{card}(\mathcal{V}_t)}{d\lambda}\right).$$

Proof. For all $t \geq 1$, we have

$$V_t = V_{t-1} + \sum_{i \in \mathcal{V}_t} x_t(i) x_t(i)^\top = V_{t-1}^{1/2} \left(I + \sum_{i \in \mathcal{V}_t} V_{t-1}^{-1/2} x_t(i) x_t(i)^\top V_{t-1}^{-1/2} \right) V_{t-1}^{1/2}$$

Thus, with $w_t(i) = V_{t-1}^{-1/2} x_t(i)$, we get

$$\det(V_t) = \det(V_{t-1}) \det\left(I + \sum_{i \in \mathcal{V}_t} w_t(i) w_t(i)^\top\right). \quad (48)$$

Regarding $\{w_t(i)\}_{i \in \mathcal{V}_t}$ as column vectors, we denote by their horizontal concatenation as $W \in \mathbb{R}^{d \times m_t}$, where $m_t = \text{card}(\mathcal{V}_t)$. We have $WW^\top = \sum_{i \in \mathcal{V}_t} w_t(i) w_t(i)^\top$. Let the eigenvalues of WW^\top be denoted

by $\lambda_1, \dots, \lambda_d$, then clearly $\det(I + WW^\top) = \prod_{k=1}^d (1 + \lambda_k)$. On the other hand, $\sum_{k=1}^d \lambda_k = \text{tr}(WW^\top) = \text{tr}(W^\top W) = \sum_{i \in \mathcal{V}_t} \|w_t(i)\|^2$. From the non-negativity of the eigenvalues of WW^\top , we obtain

$$\det(I + WW^\top) = \prod_{k=1}^d (1 + \lambda_k) \geq 1 + \sum_{k=1}^d \lambda_k = 1 + \sum_{i \in \mathcal{V}_t} \|w_t(i)\|^2 = 1 + \sum_{i \in \mathcal{V}_t} \|x_t(i)\|_{V_{t-1}^{-1}}^2. \quad (49)$$

Applying (48) and (49) recursively leads to

$$\det(V_T) \geq \det(\lambda I) \prod_{t=1}^T \left(1 + \sum_{i \in \mathcal{V}_t} \|x_t(i)\|_{V_{t-1}^{-1}}^2 \right) = \lambda^d \prod_{t=1}^T \left(1 + \sum_{i \in \mathcal{V}_t} \|x_t(i)\|_{V_{t-1}^{-1}}^2 \right). \quad (50)$$

We now turn back to bound $\sum_{i \in \mathcal{V}_t} \|x_t(i)\|_{V_{t-1}^{-1}}^2$. Since $\lambda \geq \max_{t \in [T]} \text{card}(\mathcal{V}_t)$ and $\|x_t(i)\| \leq 1$, we have

$$\sum_{i \in \mathcal{V}_t} \|x_t(i)\|_{V_{t-1}^{-1}}^2 = \min \left(1, \sum_{i \in \mathcal{V}_t} \|x_t(i)\|_{V_{t-1}^{-1}}^2 \right) \leq 2 \log \left(1 + \sum_{i \in \mathcal{V}_t} \|x_t(i)\|_{V_{t-1}^{-1}}^2 \right). \quad (51)$$

Combining (50) and (51) we get the first inequality. The second inequality is proved as in Lemma 5. \square

We are now ready to prove Theorem 5. The first part of the proof reuses the trick that we have seen several times before, notably in the proof of Theorem 1. We rearrange the UCB indices and show that the suboptimality gaps can be bounded from above by the sum of the radii of certain confidence intervals. The details will be omitted. Next, we apply Lemma 6 to bound the sum of these upper bounds and conclude.

Proof of Theorem 5. For any $z \in \mathcal{Z}_0$, choosing $\delta' = \delta/(q+1)$ in Lemma 5 gives

$$\mathbb{P} \left(\{ \exists t, \|\theta^z - \hat{\theta}_t^z\|_{V_t^z}^2 > \alpha_t \} \right) \geq \frac{\delta}{q+1}.$$

With a union bound we thus deduce that with probability at least $1 - \delta$, the inequality $\|\theta^z - \hat{\theta}_t^z\|_{V_t^z}^2 \leq \alpha_t$ holds for all t and z . This also implies that for all t and z , we have

$$|\langle \theta^z, x_t(i) \rangle - \langle \hat{\theta}_{t-1}^z, x_t(i) \rangle| \leq \alpha_t \|x_t(i)\|_{(V_{t-1}^z)^{-1}} \quad (52)$$

In the following, we will prove (47) under the condition that (52) holds, which happens with probability at least $1 - \delta$ as we just showed.

Let $t \geq 1$ and $a_t^* \in \arg \max_{a \in \mathcal{A}} \sum_{i \in \mathcal{V}} \langle \theta^{a(i)}, x_t(i) \rangle$ be an optimal action of this round. The algorithm chooses a_t that maximizes $\rho_t^a = \sum_{i \in \mathcal{V}^a} (U_t^a(i) - U_t^0(i))$, and thus

$$\sum_{i \in \mathcal{V}^{a_t}} (U_t^{a_t(i)}(i) - U_t^0(i)) \geq \sum_{i \in \mathcal{V}^{a_t^*}} (U_t^{a_t^*(i)}(i) - U_t^0(i)). \quad (53)$$

On the other hand, inequality (52) can be rewritten as

$$U_t^z(i) - 2\alpha_t \|x_t(i)\|_{(V_{t-1}^z)^{-1}} \leq \langle \theta^z, x_t(i) \rangle \leq U_t^z(i). \quad (54)$$

Combining (53) and (54) and rearranging, we get

$$\sum_{i \in \mathcal{V}} \langle \theta^{a_t^*(i)} - \theta^{a_t(i)}, x_t(i) \rangle = \sum_{i \in \mathcal{V}^{a_t} \cup \mathcal{V}^{a_t^*}} \langle \theta^{a_t^*(i)} - \theta^{a_t(i)}, x_t(i) \rangle \leq \sum_{i \in \mathcal{V}^{a_t} \cup \mathcal{V}^{a_t^*}} 2\alpha_t \|x_t(i)\|_{(V_{t-1}^{a_t(i)})^{-1}}.$$

Summing the above from $t = 1$ to T and applying the Cauchy–Schwarz inequality, we obtain

$$\begin{aligned}
\text{Reg}_T &\leq \sum_{t=1}^T \sum_{i \in \mathcal{V}^{a_t} \cup \mathcal{V}^{a_t^*}} 2\alpha_t \|x_t(i)\|_{(V_{t-1}^{a_t(i)})^{-1}} \\
&\leq 2\alpha_T \sqrt{\left(\sum_{t=1}^T \text{card}(\mathcal{V}^{a_t} \cup \mathcal{V}^{a_t^*}) \right) \left(\sum_{t=1}^T \sum_{i \in \mathcal{V}^{a_t} \cup \mathcal{V}^{a_t^*}} \|x_t(i)\|_{(V_{t-1}^{a_t(i)})^{-1}}^2 \right)} \\
&\leq 2\alpha_T \sqrt{2LT \sum_{t=1}^T \left(\sum_{z \in \mathcal{Z}} \sum_{i \in \mathcal{V}_t^z} \|x_t(i)\|_{(V_{t-1}^z)^{-1}}^2 + \sum_{i \in \mathcal{V}_t^{a_t^*} \setminus \mathcal{V}^{a_t}} \|x_t(i)\|_{(V_{t-1}^0)^{-1}}^2 \right)}.
\end{aligned}$$

Since $\lambda \geq L$, by [Lemma 6](#), for all $z \in \mathcal{Z}$ it holds that

$$\sum_{t=1}^T \sum_{i \in \mathcal{V}_t^z} \|x_t(i)\|_{(V_{t-1}^z)^{-1}}^2 \leq 2d \log \left(1 + \frac{LT}{d\lambda} \right).$$

Similarly, let us define $\tilde{V}_t^0 = \lambda I + \sum_{s=1}^t \sum_{i \in \mathcal{V}_t^{a_t^*} \setminus \mathcal{V}^{a_t}} x_s(i)x_s(i)^\top$. Using $\tilde{V}_t^0 \preceq V_t^0$ and invoking [Lemma 6](#) with $\mathcal{V}_t \leftarrow \mathcal{V}_t^{a_t^*} \setminus \mathcal{V}^{a_t}$ gives

$$\sum_{t=1}^T \sum_{i \in \mathcal{V}_t^{a_t^*} \setminus \mathcal{V}^{a_t}} \|x_t(i)\|_{(V_{t-1}^0)^{-1}}^2 \leq \sum_{t=1}^T \sum_{i \in \mathcal{V}_t^{a_t^*} \setminus \mathcal{V}^{a_t}} \|x_t(i)\|_{(\tilde{V}_{t-1}^0)^{-1}}^2 \leq 2d \log \left(1 + \frac{LT}{d\lambda} \right).$$

Putting the above three inequalities together, we get exactly [\(47\)](#). \square

C.3 Experiments for Targeted Campaign with Criteo Uplifting Dataset

In this section, we provide experiments for the contextual combinatorial uplifting model using the Criteo Uplift dataset. At each round, 100 individuals are sampled from the dataset and the problem consists in choosing 10 of them to treat so that the number of visits over the 100 individuals is maximized. Here, we compare our methods based on C2UPUCB ([Algorithm 7](#)) with approaches based on greedy and ε -greedy (see [[Lattimore and Szepesvári, 2020](#)] for a description) that are more appropriate baselines in this setup. Again, our results are averaged over 100 independent runs and the shaded areas of the plots represent the standard errors.

Algorithms. In all the considered algorithms, the learner learns a logistic model based on the feedback they have received from the individuals they select to treat.⁹ If the baseline model is unknown, the learner also learns a baseline logistic model based on the feedback received from the individuals they do not treat. In the following, we will denote respectively by \hat{g}_t^1 and \hat{g}_t^0 the conditional probability functions of the two fitted models using data up to and including round t . When the baseline model is known, we denote by g^0 its conditional probability function. The methods that we look into in this experiment mainly differ in how they select individuals in each round t , as we detail below. As before, we use the suffix (bl) to indicate modification of the respective algorithm to the case where the baseline payoffs are known.

- **Greedy (bl):** The learner chooses the 10 individuals with the largest estimated uplifts $\hat{g}_{t-1}^1(x) - g^0(x)$.

⁹Logistic models were also considered in the original Criteo dataset paper [[Diemert et al., 2018](#)] to compute uplifts.

- **ε -greedy (bl):** The learner reserves a exploration budget b where b is sampled from $B(10, \varepsilon)$ the binomial distribution with parameters 10 and ε . $10 - b$ individuals are chosen greedily as in the greedy strategy while the remaining b individuals are randomly sampled from those that have not been selected.
- **ε -greedy:** For the greedy selection part in ε -greedy, the estimated uplifts are computed by $\hat{g}_{t-1}^1(x) - \hat{g}_{t-1}^0(x)$.
- **C2UPUCB (bl):** For each individual i we define their uplifting index $\rho_t^a(i) = \hat{g}_{t-1}^1(x) + \beta \|x_t(i)\|_{(V_{t-1}^1)^{-1}} - g^0(x)$ where V_{t-1}^1 is defined following (45) with $\lambda = 1$ and β is a exploration parameter to be tuned. The algorithm consists in choosing the 10 individuals with the largest uplifting indices.
- **C2UPUCB:** It works in the same way as C2UPUCB with baseline knowledge but the uplifting index is now defined by $\rho_t^a(i) = \hat{g}_{t-1}^1(x) + \beta \|x_t(i)\|_{(V_{t-1}^1)^{-1}} - \hat{g}_{t-1}^0(x) - \beta \|x_t(i)\|_{(V_{t-1}^0)^{-1}}$; that is, it is computed as the difference between two upper confidence bounds.
- **C2UCB:** It works in the same way as C2UPUCB but the baseline is ignored, so $\rho_t^a(i) = \hat{g}_{t-1}^1(x) + \beta \|x_t(i)\|_{(V_{t-1}^1)^{-1}}$. The linear variant of this algorithm was introduced in [Qin et al., 2014] for contextual combinatorial bandits.

Note that as the reward comes from a logistic model, we use a logistic regression for the bandit model based on the idea of GLM-UCB [Filippi et al., 2010]. For the results, we select the optimal parameters from $\varepsilon \in \{0.05, 0.1, 0.2, 0.3\}$ and $\beta \in \{0.5, 1, 2, 3\}$. More details on the tuning of the parameters are provided in Appendix D.1, and the chosen optimal parameters are reported in Table 3.

Data Preprocessing. For this experiment, we normalize the data so that the features are contained in the range $[0, 1]$. Moreover, we also subsample the dataset to increase the visit rate to 0.25. In fact, we found the logistic models to perform very poorly on the original dataset (with a f1-score around 0.4) due to its small visit rate at 0.04.

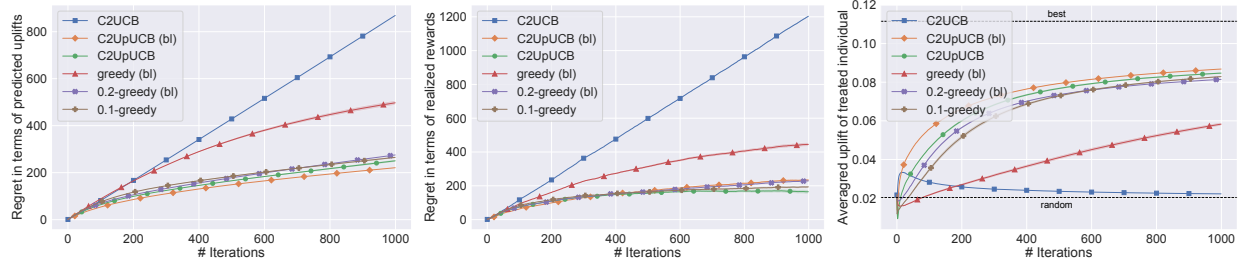
Feedback. We consider two ways to use the fitted logistic models

1. The logistic models are only used to generate the counterfactual outcomes. That is, the feedback of an individual is sampled from the Bernoulli distribution with mean predicted by either of the models only if the learner decides to treat the individual while the individual is not treated in the dataset or vice-versa.
2. The logistic models are used to generate all the outcomes.

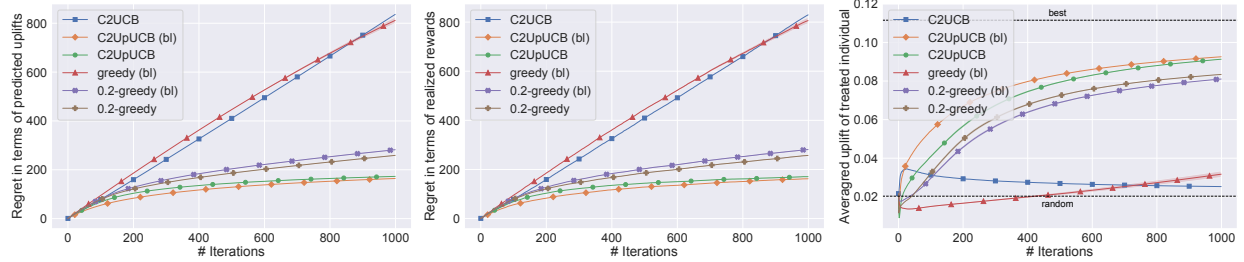
Results. In Figure 5 (Right), we plot the uplift averaged over all the individuals that have been treated up to a time horizon by an algorithm. Higher this value is better it is because it means we are selecting the individuals for which the treatment is the most effective. For comparison, we also indicate the average uplift of random selection and an oracle strategy that always chooses the 10 individuals with the highest uplifts. Both C2UPUCB and ε -greedy beat greedy and random selection, which demonstrates the benefit of considering a bandit model for this problem. ε -greedy turns out to be quite competitive, but C2UPUCB still has a slight advantage no matter whether the baseline is known or unknown.

For further comparison, we also plot the regret in terms of both the predicted uplifts (Left) and the rewards that the learner actually receives (Middle) in Figure 5.¹⁰ For the two regrets, we define the optimal action as

¹⁰We use occasionally the term predicted uplift here to remind that they are computed using logistic models fitted on the dataset and we do not have access to any type of *ground-truth* uplift.



(a) Logistic models only used to generate counterfactual outcomes.



(b) Logistic models used to generate all the outcomes.

Figure 5: Results on the targeted campaign experiment. For the top row the logistic models are only used to generate counterfactual outcomes, and in the bottom row all the outcomes are generated using the predicted conditional probabilities of the logistic models. We plot regret computed using the predicted uplifts and the realized rewards, and the predicted uplift averaged over all the treated individuals.

the one that selects the 10 individuals with the highest predicted uplifts. Then, if the logistic models correctly predict the expected rewards, the regret in terms of the predicted uplifts should be the expectation of the regret in terms of the realized rewards, and thus we should get similar curves after averaging over 100 runs. This is indeed the case in Figure 5b. Nonetheless, when the outcomes also come from true data, selecting the samples with higher predicted uplifts does not necessarily result in higher rewards, which explains the discrepancy between Figure 5a (Left) and Figure 5a (Middle). C2UpUCB and ϵ -greedy perform quite similarly here when we look at the realized rewards. Since the ground truth outcomes cannot be perfectly predict by logistic models, we also observe linear regret of all the algorithms in Figure 5a (Left), while in Figure 5b (Left), both C2UpUCB-M and C2UpUCB achieves sublinear regret and performs much better than ϵ -greedy.

In all the cases, Greedy and C2UCB perform poorly. This points out the importance of exploration (which is taken into account by bandit algorithms) and uplift modeling in this problem. Finally, to explain the surprising fact that an algorithm might perform better without knowledge of the baseline, we look into the data and find out that samples with higher baselines often have higher uplifts. Therefore, when the learner is provided with the baselines of the samples, they may avoid choosing these samples at the beginning, which results a slower exploration and a worse performance.

D Additional Experimental Details

In this appendix, we present missing experimental details from Section 7. The code implementation of the proposed methods is attached in supplementary material. All our experiments are run on amazon m4.xlarge EC2 instances. The Criteo Uplift Prediction Dataset that we use for our experiments has the ‘CC0: Public Domain’ license.

Table 3: Parameter values used for the reported experimental results. TS stands for Thompson sampling. Please refer to the text for a explanation on what is the parameter that we tune for each of the algorithms.

Non-contextual	UCB	UPUCB (bl)	UPUCB	UPUCB-nAff (bl)	UPUCB-nAff	TS
Gaussian	2	5	3	5	3	$8 \cdot 10^{-2}$
Bernoulli	$7 \cdot 10^{-7}$	$8 \cdot 10^{-5}$	$8 \cdot 10^{-5}$	$5 \cdot 10^{-7}$	$2 \cdot 10^{-6}$	$2 \cdot 10^{-7}$
Contextual	C2UCB	C2UPUCB (b)	C2UPUCB	ε -greedy (b)	ε -greedy	
Logistic for counterfactual	2	2	2	0.2	0.1	
Logistic for all	2	2	2	0.2	0.2	

D.1 Choice of Algorithm Parameters

As mentioned in the main text, we report the results for the parameters that yield the best average performance. Precisely, we conduct 100 independent runs for each experimental setup, and in each run, we fix the random seed and test all the candidate algorithms with their parameters chosen from a predefined set. The performance of a specific parameter choice is then evaluated by the sum of the mean and one standard deviation of the regrets of the final run. In fact, when the algorithm puts too much weight on exploitation, we may get lower mean regret while the standard deviation gets extremely large because the algorithm fails more frequently.

Below we explain in detail what are the parameters that we tune in the non-contextual experiments. The optimal parameters that we have found are summarized in [Table 3](#).

UCB and UPUCBs. For all UCB-type methods, we tune the exploration parameter $\beta = 1/\delta'$, which plays a crucial role in determining the widths of the confidence intervals– see (3) and (6). Since c_t^a represents the width of the confidence interval of a single variable’s payoff in UpUCB algorithms, we set the widths of the confidence intervals of UCB to $mc_t^a = m\sqrt{2\beta/N_t^a}$, i.e., the UCB indices in UCB are $U_t^a = \hat{r}_t^a + mc_t^a$, where \hat{r}_t^a is the empirical mean of action a ’s reward up to time t . In this way, we can choose exploration parameters of roughly the same order for different algorithms. In the Gaussian uplifting bandit experiment ([Section 7](#)), we pick $\beta \in [10] = \{1, \dots, 10\}$. As for the Bernoulli uplifting bandit experiment with Criteo dataset ([Section 7](#)), we select $\beta \in \{k \times 10^{-\gamma} : k \in [10], \gamma \in \{5, 6, 7\}\}$. The optimal performances are achieved with much smaller exploration parameters here because the noises in the variables are independent.

Thompson Sampling. We use Thompson sampling with Gaussian prior and Gaussian noise. The mean and the variance of the prior are set to the mean and the variance of the actions’ rewards. Therefore, Thompson sampling is actually provided with more information than UCB. Similar to before, the variance of the noise is set to $m^2\sigma^2$ since the noise in the total reward scales with m . The parameter σ^2 plays a similar role as β in UCB. In the Gaussian and Bernoulli uplifting experiments we respectively choose $\sigma^2 \in \{k \times 10^{-\gamma} : k \in [10], \gamma \in \{0, 1, 2\}\}$ and $\sigma^2 \in \{k \times 10^{-\gamma} : k \in [10], \gamma \in \{5, 6, 7\}\}$.

D.2 More Details on Bernoulli Uplifting Bandit

For sake of completeness, we report below the visit rate of the treated and the untreated individuals of the clusters (rounded to three decimal places) that we obtained for our experiments

$$\begin{aligned}\mu_{\text{treated}} &= [.001 .037 .003 .001 .003 .377 .237 .309 .071 .287 \\ &\quad .531 .044 .007 .086 .002 .019 .028 .007 .265 .013] \\ \mu_{\text{untreated}} &= [.001 .023 .002 .002 .004 .289 .206 .229 .073 .289 \\ &\quad .464 .035 .004 .052 .001 .011 .022 .004 .165 .000].\end{aligned}$$

The sizes of the clusters, which also correspond to $(L^a)_{a \in \mathcal{A}}$, are

$$\begin{aligned}&[10600 \ 2764 \ 7222 \ 11128 \ 6385 \ 1630 \ 2806 \ 1089 \ 3018 \ 4594 \\ &\quad 594 \ 7020 \ 12654 \ 2186 \ 9609 \ 5101 \ 3714 \ 4569 \ 1158 \ 2159].\end{aligned}$$

The uplifts of the actions (rounded to one decimal place) are

$$\begin{aligned}r_{\text{up}} &= [0.6 \ 39.3 \ 6.6 \ -5.0 \ -5.7 \ 143.5 \ 86.0 \ 87.1 \ -4.5 \ -6.9 \\ &\quad 39.8 \ 66.3 \ 34.7 \ 75.4 \ 4.0 \ 40.3 \ 21.1 \ 12.3 \ 115.7 \ 28.3].\end{aligned}$$