
Supplementary Material: Regret in Online Recommendation Systems

Kaito Ariu
KTH
Stockholm, Sweden
ariu@kth.se

Narae Ryu
KAIST
Daejeon, South Korea
nrryu@kaist.ac.kr

Se-Young Yun
KAIST
Daejeon, South Korea
yunseyoung@kaist.ac.kr

Alexandre Proutière
KTH
Stockholm, Sweden
alepro@kth.se

Abstract

This paper proposes a theoretical analysis of recommendation systems in an online setting, where items are sequentially recommended to users over time. In each round, a user, randomly picked from a population of m users, requests a recommendation. The decision-maker observes the user and selects an item from a catalogue of n items. Importantly, an item cannot be recommended twice to the same user. The probabilities that a user likes each item are unknown. The performance of the recommendation algorithm is captured through its regret, considering as a reference an Oracle algorithm aware of these probabilities. We investigate various structural assumptions on these probabilities: we derive for each structure regret lower bounds, and devise algorithms achieving these limits. Interestingly, our analysis reveals the relative weights of the different components of regret: the component due to the constraint of not presenting the same item twice to the same user, that due to learning the chances users like items, and finally that arising when learning the underlying structure.

Contents

1	Introduction	3
2	Related Work	4
3	Models and Preliminaries	4
3.1	Problem structures and regrets	5
3.2	Preliminaries – User arrival process	5
4	Regret Lower Bounds	6
4.1	Clustered items and statistically identical users	6
4.2	Unclassified items and statistically identical users	6
4.3	Clustered items and clustered users	7

5 Algorithms	7
5.1 Clustered items and statistically identical users	8
5.2 Unclustered items and statistically identical users	8
5.3 Clustered items and clustered users	9
6 Conclusion	10
7 Table of Notations	13
A Algorithms and experiments	15
A.1 Clustered items and statistically identical users	15
A.2 Unclustered items and statistically identical users	18
A.3 Clustered items and users	20
B Preliminaries: Properties of the user arrival process	25
C Justifying the regret definitions	27
D Fundamental limits for Model A: Proof of Theorem 1	29
D.1 Proof of Lemma 7	30
E Fundamental limits for Model B: Proof of Theorem 2	32
F Fundamental limits for Model C: Proof of Theorem 3	34
F.1 Examples	34
F.2 Proof	35
G Performance guarantees of ECT: Proof of Theorem 4	41
H Performance guarantees of ET: Proof of Theorem 5	45
I Performance guarantees of EC-UCS: Proof of Theorem 6	48
J Performance guarantees of ECB and Item Clustering: Proof of Theorems 7 and 9	57

1 Introduction

Recommendation systems [28] have over the last two decades triggered important research efforts (see, e.g., [9, 10, 21, 3] for recent works and references therein), mainly focused towards the design and analysis of algorithms with improved efficiency. These algorithms are, to some extent, all based on the principle of *collaborative filtering*: similar items should yield similar user responses, and similar users have similar probabilities of liking or disliking a given item. In turn, efficient recommendation algorithms need to learn and exploit the underlying structure tying the responses of the users to the various items together.

Most recommendation systems operate in an online setting, where items are sequentially recommended to users over time. We investigate recommendation algorithms in this setting. More precisely, we consider a system of n items and m users, where $m \geq n$ (as this is typically the case in practice). In each round, the algorithm needs to recommend an item to a *known* user, picked randomly among the m users. The response of the user is noisy: the user likes the recommended item with an a priori *unknown* probability depending on the (item, user) pair. In practice, it does not make sense to recommend an item twice to the same user (why should we recommend an item to a user who already considered or even bought the item?). We restrict our attention to algorithms that do not recommend an item twice to the same user, a constraint referred to as the *no-repetition* constraint. The objective is to devise algorithms maximizing the expected number of successful recommendations over a time horizon of T rounds.

We investigate different system structures. Specifically, we first consider the case of clustered items and statistically identical users – the probability that a user likes an item depends on the item cluster only. We then study the case of unclustered items and statistically identical users – the probability that a user likes an item depends on the item only. The third investigated structure exhibits clustered items and clustered users – the probability that a user likes an item depends on the item and user clusters only. In all cases, the structure (e.g., the clusters) is initially unknown and has to be learnt to some extent. This paper aims at answering the question: How can the structure be optimally learnt and exploited?

To this aim, we study the regret of online recommendation algorithms, defined as the difference between their expected number of successful recommendations to that obtained under an Oracle algorithm aware of the structure and of the success rates of each (item, user) pair. We are interested in regimes where n , m , and T grow large simultaneously, and $T = o(mn)$ (see §3 for details). For the aforementioned structures, we first derive non-asymptotic and problem-specific regret lower bounds satisfied by any algorithm.

- (i) For clustered items and statistically identical users, as T (and hence m) grows large, the minimal regret scales as $K \max\{\frac{\log(m)}{\log(\frac{m \log(m)}{T})}, \frac{T}{\Delta m}\}$, where K is the number of item clusters, and Δ denotes the minimum difference between the success rates of items from the optimal and sub-optimal clusters.
- (ii) For unclustered items and statistically identical users, the minimal *satisficing* regret¹ scales as $\max\{\frac{\log(m)}{\log(\frac{m \log(m)}{T})}, \frac{T}{m\varepsilon}\}$, where ε denotes the threshold defining the satisficing regret (recommending an item in the top ε percents of the items is assumed to generate no regret).
- (iii) For clustered items and users, the minimal regret scales as $\frac{m}{\Delta}$ or as $\frac{m}{\Delta} \log(T/m)$, depending on the values of the success rate probabilities.

We also devise algorithms that provably achieve these limits (up to logarithmic factors), and whose regret exhibits the right scaling in Δ or ε . We illustrate the performance of our algorithms through experiments presented in the appendix.

Our analysis reveals the relative weights of the different components of regret. For example, we can explicitly identify the regret induced by the no-repetition constraint (this constraint imposes us to select unrecommended items and induces an important learning price). We may also characterize the regret generated by the fact that the item or user clusters are initially unknown. Specifically, fully exploiting the item clusters induces a regret scaling as $K \frac{T}{\Delta m}$. Whereas exploiting user clusters has a much higher regret cost scaling as least as $\frac{m}{\Delta}$.

In our setting, deriving regret lower bounds and devising optimal algorithms cannot be tackled using existing techniques from the abundant bandit literature. This is mainly due to the no-repetition constraint, to the hidden structure, and to the specificities introduced by the random arrivals of

¹For this unstructured scenario, we will justify why considering the satisficing regret is needed.

users. Getting tight lower bounds is particularly challenging because of the non-asymptotic nature of the problem (items cannot be recommended infinitely often, and new items have to be assessed continuously). To derive these bounds, we introduce novel techniques that could be useful in other online optimization problems. The design and analysis of efficient algorithms also present many challenges. Indeed, such algorithms must include both clustering and bandit techniques, that should be jointly tuned.

Due to space constraints, we present the pseudo-codes of our algorithms, all proofs, numerical experiments, as well as some insightful discussions in the appendix.

2 Related Work

The design of recommendation systems has been framed into structured bandit problems in the past. Most of the work there consider a linear reward structure (in the spirit of the matrix factorization approach), see e.g. [9], [10], [22], [20], [21], [11]. These papers ignore the no-repetition constraint (a usual assumption there is that when a user arrives, a set of fresh items can be recommended). In [24], the authors try to include this constraint but do not present any analytical result. Furthermore, notice that the structures we impose in our models are different than that considered in the low-rank matrix factorization approach.

Our work also relates to the literature on clustered bandits. Again the no-repetition constraint is not modeled. In addition, most often, only the user clusters [6], [23] or only the item clusters are considered [18], [14]. Low-rank bandits extend clustered bandits by modeling the (item, user) success rates as a low-rank matrix, see [15], [25], still without accounting for the no-repetition constraint, and without a complete analysis (no precise regret lower bounds).

One may think of other types of bandits to model recommendation systems. However, none of them captures the essential features of our problem. For example, if we think of contextual bandits (see, e.g., [12] and references therein), where the context would be the user, it is hard to model the fact that when the same context appears several times, the set of available arms (here items) changes depending on the previous arms selected for this context. Budgeted and sleeping bandits [7], [17] model scenarios where the set of available arms changes over time, but in our problem, this set changes in a very specific way not covered by these papers. In addition, studies on budgeted and sleeping bandits do not account for any structure.

The closest related work can be found in [4] and [13]. There, the authors explicitly model the no-repetition constraint but consider user clusters only, and do not provide regret lower bounds. [3] extends the analysis to account for item clusters as well but studies a model where users in the same cluster deterministically give the same answers to items in the same cluster.

3 Models and Preliminaries

We consider a system consisting of a set $\mathcal{I} = [n] := \{1, \dots, n\}$ of items and a set $\mathcal{U} = [m]$ of users. In each round, a user, chosen uniformly at random from \mathcal{U} , needs a recommendation. The decision-maker observes the user id and selects an item to be presented to the user. Importantly an item cannot be recommended twice to a user. The user immediately rates the recommended item +1 if she likes it or 0 otherwise. This rating is observed by the decision-maker, which helps subsequent item selections.

Formally, in round t , the user $u_t \sim \text{unif}(\mathcal{U})$ requires a recommendation. If item i is recommended, the user $u_t = u$ likes the item with probability ρ_{iu} . We introduce the binary r.v. X_{iu} to indicate whether the user likes the item, $X_{iu} \sim \text{Ber}(\rho_{iu})$. Let π denote a sequential item selection strategy or algorithm. Under π , the item i_t^π is presented to the t -th user. The choice i_t^π depends on the past observations and on the identity of the t -th user, namely, i_t^π is \mathcal{F}_{t-1}^π -measurable, with $\mathcal{F}_{t-1}^\pi = \sigma(u_t, (u_s, i_s^\pi, X_{i_s^\pi u_s}), s \leq t-1)$ ($\sigma(Z)$ denotes the σ -algebra generated by the r.v. Z). Denote by Π the set of such possible algorithms. The reward of an algorithm π is defined as the expected number of positive ratings received over T rounds: $\mathbb{E}[\sum_{t=1}^T \rho_{i_t^\pi u_t}]$. We aim at devising an algorithm with maximum reward.

We are mostly interested in scenarios where (m, n, T) grow large under the constraints (i) $m \geq n$ (this is typically the case in recommendation systems), (ii) $T = o(mn)$, and (iii) $\log(m) = o(n)$.

Condition (ii) complies with the no-repetition constraint and allows some freedom in the item selection process. (iii) is w.l.o.g. as explained in [4], and is just imposed to simplify our definitions of regret (refer to Appendix C for a detailed discussion).

3.1 Problem structures and regrets

We investigate three types of systems depending on the structural assumptions made on the success rates $\rho = (\rho_{iu})_{i \in \mathcal{I}, u \in \mathcal{U}}$.

Model A. Clustered items and statistically identical users. In this case, ρ_{iu} depends on the item i only. Items are classified into K clusters $\mathcal{I}_1, \dots, \mathcal{I}_K$. When the algorithm recommends an item i for the first time, i is assigned to cluster \mathcal{I}_k with probability α_k , independently of the cluster assignments of the other items. When $i \in \mathcal{I}_k$, then $\rho_i = p_k$. We assume that both $\alpha = (\alpha_k)_{k \in [K]}$ and $p = (p_k)_{k \in [K]}$ do not depend on (n, m, T) , but are initially unknown. W.l.o.g. assume that $p_1 > p_2 \geq p_3 \geq \dots \geq p_K$. To define the regret of an algorithm $\pi \in \Pi$, we compare its reward to that of an Oracle algorithm aware of the item clusters and of the parameters p . The latter would *mostly* recommend items from cluster \mathcal{I}_1 . Due to the randomness in the user arrivals and the cluster sizes, recommending items not in \mathcal{I}_1 may be necessary. However, we define regret as if recommending items from \mathcal{I}_1 was always possible. Using our assumptions $T = o(mn)$ and $\log(m) = o(n)$, we can show that the difference between our regret and the true regret (accounting for the possible need to recommend items outside \mathcal{I}_1) is always negligible. Refer to Appendix C for a formal justification. In summary, the regret of $\pi \in \Pi$ is defined as: $R^\pi(T) = Tp_1 - \sum_{t=1}^T \mathbb{E} \left[\sum_{k=1}^K \mathbb{1}_{\{i_t^\pi \in \mathcal{I}_k\}} p_k \right]$.

Model B. Unclustered items and statistically identical users. Again here, ρ_{iu} depends on the item i only. when a new item i is recommended for the first time, its success rate ρ_i is drawn according to some distribution ζ over $[0, 1]$, independently of the success rates of the other items. ζ is arbitrary and initially unknown, but for simplicity assumed to be absolutely continuous w.r.t. Lebesgue measure. To represent ζ , we also use its inverse distribution function: for any $x \in [0, 1]$, $\mu_x := \inf\{\gamma \in [0, 1] : \mathbb{P}[\rho_i \leq \gamma] \geq x\}$. We say that an item i is within the ε -best items if $\rho_i \geq \mu_{1-\varepsilon}$. We adopt the following notion of regret: for a given $\varepsilon > 0$, $R_\varepsilon^\pi(T) = \sum_{t=1}^T \mathbb{E} \left[\max\{0, \mu_{1-\varepsilon} - \rho_{i_t^\pi}\} \right]$. Hence, we assume that recommending items within the ε -best items does not generate any regret. We also assume, as in Model A, that an Oracle policy can always recommend such items (refer to Appendix C). This notion of *satisficing* regret [29] has been used in the bandit literature to study problems with a very large number of arms (we have a large number of items). For such problems, identifying the best arm is very unlikely, and relaxing the regret definition is a necessity. Satisficing regret is all the more relevant in our problem that even if one would be able to identify the best item, we cannot recommend it (play it) more than m times (due to the no-repetition constraint), and we are actually forced to recommend sub-optimal items. A similar notion of regret is used in [4] to study recommendation systems in a setting similar to our Model B.

Model C. Clustered items and clustered users. We consider the case where both items and users are clustered. Specifically, users are classified into L clusters $\mathcal{U}_1, \dots, \mathcal{U}_L$, and when a user arrives to the system the first time, she is assigned to cluster \mathcal{U}_ℓ with probability β_ℓ , independently of the other users. There are K item clusters $\mathcal{I}_1, \dots, \mathcal{I}_K$. When the algorithm recommends an item i for the first time, it is assigned to cluster \mathcal{I}_k with probability α_k as in Model A. Now $\rho_{iu} = p_{k\ell}$ when $i \in \mathcal{I}_k$ and $u \in \mathcal{U}_\ell$. Again, we assume that $p = (p_{k\ell})_{k,\ell}$, $\alpha = (\alpha_k)_{k \in [K]}$ and $\beta = (\beta_\ell)_{\ell \in [L]}$ do not depend on (n, m, T) . For any ℓ , let $k_\ell^* = \arg \max_k p_{k\ell}$ be the best item cluster for users in \mathcal{U}_ℓ . We assume that k_ℓ^* is unique. In this scenario, we assume that an Oracle algorithm, aware of the item and user clusters and of the parameters p , would only recommend items from cluster k_ℓ^* to a user in \mathcal{U}_ℓ (refer to Appendix C). The regret of an algorithm $\pi \in \Pi$ is hence defined as: $R^\pi(T) = T \sum_\ell \beta_\ell p_{k_\ell^* \ell} - \sum_{t=1}^T \mathbb{E} \left[\sum_{k,\ell} \mathbb{1}_{\{u_t \in \mathcal{U}_\ell, i_t^\pi \in \mathcal{I}_k\}} p_{k\ell} \right]$.

3.2 Preliminaries – User arrival process

The user arrival process is out of the decision maker's control and strongly impacts the performance of the recommendation algorithms. To analyze the regret of our algorithms, we will leverage the following results. Let $N_u(T)$ denote the number of requests of user u up to round T . From the

literature on "Balls and Bins process", see e.g. [27], we know that if $\bar{n} := \mathbb{E}[\max_{u \in \mathcal{U}} N_u(T)]$, then

$$\bar{n} = \begin{cases} \frac{\log(m)}{\log(\frac{m \log(m)}{T})} (1 + o(1)) & \text{if } T = o(m \log(m)), \\ \log(m)(d_c + o(1)) & \text{if } T = cm \log(m), \\ \frac{T}{m} (1 + o(1)) & \text{if } T = \omega(m \log(m)), \end{cases}$$

where d_c is a constant depending on c only. We also establish the following concentration result controlling the tail of the distribution of $N_u(T)$ (refer to Appendix B):

Lemma 1. Define $\bar{N} = \frac{4 \log(m)}{\log(\frac{m \log(m)}{T} + e)} + \frac{e^2 T}{m}$. Then, $\forall u \in \mathcal{U}$, $\mathbb{E}[\max\{0, N_u(T) - \bar{N}\}] \leq \frac{1}{(e-1)m}$.

The quantities \bar{n} and \bar{N} play an important role in our regret analysis.

4 Regret Lower Bounds

In this section, we derive regret lower bounds for the three envisioned structures. Interestingly, we are able to quantify the minimal regret induced by the specific features of the problem: (i) the no-repetition constraint, (ii) the unknown success probabilities, (iii) the unknown item clusters, (iv) the unknown user clusters. The proofs of the lower bounds are presented in Appendices D-E-F.

4.1 Clustered items and statistically identical users

We denote by $\Delta_k = p_1 - p_k$ the gap between the success rates of items from the best cluster and of items from cluster \mathcal{I}_k , and introduce the function: $\phi(k, m, p) = \frac{1 - e^{-m\gamma(p_1, p_k)}}{8(1 - e^{-\gamma(p_1, p_k)})}$, where $\gamma(p, q) = \text{kl}(p, q) + \text{kl}(q, p)$ and $\text{kl}(p, q) = p \log \frac{p}{q} + (1 - p) \log \frac{1-p}{1-q}$. Using the fact that $\text{kl}(p, q) \leq (p - q)^2 / (q(1 - q))$, we can easily show that as m grows large, $\phi(k, m, p)$ scales as $\eta / (16\Delta_k^2)$ when Δ_k is small, where $\eta := \min_k p_k(1 - p_k)$.

We derive problem-specific regret lower bounds, and as in the classical stochastic bandit literature, we introduce the notion of *uniformly good* algorithm. π is uniformly good if its expected regret $R^\pi(T)$ is $O(\max\{\sqrt{T}, \frac{\log(m)}{\log(\frac{m \log(m)}{T} + e)}\})$ for all possible system parameters (p, α) when T, m, n grow large with $T = o(nm)$ and $m \geq n$. As shown in the next section, uniformly good algorithms exist.

Theorem 1. Let $\pi \in \Pi$ be an arbitrary algorithm. The regret of π satisfies: for all $T \geq 1$ such that $m \geq c/\Delta_2^2$ (for some constant c large enough), $R^\pi(T) \geq \max\{R_{\text{nr}}(T), R_{\text{ic}}(T)\}$, where $R_{\text{nr}}(T)$ and $R_{\text{ic}}(T)$, the regrets due to the no-repetition constraint and to the unknown item clusters, respectively, are defined by $R_{\text{nr}}(T) := \bar{n} \sum_{k \neq 1} \alpha_k \Delta_k$ and $R_{\text{ic}}(T) := \frac{T}{m} \sum_{k \neq 1} \alpha_k \phi(k, m, p) \Delta_k$.

Assume that π is uniformly good, then we have²: $R^\pi(T) \gtrsim R_{\text{sp}}(T) := \log(T) \sum_{k \neq 1} \frac{\Delta_k}{2\text{kl}(p_k, p_1)}$, where $R_{\text{sp}}(T)$ refers to the regret due to the unknown success probabilities.

From the above theorem, analyzing the way $R_{\text{nr}}(T)$, $R_{\text{ic}}(T)$, and $R_{\text{sp}}(T)$ scale, we can deduce that:

- (i) When $T = o(m \log(m))$, the regret arises mainly due to either the no-repetition constraint or the need to learn the success probabilities, and it scales at least as $\max\{\frac{\log(m)}{\log(\frac{m \log(m)}{T})}, \log(T)\}$.
- (ii) When $T = cm \log(m)$, the three components of the regret lower bound scales in the same way, and the regret scales at least as $\log(T)$.
- (iii) When $T = \omega(m \log(m))$, the regret arises mainly due to either the no-repetition constraint or the need to learn the item clusters, and it scales at least as $\frac{T}{m}$.

4.2 Unclustered items and statistically identical users

In this scenario, the regret is induced by the no-repetition constraint, and by the fact the success rate of an item when it is first selected and the distribution ζ are unknown. These two sources of regret lead to the terms $R_{\text{nr}}(T)$ and $R_i(T)$, respectively, in our regret lower bound.

²We write $a \gtrsim b$ if $\liminf_{T \rightarrow \infty} a/b \geq 1$.

Theorem 2. Assume that the density of ζ satisfies, for some $C > 0$, $\zeta(\mu) \leq C$ for all $\mu \in [0, 1]$. Let $\pi \in \Pi$ be an arbitrary algorithm. Then its satisficing regret satisfies: for all $T \geq 1$ such that $m \geq c/\varepsilon^2$ (for some constant $c \geq 1$ large enough), $R_\varepsilon^\pi(T) \geq \max\{R_{\text{nr}}(T), R_i(T)\}$, where $R_{\text{nr}}(T) := \bar{n} \int_0^{\mu_{1-\varepsilon}} (\mu_{1-\varepsilon} - \mu)\zeta(\mu)d\mu$ and $R_i(T) := \frac{T}{m} \frac{\frac{(1-\varepsilon)^2}{2C} (1-\frac{\varepsilon C}{1-\varepsilon})^2}{\min\{1, (1+C)\varepsilon\} + 1/m}$.

4.3 Clustered items and clustered users

To state regret lower bounds in this scenario, we introduce the following notations. For any $\ell \in [L]$, let $\Delta_{k\ell} = p_{k_\ell^* \ell} - p_{k\ell}$ be the gap between the success rates of items from the best cluster $\mathcal{I}_{k_\ell^*}$ and of items from cluster \mathcal{I}_k . We also denote by $\mathcal{R}_\ell = \{r \in [L] : k_\ell^* \neq k_r^*\}$. We further introduce the functions:

$$\phi(k, \ell, m, p) = \frac{1 - e^{-m\gamma(p_{k_\ell^* \ell}, p_{k\ell})}}{8 \left(1 - e^{-\gamma(p_{k_\ell^* \ell}, p_{k\ell})}\right)} \quad \text{and} \quad \psi(\ell, k, T, m, p) = \frac{1 - e^{-\frac{T}{m}\gamma(p_{k_\ell^* \ell}, p_{k\ell})}}{8 \left(1 - e^{-\gamma(p_{k_\ell^* \ell}, p_{k\ell})}\right)}.$$

Compared to the case of clustered items and statistically identical users, this scenario requires the algorithm to actually learn the user clusters. To discuss how this induces additional regret, assume that the success probabilities p are known. Define $\mathcal{L}_\perp = \{(\ell, \ell') \in [L]^2 : p_{k_\ell^* \ell} \neq p_{k_{\ell'}^* \ell'}\}$, the set of pairs of user clusters whose best item clusters differ. If $\mathcal{L}_\perp \neq \emptyset$, then there isn't a single optimal item cluster for all users, and when a user u first arrives, we need to learn its cluster. If p is known, this classification generates at least a constant regret (per user) – corresponding to the term $R_{\text{uc}}(T)$ in the theorem below. For specific values of p , we show that this classification can even generate a regret scaling as $\log(T/m)$ (per user). This happens when $\mathcal{L}^\perp(\ell) = \{\ell' \neq \ell : k_\ell^* \neq k_{\ell'}^*, p_{k_\ell^* \ell} = p_{k_{\ell'}^* \ell'}\}$ is not empty – refer to Appendix F for examples. In this case, we cannot distinguish users from \mathcal{U}_ℓ and $\mathcal{U}_{\ell'}$ by just presenting items from $\mathcal{I}_{k_\ell^*}$ (the greedy choice for users in \mathcal{U}_ℓ). The corresponding regret term in the theorem below is $R'_{\text{uc}}(T)$. To formalize this last regret component, we define uniformly good algorithms as follows. An algorithm is uniformly good if for any user u , $R_u^\pi(N) = o(N^\alpha)$ as N grows large for all $\alpha > 0$, where $R_u^\pi(N)$ denotes the accumulated expected regret under π for user u when the latter has arrived N times.

Theorem 3. Let $\pi \in \Pi$ be an arbitrary algorithm. Then its regret satisfies: for all $T \geq 2m$ such that $m \geq c/\min_{k, \ell} \Delta_{k\ell}^2$ (for some constant c large enough), $R^\pi(T) \geq \max\{R_{\text{nr}}(T), R_{\text{ic}}(T), R_{\text{uc}}(T)\}$, where $R_{\text{nr}}(T)$, $R_{\text{ic}}(T)$, and $R_{\text{uc}}(T)$ are regrets due to the no-repetition constraint, to the unknown item clusters, and to the unknown user clusters respectively, defined by:

$$\begin{cases} R_{\text{nr}}(T) := \bar{n} \sum_\ell \beta_\ell \sum_{k \neq k_\ell^*} \alpha_k \Delta_{k\ell}, \\ R_{\text{ic}}(T) := \frac{T}{m} \sum_\ell \beta_\ell \sum_{k \neq k_\ell^*} \alpha_k \phi(k, \ell, m, p) \Delta_{k\ell}, \\ R_{\text{uc}}(T) := m \sum_{\ell \in [L]} \beta_\ell \frac{\sum_{k \in \mathcal{R}_\ell} \Delta_{k\ell} \psi(\ell, k, T, m, p)}{K}. \end{cases}$$

In addition, when $T = \omega(m)$, if π is uniformly good, $R^\pi(T) \gtrsim R'_{\text{uc}}(T) := c(\beta, p)m \log(T/m)$ where $c(\beta, p) = \inf_{n \in \mathcal{F}} \sum_\ell \beta_\ell \sum_{k \neq k_\ell^*} \Delta_{k\ell} n_{k\ell}$ with $\mathcal{F} = \{n \geq 0 : \forall \ell, \forall \ell' \in \mathcal{L}^\perp(\ell), \sum_{k \neq k_\ell^*} \text{kl}(p_{k\ell}, p_{k\ell'}) n_{k\ell} \geq 1\}$.

Note that we do not include in the lower bound the term $R_{\text{sp}}(T)$ corresponding to the regret induced by the lack of knowledge of the success probabilities. Indeed, it would scale as $\log(T)$, and this regret would be negligible compared to $R_{\text{uc}}(T)$ (remember that $T = o(m^2)$), should $\mathcal{L}_\perp \neq \emptyset$. Under the latter condition, the main component of regret is for any time horizon is due to the unknown user clusters. When $\mathcal{L}_\perp \neq \emptyset$, the regret scales at least as m if for all ℓ , $\mathcal{L}^\perp(\ell) = \emptyset$, and $m \log(T/m)$ otherwise.

5 Algorithms

This section presents algorithms for our three structures and an analysis of their regret. The detailed pseudo-codes of our algorithms and numerical experiments are presented in Appendix A. The proofs of the regret upper bounds are postponed to Appendices G-H-I.

5.1 Clustered items and statistically identical users

To achieve a regret scaling as in our lower bounds, the structure needs to be exploited. Even without accounting for the no-repetition constraint, the KL-UCB algorithm would, for example, yield a regret scaling as $\frac{n}{\Delta_2} \log(T)$. Now we could first sample T/m items and run KL-UCB on this restricted set of items – this would yield a regret scaling as $\frac{T}{m\Delta_2} \log(T)$, without accounting for the no-repetition constraint. Our proposed algorithm, Explore-Cluster-and-Test (ECT), achieves a better regret scaling and complies with the no-repetition constraint. Refer to Appendix A for numerical experiments illustrating the superiority of ECT.

The Explore-Cluster-and-Test algorithm. ECT proceeds in the following phases:

(a) **Exploration phase.** This first phase consists in gathering samples for a subset \mathcal{S} of randomly selected items so that the success probabilities and the clusters of these items are learnt accurately. Specifically, we pick $|\mathcal{S}| = \lfloor \log(T)^2 \rfloor$ items, and for each of these items, gather roughly $\log(T)$ samples.

(b) **Clustering phase.** We leverage the information gathered in the exploration phase to derive an estimate $\hat{\rho}_i$ of the success probability ρ_i for item $i \in \mathcal{S}$. These estimates are used to cluster items, using an appropriate version of the K-means algorithm. In turn, we extract from this phase, accurate estimates \hat{p}_1 and \hat{p}_2 of the success rates of items in the two best item clusters, and a set $\mathcal{V} \subset \mathcal{S}$ of items believed to be in the best cluster: $\mathcal{V} := \{i \in \mathcal{S} : \hat{\rho}_i > (\hat{p}_1 + \hat{p}_2)/2\}$.

(c) **Test phase.** The test phase corresponds to an exploitation phase. Whenever this is possible (the no-repetition constraint is not violated), items from \mathcal{V} are recommended. When an item outside \mathcal{V} has to be selected due to the no-repetition constraint, we randomly sample and recommend an item outside \mathcal{V} . This item is appended to \mathcal{V} . To ensure that any item i in the (evolving) set \mathcal{V} is from the best cluster with high confidence, we keep updating its empirical success rate $\hat{\rho}_i$, and periodically test whether $\hat{\rho}_i$ is close enough from \hat{p}_1 . If this is not the case, i is removed from \mathcal{V} .

In all phases, ECT is designed to comply with the no-repetition constraint: for example, in the exploration phase, when the user arrives, if we cannot recommend an item from \mathcal{S} due to the constraint, we randomly select an item not violating the constraint. In the analysis of ECT regret, we upper bound the regret generated in rounds where a random item selection is imposed. Observe that ECT does not depend on any parameter (except for the choice of the number of items initially explored in the first phase).

Theorem 4. *We have: $R^{\text{ECT}}(T) = \mathcal{O}\left(\frac{2\bar{N}}{\alpha_1} \sum_{k=2}^K \frac{\alpha_k(p_1 - p_k)}{(p_1 - p_2)^2} + (\log T)^3\right)$.*

The regret lower bound of Theorem 1 states that for any algorithm π , $R^\pi(T) = \Omega(\bar{N})$, and if π is uniformly good $R^\pi(T) = \Omega(\max\{\bar{N}, \log(T)\})$. Thus, in view of the above theorem, ECT is order-optimal if $\bar{N} = \Omega((\log T)^3)$, and order-optimal up to an $(\log T)^2$ factor otherwise. Furthermore, note that when $R_{\text{ic}}(T) = \Omega(\frac{T}{\Delta_2 m})$ is the leading term in our regret lower bound, ECT regret has also the right scaling in Δ_2 : $R^{\text{ECT}}(T) = \mathcal{O}(\frac{T}{\Delta_2 m})$.

5.2 Unclustered items and statistically identical users

When items are not clustered, we propose ET (Explore-and-Test), an algorithm that consists of two phases: an exploration phase that aims at estimating the threshold level $\mu_{1-\varepsilon}$, and a test phase where we apply to each item a sequential test to determine whether the item is above the threshold.

The Explore-and-Test algorithm. The ET algorithm proceeds as follows.

(a) **Exploration phase.** In this phase, we randomly select of set \mathcal{S} consisting of $\lfloor \frac{8}{\varepsilon^2} \log T \rfloor$ items and recommend each selected item to $\lfloor 4^2 \log T \rfloor$ users. For each item $i \in \mathcal{S}$, we compute its empirical success rate $\hat{\rho}_i$. We then estimate $\mu_{1-\frac{\varepsilon}{2}}$ by $\hat{\mu}_{1-\frac{\varepsilon}{2}}$ defined so that: $\frac{\varepsilon}{2} |\mathcal{S}| = |\{i \in \mathcal{S} : \hat{\rho}_i \geq \hat{\mu}_{1-\frac{\varepsilon}{2}}\}|$. We also initialize the set \mathcal{V} of candidate items to exploit as $\mathcal{V} = \{i \in \mathcal{S} : \hat{\rho}_i \geq \hat{\mu}_{1-\frac{\varepsilon}{2}}\}$.

(b) **Test phase.** In this phase, we recommend items in \mathcal{V} , and update the set \mathcal{V} . Specifically, when a user u arrives, we recommend the item $i \in \mathcal{V}$ that has been recommended the least recently among items that would not break the no-repetition constraint. If no such items exist in \mathcal{V} , we randomly

recommend an item outside \mathcal{V} and add it to \mathcal{V} .

Now to ensure that items in \mathcal{V} are above the threshold, we perform the following sequential test, which is reminiscent of sequential tests used in optimal algorithms for infinite bandit problems [2]. For each item, the test is applied when the item has been recommended for the $\lfloor 2^\ell \log \log_2(2^e m^2) \rfloor$ times for any positive integer ℓ . For the ℓ -th test, we denote by $\bar{\rho}^{(\ell)}$ the real number such that $\text{kl}(\bar{\rho}^{(\ell)}, \hat{\mu}_{1-\frac{\varepsilon}{2}}) = 2^{-\ell}$. If $\bar{\rho}^{(\ell)} \leq \hat{\mu}_{1-\frac{\varepsilon}{2}}$, the item is removed from \mathcal{V} .

Theorem 5. *Assume that the density of ζ satisfies $\zeta(\mu) \leq C$ for all $\mu \in [0, 1]$.*

For any $\varepsilon \geq C \sqrt{\frac{\pi}{2 \log T}}$, we have: $R_\varepsilon^{\text{ET}}(T) = \mathcal{O}\left(\bar{N} \frac{\log(1/\varepsilon) \log \log(m)}{\varepsilon} + \frac{(\log T)^2}{\varepsilon^2}\right)$.

In view of Theorem 2, the regret of any algorithm scales at least as $\Omega(\frac{\bar{N}}{\varepsilon})$. Hence, the above theorem states that ET is order-optimal at least when $\bar{N} = \Omega((\log T)^2)$.

5.3 Clustered items and clustered users

The main challenge in devising an algorithm in this setting stems from the fact that we do not control the user arrival process. In turn, clustering users with low regret is delicate. We present Explore-Cluster with Upper Confidence Sets (EC-UCS), an algorithm that essentially exhibits the same regret scaling as our lower bound. The idea behind the design of EC-UCS is as follows. We estimate the success rates $(p_{k\ell})_{k,\ell}$ using small subsets of items and users. Then based on these estimates, each user is optimistically associated with a UCS, *Upper Confidence Set*, a set of clusters the user may likely belong to. The UCS of a user then shrinks as the number of requests made by this user increases (just as the UCB index of an arm in bandit problems gets closer to its average reward). The design of our estimation procedure and of the various UCS is made so as to get an order-optimal algorithm. In what follows, we assume that $m^2 \geq T(\log T)^3$ and $T \geq m \log(T)$.

The Explore-Cluster-with-Upper Confidence Sets algorithm.

(a) **Exploration and item clustering phase.** The algorithm starts by collecting data to infer the item clusters. It randomly selects a set \mathcal{S} consisting of $\min\{n, \lfloor \frac{m}{(\log T)^2} \rfloor\}$ items. For the $10m$ first user arrivals, it recommends items from \mathcal{S} uniformly at random. These $10m$ recommendations and the corresponding user responses are recorded in the dataset \mathcal{D} . From the dataset \mathcal{D} , the item clusters are extracted using a spectral algorithm (see Algorithm 4 in the appendix). This algorithm is taken from [33], and considers the *indirect edges* between items created by users. Specifically, when a user appears more than twice in \mathcal{D} , she creates an indirect edge between the items recommended to her for which she provided the same answer (1 or 0). Items with indirect edges are more likely to belong to the same cluster. The output of this phase is a partition of \mathcal{S} into item clusters $\hat{I}_1, \dots, \hat{I}_K$. We can show that with an exploration budget of $10m$, w.h.p. at least $m/2$ indirect edges are created and that in turn, the spectral algorithm does not make any clustering errors w.p. at least $1 - \frac{1}{T}$.

(b) **Exploration and user clustering phase.** To the $(10 + \log(T))m$ next user arrivals, EC-UCS clusters a subset of users using a Nearest-Neighbor algorithm. The algorithm selects a subset \mathcal{U}^* of users to cluster, and recommendations to the remaining users will be made depending some distance to the inferred clusters in \mathcal{U}^* . Users from all clusters must be present in \mathcal{U}^* . To this aim, EC-UCS first randomly selects a subset \mathcal{U}_0 of $\lfloor m/\log(T) \rfloor$ users from which it extracts the set \mathcal{U}^* of $\lfloor \log(T)^2 \rfloor$ users who have been observed the most. The extraction and the clustering of \mathcal{U}^* is made several times until the $\lfloor (10 + \log(T))m \rfloor$ -th user arrives so as to update and improve the user clusters. From these clusters, we deduce estimates $\hat{p}_{k\ell}$ of the success probabilities.

(c) **Recommendations based on Optimistic Assignments.** After the $10m$ -th arrivals, recommendations are made based on the estimated $\hat{p}_{k\ell}$'s. For user $u_t \notin \mathcal{U}_0$, the item selection further depends on the \hat{p}_{ku_t} 's, the empirical success rates of user u_t for items in the various clusters. A greedy recommendation for u_t would consist in assigning u_t to cluster ℓ minimizing $\|\hat{p}_{\cdot\ell} - \hat{p}_{\cdot u_t}\|$ over ℓ , and then in picking an item from cluster \hat{I}_k with maximal $\hat{p}_{k\ell}$. Such a greedy recommendation would not work as when u_t has not been observed many times, the cluster she belongs to remains uncertain. To address this issue, we apply the Optimism in Front of Uncertainty principle often used in bandit algorithms to foster exploration. Specifically, we build a set $\mathcal{L}(u_t)$ of clusters u_t is likely to belong to. $\mathcal{L}(u_t)$ is referred to as the Upper Confidence Set of u_t . As we get more observations of u_t , this set shrinks. Specifically, we let $x_{k\ell} = \max\{|\hat{p}_{k\ell} - \hat{p}_{ku_t}| - \varepsilon, 0\}$, for some well defined $\varepsilon > 0$ (essentially scaling as $\sqrt{\log \log(T)/\log(T)}$, see Appendix A for details), and

define $\mathcal{L}(u_t) = \{\ell \in [L] : \sum_k x_k^2 n_{k u_t} < 2K \log(n_{u_t})\}$ (n_{u_t} is the number of time u_t has arrived, and $n_{k u_t}$ is the number of times u_t has been recommended an item from cluster \hat{I}_k). After optimistically composing the set $\mathcal{L}(u_t)$, u_t is assigned to cluster ℓ chosen uniformly at random in $\mathcal{L}(u_t)$, and recommended an item from cluster \hat{I}_k with maximal $\hat{p}_{k\ell}$.

Theorem 6. For any ℓ , let σ_ℓ be the permutation of $[K]$ such that $p_{\sigma_\ell(1)\ell} > p_{\sigma_\ell(2)\ell} \geq \dots \geq p_{\sigma_\ell(K)\ell}$. Let $\mathcal{R}_\ell = \{r \in [L] : k_\ell^* \neq k_r^*\}$, $\mathcal{S}_{\ell r} = \{k \in [K] : p_{k\ell} \neq p_{kr}\}$, $y_{\ell r} = \min_{k \in \mathcal{S}_{\ell r}} |p_{k\ell} - p_{kr}|$, $\delta = \min_\ell (p_{\sigma_\ell(1)\ell} - p_{\sigma_\ell(2)\ell})$, and $\phi(x) := x/\log(1/x)$. Then, we have:

$$R^{\text{EC-UCS}}(T) = \mathcal{O} \left(m \sum_\ell \beta_\ell (p_{\sigma_\ell(1)\ell} - p_{\sigma_\ell(K)\ell}) \left(\max \left(\frac{K^3 \log K}{\phi(\min(y_{\ell r}, \delta))^2}, \frac{\sqrt{K}}{\min_\ell \beta_\ell} \right) + \sum_{r \in \mathcal{R}_\ell \setminus \mathcal{L}^\perp(\ell)} \frac{K^2 \log K}{\phi(|p_{k_\ell^* r} - p_{k_\ell^* \ell}|^2)} + \sum_{k \in \mathcal{S}_{\ell r}} \sum_{r \in \mathcal{L}^\perp(\ell)} \frac{K \log \bar{N}}{|\mathcal{S}_{\ell r}| |p_{k\ell} - p_{kr}|^2} \right) \right).$$

EC-UCS blends clustering and bandit algorithms, and its regret analysis is rather intricate. The above theorem states that remarkably, the regret of the EC-UCS algorithm matches our lower bound order-wise. In particular, the algorithm manages to get a regret (i) scaling as m whenever it is possible, i.e., when $\mathcal{L}^\perp(\ell) = \emptyset$ for all ℓ , (ii) scaling as $m \log(\bar{N})$ otherwise.

In Appendix A.3, we present ECB, a much simpler algorithm than EC-UCS, but whose regret upper bound, derived in Appendix J, always scales as $m \log(\bar{N})$.

6 Conclusion

This paper proposes and analyzes several models for online recommendation systems. These models capture both the fact that items cannot repeatedly be recommended to the same users and some underlying user and item structure. We provide regret lower bounds and algorithms approaching these limits for all models. Many interesting and challenging questions remain open. We may, for example, investigate other structural assumptions for the success probabilities (e.g. soft clusters), and adapt our algorithms. We may also try to extend our analysis to the very popular linear reward structure, but accounting for no-repetition constraint.

Broader Impact

This work, although mostly theoretical, may provide guidelines and insights towards an improved design of recommendation systems. The benefits of such improved design could be to increase user experience with these systems, and to help companies to improve their sales strategies through differentiated recommendations. The massive use of recommendation systems and its potential side effects have recently triggered a lot of interest. We must remain aware of and investigate such effects. These include: opinion polarization, a potential negative impact on users' behavior and their willingness to pay, privacy issues.

Acknowledgements

K. Ariu was supported by the Nakajima Foundation Scholarship. S. Yun and N. Ryu were supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT)(No.2019-0-00075, Artificial Intelligence Graduate School Program(KAIST)). A. Proutiere's research is supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

References

- [1] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, pages 235–256, 2002.
- [2] Thomas Bonald and Alexandre Proutiere. Two-target algorithms for infinite-armed bandits with bernoulli rewards. In *Advances in Neural Information Processing Systems 26*, pages 2184–2192. 2013.

- [3] Guy Bresler and Mina Karzand. Regret bounds and regimes of optimality for user-user and item-item collaborative filtering. In *2018 Information Theory and Applications Workshop (ITA)*, pages 1–37, 2018.
- [4] Guy Bresler, George H Chen, and Devavrat Shah. A latent source model for online collaborative filtering. In *Advances in Neural Information Processing Systems*, pages 3347–3355, 2014.
- [5] Sebastien Bubeck, Vianney Perchet, and Philippe Rigollet. Bounded regret in stochastic multi-armed bandits. In *Proceedings of the 26th Annual Conference on Learning Theory*, pages 122–134, 2013.
- [6] Loc Bui, Ramesh Johari, and Shie Mannor. Clustered bandits. *arXiv preprint arXiv:1206.4169*, 2012.
- [7] Richard Combes, Chong Jiang, and Rayadurgam Srikant. Bandits with budgets: Regret lower bounds and optimal algorithms. In *Proceedings of the 2015 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, pages 245–257, 2015.
- [8] Aurélien Garivier, Pierre Ménard, and Gilles Stoltz. Explore first, exploit next: The true shape of regret in bandit problems. *Mathematics of Operations Research*, 2018.
- [9] Claudio Gentile, Shuai Li, and Giovanni Zappella. Online clustering of bandits. In *Proceedings of the 31th International Conference on Machine Learning*, pages 757–765, 2014.
- [10] Claudio Gentile, Shuai Li, Purushottam Kar, Alexandros Karatzoglou, Giovanni Zappella, and Evans Etrúe. On context-dependent clustering of bandits. In *Proceedings of the 34th International Conference on Machine Learning*, pages 1253–1262, 2017.
- [11] Aditya Gopalan, Odalric-Ambrym Maillard, and Mohammadi Zaki. Low-rank bandits with latent mixtures. *arXiv preprint arXiv:1609.01508*, 2016.
- [12] Botao Hao, Tor Lattimore, and Csaba Szepesvari. Adaptive exploration in linear contextual bandit. *arXiv preprint arXiv:1910.06996*, 2019.
- [13] Reinhard Heckel and Kannan Ramchandran. The sample complexity of online one-class collaborative filtering. In *Proceedings of the 34th International Conference on Machine Learning*, pages 1452–1460, 2017.
- [14] Matthieu Jedor, Vianney Perchet, and Jonathan Louedec. Categorized bandits. In *Advances in Neural Information Processing Systems*, pages 14399–14409, 2019.
- [15] Kwang-Sung Jun, Rebecca Willett, Stephen Wright, and Robert Nowak. Bilinear bandits with low-rank structure. In *Proceedings of the 36th International Conference on Machine Learning*, pages 3163–3172, 2019.
- [16] O. Kallenberg. *Random Measures, Theory and Applications*. Probability Theory and Stochastic Modelling. Springer International Publishing, 2017. ISBN 9783319415987.
- [17] Robert D. Kleinberg, Alexandru Niculescu-Mizil, and Yogeshwer Sharma. Regret bounds for sleeping experts and bandits. In *Proceedings of the 21st Annual Conference on Learning Theory*, pages 425–436, 2008.
- [18] Joon Kwon, Vianney Perchet, and Claire Vernade. Sparse stochastic bandits. In *Proceedings of the 30th Conference on Learning Theory*, pages 1269–1270, 2017.
- [19] Tze Leung Lai and Herbert Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- [20] Shuai Li and Shengyu Zhang. Online clustering of contextual cascading bandits. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [21] Shuai Li, Alexandros Karatzoglou, and Claudio Gentile. Collaborative filtering bandits. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, pages 539–548, 2016.
- [22] Shuai Li, Wei Chen, and Kwong-Sak Leung. Improved algorithm on online clustering of bandits. *arXiv preprint arXiv:1902.09162*, 2019.
- [23] Odalric-Ambrym Maillard and Shie Mannor. Latent bandits. In *Proceedings of the 31th International Conference on Machine Learning*, pages 136–144, 2014.
- [24] Jérémie Mary, Romaric Gaudel, and Philippe Preux. Bandits and recommender systems. In *International Workshop on Machine Learning, Optimization and Big Data*, pages 325–336, 2015.

- [25] Jonas W Mueller, Vasilis Syrgkanis, and Matt Taddy. Low-rank bandit methods for high-dimensional dynamic pricing. In *Advances in Neural Information Processing Systems*, pages 15442–15452, 2019.
- [26] Jungseul Ok, Se-Young Yun, Alexandre Proutiere, and Rami Mochaourab. Collaborative clustering: Sample complexity and efficient algorithms. In *International Conference on Algorithmic Learning Theory*, pages 288–329, 2017.
- [27] Martin Raab and Angelika Steger. Balls into Bins - A Simple and Tight Analysis. In *Proceedings of the Second International Workshop on Randomization and Approximation Techniques in Computer Science*, pages 159–170, 1998.
- [28] Paul Resnick and Hal R Varian. Recommender systems. *Communications of the ACM*, 40(3):56–58, 1997.
- [29] Daniel Russo and Benjamin Van Roy. Satisficing in time-sensitive bandit learning. *arXiv preprint arXiv:1803.02855*, 2018.
- [30] Joel A Tropp et al. An introduction to matrix concentration inequalities. *Foundations and Trends in Machine Learning*, 8(1-2):1–230, 2015.
- [31] Alexandre B Tsybakov. *Introduction to nonparametric estimation*. Springer Science & Business Media, 2008.
- [32] Se-Young Yun and Alexandre Proutiere. Optimal cluster recovery in the labeled stochastic block model. In *Advances in Neural Information Processing Systems*, pages 965–973, 2016.
- [33] Se-Young Yun, Alexandre Proutiere, et al. Streaming, memory limited algorithms for community detection. In *Advances in Neural Information Processing Systems*, pages 3167–3175, 2014.

7 Table of Notations

Notations common to all models	
n	Number of items
m	Number of users
\mathcal{I}	Set of items
\mathcal{U}	Set of users
u_t	User requesting recommendation at round t
T	Time horizon
X_{iu}	Binary random variable to indicate whether user u likes the item i
$\rho = (\rho_{iu})_{i \in \mathcal{I}, u \in \mathcal{U}}$	Probability that the user u likes the item i
π	Algorithm for sequential item selection
Π	Set of all algorithms for sequential item selection
i_t^π	Item selected at round t under π
\mathcal{F}_{t-1}^π	σ -algebra generated by $(u_s, (u_s, i_s^\pi, X_{i_s^\pi u_s}), s \leq t-1)$
\bar{n}	Term $\mathbb{E}[\max_{u \in \mathcal{U}} N_u(T)]$
\bar{N}	Term $\frac{4 \log(m)}{\log\left(\frac{m \log(m)}{T} + e\right)} + \frac{e^2 T}{m}$
Generic notations	
\hat{a}	Estimated value of a
$\sigma(A)$	σ -algebra generated by A
$\text{kl}(p, q)$	Kullback–Leibler divergence from Bernoulli random variable with parameter p to that with parameter q
\gtrsim	We write $a \gtrsim b$ if $\liminf_{T \rightarrow \infty} a/b \geq 1$

Table 1: Table of notations common to all models

Model A: Clustered items and statistically identical users	
\mathcal{I}_k	Set of items in the item cluster k
$\alpha = (\alpha_k)_{k \in [K]}$	Probability that an item is assigned to the item cluster k
K	Number of item clusters
Δ	Minimum difference between the success rates of items in optimal cluster and of items in sub-optimal cluster
$p = (p_k)_{k \in [K]}$	Probability that the user likes the item $i \in \mathcal{I}_k$
$R^\pi(T)$	Regret of an algorithm π
Δ_k	Term $p_1 - p_k$
$\phi(k, m, p)$	Term $\frac{1 - e^{-m\gamma(p_1, p_k)}}{8(1 - e^{-\gamma(p_1, p_k)})}$
$\gamma(p, q)$	Term $\text{kl}(p, q) + \text{kl}(q, p)$
η	Term $\min_k p_k (1 - p_k)$
\mathcal{S}	Set of initially sampled items
\mathcal{V}	Set of items believed to be in the best cluster

Table 2: Table of notations: Model A

Model B: Unclustered items and statistically identical users	
ζ	Distribution over $[0, 1]$
μ_x	Term $\inf\{\gamma \in [0, 1] : \mathbb{P}[\rho_i \leq \gamma] \geq x\}$
ε	Constant that specifies the ε -best items
$R_\varepsilon^\pi(T)$	Satisficing regret of algorithm π with a given $\varepsilon > 0$
C	Constant that regularizes the distribution $\zeta(\mu)$
S	Set of initially sampled items
\mathcal{V}	Set of items believed to be in the best cluster

Table 3: Table of notations: Model B

Model C: Clustered items and clustered users	
\mathcal{I}_k	Set of items in the item cluster k
$\alpha = (\alpha_k)_{k \in [K]}$	Probability that an item is assigned to the item cluster k
\mathcal{U}_ℓ	Set of users in the user cluster ℓ
$\beta = (\beta_\ell)_{\ell \in [L]}$	Probability that a user is assigned to the user cluster ℓ
K	Number of item clusters
L	Number of user clusters
$p = (p_{k\ell})_{k \in [K], \ell \in [L]}$	Probability that the user u likes the item i such that $i \in \mathcal{I}_k$ and $u \in \mathcal{U}_\ell$
k_ℓ^*	Term $\arg \max_k p_{k\ell}$
$\Delta_{k\ell}$	Term $p_{k_\ell^* \ell} - p_{k\ell}$
Δ	Minimum difference between the success rates of items in optimal cluster and of items in sub-optimal cluster
δ_ℓ	Term $\min_{k: \Delta_{k\ell} > 0} \Delta_{k\ell}$
$\phi(k, \ell, m, p)$	Term $\frac{1-e^{-\frac{T}{m} \gamma(p_{k_\ell^* \ell}, p_{k\ell})}}{8 \left(1-e^{-\gamma(p_{k_\ell^* \ell}, p_{k\ell})}\right)}$
$\psi(\ell, k, T, m, p)$	Term $\frac{1-e^{-\frac{T}{m} \gamma(p_{k_\ell^* \ell}, p_{k\ell})}}{8 \left(1-e^{-\gamma(p_{k_\ell^* \ell}, p_{k\ell})}\right)}$
\mathcal{L}_\perp	Set $\{(\ell, \ell') \in [L]^2 : p_{k_\ell^* \ell} \neq p_{k_{\ell'}^* \ell'}\}$
$\mathcal{L}^\perp(\ell)$	Set $\{\ell' \neq \ell : k_\ell^* \neq k_{\ell'}^*, p_{k_\ell^* \ell} = p_{k_{\ell'}^* \ell'}\}$
$R_u^\pi(N)$	Accumulated expected regret under π for user u when the user has arrived N times
$R^\pi(T)$	Regret of an algorithm π
S	Set of initially sampled items
\mathcal{U}_0	Set of initially sampled users
\mathcal{U}^*	Set of $(\log T)^2$ users in \mathcal{U}_0 who have been arrived the most
$\mathcal{L}(u_t)$	Upper Confidence Set of the user u_t
σ_ℓ	Permutation of $[K]$ such that $p_{\sigma_\ell(1)\ell} > p_{\sigma_\ell(2)\ell} \geq \dots \geq p_{\sigma_\ell(K)\ell}$
\mathcal{R}_ℓ	Set $\{r \in [L] : k_\ell^* \neq k_r^*\}$
$S_{\ell r}$	Set $\{k \in [K] : p_{k\ell} \neq p_{kr}\}$
$y_{\ell r}$	Term $\min_{k \in S_{\ell r}} p_{k\ell} - p_{kr} $
δ	Term $\min_\ell (p_{\sigma_\ell(1)\ell} - p_{\sigma_\ell(2)\ell})$
ε	Term $K \sqrt{\frac{8Km}{t}} \log \frac{t}{m}$ (Updated only when the user clustering is executed)

Table 4: Table of notations: Model C

A Algorithms and experiments

In this section, we present the detailed pseudo-codes of our algorithms. We also illustrate the performance of these algorithms numerically.

A.1 Clustered items and statistically identical users

Algorithm 1 Explore-Cluster-and-Test

Input: T, K

1. Exploration
 Sample a set \mathcal{S} of $\lfloor (\log T)^2 \rfloor$ items (uniformly at random)
 Recommend each item in \mathcal{S} to $\lfloor \log T \rfloor$ users
 (when this is not possible due to the no-repetition constraint) recommend a random feasible item.
 $T_0 \leftarrow$ round where the exploration phase ends

2. Clustering
 $\hat{\rho}_i \leftarrow$ the empirical average of i for all $i \in \mathcal{S}$
 $Q_i \leftarrow \{j \in \mathcal{S} : |\hat{\rho}_i - \hat{\rho}_j| \leq (\log T)^{-\frac{1}{4}}\}$ for all $i \in \mathcal{S}$
 $M \leftarrow \emptyset$
for $k = 1$ **to** K **do**
 $i_k \leftarrow \arg \max_{j \in \mathcal{S}} |Q_j \setminus \cup_{\ell=1}^{k-1} Q_{i_\ell}|$
 if $|Q_{i_k}| < \log T$ **then**
 break
 end if
 $M \leftarrow M \cup \{i_k\}$
end for
 $i_1 \leftarrow \arg \max_{i \in M} \hat{\rho}_i$ and $\hat{p}_1 \leftarrow \hat{\rho}_{i_1}$
 $i_2 \leftarrow \arg \max_{i \in M \setminus \{i_1\}} \hat{\rho}_i$ and $\hat{p}_2 \leftarrow \hat{\rho}_{i_2}$

3. Test
 $\Delta_0 \leftarrow \hat{p}_1 - \hat{p}_2$
 $\mathcal{V}, \mathcal{V}_0 \leftarrow \{i \in \mathcal{S} : \hat{\rho}_i > (\hat{p}_1 + \hat{p}_2)/2\}$
for $t = (T_0 + 1)$ **to** T **do**
 Recommend item from \mathcal{V} with the highest empirical average if possible, otherwise randomly recommend item i from $\mathcal{T} \setminus \mathcal{V}_0$ and add i to \mathcal{V} and \mathcal{V}_0
 if the number of times i has been recommended is a multiple of $\lfloor \frac{2 \log 3}{\Delta_0^2} \rfloor$ and $\hat{\rho}_i < (\hat{p}_1 + \hat{p}_2)/2$ **then**
 Remove i from \mathcal{V}
 end if
end for

Numerical experiments. We illustrate the performance of ECT in the following scenario: $K = 2$ item clusters, $n = 3000$ items, $m = 5000$ users, $p_1 = 0.7$, $p_2 = 0.2$, $\alpha_1 = \alpha_2 = 0.5$.

We compare the performance of ECT to two naive algorithms:

(i) B-KLUCB [7]: This algorithm was proposed for budgeted bandits. Here the budget per arm is m . The algorithm ranks the arms (the items) according to their KL-UCB indexes, and selects the available item (accounting for the no-repetition constraint) with the highest index.

(ii) B-KLUCB with sampling: The algorithm first samples $\lfloor (T/m) \log T \rfloor$ items randomly, and play B-KLUCB only for these items. When none of these items can be played, the algorithm plays a randomly selected item (as in ECT).

Figure 1 plots the regret vs time for the 3 algorithms, for a time horizon $T = 20000$. The regret is averaged over 200 runs for ECT and B-KLUCB with sampling, and 20 runs for B-KL-UCB (we do not need more runs since there is no randomness induced by the initial item sampling procedure). For ECT, the number of items initially sampled is $|\mathcal{S}| = \lfloor 0.225 \log(T)^2 \rfloor$, whereas for B-KLUCB, it is 39 (this number is optimized so as to get the best performance – refer to Figure 3 for a sensitivity analysis of the regret depending on the number of items initially sampled).

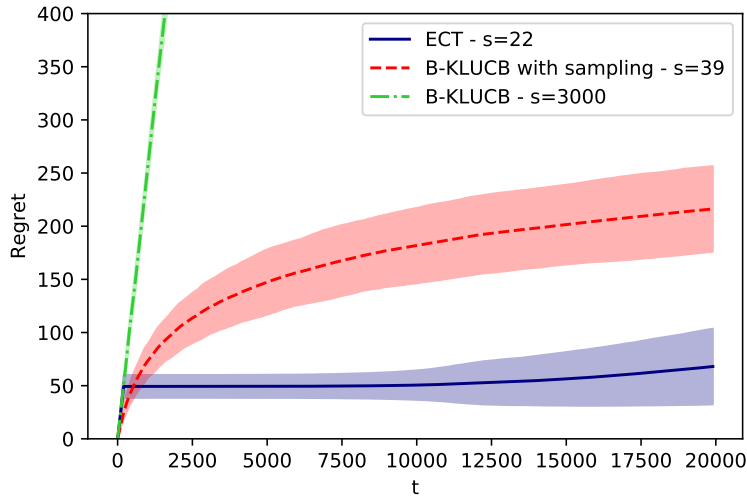


Figure 1: Regret vs time averaged over 200 instances for ECT and B-KLUCB with sampling, and over 20 instances for B-KLUCB. $T = 20000$, $n = 3000$, $m = 5000$, $p_1 = 0.7$, $p_2 = 0.2$, $\alpha_1 = \alpha_2 = 0.5$. Shaded areas correspond to one standard deviations.




Figure 2: Regret vs time averaged over 100 instances for each algorithm. $T = 100000$, $n = 3000$, $m = 5000$, $p_1 = 0.7$, $p_2 = 0.2$, $\alpha_1 = \alpha_2 = 0.5$. Shaded areas correspond to one standard deviations.

Figure 2 compares the regret of ECT to that obtained under B-KLUCB with sampling when $T = 100000$. ECT initially samples 29 items. For B-KLUCB with sampling, we have tested two different numbers of items initially sampled, namely 29 and 230. After round 20000, ECT starts playing items that have not been used in the exploration phase. To keep regret low, ECT hence relies on sequential tests. The regret curve of ECT shows that these tests perform very well. This contrasts with B-KLUCB with sampling: when the number of initially sampled items is 29, as for ECT, after round 20000, new items must be selected, and B-KLUCB performs very poorly (the regret rapidly grows).

Finally, we assess the sensitivity of ECT and B-KLUCB with sampling w.r.t. the number of initially sampled items. Figure 3 plots the regret after $T = 20000$ rounds depending on this number. Again




Figure 3: Regret at $T = 20000$ vs number of items initially sampled $|\mathcal{S}|$ for ECT and B-KLUCB with sampling, averaged over 200 instances. $T = 20000$, $n = 3000$, $m = 5000$, $p_1 = 0.7$, $p_2 = 0.2$, $\alpha_1 = \alpha_2 = 0.5$. One standard deviations are shown in the error bars.

we average over 200 runs. ECT is not very sensitive to the number of sampled items; B-KLUCB is, on the other hand, very sensitive. For ECT, this provides further evidence that the sequential tests applied to items not used in the exploration phase are very efficient.

A.2 Unclustered items and statistically identical users

Algorithm 2 Explore-and-Test

Input: T, K

1. Exploration

Sample a set \mathcal{S} of $\lfloor \frac{8^2}{\varepsilon^2} \log T \rfloor$ items (uniformly at random)

Recommend each item in \mathcal{S} to $\lfloor 4^2 \log T \rfloor$ users

(when this is not possible due to the no-repetition constraint) recommend a random feasible item

$\hat{\rho}_i \leftarrow$ the empirical average of i for all $i \in \mathcal{S}$

$T_0 \leftarrow$ round where the exploration phase ends

2. Test

Compute $\hat{\mu}_{1-\frac{\varepsilon}{2}}$ s.t. $\frac{\varepsilon}{2}|\mathcal{S}| = |\{i \in \mathcal{S} : \hat{\rho}_i \geq \hat{\mu}_{1-\frac{\varepsilon}{2}}\}|$

$\mathcal{V} \leftarrow \{i \in \mathcal{S} : \hat{\rho}_i \geq \hat{\mu}_{1-\frac{\varepsilon}{2}}\}$

Reset the reward observation history

for $t = (T_0 + 1)$ **to** T **do**

Recommend an item i that was recommended the least recently among items in \mathcal{V} . If items in \mathcal{V} cannot be selected, recommend an item i , randomly selected from the set of unrecommended items, and add i to \mathcal{V}

if the number of times i has been recommended is exactly $\lfloor 2^\ell \log \log_2(2^e m^2) \rfloor$ for some positive integer ℓ **then**

Compute $\bar{\rho}^{(\ell)} (\leq \hat{\mu}_{1-\frac{\varepsilon}{2}})$ s.t. $\text{kl}(\bar{\rho}^{(\ell)}, \hat{\mu}_{1-\frac{\varepsilon}{2}}) = 2^{-\ell}$

if $\hat{\rho}_i \leq \bar{\rho}^{(\ell)}$ **then**

$\mathcal{V} \leftarrow \mathcal{V} \setminus \{i\}$

end if

end if

end for

Numerical experiments. Consider a system with $n = 700$ items, and $m = 1300$ users. The time horizon is $T = 100000$. Assume that the distribution ζ is uniform over the interval $[0.1, 0.9]$, and let us target items within the 30% best items, i.e., $\varepsilon = 0.3$. In Figure 4, we compare the satisficing regret averaged over 100 runs of the ET algorithm with $|\mathcal{S}| = 65$ items used in the exploration phase, to that achieved under B-KLUCB with sampling. Since under any algorithm, one needs to use at least T/m items, the number of items sampled under B-KLUCB with sampling is chosen as $\min((T/m) \log T, n)$, which is equal to $n = 700$ in our setting. Figure 4 illustrates the efficiency of the sequential tests used under ET.

Next, we assess the sensitivity of ET and B-KLUCB w.r.t. the number of initially sampled items. Figure 5 compares the satisficing regret after $T = 100000$ depending on this number. The values are averaged over 20 runs. ET seems robust to the number of sampled items. B-KLUCB is, however, very sensitive to the number of sampled items and shows larger regret than that of ET. This result presents further evidence that the sequential testing procedures used in ET are efficient.




Figure 4: Satisficing regret vs time averaged over 100 runs of ET and B-KLUCB with sampling. $n = 700, m = 1300, T = 100000, \varepsilon = 0.3$. ζ is uniform on $[0.1, 0.9]$. Shaded areas correspond to one standard deviations.

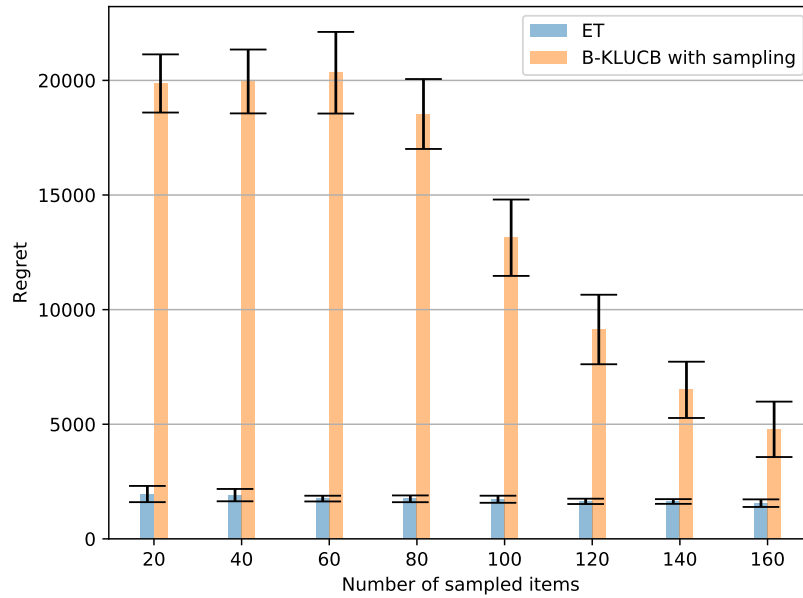


Figure 5: Satisficing regret at $T = 100000$ vs number of items initially sampled $|\mathcal{S}|$ for ET and B-KLUCB with sampling, averaged over 20 instances. $n = 700, m = 1300, T = 100000, \varepsilon = 0.3$. ζ is uniform on $[0.1, 0.9]$. One standard deviations are shown in the error bars.

A.3 Clustered items and users

Here we start by providing the description of our order-optimal algorithm, EC-UCS. We then present ECB (Explore-Cluster-Bandit), a much simpler algorithm but with lower performance guarantees.

A.3.1 The EC-UCS algorithm

We present the pseudo-code of the EC-UCS algorithm in Algorithm 3. The algorithm calls spectral clustering algorithms whose pseudo-codes are provided in Algorithms 4-5-6.

Algorithm 3 Explore-and-Cluster-with-Upper-Confidence-Sets (EC-UCS)

Input: T, K, L

1. Exploration for Item Clustering
 $n_0 \leftarrow \min\{n, \lfloor m/(\log T)^2 \rfloor\}$
 Sample a set \mathcal{S} of n_0 items.
for $1 \leq t \leq 10m$ **do**
 Recommend items from \mathcal{S} randomly (when this is not possible due to the no-repetition constraint, recommend a random feasible item). Record the user responses in the dataset \mathcal{D} .
end for

2. Item Clustering
 Run Algorithm 4 with input $\mathcal{S}, \mathcal{D}, K$ and output $\hat{I}_1, \dots, \hat{I}_K$.

3. Exploitation
 $\mathcal{U}_0 \leftarrow$ a set of randomly chosen $\lfloor \frac{m}{\log T} \rfloor$ users
 $\hat{\rho}_u = (\hat{\rho}_{uk})_{k \in [K]} \leftarrow$ the empirical average of u for each item cluster \hat{I}_k for all $u \in \mathcal{U}$
for $10m < t \leq T$ **do**
 if $t \leq \lfloor (10 + \log T)m \rfloor$ and $t = (9 + 2^i)m + 1$ for some non-negative integer i **then**
 $\mathcal{U}^* \leftarrow$ a set of $\lfloor (\log T)^2 \rfloor$ users in \mathcal{U}_0 who have been observed the most
 $\epsilon \leftarrow K \sqrt{\frac{8Km}{t} \log \frac{t}{m}}$
 $Q_u \leftarrow \{v \in \mathcal{U}^* : \|\hat{\rho}_u - \hat{\rho}_v\| \leq \epsilon\}$ for all $u \in \mathcal{U}^*$
 for $\ell = 1$ **to** L **do**
 $u_\ell \leftarrow \arg \max_{v \in \mathcal{U}^*} |Q_v \setminus \cup_{r=1}^{\ell-1} Q_{u_r}|$ and $\hat{p}_\ell \leftarrow \hat{\rho}_{u_\ell}$
 $L_0 \leftarrow \ell$
 if $\cup_{r=1}^{\ell} Q_{i_r} = \mathcal{U}^*$ **then**
 break
 end if
 end for
 end if
 if $u_t \in \mathcal{U}_0$ and $t \leq \lfloor (10 + \log T)m \rfloor$ **then**
 Recommend an item in a round-robin fashion from $\hat{I}_1, \dots, \hat{I}_K$ for each user
 else
 $x_{k\ell} \leftarrow \max\{|\hat{p}_{k\ell} - \hat{\rho}_{ku_t}| - \epsilon, 0\}$
 $\mathcal{L}(u_t) \leftarrow \{\ell \in [L_0] : \sum_{k=1}^K n_{ku_t} x_{k\ell}^2 < 2K \log n_{u_t}\}$
 if $|\mathcal{L}(u_t)| \geq 1$ **then**
 Recommend an item uniformly at random from $(\hat{I}_{k_\ell^*})_{\ell \in \mathcal{L}(u_t)}$
 else
 Recommend an item uniformly at random from $\hat{I}_1, \dots, \hat{I}_K$
 end if
 end if
 Update $\hat{\rho}_{u_t}$
end for

Algorithm 4 Spectral Item Clustering by indirect edges (inspired by Algorithm 1 in [33])

Input: $\mathcal{S}, \mathcal{D}, K$
 $A \leftarrow \mathbf{0} \in \mathbb{R}^{\mathcal{S} \times \mathcal{S}}, s \leftarrow 0$
for $u \in \mathcal{U}$ **do**
 if u has received recommendations at least two times **then**
 $s \leftarrow s + 1$
 if user u gives positive responses to both i and j that are the first two recommended items to user u **then**
 $A_{ij} \leftarrow A_{ij} + 1, A_{ji} \leftarrow A_{ji} + 1$
 end if
 end if
end for
Run Algorithm 5 with input A, s, K and output $\hat{I}_1, \dots, \hat{I}_K$.
Output: $\hat{I}_1, \dots, \hat{I}_K$

Algorithm 5 Spectral Partitioning+ (an improved version of Algorithm 2 in [26])

Input: Observation matrix A, s, K
1. Spectral Decomposition
Run Algorithm 6, with input A, K and output $(S_k)_{k=1, \dots, K}$.
2. Improvement
 $\hat{p}(i, j) \leftarrow \frac{\sum_{v \in S_i} \sum_{v' \in S_j} A_{v, v'}}{|S_i|s}$ for all $1 \leq i, j \leq K$
 $S_k^{(0)} \leftarrow S_k$ for all $1 \leq k \leq K$
for $t = 1$ **to** $\lfloor \log n_0 \rfloor$ **do**
 $S_k^{(t)} \leftarrow \emptyset$ for all $1 \leq k \leq K$
 for $v \in \mathcal{S}$ **do**
 $i^* \leftarrow \arg \max_{1 \leq i \leq K}$
 $\left\{ \sum_{k=0}^K (\sum_{w \in S_k^{(t-1)}} A_{vw}) \log \hat{p}(i, k) \right\}$
 where $\sum_{w \in S_0^{(t-1)}} A_{vw} := s - \sum_{k=1}^K \sum_{w \in S_k^{(t-1)}} A_{vw}$
 and $\hat{p}(i, 0) := 1 - \sum_{k=1}^K \hat{p}(i, k)$ (ties are broken uniformly at random)
 $S_{i^*}^{(t)} \leftarrow S_{i^*}^{(t-1)} \cup \{v\}$
 end for
end for
 $\hat{I}_k \leftarrow S_k^{(\log n)}$ for all $1 \leq k \leq K$
Output: $\hat{I}_1, \dots, \hat{I}_K$

Algorithm 6 Spectral Decomposition (Algorithm 3 in [26])

Input: Observation matrix A , K
 $\hat{A} \leftarrow$ rank- K approximation of A
 $\tilde{p} \leftarrow \frac{\sum_{v,w \in \mathcal{I}} A_{vw}}{|\mathcal{S}|(|\mathcal{S}|-1)}$
for $x = 1$ to $\lfloor \log n_0 \rfloor$ **do**
 $Q_v^{(x)} \leftarrow \left\{ w \in \mathcal{S} : \|\hat{A}_w - \hat{A}_v\|^2 \leq x \frac{\tilde{p}}{100} \right\}$ for all $v \in \mathcal{S}$
 $T_l^{(x)} \leftarrow \emptyset$ for all $l \in [K]$
 for $k = 1$ to K **do**
 $v_k^* \leftarrow \arg \max_{v \in \mathcal{I}} \left| Q_v^{(x)} \setminus \bigcup_{i=1}^{k-1} T_i^{(x)} \right|$
 $T_k^{(x)} \leftarrow Q_{v_k^*}^{(x)} \setminus \bigcup_{i=1}^{k-1} T_i^{(x)}$
 $\xi_k^{(x)} \leftarrow \sum_{v \in T_k^{(x)}} \frac{\hat{A}_v}{|T_k^{(x)}|}$
 end for
 for $v \in \mathcal{S} \setminus \bigcup_{k=1}^K T_k^{(x)}$ **do**
 $k^* \leftarrow \arg \min_{1 \leq k \leq K} \|\hat{A}_v - \xi_k^{(x)}\|^2$
 $T_{k^*}^{(x)} \leftarrow T_{k^*}^{(x)} \cup \{v\}$
 end for
 $r_x \leftarrow \sum_{k=1}^K \sum_{v \in T_k^{(x)}} \|\hat{A}_v - \xi_k^{(x)}\|^2$
end for
 $x^* \leftarrow \arg \min_x r_x$
 $S_k \leftarrow T_k^{(x^*)}$ for all $k \in [K]$
Output: S_1, \dots, S_K

A.3.2 The ECB algorithm

The ECB algorithm presented in Algorithm 7. ECB achieves a regret scaling as $O(m \log(\bar{N}))$ for all p (ECB treats each user independently, and does not transfer the information gathered across users). The algorithm proceeds as follows.

(a)-(b) **Exploration and clustering phases.** (b) These phases are identical to those of EC-UCS. The algorithm outputs item cluster estimates $\hat{I}_1, \dots, \hat{I}_K$. We can show that with an exploration budget of $10m$, the spectral algorithm does not make any clustering errors w.p. at least $1 - \frac{1}{T}$.

(c) **Bandit phase.** The last phase consists in just applying m (one for each user) UCB1 algorithms [1] with the set of arms $1, \dots, K$. There, selecting arm k means recommending an item from \hat{I}_k , accounting for the no-repetition constraint (which is possible w.h.p. since $m^2 \geq T(\log T)^3$).

ECB calls the clustering algorithm presented in Algorithm 4, that first constructs an item adjacency matrix (using indirect edges from users), and then applies the spectral clustering algorithm, Algorithm 5, to output K item clusters. Note that Algorithm 5 further calls the spectral decomposition algorithm, shown in Algorithm 6.

We have the following performance guarantee on ECB (the proof is presented in Appendix J.1):

Theorem 7. *When $m^2 \geq T(\log T)^3$, the regret of ECB satisfies:*

$$R^{\text{ECB}}(T) = \mathcal{O} \left(\sum_{\ell=1}^L \beta_\ell m \sum_{k \neq k_\ell^*} \frac{\log(\bar{N})}{p_{k_\ell^* \ell} - p_{k \ell}} \right).$$

Algorithm 7 Explore-Cluster-Bandit

Input: T, K **1. Exploration for Item Clustering** $n_0 \leftarrow \min\{n, \lfloor m/(\log T)^2 \rfloor\}$ Sample a set \mathcal{S} of n_0 items.**for** $1 \leq t \leq 10m$ **do**Recommend items from \mathcal{S} randomly (when this is not possible due to the no-repetition constraint, recommend a random feasible item). Record the user responses in the dataset \mathcal{D} .**end for****2. Item Clustering**Run Algorithm 4 with input $\mathcal{S}, \mathcal{D}, K$ and output $\hat{I}_1, \dots, \hat{I}_K$.**3. Run UCB1 [1] for each user****for** $10m < t \leq T$ **do** $k^* \leftarrow \arg \max_{k \in [K]} \hat{\rho}_k^{u_t} + \sqrt{\frac{2 \ln(t)}{N_k^{u_t}(t)}}$, where $\hat{\rho}_k^{u_t}$ is the empirical average reward of user u_t for items in \hat{I}_k and $N_k^{u_t}(t)$ is the total number of samples of user u_t for items in \hat{I}_k .Recommend an item from \hat{I}_{k^*} randomly (when this is not possible due to the no-repetition constraint) recommend a random feasible item.**end for**

A.3.3 Numerical Experiment

Consider a system with $n = 2000$ items and $m = 5000$ users. The time horizon is $T = 800000$. The statistical parameters of $p_{k\ell}$ are given in Table 5. $\alpha_1 = \alpha_2 = \frac{1}{2}$ and $\beta_1 = \beta_2 = \frac{1}{2}$. With this parameter setting, for each $\ell = 1, 2$, $\mathcal{L}^\perp(\ell) = \emptyset$. From Theorems 6 and 7, we know that the regret of EC-UCS is $R^{\text{EC-UCS}}(T) = \mathcal{O}(m)$ whereas that of ECB is $R^{\text{ECB}}(T) = \mathcal{O}(m \log(\bar{N}))$. Hence, we expect that EC-UCS to outperform ECB. Figure 6 shows the regret evolution over time of EC-UCS algorithm and ECB algorithm after the item clustering phase. The curves are averaged over 10 instances. The rate at which the regret of EC-UCS increases is rapidly decreasing. This is not the case for that of the regret of ECB.

Next, we assess the regret of the two algorithms after T rounds as a function of T . We consider a system with $n = 3000$ items and $m = 5000$ users. $\alpha_1 = \alpha_2 = \frac{1}{2}$ and $\beta_1 = \beta_2 = \frac{1}{2}$. The statistical parameters of $p_{k\ell}$ are the same as in the previous system. Figure 7 presents the results. Here, the regrets are averages over 20 runs. The regret of ECB clearly increases with T , while the regret of EC-UCS does not seem to be sensitive to T . Overall, our results confirm our theoretical results, at least on simple examples.

	$k = 1$	$k = 2$
$\ell = 1$	0.2	0.8
$\ell = 2$	0.8	0.2

Table 5: The values of $(p_{k\ell})$.




Figure 6: Regret vs time after the item clustering phase, averaged over 10 runs of EC-UCS and ECB. $T = 800000$, $n = 2000$, $m = 5000$, $\alpha_1 = \alpha_2 = 0.5$, $\beta_1 = \beta_2 = 0.5$. One standard deviations are shown as the shaded areas.

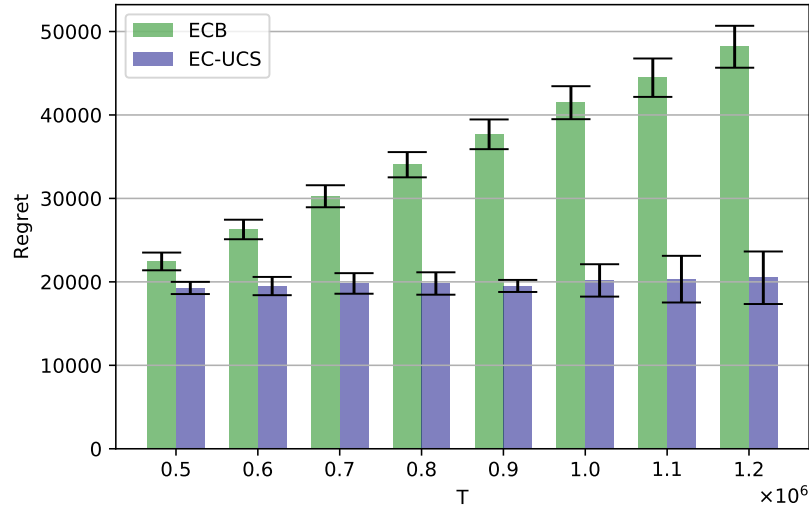


Figure 7: Regret at round T vs budget T after the item clustering phase, averaged over 20 runs of EC-UCS and ECB. $n = 3000$, $m = 5000$, $\alpha_1 = \alpha_2 = 0.5$, $\beta_1 = \beta_2 = 0.5$. One standard deviations are shown as the shaded areas.

A.4 Experimental set-up

The simulations were performed on a desktop computer with Intel Core i7-8700B 3.2 GHz CPU and 32 GB RAM.

B Preliminaries: Properties of the user arrival process

This section presents several preliminary results on the user arrival process, extensively used throughout the proofs of the main theorems. Here we also provide the proof of Lemma 1.

Lemma 2 (Chernoff-Hoeffding theorem). *Let X_1, \dots, X_n be i.i.d. Bernoulli random variables with mean ν . Then, for any $\delta > 0$,*

$$\begin{aligned} \mathbb{P}\left(\frac{1}{n}\sum_{i=1}^n X_i \geq \nu + \delta\right) &\leq \exp(-n\text{kl}(\nu + \delta, \nu)) \\ \mathbb{P}\left(\frac{1}{n}\sum_{i=1}^n X_i \leq \nu - \delta\right) &\leq \exp(-n\text{kl}(\nu - \delta, \nu)) \end{aligned}$$

Lemma 3 (Pinsker's inequality [31]). *For any $0 \leq p, q \leq 1$, $2(p - q)^2 \leq \text{kl}(p, q)$.*

Lemma 4. *For any $0 \leq p, q \leq 1$, $\text{kl}(p, q) \geq p \log \frac{p}{q} + (q - p)$.*

Proof. This follows from $(1 - p) \log \frac{1-p}{1-q} \geq q - p$. □

Proof of Lemma 1. This lemma is quoted below for convenience:

Lemma 1 (restated). *For every user $u \in \mathcal{U}$, we have*

$$\mathbb{E}\left[\max\left\{0, N_u(T) - \bar{N}\right\}\right] \leq \frac{1}{m(e-1)}$$

where

$$\bar{N} = \frac{4 \log m}{\log\left(\frac{m \log m}{T} + e\right)} + \frac{e^2 T}{m}.$$

Proof. Since u arrives with probability $\frac{1}{m}$ in each round, the probability that u arrives for more than $\bar{N} + x$ times in the T first round is,

$$\begin{aligned} \mathbb{P}(N_u(T) \geq \bar{N} + x) &\stackrel{(a)}{\leq} \exp\left(-T \text{kl}\left(\frac{\bar{N} + x}{T}, \frac{1}{m}\right)\right) \\ &\stackrel{(b)}{\leq} \exp\left(-(\bar{N} + x) \log\left(\frac{m(\bar{N} + x)}{T}\right) - \frac{T}{m} + (\bar{N} + x)\right) \\ &\stackrel{(c)}{\leq} \exp\left(-\frac{\bar{N} + x}{2} \log\left(\frac{m(\bar{N} + x)}{T}\right)\right) \\ &\stackrel{(d)}{\leq} \exp\left(-\frac{\bar{N}}{2} \log\left(\frac{m\bar{N}}{T}\right)\right) \exp(-x) \end{aligned}$$

where (a) follows from Lemma 2, (b) from Lemma 4, (c) is obtained from the fact that $\frac{m\bar{N}}{T} \geq e^2$, and (d) holds since $x \geq 0$ and $\log\left(\frac{m\bar{N}}{T}\right) \geq 2$. We deduce that:

$$\begin{aligned} \mathbb{E}\left[\max\left\{0, N_u(T) - \bar{N}\right\}\right] &= \sum_{x=1}^{\infty} \mathbb{P}(N_u(T) \geq \bar{N} + x) \\ &\leq \sum_{x=1}^{\infty} \exp\left(-\frac{\bar{N}}{2} \log\left(\frac{m\bar{N}}{T}\right)\right) \exp(-x) \\ &= \frac{\exp\left(-\frac{\bar{N}}{2} \log\left(\frac{m\bar{N}}{T}\right)\right)}{e-1}. \end{aligned} \tag{1}$$

To conclude the proof, we compute an upper bound of (1) for two cases: $\frac{T}{m} \geq \frac{\log m}{e}$ and $\frac{T}{m} \leq \frac{\log m}{e}$.

If $\frac{T}{m} \geq \frac{\log m}{e}$, then $\frac{\bar{N}}{2} > \log m$ and $\log\left(\frac{m\bar{N}}{T}\right) \geq 2$. Thus,

$$\exp\left(-\frac{\bar{N}}{2} \log\left(\frac{m\bar{N}}{T}\right)\right) < \frac{1}{m^2}. \quad (2)$$

We now consider the case where $\frac{T}{m} \leq \frac{\log m}{e}$. When we define $f(x) = \log\left(\frac{4x}{\log(x+e)}\right) - \frac{\log(x+e)}{2}$, one can easily check that $f(e) \geq 0$ and $f'(x) \geq 0$ for all $x \geq e$. Therefore, since $\frac{m \log m}{T} \geq e$, we can deduce that

$$\frac{2 \log m}{\log\left(\frac{m \log m}{T} + e\right)} \log\left(\frac{4m \log m}{T \log\left(\frac{m \log m}{T} + e\right)}\right) \geq \log m$$

which directly implies that

$$\exp\left(-\frac{\bar{N}}{2} \log\left(\frac{m\bar{N}}{T}\right)\right) < \frac{1}{m}. \quad (3)$$

Lemma 1 is obtained by combining (1), (2) and (3). \square

The following Lemma characterizes the lower tail of the number of user arrivals.

Lemma 5. *For every user $u \in \mathcal{U}$, we have*

$$\mathbb{P}\left(N_u(T) \leq \frac{T}{2m}\right) \leq \exp\left(-\frac{T}{2m} (1 - \log(2))\right).$$

Proof.

$$\begin{aligned} \mathbb{P}\left(N_u(T) \leq \frac{T}{2m}\right) &\stackrel{(a)}{\leq} \exp\left(-T \text{kl}\left(\frac{1}{2m}, \frac{1}{m}\right)\right) \\ &\stackrel{(b)}{\leq} \exp\left(-T \left(\frac{1}{2m} \log(1/2) + \left(\frac{1}{m} - \frac{1}{2m}\right)\right)\right) \\ &= \exp\left(-\frac{T}{2m} (1 - \log(2))\right), \end{aligned}$$

where (a) is from Lemma 2 and (b) is from Lemma 4. \square

The next lemma is instrumental in the performance analysis of the EC-UCS and ECB algorithms (for systems with clustered items and users).

Lemma 6. *With probability $1 - \frac{1}{T}$, at least $\frac{m}{2}$ users arrive at least two times among the first $10m$ arrivals.*

Proof. We denote by $N_u(10m)$ the number of times user u arrives in the $10m$ first arrivals. For any set $A \subset \mathcal{U}$, let $N(A, 10m)$ denote the total number of arrivals of users in A among the first $10m$ arrivals.

We write the probability that less than $m/2$ users arrive twice in the $10m$ first arrivals as:

$$\begin{aligned} \mathbb{P}\left[\sum_u \mathbf{1}_{\{N_u(10m) \geq 2\}} < m/2\right] &= \mathbb{P}\left[\sum_u \mathbf{1}_{\{N_u(10m) < 2\}} \geq m/2\right] \\ &\leq \mathbb{P}[\exists A : |A| = \frac{m}{2}, \forall u \in A, N_u(10m) < 2] \\ &\leq \mathbb{P}[\exists A : |A| = \frac{m}{2}, N(A, 10m) \leq \frac{m}{2}] \\ &\leq \binom{m}{m/2} \exp\left(-20m \left(\frac{1}{2} - \frac{1}{10}\right)^2\right) \\ &\leq (2e)^{m/2} \exp\left(-\frac{9}{5}m\right) \\ &\leq e^{-0.8m} \leq \frac{1}{T^2}. \end{aligned} \quad (4)$$

\square

C Justifying the regret definitions

In this section, we justify our definitions of regret for Models A, B and C. In these models, we define regret as if an Oracle policy would always be able to select for any user u an item from the best cluster for Models A and C for this user, or an ε -best item for Model B, even under the no-repetition constraint. In fact, the definition of the true regret should account for the no-repetition constraint and in turn for the fact that an Oracle policy may be obliged to select items that do not belong to the best cluster for the user u , because user u may arrive too often before the time horizon T or because the size of this cluster is too small (remember that when an item, selected for the first time, is randomly assigned to a cluster). We prove here that the difference between our notion of regret and the true regret is actually negligible when compared to any of the terms involved in our regret lower bounds, namely T/m and \bar{n} .

To establish this claim, recall the assumptions made on (n, m, T) . $T = o(nm)$, $m \geq n$, and $\log(m) = o(n)$ (and as a consequence, for any $\beta > 0$ independent of (n, m, T) , $m = o(e^{\beta n})$).

For illustrative purposes, we prove our claim for Model A (the same result holds for Models B and C). Let Z_u denote the (random) number of items in the best cluster for user u . It is easy to show that the difference between our notion of regret $R^\pi(T)$ and the true regret $R_{true}^\pi(T)$ satisfies:

$$|R^\pi(T) - R_{true}^\pi(T)| \leq \sum_{u \in \mathcal{U}} \mathbb{E}[\max\{N_u(T) - Z_u, 0\}].$$

In Model A, Z_u is the size of cluster \mathcal{I}_1 . The average size of \mathcal{I}_1 is $\mathbb{E}[Z_u] = \alpha_1 n$. Let $\varepsilon > 0$ be such that $\varepsilon < \alpha_1$ and $\frac{\varepsilon n m}{T} \geq e^2$. Using the same arguments as those leading to Lemma 1, we obtain the following concentration result:

$$\mathbb{E}\left[\max\left\{0, N_u(T) - \varepsilon n\right\}\right] \leq \frac{1}{(e-1)} \exp\left(-\frac{\varepsilon n}{2} \log\left(\frac{\varepsilon n m}{T}\right)\right).$$

In addition, we also have as a direct application of Chernoff-Hoeffding inequality presented in Lemma 2:

$$\mathbb{P}[Z_u \leq \varepsilon n] \leq \exp(-n \text{kl}(\varepsilon, \alpha_1)).$$

From the two above inequalities, we deduce that:

$$\begin{aligned} \mathbb{E}[\max\{N_u(T) - Z_u, 0\}] &\leq \frac{T}{m} \mathbb{P}[Z_u \leq \varepsilon n] + \mathbb{E}\left[\max\left\{0, N_u(T) - \varepsilon n\right\}\right] \\ &\leq \frac{T}{m} \exp(-n \text{kl}(\varepsilon, \alpha_1)) + \exp\left(-\frac{\varepsilon n}{2} \log\left(\frac{\varepsilon n m}{T}\right)\right). \end{aligned}$$

We conclude that:

$$|R^\pi(T) - R_{true}^\pi(T)| \leq \underbrace{T \exp(-n \text{kl}(\varepsilon, \alpha_1))}_{=A(n, m, T)} + \underbrace{m \exp\left(-\frac{\varepsilon n}{2} \log\left(\frac{\varepsilon n m}{T}\right)\right)}_{=B(n, m, T)}.$$

Next we verify that $\max\{A(n, m, T), B(n, m, T)\} = o(\min\{\frac{T}{m}, \bar{n}\})$. This will be enough to justify our definition of regret, since our regret lower bounds are all larger than $\min\{\frac{T}{m}, \bar{n}\}$.

(i) Let us check that $\max\{A(n, m, T), B(n, m, T)\} = o(\frac{T}{m})$. Indeed, for $A(n, m, T)$, we have $m = o(e^{n \text{kl}(\varepsilon, \alpha_1)})$; as for $B(n, m, T)$, we have $m^2/T = o(e^{\varepsilon n})$ which results from $m = o(e^{\varepsilon n/2})$.

(ii) Let us check that $\max\{A(n, m, T), B(n, m, T)\} = o(\bar{n})$. We consider the three regimes for \bar{n} :

1. When $T = \omega(m \log(m))$. Then $\bar{n} = T/m(1 + o(1))$ and we conclude as in (i).
2. When $T = \Theta(m \log(m))$. To simplify, we just prove the statement for $T = cm \log(m)$. Then $\bar{n} = \log(m)(d_c + o(1))$ and $A(n, m, T) = m \log(m) \exp(-n \text{kl}(\varepsilon, \alpha_1))$. We hence conclude that $A(n, m, T) = \bar{n}$ since $m = o(e^{n \text{kl}(\varepsilon, \alpha_1)})$. Now $B(n, m, T) = m \exp(-\frac{\varepsilon n}{2} \log(\frac{\varepsilon n}{c \log(m)}))$, and $B(n, m, T) = o(\bar{n})$ is a consequence of $m = o(e^{\frac{\varepsilon n}{2}})$.
3. When $T = o(m \log(m))$. Then $\bar{n} = \frac{\log(m)}{\log(m \log(m)/T)}(1 + o(1))$. $A(n, m, T) = o(\bar{n})$ is equivalent to $X = o(e^{n \text{kl}(\varepsilon, \alpha_1)})$ where $X = \frac{T}{\log(m)} \log(\frac{m \log(m)}{T})$. Now $X = o(m)$ since

$T = o(m \log(m))$, and thus $A(n, m, T) = o(\bar{n})$. Finally, $B(n, m, T) = o(\bar{n})$ is equivalent to $Y = o(\exp(\frac{\epsilon n}{2} \log(\epsilon n m / T)))$ where $Y = \frac{m \log(m \log(m) / T)}{\log(m)}$. Since $T = o(nm)$, this would be implied by $Y = o(\exp(\frac{\epsilon n}{2}))$. However, since $T \geq 1$, $Y \leq \frac{m \log(m \log(m))}{\log(m)} = m(1 + o(1))$. We conclude that $B(n, m, T) = o(\bar{n})$ since $m = o(e^{\frac{\epsilon n}{2}})$.

D Fundamental limits for Model A: Proof of Theorem 1

Proof of $R^\pi(T) \geq R_{\text{ic}}(T)$: Let $\pi \in \Pi$. Assume that the success probabilities p_k 's are known. Further simplify the problem by relaxing the no-repetition constraint: instead, we just impose that any item cannot be selected more than m times. The algorithm π has an expected regret larger than an optimal algorithm for the problem where p_k 's are known, and the no-repetition constraint is relaxed as explained above. Denote by τ this optimal algorithm. Next we establish a regret lower bound for τ . To this aim, we first consider that the cluster ids of the items are drawn before the first round. That way, all the n items belong to a cluster even before the first round. Let J_1, \dots, J_K denote the sizes of the various clusters of items. Hence, we work conditioning on the cluster ids of the items. Now, define $\mathbb{E}^\tau[N_k]$ as the expected number of times an item of cluster \mathcal{I}_k is selected under τ (of course, $N_k \leq m$) until round T . More precisely, denote by $(n_1, n_2, \dots, n_{J_k})$ the random variables representing the numbers of times the first, the second, the third, etc. items in \mathcal{I}_k are selected under τ . Then by definition: $\mathbb{E}^\tau[N_k] = \mathbb{E}^\tau[\sum_{i=1}^{J_k} n_i]/J_k$. We prove:

Lemma 7. *We have for any $m \geq c/\Delta_2^2$ and for any $k \neq 1$: $\frac{\phi(k, m, p)}{m} \leq \frac{\mathbb{E}^\tau[N_k]}{\mathbb{E}^\tau[N_1]} \leq 1$.*

The constant c in the above lemma is the same as that in Theorem 1. It is chosen such that if $m \geq c/\Delta_2^2$, then, for any k , $\phi(k, m, p) \leq m$ and thus, the first inequality of Lemma 7 makes sense. Remember that $\phi(k, m, p)$ scales as $1/\Delta_k^2$ (refer to the remark above Theorem 1), and hence such a choice for c is possible.

Lemma 7 is proved at the end of this section. Assume for now that it holds. We complete the proof by deriving a lower bound of the optimal algorithm τ , $R^\tau(T)$. When the item clusters are fixed, the expected conditional regret of τ is: $R^\tau(T|\mathcal{I}_1, \dots, \mathcal{I}_K) = \sum_{k \neq 1} J_k \mathbb{E}^\tau[N_k] \Delta_k$. Hence we have:

$$\begin{aligned} R^\tau(T|\mathcal{I}_1, \dots, \mathcal{I}_K) &= T \frac{\sum_{k \neq 1} J_k \mathbb{E}^\tau[N_k] \Delta_k}{\sum_k J_k \mathbb{E}^\tau[N_k]} \stackrel{(a)}{\geq} \frac{T}{m} \frac{\sum_{k \neq 1} J_k \mathbb{E}^\tau[N_1] \phi(k, m, p) \Delta_k}{\sum_k J_k \mathbb{E}^\tau[N_1]} \\ &= \frac{T}{m} \sum_{k \neq 1} \frac{J_k}{n} \phi(k, m, p) \Delta_k, \end{aligned}$$

where (a) stems from Lemma 7 (in the numerator, we use $\mathbb{E}^\tau[N_k] \geq \mathbb{E}^\tau[N_1] \frac{\phi(k, m, p)}{m}$, and in the denominator $\mathbb{E}^\tau[N_k] \leq \mathbb{E}^\tau[N_1]$). Taking the expectation of the above inequality (noting that $\mathbb{E}[J_k/n] = \alpha_k$), we conclude that: $R^\pi(T) \geq R^\tau(T) \geq \frac{T}{m} \sum_{k \neq 1} \alpha_k \phi(k, m, p) \Delta_k$.

Proof of $R^\pi(T) \geq R_{\text{nr}}(T)$: Observe that under the no-repetition constraint, the number of items that π will select is greater than $\max_u N_u(T)$. Now when an item is selected for the first time, by assumption, it belongs to the sub-optimal cluster \mathcal{I}_k with probability α_k , in which case this initial selection induces an expected regret $\Delta_k = p_1 - p_k$. Hence $R^\pi(T) \geq \mathbb{E}[\max_u N_u(T)] \sum_{k \neq 1} \alpha_k \Delta_k = \bar{n} \sum_{k \neq 1} \alpha_k \Delta_k$. \square

Proof of $R^\pi(T) \gtrsim R_{\text{sp}}(T)$: To prove this asymptotic lower bound, we consider a simpler problem: the algorithm knows the item clusters $(\mathcal{I}_k)_{k \in [K]}$. Then, the problem reduces to a K -armed Bernoulli bandit problem with unknown parameters $(p_k)_{k \in [K]}$.

The proof then proceeds using a classical change-of-measure argument as in [19]. We present this argument for completeness. Assume that the algorithm π is uniformly good. Pick any $k \in [K]$ s.t. $k \neq 1$. Let $p = (p_k)_{k \in [K]}$ be original parameters and let $(p'_k)_{k \in [K]}$ be perturbed parameters where $\forall k' \neq k$, $p'_{k'} = p_{k'}$ and $p'_k = p_1 + \varepsilon$ with some constant $\varepsilon > 0$. We denote $N_k(T) := \sum_{t=1}^T \mathbb{1}\{i_t^\pi \in \mathcal{I}_k\}$. We use $\mathbb{E}^\pi[\cdot]$ (or $\mathbb{E}[\cdot]$) and $\mathbb{E}_{p'}[\cdot]$ to denote the expectation under the original model and under the perturbed model, respectively.

Let L_T be the log-likelihood defined as:

$$L_T := \sum_{t=1}^T \mathbb{1}\{i_t^\pi \in \mathcal{I}_k\} \left(\mathbb{1}\{X_{i_t^\pi u_t} = 1\} \log\left(\frac{p_k}{p_1 + \varepsilon}\right) + \mathbb{1}\{X_{i_t^\pi u_t} = 0\} \log\left(\frac{1 - p_k}{1 - (p_1 + \varepsilon)}\right) \right).$$

Taking the expectation under p , we have

$$\begin{aligned} \mathbb{E}_p[L_T] &= \mathbb{E}_p[N_k(T)]\text{kl}(p_k, p_1 + \varepsilon) \stackrel{(a)}{\geq} \text{kl}\left(\frac{\mathbb{E}_p[N_k(T)]}{T}, \frac{\mathbb{E}_{p'}[N_k(T)]}{T}\right) \\ &\stackrel{(b)}{\geq} \left(1 - \frac{\mathbb{E}_p[N_k(T)]}{T}\right) \log\left(\frac{T}{T - \mathbb{E}_{p'}[N_k(T)]}\right) - \log 2. \end{aligned}$$

where (a) stems from the *data processing inequality*, see [8], and (b) is from the fact that for all $(x, y) \in [0, 1]^2$, $\text{kl}(x, y) \geq (1-x) \log \frac{1}{1-y} - \log 2$. By the uniform goodness assumption, with some constant $C > 0$, we have:

$$\mathbb{E}_p[N_k(T)] \lesssim C\sqrt{T} \quad \text{and} \quad T - \mathbb{E}_{p'}[N_k(T)] \lesssim C\sqrt{T}.$$

Hence:

$$\frac{1}{\log(T)} \log \frac{T}{T - \mathbb{E}_{p'}[N_k(T)]} \gtrsim \frac{1}{\log(T)} \log \frac{T}{C\sqrt{T}} \gtrsim \frac{1}{2}.$$

The inequality $\mathbb{E}_p[N_k(T)]\text{kl}(p_k, p_1 + \varepsilon) \gtrsim \frac{1}{2}$ holds for any $\varepsilon > 0$. Therefore, we have:

$$\mathbb{E}_p[N_k(T)]\text{kl}(p_k, p_1) \gtrsim \frac{1}{2}.$$

Thus, we get the regret lower bound:

$$R^\pi(T) = \sum_{k \neq 1} \Delta_k \mathbb{E}_p[N_k(T)] \gtrsim \left(\sum_{k \neq 1} \frac{\Delta_k}{2\text{kl}(p_k, p_1)} \right) \log(T).$$

This concludes the proof of Theorem 1. \square

D.1 Proof of Lemma 7

To establish the lemma, we build, from the optimal algorithm τ , an algorithm ζ that can be applied to a 2-armed bandit problem with known expected rewards p_1 and p_k . We then provide a connection between the regret of ζ in the 2-armed bandit problem and $\mathbb{E}^\tau[N_k]$. We conclude the proof by establishing a regret lower bound for ζ , using similar techniques as in [5].

2-armed bandit problem with known rewards and the algorithm ζ . Consider a 2-armed bandit problem with Bernoulli arms 1 and k of means p_1 and p_k . The means are known but the arm with the highest mean is unknown. That is to say that the expected reward of arm 1 can be either p_1 and p_k . For this bandit problem, we build an algorithm, denoted by ζ , based on the algorithm τ .

1. Pick i_1 and i_k uniformly at random in the clusters \mathcal{I}_1 and \mathcal{I}_k . We run τ for T rounds. When τ selects item i_1 (resp. i_k), then ζ also selects arm 1 (resp. k).
2. We repeat the above Step 1, to determine the arm selections made by ζ . At the beginning of the successive episodes of T rounds, the items i_1 and i_k are again chosen uniformly at random in the clusters \mathcal{I}_1 and \mathcal{I}_k , independently of the choices made in earlier episodes.

Regret of ζ and its connection to $\mathbb{E}^\tau[N_k]$. Consider an episode of T rounds for τ . Let i_k denote the item selected from \mathcal{I}_k in the design of ζ , the expected regret accumulated by ζ in this episode is $\Delta_k \mathbb{E}^\tau[n_{i_k}]$, where n_{i_k} is the number of times τ selects i_k in the episode. Since i_k is chosen uniformly at random, and by definition of $\mathbb{E}^\tau[N_k]$, we actually have $\mathbb{E}^\tau[n_{i_k}] = \mathbb{E}^\tau[N_k]$, which connects the regret of ζ and $\mathbb{E}^\tau[N_k]$.

Next, assume that we stop the algorithm ζ after κ episodes of T rounds, where κ is the first episode where ζ has made more than m selections. κ is a random variable, and Wald's first lemma implies that the expected regret R^ζ accumulated by ζ before we stop playing is:

$$R^\zeta = \Delta_k \mathbb{E}[\kappa] \mathbb{E}^\tau[N_k].$$

Since $\max\{\mathbb{E}^\tau[N_1], \mathbb{E}^\tau[N_k]\} \leq m$, we have $\mathbb{E}[\kappa] \mathbb{E}^\tau[N_1] \leq 2m$. Hence:

$$R^\zeta = \Delta_k \mathbb{E}[\kappa] \mathbb{E}^\tau[N_k] \leq 2m \Delta_k \frac{\mathbb{E}^\tau[N_k]}{\mathbb{E}^\tau[N_1]}. \quad (5)$$

By construction, R^ζ corresponds to the expected regret of our algorithm ζ for a number of rounds larger than m in the 2-arm bandit problem with known average rewards. The proof of Lemma 7 is completed by establishing a lower bound on R^ζ .

Regret lower bound of R^ζ . We prove the following lemma, which combined with (5) yields Lemma 7.

Lemma 8. *We have: $R^\zeta \geq 2\Delta_k\phi(k, m, p)$.*

Proof of Lemma 8. The proof is similar to that of Theorem 6 in [5]. The following lemma by [31, 5] is the essential ingredient of the proof:

Lemma 9. *Let P_0 and P_1 be two probability measures on a measurable space (Ω, \mathcal{F}) , with P_0 is absolutely continuous with respect to P_1 . Then, for any \mathcal{F} -measurable function $\Psi : \Omega \rightarrow \{0, 1\}$, we have:*

$$\mathbb{P}_{P_0}(\Psi(\omega) = 1) + \mathbb{P}_{P_1}(\Psi(\omega) = 0) \geq \frac{1}{2} \exp(-\text{KL}(P_0, P_1)).$$

Consider $a_t \in \{1, k\}$, defined as the t -th arm selection made by ζ . This selection happens in the m_t -th round of an episode of T rounds for the algorithm τ . At the beginning of this episode, in the design of ζ , items i_1 and i_k have been selected, and in this m_t -th round, τ selects either i_1 or i_k . The decisions made under τ in this episode depend on the observations made in this episode only, and this remark holds for ζ as well. We define by \mathcal{F} the σ -algebra generated by the observations made before the m_t -th round in the episode. To build \mathcal{F} , we assume that each time i_1 or i_k is selected, then a sample of the reward of both items is observed.

With the above definitions, we have $\{a_t = 1\} \in \mathcal{F}$, and $\{a_t = k\} \in \mathcal{F}$. Next consider the following two probability measures on \mathcal{F} : P_0 corresponds to the observations made in the original model (with the true item clusters), and P_1 to the observations made assuming that items i_1 and i_k are swapped: the average reward of i_1 is p_k and that of i_k is p_1 . P_0 and P_1 differ only when it comes to observations made in rounds where items i_1 and i_k are selected. At round m_t , we know that we have had at most $t - 1$ such rounds. We deduce that:

$$\text{KL}(P_0, P_1) \leq (t - 1)(\text{kl}(p_1, p_k) + \text{kl}(p_k, p_1)).$$

Applying Lemma 9, we get:

$$\mathbb{P}_{P_0}(a_t = k) + \mathbb{P}_{P_1}(a_t = 1) \geq \frac{1}{2} \exp(-(t - 1)(\text{kl}(p_1, p_k) + \text{kl}(p_k, p_1))) = \frac{1}{2} e^{-(t-1)\gamma(p_1, p_k)}. \quad (6)$$

Observe that $\mathbb{P}_{P_0}(a_t = k) + \mathbb{P}_{P_1}(a_t = 1)$ is the expected instantaneous regret of ζ for its t -th arm selection. Hence, we have:

$$\begin{aligned} R^\zeta &\geq \frac{\Delta_k}{4} \sum_{t=1}^m e^{-(t-1)\gamma(p_1, p_k)} \\ &= \frac{\Delta_k}{4} \frac{1 - e^{-m\gamma(p_1, p_k)}}{1 - e^{-\gamma(p_1, p_k)}} \\ &= 2\Delta_k\phi(k, m, p), \end{aligned}$$

where the first inequality stems from the fact that R^ζ is the regret accumulated over more than m rounds, and the second inequality is from (6). \square

E Fundamental limits for Model B: Proof of Theorem 2

We apply the same strategy as in the case of clustered items. Let π be an arbitrary algorithm.

Proof of $R^\pi(T) \geq R_{\text{nr}}(T)$: Using the same reasoning as for Model A, the algorithm needs to sample at least $\max_u N_u(T)$, and a new item generated a satisficing regret equal to $\int_0^{\mu_{1-\varepsilon}} (\mu_{1-\varepsilon} - \mu)\zeta(\mu)d\mu$. Hence, we get:

$$R^\pi(T) \geq \mathbb{E}[\max_u N_u(T)] \int_0^{\mu_{1-\varepsilon}} (\mu_{1-\varepsilon} - \mu)\zeta(\mu)d\mu = R_{\text{nr}}(T).$$

Proof of $R^\pi(T) \geq R_i(T)$: For the term $R_i(T)$, assume that ζ is known. With this knowledge, we denote τ an optimal algorithm. We denote by $\mathbb{E}^\tau[N_\mu]$ the expected number of rounds an item with success rate μ is selected under τ . Formally, the algorithm τ induced the two following random counting measures on the interval $[0, 1]$: (i) Ξ counts the number of the items whose parameter is $\mu \in [0, 1]$ seen by the algorithm τ ('seen' means selected at least once), (ii) Υ counts the number of times the algorithm τ selects items whose parameter is $\mu \in [0, 1]$. Now the intensity measures [16] γ and ω of Ξ and Υ are absolutely continuous w.r.t. ζ , and in addition, ω is absolutely continuous w.r.t. γ . Denote by d_γ and d_ω the densities of γ and ω w.r.t. ζ . Then, $\mathbb{E}^\tau[N_\mu]$ is defined by $d_\omega(\mu)/d_\gamma(\mu)$. In the remark at the end of this proof, we make these definitions and the expression of the regret of τ explicit in the case where ζ is constant over intervals of $[0, 1]$. Our proof could actually directly use a sequence of such discretizations, and then concludes by monotone limits.

Now the regret of an algorithm π satisfies:

$$\begin{aligned} R_\varepsilon^\pi(T) &\geq T \frac{\int_0^{\mu_{1-\varepsilon}} \mathbb{E}^\tau[N_\mu](\mu_{1-\varepsilon} - \mu)\zeta(\mu)d\mu}{\int_0^1 \mathbb{E}^\tau[N_\mu]\zeta(\mu)d\mu} \\ &\geq T \frac{\int_0^{\mu_{1-\varepsilon}} \mathbb{E}^\tau[N_\mu](\mu_{1-\varepsilon} - \mu)\zeta(\mu)d\mu}{\min\{1, (1+C)\varepsilon\}m + \int_0^{\mu_{1-\varepsilon}-\varepsilon} \mathbb{E}^\tau[N_\mu]\zeta(\mu)d\mu} \\ &\geq \frac{T}{m} \frac{\frac{(1-\varepsilon)^2}{2C} \left(1 - \frac{\varepsilon C}{1-\varepsilon}\right)^2}{\min\{1, (1+C)\varepsilon\} + 1/m}, \end{aligned}$$

where we use the fact that $\varepsilon \geq \frac{1}{\sqrt{m}}$ and $1 \leq \mathbb{E}^\tau[N_\mu] \leq m$ for all μ .

To complete the proof of the theorem, we just establish the following inequality:

$$T \frac{\int_0^{\mu_{1-\varepsilon}} \mathbb{E}^\tau[N_\mu](\mu_{1-\varepsilon} - \mu)\zeta(\mu)d\mu}{\min\{1, (1+C)\varepsilon\}m + \int_0^{\mu_{1-\varepsilon}-\varepsilon} \mathbb{E}^\tau[N_\mu]\zeta(\mu)d\mu} \geq T \frac{\frac{(1-\varepsilon)^2}{2C} \left(1 - \frac{\varepsilon C}{1-\varepsilon}\right)^2}{\min\{1, (1+C)\varepsilon\}m + 1}.$$

Let $\psi_i = \int_{\mu_{1-\varepsilon} - (i+1)\varepsilon}^{\mu_{1-\varepsilon} - i\varepsilon} \mathbb{E}^\tau[N_\mu]\zeta(\mu)d\mu$. Then,

$$\psi_i \in \left[\int_{\mu_{1-\varepsilon} - (i+1)\varepsilon}^{\mu_{1-\varepsilon} - i\varepsilon} \zeta(\mu)d\mu, \infty \right),$$

since $\mathbb{E}^\tau[N_\mu] \geq 1$ for all μ . We have:

$$T \frac{\int_0^{\mu_{1-\varepsilon}} \mathbb{E}^\tau[N_\mu](\mu_{1-\varepsilon} - \mu)\zeta(\mu)d\mu}{\min\{1, (1+C)\varepsilon\}m + \int_0^{\mu_{1-\varepsilon}-\varepsilon} \mathbb{E}^\tau[N_\mu]\zeta(\mu)d\mu} \geq T \frac{\sum_{i=1}^{\lfloor \mu_{1-\varepsilon}/\varepsilon \rfloor} i\varepsilon\psi_i}{\min\{1, (1+C)\varepsilon\}m + \sum_{i=1}^{\lfloor \mu_{1-\varepsilon}/\varepsilon \rfloor} \psi_i}. \quad (7)$$

As the derivate of $\frac{a+bx}{c+dx}$ is $\frac{bc-da}{(c+dx)^2}$, $\frac{\sum_{i=1}^{\lfloor \mu_{1-\varepsilon}/\varepsilon \rfloor} i\varepsilon\psi_i}{\min\{1, (1+C)\varepsilon\}m + \sum_{i=1}^{\lfloor \mu_{1-\varepsilon}/\varepsilon \rfloor} \psi_i}$ is either an increasing function or a decreasing function of ψ_i when all other ψ_j 's are fixed. Therefore, the r.h.s. of (7) can be optimized only when the ψ_i 's are at extreme points, either $\int_{\mu_{1-\varepsilon} - (i+1)\varepsilon}^{\mu_{1-\varepsilon} - i\varepsilon} \zeta(\mu)d\mu$ or ∞ .

When $\psi_i = \int_{\mu_{1-\varepsilon} - (i+1)\varepsilon}^{\mu_{1-\varepsilon} - i\varepsilon} \zeta(\mu) d\mu$ for all i ,

$$\begin{aligned} T \frac{\sum_{i=1}^{\lfloor \mu_{1-\varepsilon}/\varepsilon \rfloor} i\varepsilon\psi_i}{\min\{1, (1+C)\varepsilon\}m + \sum_{i=1}^{\lfloor \mu_{1-\varepsilon}/\varepsilon \rfloor} \psi_i} &\geq T \frac{\sum_{i=1}^{\lfloor \mu_{1-\varepsilon}/\varepsilon \rfloor} i\varepsilon \int_{\mu_{1-\varepsilon} - (i+1)\varepsilon}^{\mu_{1-\varepsilon} - i\varepsilon} \zeta(\mu) d\mu}{\min\{1, (1+C)\varepsilon\}m + 1} \\ &\geq T \frac{\frac{(1-\varepsilon)^2}{2C} \left(1 - \frac{\varepsilon C}{1-\varepsilon}\right)^2}{\min\{1, (1+C)\varepsilon\}m + 1}. \end{aligned}$$

When $\psi_i = \infty$ for some $i \geq 1$, we have

$$T \frac{\sum_{i=1}^{\lfloor \mu_{1-\varepsilon}/\varepsilon \rfloor} i\varepsilon\psi_i}{\min\{1, (1+C)\varepsilon\}m + \sum_{i=1}^{\lfloor \mu_{1-\varepsilon}/\varepsilon \rfloor} \psi_i} \geq \varepsilon T.$$

Thus, we have

$$\begin{aligned} T \frac{\int_0^{\mu_{1-\varepsilon}} \mathbb{E}^\tau[N_\mu](\mu_{1-\varepsilon} - \mu)\zeta(\mu) d\mu}{\min\{1, (1+C)\varepsilon\}m + \int_0^{\mu_{1-\varepsilon} - \varepsilon} \mathbb{E}^\tau[N_\mu]\zeta(\mu) d\mu} &\geq T \min \left\{ \frac{\frac{(1-\varepsilon)^2}{2C} \left(1 - \frac{\varepsilon C}{1-\varepsilon}\right)^2}{\min\{1, (1+C)\varepsilon\}m + 1}, \varepsilon \right\} \\ &= T \frac{\frac{(1-\varepsilon)^2}{2C} \left(1 - \frac{\varepsilon C}{1-\varepsilon}\right)^2}{\min\{1, (1+C)\varepsilon\}m + 1}, \end{aligned}$$

where the last equation stems from the assumption $m \geq \frac{c}{\varepsilon^2}$. This concludes the proof. \square

Remark. Assume that there exists a finite set of non-overlapping intervals of $[0, 1]$ and covering $[0, 1]$ such that the density of ζ is constant over each of these intervals. We denote by ζ_i the probability that when a new item is selected, its parameter lies in the i -th interval. Further assume that the satisficing regret of an item with parameter in the i -th interval does depend on the parameter, and is equal to Δ_i . Under the algorithm τ , let $\mathbb{E}^\tau[P_i]$ denote the expected number of items seen (selected at least once) by τ and whose parameter is in the i -th interval, and let $\mathbb{E}^\tau[N_i]$ the expected number of times τ selects an item with parameter in the i -th interval. Then, the equivalent of $\mathbb{E}^\tau[N_\mu]$ is η_i , the expected number of times an item with parameter in the i -th interval is selected. η_i is defined as $\eta_i = \mathbb{E}^\tau[N_i]/\mathbb{E}^\tau[P_i]$. Observe that since the item parameters of newly selected items are i.i.d. with distribution ζ_i , $\mathbb{E}^\tau[P_i]$ is proportional to ζ_i . This is just a consequence of the general Wald lemma: indeed, if $X_{k,i}$ is the binary r.v. indicating whether the k -th item seen by the algorithm has a parameter in the i -th interval, and if κ denotes the random number of items seen by the algorithm within the time horizon T , then Wald's equation holds if $\mathbb{E}^\tau[X_{k,i} 1_{\kappa \geq k}] = \mathbb{E}[X_{k,i}] \mathbb{P}^\tau[\kappa \geq k]$. This is true in our case since the event $\{\kappa \geq k\}$ corresponds to the fact that τ decides to sample the k -th item, and this decision is solely based on observations made on the $(k-1)$ first items. Finally, the regret of τ is:

$$R^\tau(T) = \sum_i \Delta_i \mathbb{E}^\tau[N_i] = T \frac{\sum_i \Delta_i \eta_i \zeta_i}{\sum_i \eta_i \zeta_i}.$$

In the above formula we used the fact that $T = \sum_i \mathbb{E}^\tau[N_i]$. Note that we obtained a discrete version of $T \frac{\int_0^{\mu_{1-\varepsilon}} \mathbb{E}^\tau[N_\mu](\mu_{1-\varepsilon} - \mu)\zeta(\mu) d\mu}{\int_0^1 \mathbb{E}^\tau[N_\mu]\zeta(\mu) d\mu}$.

F Fundamental limits for Model C: Proof of Theorem 3

We start this section by illustrating the various terms involved in the regret lower bound in Theorem 3. We then prove the theorem.

F.1 Examples

Let π be a uniformly good algorithm. Then under the conditions of Theorem 3, we have: $R^\pi(T) \geq \max\{R_{\text{nr}}(T), R_{\text{ic}}(T), R_{\text{uc}}(T)\}$ and $R^\pi(T) \gtrsim R'_{\text{uc}}(T)$. We exemplify the scalings of these terms below, with a particular emphasis on those due to the need of learning user clusters, $R_{\text{uc}}(T)$ and $R'_{\text{uc}}(T)$.

Case 1. We consider the case $L = 2$ and $K = 2$ with a following parameter set in Table 6.

	$k = 1$	$k = 2$
$\ell = 1$	0.8	0.6
$\ell = 2$	0.8	0.9

Table 6: Values of $(p_{k\ell})$.

For this parameter, $\mathcal{L}_\perp = \{(1, 2)\}$ and $\mathcal{L}^\perp(1) = \{2\}$. We have: $R_{\text{nr}} = \Omega(\bar{n})$, $R_{\text{ic}} = \Omega(\frac{T}{m})$ and $R_{\text{uc}} = \Omega(m)$. Furthermore, when $T = \omega(m)$,

$$R^\pi(T) \gtrsim \frac{\beta_1(0.8 - 0.6)}{\text{kl}(0.6, 0.9)} m \log\left(\frac{T}{m}\right) = \Omega\left(m \log\left(\frac{T}{m}\right)\right).$$

Case 2. We consider the case $L = 2$ and $K = 2$ with a following parameter set in Table 7.

	$k = 1$	$k = 2$
$\ell = 1$	0.8	0.6
$\ell = 2$	0.8	0.7

Table 7: Values of $(p_{k\ell})$.

For this parameter, $\mathcal{L}_\perp = \emptyset$ and $\forall \ell \in [L]$, $\mathcal{L}^\perp(\ell) = \emptyset$. We have: $R_{\text{nr}} = \Omega(\bar{n})$, $R_{\text{ic}} = \Omega(\frac{T}{m})$, $R_{\text{uc}} = 0$ and $c(\beta, p) = 0$.

Case 3. We consider the case $L = 2$ and $K = 2$ with a following parameter set in Table 8.

	$k = 1$	$k = 2$
$\ell = 1$	0.8	0.9
$\ell = 2$	0.8	0.85

Table 8: Values of $(p_{k\ell})$.

In this case, $\mathcal{L}_\perp = \emptyset$ and $\forall \ell \in [L]$, $\mathcal{L}^\perp(\ell) = \emptyset$. We have: $R_{\text{nr}} = \Omega(\bar{n})$, $R_{\text{ic}} = \Omega(\frac{T}{m})$, $R_{\text{uc}} = 0$ and $c(\beta, p) = 0$.

Case 4. We consider the case $L = 2$ and $K = 3$ with a following parameter set in Table 9.

	$k = 1$	$k = 2$	$k = 3$
$\ell = 1$	0.9	0.8	0.7
$\ell = 2$	0.7	0.8	0.9

Table 9: Values of $(p_{k\ell})$.

In this case, $\mathcal{L}_\perp = \{(1, 2)\}$ and $\forall \ell \in [L]$, $\mathcal{L}^\perp(\ell) = \emptyset$. We have: $R_{\text{nr}} = \Omega(\bar{n})$, $R_{\text{ic}} = \Omega(\frac{T}{m})$, $R_{\text{uc}} = \Omega(m)$ and $c(\beta, p) = 0$.

E.2 Proof

Proof of $R^\pi(T) \geq \max\{R_{\text{nr}}(T), R_{\text{ic}}(T)\}$. The proof is the same as in that of Theorem 1.

Proof of $R^\pi(T) \geq R_{\text{uc}}(T)$. We first give a simple proof that learning user clusters induces a regret scaling as m , i.e., that $R^\pi(T) = \Omega(m)$. When a user first arrives, we do not know her cluster, and hence we have to recommend an item from a cluster picked randomly. This selection induces an average regret at least equal to $\min_\ell \{\beta_\ell\} \Delta$ when $\mathcal{L}_\perp \neq \emptyset$. Since the number of users that arrive at least once is in expectation larger than $m(1 - (1 - \frac{1}{m})^T) \geq m(1 - e^{-T/m})$, we get that: $R^\pi(T) \geq \min_\ell \{\beta_\ell\} \Delta m(1 - e^{-T/m})$.

To get the right constant in the regret lower bound, we need to develop a more involved argument. Assume that the item clusters \mathcal{I}_k 's and the success rates p are known. With this knowledge, we denote by τ an optimal algorithm. We derive a regret lower bound for τ . Define $N_k^u := \sum_{t=1}^T \mathbb{1}_{\{i_t^\tau \in \mathcal{I}_k, u_t = u\}}$ as the number of times user u is presented items in cluster \mathcal{I}_k (under π). Fix user clusters $\mathcal{U}_1, \dots, \mathcal{U}_L$. Similar to the proof of Lemma 7, we will prove that:

Lemma 10. *We have for all $T \geq 2m$, for any $\ell \in [L]$, for any $k \in [K]$ such that $\exists r \in \mathcal{R}_\ell$ such that $k = k_r^*$: $\frac{\psi(\ell, k, T, m, p)}{(T/m)} \leq \frac{\mathbb{E}^\tau[N_k^u | u \in \mathcal{U}_\ell]}{\mathbb{E}^\tau[N_{k_\ell^*}^u | u \in \mathcal{U}_\ell]} \leq 1$, where $\psi(\ell, k, T, m, p) = \frac{1 - e^{-\frac{T}{m} \gamma(p_{k_\ell^*}^\ell, p_{k\ell}})}{8 \left(1 - e^{-\gamma(p_{k_\ell^*}^\ell, p_{k\ell})}\right)}$.*

For fixed user clusters, the expected conditional regret $R^\tau(T | \mathcal{U}_1, \dots, \mathcal{U}_L)$ is: $R^\tau(T | \mathcal{U}_1, \dots, \mathcal{U}_L) = \sum_{\ell \in [L]} |\mathcal{U}_\ell| \sum_{k \in \mathcal{R}_\ell} \mathbb{E}^\tau[N_k^u | u \in \mathcal{U}_\ell] \Delta_{k\ell}$. Therefore, we have:

$$\begin{aligned} R^\tau(T | \mathcal{U}_1, \dots, \mathcal{U}_L) &= \sum_{\ell \in [L]} |\mathcal{U}_\ell| \frac{\frac{T}{m} \sum_{k \in \mathcal{R}_\ell} \mathbb{E}^\tau[N_k^u | u \in \mathcal{U}_\ell] \Delta_{k\ell}}{\sum_{k \in [K]} \mathbb{E}^\tau[N_k^u | u \in \mathcal{U}_\ell]} \\ &\stackrel{(a)}{\geq} \sum_{\ell \in [L]} |\mathcal{U}_\ell| \frac{\frac{T}{m} \sum_{k \in \mathcal{R}_\ell} \mathbb{E}^\tau[N_{k_\ell^*}^u | u \in \mathcal{U}_\ell] \Delta_{k\ell} \psi(\ell, k, T, m, p)}{\sum_{k \in [K]} \mathbb{E}^\tau[N_{k_\ell^*}^u | u \in \mathcal{U}_\ell] \frac{T}{m}} \\ &= \sum_{\ell \in [L]} |\mathcal{U}_\ell| \frac{\sum_{k \in \mathcal{R}_\ell} \Delta_{k\ell} \psi(\ell, k, T, m, p)}{K}, \end{aligned}$$

where for (a) we used Lemma 10 ($\psi(\ell, k, T, m, p) \mathbb{E}^\tau[N_{k_\ell^*}^u] / (T/m) \leq \mathbb{E}^\tau[N_k^u | u \in \mathcal{U}_\ell]$ in the numerator, and $\mathbb{E}^\tau[N_k^u | u \in \mathcal{U}_\ell] \leq \mathbb{E}^\tau[N_{k_\ell^*}^u | u \in \mathcal{U}_\ell]$ in the denominator). Taking the expectation over $\mathcal{U}_1, \dots, \mathcal{U}_L$, we have (since $\mathbb{E}[|\mathcal{U}_\ell|] = \beta_\ell m$):

$$R^\tau(T) \geq m \sum_{\ell \in [L]} \beta_\ell \frac{\sum_{k \in \mathcal{R}_\ell} \Delta_{k\ell} \psi(\ell, k, T, m, p)}{K}.$$

Proof of Lemma 10. Consider a 2-armed bandit problem with Bernoulli arms k_ℓ^* and k of means $p_{k_\ell^*}$ and p_k . The means are known but the arm with the highest mean is unknown. For this bandit problem, we build an algorithm ζ based on the decisions by τ .

A valid algorithm ζ based on τ .

1. Pick a user u uniformly at random in the cluster \mathcal{U}_ℓ . We run τ for T rounds. When user u comes to the system and τ selects the item in $\mathcal{I}_{k_\ell^*}$ (resp. \mathcal{I}_k), then ζ also selects the arm k_ℓ^* (resp. k). We call this procedure as an episode.
2. We repeat the above Step 1, to determine the arm selections made by ζ . At the beginning of the successive episodes of T rounds, the user u are again chosen uniformly at random in the clusters \mathcal{I}_ℓ , independently of the choices made in earlier episodes.

ζ is a valid algorithm as the decision by τ is based on past observations. We stop ζ after κ episodes of T rounds, where κ is the first episode where ζ has made more than $\frac{T}{m}$ selections. By Wald's first lemma, the expected regret accumulated by ζ before we stop playing is:

$$R^\zeta = \Delta_{k\ell} \mathbb{E}^\tau[\kappa] \mathbb{E}^\tau[N_k^u | u \in \mathcal{U}_\ell].$$

Since $\mathbb{E}^\tau[N_k^u | u \in \mathcal{U}_\ell] + \mathbb{E}^\tau[N_{k_\ell^*}^u | u \in \mathcal{U}_\ell] \leq \frac{T}{m}$, We have:

$$\begin{aligned} \mathbb{E}[\kappa] \mathbb{E}^\tau[N_{k_\ell^*}^u | u \in \mathcal{U}_\ell] &\leq \underbrace{\frac{T}{m}}_{\text{Stopping criteria of } \zeta} + \underbrace{\mathbb{E}[\sum_k N_k^u | u \in \mathcal{U}_\ell]}_{\text{expected number of drawing the user } u \text{ in a single episode}} \\ &= \frac{2T}{m} \end{aligned}$$

We will also prove a lower bound on R^ζ :

Lemma 11. *We have for all $T \geq 2m$: $R^\zeta \geq 2\Delta_{k\ell}\psi(\ell, k, T, m, p)$.*

Combining this lemma with Lemma 10 concludes the proof of $R^\pi(T) \geq R_{\text{uc}}(T)$.

Proof of Lemma 11. When the algorithm decides which arm to choose at time t , we assume that the algorithm has access to the rewards of both arms up to time $t-1$. This is a simpler problem than the original 2-armed bandit problem. Hence the regret in the original problem is higher than that in the simpler problem. Let θ denote the distribution of the rewards of both arms, when the average reward of the first arm is $p_{k_\ell^*}$ and that of the second arm is p_k . Let θ' denote the distribution of the rewards of both arms, when arms are swapped: the average reward of the first arm is p_k and that of the second arm is $p_{k_\ell^*}$. Let $\theta^{\otimes(t-1)}$ be a product measure for the reward observations up to time $t-1$ under the measure θ . Let k_t be the arm selected at time t . From Lemma 9, we have, for each $t \geq 2$,

$$\begin{aligned} \mathbb{P}_\theta(k_t = k) + \mathbb{P}_{\theta'}(k_t = k_\ell^*) &\geq \frac{1}{2} \exp(-\text{KL}(\theta^{\otimes(t-1)}, \theta'^{\otimes(t-1)})) \\ &= \frac{1}{2} \exp(-(t-1)(\text{kl}(p_{k_\ell^* \ell}, p_{k\ell}) + \text{kl}(p_{k\ell}, p_{k_\ell^* \ell}))) = \frac{1}{2} e^{-(t-1)\gamma(p_{k_\ell^* \ell}, p_{k\ell})}. \end{aligned} \quad (8)$$

Note that (8) also holds for $t=1$ by the symmetry. We have,

$$\begin{aligned} R^\zeta &\geq \frac{\Delta_{k\ell}}{2} \sum_{t=1}^{T/m} (\mathbb{P}_\theta(k_t = k) + \mathbb{P}_{\theta'}(k_t = k_\ell^*)) \\ &\geq \frac{\Delta_{k\ell}}{4} \sum_{t=1}^{T/m} e^{-(t-1)\gamma(p_{k_\ell^* \ell}, p_{k\ell})} \\ &= \frac{\Delta_{k\ell}}{4} \frac{1 - e^{-\frac{T}{m}\gamma(p_{k_\ell^* \ell}, p_{k\ell})}}{1 - e^{-\gamma(p_{k_\ell^* \ell}, p_{k\ell})}} \\ &= 2\Delta_{k\ell}\psi(\ell, k, T, m, p), \end{aligned}$$

where the first inequality stems from the fact that R^ζ is the regret accumulated over more than $\frac{T}{m}$ rounds, and the second inequality is from (8). \square

Proof of $R^\pi(T) \gtrsim c(\beta, p)m \log(T/m)$. We define:

$$\Theta := \{p \in [0, 1]^{K \times L} : \forall \ell \in [L], \exists \Gamma \text{ (permutation of } [K]) \text{ s.t. } p_{\Gamma(1)\ell} > p_{\Gamma(2)\ell} \geq \dots \geq p_{\Gamma(K)\ell}\}$$

as a set of all possible problems. We denote $k_\ell^* := \arg \max_k p_{k\ell}$ as the index of the best item cluster for users in the cluster \mathcal{U}_ℓ and $\Delta_{k\ell} := p_{k_\ell^* \ell} - p_{k\ell} \forall k \in [K], \forall \ell \in [L]$. Consider an arbitrary algorithm π . We define the regret of a single user u as:

$$\begin{aligned} R_u^\pi(T) &:= \frac{T}{m} \sum_\ell \beta_\ell p_{k_\ell^* \ell} - \sum_{t=1}^T \mathbb{E} \left[\sum_{k, \ell} \mathbb{1}_{\{u \in \mathcal{U}_\ell, u_t = u, i_t^\pi \in \mathcal{I}_k\}} p_{k\ell} \right] \\ &= \sum_\ell \beta_\ell \sum_{k \neq k_\ell^*} \Delta_{k\ell} \mathbb{E}[N_k^u | u \in \mathcal{U}_\ell], \end{aligned}$$

where $N_k^u := \sum_{t=1}^T \mathbb{1}_{\{i_t^\pi \in \mathcal{I}_k, u_t = u\}}$ is the number of times user u is presented an item of cluster \mathcal{I}_k (under π). Remember that u is random and belongs to \mathcal{U}_ℓ with probability β_ℓ . We further define the

conditional regret of a single user $u \in \mathcal{U}_\ell$ given $N_u(T) = N$ and $u \in \mathcal{U}_\ell$ as:

$$\begin{aligned} R_u^\pi(T)_{N,\ell} &:= N p_{k_\ell^* \ell} - \sum_{t=1}^T \mathbb{E} \left[\sum_k \mathbb{1}_{\{u_t=u, i_t^\pi \in \mathcal{I}_k\}} p_{k\ell} \mid N_u(T) = N, u \in \mathcal{U}_\ell \right] \\ &= \sum_{k \neq k_\ell^*} \Delta_{k\ell} \mathbb{E}[N_k^u \mid u \in \mathcal{U}_\ell, N_u(T) = N]. \end{aligned}$$

Note that we have $R^\pi(T) = \sum_{u \in \mathcal{U}} R_u^\pi(T) = m R_u^\pi(T)$ and $R_u^\pi(T) = \sum_{N \geq 1}^T \mathbb{P}(N_u(T) = N) \sum_{\ell \in [L]} \beta_\ell R_u^\pi(T)_{N,\ell}$.

Assume that π is uniformly good. This means that if for all problem $p \in \Theta$, as $N \rightarrow \infty$, the conditional regret $R_u^\pi(T)_{N,\ell}$ satisfies:

$$\forall \ell \in [L], \forall 0 < \alpha < 1, \quad R_u^\pi(T)_{N,\ell} = o((N)^\alpha). \quad (9)$$

The existence of uniformly good algorithms is guaranteed because applying the classical algorithms (e.g, UCB1) to each user satisfy indeed is uniformly good (this is proved for the ECB algorithm using Theorem 9 presented in Appendix J).

We state our claim in the following theorem, providing a lower bound of the regret of a single user:

Theorem 8. *For any uniformly good algorithm $\pi \in \Pi$, for any $p \in \Theta$, when $T = \omega(m)$, we have: for any $u \in \mathcal{U}$,*

$$\liminf_{T \rightarrow \infty} \frac{R_u^\pi(T)}{\log(T/m)} \geq c(\beta, p),$$

where $c(\beta, p)$ is the value of the following optimization problem:

$$\begin{aligned} \inf_{n=(n_{k\ell}) \geq 0} \sum_{\ell \in [L]} \beta_\ell \sum_{k \neq k_\ell^*} \Delta_{k\ell} n_{k\ell} \quad (10) \\ \text{s.t.} \quad \forall \ell \in [L], \quad \forall \ell' \in \mathcal{L}^\perp(\ell), \quad \sum_{k \neq k_\ell^*} \text{kl}(p_{k\ell}, p_{k\ell'}) n_{k\ell} \geq 1. \end{aligned}$$

In the above theorem, we can interpret $n_{k\ell}$ as

$$n_{k\ell} \stackrel{T \rightarrow \infty}{\sim} \frac{\mathbb{E}[N_k^u \mid u \in \mathcal{U}_\ell, N_u(T) = T/m]}{\log(T/m)}.$$

Proof of Theorem 8: Case $K = 2, L = 2$. To illustrate the idea behind the proof, we address the simple case with two item and user clusters. We define the values of $(p_{k\ell})$ as in Table F.2, where $\mu_1 > \mu_2$ and $\mu_1 < \mu'_2$, so that $\mathcal{L}^\perp(1) = \{2\}$.

	$k = 1$	$k = 2$
$\ell = 1$	μ_1	μ_2
$\ell = 2$	μ_1	μ'_2

Table 10: The values of $(p_{k\ell})$. $\mu_1 > \mu_2$ and $\mu_1 < \mu'_2$.

The proof is in two steps. In the first step we derive a lower bound of the conditional regret, and in the second step, we de-condition using properties of the user arrival process.

Step 1. In this step, we condition on $N_u(T) = N$ and $u \in \mathcal{U}_1$. All expectations and probabilities are conditioned with respect by these events. We apply a classical change-of-measure argument. Let p denote the original model. We build a perturbed model p' obtained from p by just swapping the ideas of the user clusters. Let \mathbb{P}_p and \mathbb{E}_p (resp. $\mathbb{P}_{p'}$ and $\mathbb{E}_{p'}$) be the probability measure and the expectation under p (resp. p'), respectively. We compute the log-likelihood ratio of the observations for user u generated under p and p' as:

$$\begin{aligned} L_T &:= \sum_{t=1}^T \mathbb{1}_{\{u_t=u\}} \left(\mathbb{1}_{\{u \in \mathcal{U}_1, i_t^\pi \in \mathcal{I}_2\}} \left(\mathbb{1}_{\{X_{i_t^\pi u} = +1\}} \log \frac{\mu_2}{\mu'_2} + \mathbb{1}_{\{X_{i_t^\pi u} = 0\}} \log \frac{1 - \mu_2}{1 - \mu'_2} \right) \right. \\ &\quad \left. + \mathbb{1}_{\{u \in \mathcal{U}_2, i_t^\pi \in \mathcal{I}_2\}} \left(\mathbb{1}_{\{X_{i_t^\pi u} = +1\}} \log \frac{\mu'_2}{\mu_2} + \mathbb{1}_{\{X_{i_t^\pi u} = 0\}} \log \frac{1 - \mu'_2}{1 - \mu_2} \right) \right). \end{aligned}$$

For any measurable random variable $Z \in [0, 1]$, we have:

$$\mathbb{E}_p[L_T] = \mathbb{E}_p[N_2^u] \text{kl}(\mu_2, \mu'_2) \stackrel{(a)}{\geq} \text{kl}(\mathbb{E}_p[Z], \mathbb{E}_{p'}[Z]),$$

where (a) stems from the data-processing inequality (cf. [8]). Taking $Z = N_2^u/N$, we have:

$$\begin{aligned} \text{kl}(\mathbb{E}_p[Z], \mathbb{E}_{p'}[Z]) &= \text{kl}\left(\frac{\mathbb{E}_p[N_2^u]}{N}, \frac{\mathbb{E}_{p'}[N_2^u]}{N}\right) \\ &\geq \left(1 - \frac{\mathbb{E}_{p'}[N_2^u]}{N}\right) \log\left(\frac{N}{N - \mathbb{E}_{p'}[N_2^u]}\right) - \log 2, \end{aligned}$$

where for the last inequality, we used that for all $(x, y) \in [0, 1]^2$,

$$\text{kl}(x, y) \geq (1 - x) \log \frac{1}{1 - y} - \log 2.$$

As the algorithm π is uniformly good, $N - \mathbb{E}_{p'}(N_2^u) = o(N^\alpha)$ for all $\alpha \in (0, 1)$ as $N \rightarrow \infty$. Therefore, for all $\alpha \in (0, 1)$,

$$\begin{aligned} \liminf_{N \rightarrow \infty} \frac{1}{\log N} \log \frac{N}{N - \mathbb{E}_{p'}(N_2^u)} &\geq \liminf_{N \rightarrow \infty} \frac{1}{\log N} \log \frac{N}{N^\alpha} \\ &= 1 - \alpha. \end{aligned}$$

Furthermore, $\lim_{N \rightarrow \infty} \frac{\mathbb{E}_p[N_2^u]}{N} = 0$ as π is uniformly good. Therefore, we have:

$$\liminf_{N \rightarrow \infty} \frac{\mathbb{E}_p[N_2^u]}{\log N} \geq \frac{1}{\text{kl}(\mu_2, \mu'_2)}. \quad (11)$$

Step 2. De-conditioning. In view of Lemma 5, we have:

$$\mathbb{P}\left(N_u(T) \leq \frac{T}{2m}\right) \leq \exp\left(-\frac{T}{2m}(1 - \log(2))\right).$$

From the above inequality and (11), we deduce:

$$\begin{aligned} \liminf_{T \rightarrow \infty} \frac{R_u^\pi(T)}{\log(T/m)} &\geq \liminf_{T \rightarrow \infty} \frac{\sum_{N=1}^T \mathbb{P}(N_u(T) = N) \beta_1 R_u^\pi(T)_{N,1}}{\log(T/m)} \\ &\geq \liminf_{T \rightarrow \infty} \frac{\sum_{N=\frac{T}{2m}}^T \mathbb{P}(N_u(T) = N) \beta_1 R_u^\pi(T)_{\frac{T}{2m},1}}{\log(T/m)} \\ &\geq \liminf_{T \rightarrow \infty} \beta_1 (\mu_1 - \mu_2) \sum_{N=\frac{T}{2m}}^T \mathbb{P}(N_u(T) = N) \frac{\mathbb{E}_p[N_2^u | u \in \mathcal{U}_1, N_u(T) = T/(2m)]}{\log(T/m)} \\ &\stackrel{(a)}{\geq} \liminf_{T \rightarrow \infty} \beta_1 (\mu_1 - \mu_2) \left(1 - \exp\left(-\frac{T}{2m}(1 - \log 2)\right)\right) \frac{\log(T/(2m))}{\text{kl}(\mu_2, \mu'_2) \log(T/m)} \\ &\stackrel{(b)}{=} \frac{\beta_1 (\mu_1 - \mu_2)}{\text{kl}(\mu_2, \mu'_2)}, \end{aligned}$$

where for (a) we used Lemma 5 with (11) and for (b) we used $T = \omega(m)$. This concludes the proof of the case $K = 2$ and $L = 2$. \square

Proof of Theorem 8: General case. We consider a simpler problem: the algorithm knows the values of $(p_{k\ell})$ and β_ℓ . Take $(\ell, \ell') \in [L]^2$ such that $\ell' \in \mathcal{L}^\perp(\ell)$. As in the case of two user and item clusters, we will prove that (this is done at the end of this proof):

$$\liminf_{N \rightarrow \infty} \frac{\sum_{k \neq k_\ell^*} \mathbb{E}_p[N_k^u | u \in \mathcal{U}_\ell, N_u(T) = N] \text{kl}(p_{k\ell}, p_{k\ell'})}{\log N} \geq 1. \quad (12)$$

This inequality holds for any possible (ℓ, ℓ') such that $\ell' \in \mathcal{L}^\perp(\ell)$. Therefore, for all $\ell \in [L]$, for all $\ell' \in \mathcal{L}^\perp(\ell)$

$$\liminf_{N \rightarrow \infty} \frac{\sum_{k \neq k_\ell^*} \mathbb{E}_p[N_k^u \mid u \in \mathcal{U}_\ell, N_u(T) = N] \text{kl}(p_{k\ell}, p_{k\ell'})}{\log N} \geq 1.$$

Then, we have:

$$\begin{aligned} & \liminf_{T \rightarrow \infty} \frac{R_u^\pi(T)}{\log(T/m)} \\ & \geq \liminf_{T \rightarrow \infty} \sum_{\ell \in [L]} \beta_\ell \sum_{N=1}^T \mathbb{P}(N_u(T) = N) \frac{R_u^\pi(T)_{N,\ell}}{\log(T/m)} \\ & \geq \liminf_{T \rightarrow \infty} \sum_{\ell \in [L]} \beta_\ell \sum_{N=\frac{T}{2m}}^T \mathbb{P}(N_u(T) = N) \frac{\sum_{k \neq k_\ell^*} \Delta_{k\ell} \mathbb{E}[N_k^u \mid u \in \mathcal{U}_\ell, N_u(T) = T/(2m)]}{\log(T/m)} \\ & \stackrel{(a)}{\geq} \liminf_{T \rightarrow \infty} \left(1 - \exp\left(-\frac{T}{2m}(1 - \log 2)\right)\right) \frac{\sum_{\ell \in [L]} \beta_\ell \sum_{k \neq k_\ell^*} \Delta_{k\ell} \mathbb{E}[N_k^u \mid u \in \mathcal{U}_\ell, N_u(T) = T/(2m)]}{\log(T/m)} \\ & = \liminf_{T \rightarrow \infty} \left(1 - \exp\left(-\frac{T}{2m}(1 - \log 2)\right)\right) \frac{\sum_{\ell \in [L]} \beta_\ell \sum_{k \neq k_\ell^*} \Delta_{k\ell} \mathbb{E}[N_k^u \mid u \in \mathcal{U}_\ell, N_u(T) = T/(2m)]}{\log(T/(2m)) + \log 2} \\ & \stackrel{(b)}{=} \liminf_{T \rightarrow \infty} \frac{\sum_{\ell \in [L]} \beta_\ell \sum_{k \neq k_\ell^*} \Delta_{k\ell} \mathbb{E}[N_k^u \mid u \in \mathcal{U}_\ell, N_u(T) = T/(2m)]}{\log(T/(2m))}, \end{aligned}$$

where (a) is from Lemma 5 and (b) is from $T = \omega(m)$. Thus, we have:

$$\liminf_{T \rightarrow \infty} \frac{R_u^\pi(T)}{\log(T/m)} \geq c(\beta, p),$$

where $c(\beta, p)$ is the value of the following optimization problem:

$$\begin{aligned} & \inf_{n=(n_{k\ell}) \geq 0} \sum_{\ell \in [L]} \beta_\ell \sum_{k \neq k_\ell^*} \Delta_{k\ell} n_{k\ell} \\ & \text{s.t. } \forall \ell \in [L], \quad \forall \ell' \in \mathcal{L}^\perp(\ell), \quad \sum_{k \neq k_\ell^*} \text{kl}(p_{k\ell}, p_{k\ell'}) n_{k\ell} \geq 1. \end{aligned}$$

Proof of the inequality (12). Again, we use a change-of-measure argument. Let p and p' be a original model and a model with the indices of user clusters (ℓ, ℓ') are swapped from the original model, respectively. Let \mathbb{P}_p and \mathbb{E}_p (resp. $\mathbb{P}_{p'}$ and $\mathbb{E}_{p'}$) be the probability measure and the expectation under p (resp. p'), respectively. We define our log-likelihood ratio as:

$$\begin{aligned} L_T := & \sum_{t=1}^T \sum_{k \neq k_\ell^*} \mathbb{1}_{\{u_t = u\}} \left(\mathbb{1}_{\{u \in \mathcal{U}_\ell, i_t^\pi \in \mathcal{I}_k\}} \left(\mathbb{1}_{\{X_{i_t^\pi} u = +1\}} \log \frac{p_{k\ell}}{p_{k\ell'}} + \mathbb{1}_{\{X_{i_t^\pi} u = 0\}} \log \frac{1 - p_{k\ell}}{1 - p_{k\ell'}} \right) \right. \\ & \left. + \mathbb{1}_{\{u \in \mathcal{U}_{\ell'}, i_t^\pi \in \mathcal{I}_k\}} \left(\mathbb{1}_{\{X_{i_t^\pi} u = +1\}} \log \frac{p_{k\ell'}}{p_{k\ell}} + \mathbb{1}_{\{X_{i_t^\pi} u = 0\}} \log \frac{1 - p_{k\ell'}}{1 - p_{k\ell}} \right) \right). \end{aligned}$$

Taking the conditional expectation $\mathbb{E}_p[\cdot \mid u \in \mathcal{U}_\ell, N_u(T) = N]$, we have:

$$\begin{aligned} \mathbb{E}_p[L_T \mid u \in \mathcal{U}_\ell, N_u(T) = N] & = \sum_{k \neq k_\ell^*} \mathbb{E}_p[N_k^u \mid u \in \mathcal{U}_\ell, N_u(T) = N] \text{kl}(p_{k\ell}, p_{k\ell'}) \\ & \stackrel{(a)}{\geq} \text{kl} \left(\frac{\mathbb{E}_p[N_{k_\ell^*}^u]}{N}, \frac{\mathbb{E}_{p'}[N_{k_\ell^*}^u]}{N} \right) \\ & \stackrel{(b)}{\geq} \left(1 - \frac{\mathbb{E}_p[N_{k_\ell^*}^u]}{N}\right) \log \left(\frac{N}{N - \mathbb{E}_{p'}[N_{k_\ell^*}^u]} \right) - \log 2, \end{aligned}$$

where for (a), we used the data processing inequality by [8] and for (b), we used that for all $(x, y) \in [0, 1]^2$,

$$\text{kl}(x, y) \geq (1 - x) \log \frac{1}{1 - y} - \log 2.$$

As the algorithm π is uniformly good, $\mathbb{E}_p[N_{k_{\ell'}^*}^u] = o(N^\alpha)$ and $N - \mathbb{E}_{p'}(N_{k_{\ell'}^*}^u) = o(N^\alpha)$ for all $\alpha \in (0, 1)$ as $N \rightarrow \infty$. Therefore, for all $\alpha \in (0, 1)$, we have:

$$\begin{aligned} \liminf_{N \rightarrow \infty} \frac{1}{\log N} \log \frac{N}{N - \mathbb{E}_{p'}[N_{k_{\ell'}^*}^u]} &\geq \liminf_{N \rightarrow \infty} \frac{1}{\log N} \log \frac{N}{N^\alpha} \\ &= 1 - \alpha \end{aligned}$$

and

$$\liminf_{N \rightarrow \infty} \frac{\mathbb{E}_p[N_{k_{\ell'}^*}^u]}{N} = 0.$$

Therefore, we have:

$$\liminf_{N \rightarrow \infty} \frac{\sum_{k \neq k_{\ell'}^*} \mathbb{E}_p[N_k^u \mid u \in \mathcal{U}_\ell, N_u(T) = N] \text{kl}(p_{k\ell}, p_{k_{\ell'}\ell'})}{\log(N)} \geq 1.$$

This concludes the proof of Theorem 8. □

G Performance guarantees of ECT: Proof of Theorem 4

The proof consists in several parts. First we study the initial sampling procedure (at the beginning of the exploration phase). We then upper bound the regret induced by the exploration phase. We analyze the performance of the clustering part of the algorithm, and finally upper bound the regret generated during the test phase.

Item sampling procedure. Let $\tilde{\mathcal{I}}_k = \mathcal{S} \cap \mathcal{I}_k$ be the set of items from \mathcal{I}_k that are sampled.

Then, for $\epsilon_1 = \sqrt{\frac{\log TK}{2(\min_k \alpha_k)^2 (\log T)^2}}$,

$$\begin{aligned} \mathbb{P}\left(|\tilde{\mathcal{I}}_k| \leq \alpha_k(1 - \epsilon_1)(\log T)^2\right) &\stackrel{(a)}{\leq} \exp\left(-(\log T)^2 \text{kl}((1 - \epsilon_1)\alpha_k, \alpha_k)\right) \\ &\stackrel{(b)}{\leq} \exp\left(-(\log T)^2 2\epsilon_1^2 \alpha_k^2\right) \\ &\leq \frac{1}{TK}. \end{aligned}$$

where (a) is from Chernoff-Hoeffding bound (b) is from Pinsker's inequality.

Hence, the event $\mathcal{A}_1 = \{|\tilde{\mathcal{I}}_k| \geq \alpha_k(1 - \epsilon_1)(\log T)^2 \text{ for } 1 \leq k \leq K\}$ holds with probability at least $1 - \frac{1}{T}$. As a consequence, the expected regret due the event \mathcal{A}_1^c is $O(1)$. Thus, we can assume that the event \mathcal{A}_1 holds throughout the remaining of the proof.

Exploration phase. In this phase, we wish to recommend each item in \mathcal{S} for $\log T$ times. We prove that this exploration phase takes around $(\log T)^3$ rounds (and this is the regret it generates). Let us consider a user u . This user can make the exploration phase longer if it arrives more than $(\log T)^2$ during the $(\log T)^3$ first rounds. We have:

$$\begin{aligned} \mathbb{P}\left(N_u((\log T)^3) \geq (\log T)^2\right) &\leq \exp\left(-(\log T)^3 \text{kl}\left(\frac{1}{\log T}, \frac{1}{m}\right)\right) \\ &\leq \exp\left(-2(\log T)^3 \left(\frac{1}{\log T} - \frac{1}{m}\right)^2\right) \\ &\stackrel{(a)}{\leq} \exp\left(-2\left(\log T - \frac{2}{\log T}\right)\right) \end{aligned} \tag{13}$$

where (a) is obtained from $m > (\log T)^3$ (remember that $T = o(m^2)$).

We deduce the probability that the duration of exploration phase T_{exp} exceeds $(\log T)^3$,

$$\begin{aligned} \mathbb{P}(T_{\text{exp}} \geq (\log T)^3) &= \mathbb{P}(\exists u \text{ s.t. } N_u((\log T)^3) \geq (\log T)^2) \\ &\leq \frac{(\log T)^3}{T^2} \exp\left(\frac{4}{\log T}\right) \end{aligned}$$

where (a) is obtained from the union bound and (13).

Now the expected time taken in the exploration phase is,

$$\begin{aligned} \mathbb{E}[T_{\text{exp}}] &= (\log T)^3 + \mathbb{E}[T_{\text{exp}} | T_{\text{exp}} \geq (\log T)^3] \mathbb{P}(T_{\text{exp}} \geq (\log T)^3) \\ &\leq (\log T)^3 + \frac{(\log T)^3}{T} \exp\left(\frac{4}{\log T}\right) \\ &= \mathcal{O}((\log T)^3). \end{aligned}$$

Therefore, we can conclude that the expected regret that occurs in the exploration phase is $\mathcal{O}((\log T)^3)$.

Clustering phase. The performance of the clustering phase can be analyzed using the same arguments as in the proof of Theorem 6 in [32]. To simplify the notation, let $\epsilon = (\log T)^{-\frac{1}{4}}$. Recall that $Q_i = \{j \in \mathcal{S} : |\hat{\rho}_i - \hat{\rho}_j| \leq \epsilon\}$ for all $i \in \mathcal{S}$. We also define a set \mathcal{B}_k for $1 \leq k \leq K$ as:

$$\mathcal{B}_k = \{i \in \mathcal{S} : |p_k - \hat{\rho}_i| \leq \frac{1}{2}\epsilon\}.$$

This set has the following properties:

- (i) $|\mathcal{B}_k| = \Omega((\log T)^2)$ with probability at least $1 - \frac{1}{T}$. This follows from the following argument.

$$\begin{aligned} \mathbb{P}(i \in \mathcal{B}_k) &\geq \mathbb{P}(i \in \mathcal{B}_k | i \in \tilde{\mathcal{I}}_k) \mathbb{P}(i \in \tilde{\mathcal{I}}_k) \\ &\stackrel{(a)}{\geq} \alpha_k (1 - \epsilon_1) \mathbb{P}\left(|p_k - \hat{\rho}_i| \leq \frac{1}{2}\epsilon \mid i \in \tilde{\mathcal{I}}_k\right) \\ &\stackrel{(b)}{\geq} \alpha_k (1 - \epsilon_1) \left(1 - 2 \exp\left(-\frac{\sqrt{\log T}}{2}\right)\right), \end{aligned}$$

where (a) follows from the assumption that \mathcal{A}_1 holds and (b) stems from Chernoff-Hoeffding's bound. Let $r = \alpha_k (1 - \epsilon_1) \left(1 - 2 \exp\left(-\frac{\sqrt{\log T}}{2}\right)\right)$. Then,

$$\begin{aligned} \mathbb{P}\left(|\mathcal{B}_k| < \left(r - \frac{1}{\sqrt{2 \log T}}\right) (\log T)^2\right) &\leq \exp\left(-(\log T)^2 \text{kl}\left(r - \frac{1}{\sqrt{2 \log T}}, r\right)\right) \\ &\leq \exp\left(-2(\log T)^2 \left(\frac{1}{\sqrt{2 \log T}}\right)^2\right) \\ &\leq \frac{1}{T}. \end{aligned}$$

Therefore, $|\mathcal{B}_k| = \Omega((\log T)^2)$ with probability at least $1 - \frac{1}{T}$.

- (ii) $|\mathcal{S} \setminus (\cup_{k=1}^K \mathcal{B}_k)| = \mathcal{O}(\log T)$ with probability at least $1 - \frac{1}{T}$. To show this, we use a similar argument as in (i):

$$\begin{aligned} \mathbb{P}(i \in \mathcal{S} \setminus (\cup_{k=1}^K \mathcal{B}_k)) &\leq \sum_{k=1}^K \mathbb{P}(i \in \mathcal{I}_k) \mathbb{P}(i \in \mathcal{S} \setminus (\cup_{k=1}^K \mathcal{B}_k) | i \in \mathcal{I}_k) \\ &\leq \sum_{k=1}^K \mathbb{P}(i \in \mathcal{I}_k) \mathbb{P}\left(|p_k - \hat{\rho}_i| > \frac{1}{2}\epsilon \mid i \in \mathcal{I}_k\right) \\ &\leq \sum_{k=1}^K \mathbb{P}(i \in \mathcal{I}_k) 2 \exp\left(-\frac{\sqrt{\log T}}{2}\right) \\ &= 2 \exp\left(-\frac{\sqrt{\log T}}{2}\right) \end{aligned}$$

Then, the probability that the size of $|\mathcal{S} \setminus (\cup_{k=1}^K \mathcal{B}_k)|$ is greater than $\log T$ is,

$$\begin{aligned} &\mathbb{P}\left(|\mathcal{S} \setminus (\cup_{k=1}^K \mathcal{B}_k)| \geq \log T\right) \\ &\leq \exp\left(-(\log T)^2 \text{kl}\left(\frac{1}{\log T}, 2 \exp\left(-\frac{\sqrt{\log T}}{2}\right)\right)\right) \\ &\stackrel{(a)}{\leq} \exp\left(-(\log T) \left(\frac{\sqrt{\log T}}{4} + 2 \exp\left(-\frac{\sqrt{\log T}}{2}\right) - \log 2 - 1\right)\right) \\ &\leq \frac{1}{T}, \end{aligned}$$

where (a) is obtained from Lemma 4 when $\log T \geq 2^4$.

- (iii) If $|\mathcal{B}_k \cap Q_i| \geq 1$, then $|\mathcal{B}_j \cap Q_i| = 0$ for all j, k such that $|p_k - p_j| = \Theta(1)$. Because $|\hat{\rho}_i - \hat{\rho}_l| \geq |p_k - p_j| - |p_k - \hat{\rho}_i| - |p_j - \hat{\rho}_l| \geq |p_k - p_j| - \epsilon$ for $i \in \mathcal{B}_k$ and $l \in \mathcal{B}_j$.
(iv) $\mathcal{B}_k \subset Q_i$ for all $i \in \mathcal{B}_k$, since $|\hat{\rho}_i - \hat{\rho}_j| \leq |\hat{\rho}_i - p_k| + |\hat{\rho}_j - p_k| \leq \epsilon$ for all $j \in \mathcal{B}_k$.

From properties (iii) and (iv), there exists an item $i \in (\cup_{k=1}^K \mathcal{B}_k) \setminus (\cup_{\ell=1}^{k-1} Q_{i_\ell})$ such that $|Q_i \setminus (\cup_{\ell=1}^{k-1} Q_{i_\ell})| \geq m_k$ where m_k is the k -th largest value among $\{|\mathcal{B}_1|, \dots, |\mathcal{B}_{K_0}|\}$ for $K_0 = |\{p_k : 1 \leq k \leq K\}|$. Here, $m_k = \Omega((\log T^2)^2)$ from property (i).

We also have $|Q_v| = \mathcal{O}(\log T)$ for v such that $|Q_v \cap (\cup_{k=1}^K \mathcal{B}_k)| = 0$ from property (ii). Thus, the item v cannot be chosen as i_k .

We can conclude that $|p_k - \hat{p}_k| \leq \epsilon$ for $k = 1, 2$ with probability $1 - 2/T$, since $|\hat{\rho}_i - p_k| \leq \epsilon$ when $|Q_i \cap \mathcal{B}_k| \geq 1$. Hence as before for event \mathcal{A}_1 , we can assume that $\mathcal{A}_2 = \{|p_k - \hat{p}_k| \leq \epsilon \text{ for } k = 1, 2\}$ holds in the remaining of the proof.

Test phase. After n recommendations of an item i from $\mathcal{I}_k \neq \mathcal{I}_1$, the probability that i passes the test is,

$$\begin{aligned} \mathbb{P}(\hat{\rho}_i > \hat{p}_1 - \Delta_0/2) &= \mathbb{P}\left(\hat{\rho}_i > \frac{1}{2}(\hat{p}_1 + \hat{p}_2)\right) \\ &\stackrel{(a)}{\leq} \mathbb{P}\left(\hat{\rho}_i > \frac{1}{2}(p_1 + p_2) - \epsilon\right) \\ &\leq \exp\left(-n \text{kl}\left(\frac{1}{2}(p_1 + p_2) - \epsilon, p_k\right)\right) \\ &\leq \exp\left(-2n \left(\frac{1}{2}(p_1 + p_2) - p_k - \epsilon\right)^2\right) \end{aligned}$$

where (a) is obtained from the assumption that \mathcal{A}_2 holds.

To simplify the notation, let $x = \lfloor \frac{2 \log 3}{\Delta_0^2} \rfloor$. Since we test the item after every x recommendations, we have at most m/x tests for each item. Therefore, the expected number of times a sub-optimal item j is recommended is:

$$\begin{aligned} \mathbb{E}[N(j)] &= x + \sum_{i=1}^{m/x} x \mathbb{P}(\hat{\rho}_k > \hat{p}_1 - \Delta_0/2 \text{ for } i\text{-th test}) \\ &\leq x + \sum_{i=1}^{m/x} x \exp\left(-2ix \left(\frac{1}{2}(p_1 + p_2) - p_k - \epsilon\right)^2\right) \\ &\leq x + \frac{2}{(p_1 + p_2 - 2p_k - 2\epsilon)^2} \\ &\leq \frac{4 \log 3}{(p_1 - p_2)^2} + \frac{2}{(p_1 + p_2 - 2p_k - 2\epsilon)^2}. \end{aligned} \tag{14}$$

Furthermore, the probability that item $i \in I_1$ is not removed until the last test is,

$$\begin{aligned} \mathbb{P}\left(\bigcap_{t=1}^{m/x} \{\hat{\rho}_i > \hat{p}_1 - \Delta_0/2 \text{ for } t\text{-th test}\}\right) &\geq 1 - \sum_{t=1}^{m/x} \mathbb{P}(\hat{\rho}_i \leq \hat{p}_1 - \Delta_0/2 \text{ for } t\text{-th test}) \\ &= 1 - \sum_{t=1}^{m/x} \mathbb{P}\left(\hat{\rho}_i < \frac{1}{2}(p_1 + p_2) - \epsilon\right) \\ &\geq 1 - \sum_{t=1}^{m/x} \exp\left(-t \text{kl}\left(\frac{1}{2}(p_1 + p_2) - \epsilon, p_1\right)\right) \\ &\geq 1 - \sum_{t=1}^{m/x} \exp\left(-\frac{1}{2}tx (p_1 - p_2 + 2\epsilon)^2\right) \\ &\geq 1 - \frac{\exp\left(-\frac{1}{2}x (p_1 - p_2 + 2\epsilon)^2\right)}{1 - \exp\left(-\frac{1}{2}x (p_1 - p_2 + 2\epsilon)^2\right)} \\ &\geq \frac{1}{2}. \end{aligned} \tag{15}$$

By (15), if we assume that user arrives for \bar{N} times at most, we need at most $2\bar{N}$ optimal items in \mathcal{V} in expectation. Thus, the required number of new samples from $\mathcal{I} \setminus \mathcal{V}_0$ is less than $\frac{2\bar{N}}{\alpha_1}$. Therefore,

from (14), the expected regret that occurs in the test phase under the assumption that every user arrives for less than \bar{N} times is,

$$\mathcal{O}\left(\frac{2\bar{N}}{\alpha_1} \sum_{k=2}^K \alpha_k (p_1 - p_k) \left(\frac{4 \log 3}{(p_1 - p_2)^2} + \frac{2}{(p_1 + p_2 - 2p_k)^2} \right)\right).$$

On the other hand, the regret due to more than \bar{N} arrivals of users is $\frac{m}{m(e-1)} = O(1)$ by Lemma 1.

Finally, the expected regret of ECT satisfies:

$$R^{ECT}(T) = \mathcal{O}\left((\log T)^3 + \frac{2\bar{N}}{\alpha_1} \sum_{k=2}^K \left[\alpha_k (p_1 - p_k) \left(\frac{4 \log 3}{(p_1 - p_2)^2} + \frac{2}{(p_1 + p_2 - 2p_k)^2} \right) \right]\right).$$

□

H Performance guarantees of ET: Proof of Theorem 5

Recall that μ_x is the expected reward such that $\mathbb{P}(\rho_i \leq \mu_x) = x$, and that we are interested in the satisficing regret defined by:

$$R_\varepsilon^\pi(T) = \mathbb{E}^\pi \left(\sum_{t=1}^T \max\{0, \mu_{1-\varepsilon} - \rho_{i_t}^\pi\} \right).$$

We consider the case where $\varepsilon \geq C \sqrt{\frac{\pi}{2 \log T}}$. Further recall that we assume that $\zeta(x) \leq C$ for all $x \in [0, 1]$.

To prove Theorem 5, we first analyze the performance of the exploration phase, and in particular show that $\hat{\mu}_{1-\frac{\varepsilon}{2}}$ is very close to $\mu_{1-\varepsilon}$. We then study the regret generated during the test phase.

Exploration Phase. We first derive an upper and a lower bound of $\hat{\mu}_{1-\frac{\varepsilon}{2}}$. Here, we use the fact that for all $i \in \mathcal{S}$, the $\hat{\rho}_i$'s are i.i.d. random variables.

From Chernoff bound and Pinsker's inequality,

$$\begin{aligned} \mathbb{P}(\hat{\rho}_i \geq \mu_{1-\frac{\varepsilon}{4}}) &\leq \frac{\varepsilon}{4} + \int_0^{\mu_{1-\frac{\varepsilon}{4}}} \exp(-2^4 \log(T) \text{kl}(\mu_{1-\frac{\varepsilon}{4}}, \mu)) \zeta(\mu) d\mu \\ &\leq \frac{\varepsilon}{4} + \int_0^{\mu_{1-\frac{\varepsilon}{4}}} \exp(-2^5 (\mu_{1-\frac{\varepsilon}{4}} - \mu)^2 \log(T)) \zeta(\mu) d\mu \\ &\leq \frac{\varepsilon}{4} + \int_0^\infty C \exp(-2^5 x^2 \log(T)) dx \\ &\leq \frac{\varepsilon}{4} + \frac{\varepsilon}{8} = \frac{3\varepsilon}{8}, \end{aligned} \tag{16}$$

where for the last inequality, we use the Gaussian integral $\int_{-\infty}^\infty e^{-x^2} dx = \sqrt{\pi}$. When $\varepsilon \geq C \sqrt{\frac{\pi}{2 \log T}}$,

$$\begin{aligned} \int_0^\infty C \exp(-2^5 x^2 \log(T)) dx &= \frac{1}{2} \int_{-\infty}^\infty \frac{C}{\sqrt{2^5 \log(T)}} \exp(-x^2) dx \\ &= \frac{1}{8} C \sqrt{\frac{\pi}{2 \log T}} \leq \frac{\varepsilon}{8}. \end{aligned}$$

Similarly,

$$\begin{aligned} \mathbb{P}(\hat{\rho}_i \leq \mu_{1-\frac{3\varepsilon}{4}}) &\leq 1 - \frac{3\varepsilon}{4} + \int_{\mu_{1-\frac{3\varepsilon}{4}}}^1 e^{-2^4 \log(T) \text{kl}(\mu_{1-\frac{3\varepsilon}{4}}, \mu)} \zeta(\mu) d\mu \\ &\leq 1 - \frac{3\varepsilon}{4} + \int_{\mu_{1-\frac{3\varepsilon}{4}}}^1 e^{-2^5 (\mu_{1-\frac{3\varepsilon}{4}} - \mu)^2 \log(T)} \zeta(\mu) d\mu \\ &\leq 1 - \frac{3\varepsilon}{4} + \int_0^\infty C \exp(-2^5 \mu^2 \log(T)) d\mu \\ &< 1 - \frac{3\varepsilon}{4} + \frac{\varepsilon}{8} \leq 1 - \frac{5\varepsilon}{8}. \end{aligned} \tag{17}$$

From the Chernoff-Hoeffding and (16),

$$\begin{aligned} \mathbb{P} \left(\left| \{i \in \mathcal{S} : \hat{\rho}_i \geq \mu_{1-\frac{\varepsilon}{4}}\} \right| \geq \frac{\varepsilon}{2} |\mathcal{S}| \right) &\leq \exp \left(-|\mathcal{S}| \text{kl} \left(\frac{\varepsilon}{2}, \frac{3\varepsilon}{8} \right) \right) \\ &\leq \exp \left(-2|\mathcal{S}| \frac{\varepsilon^2}{8^2} \right) \leq \frac{1}{T^2}. \end{aligned} \tag{18}$$

From the Chernoff-Hoeffding and (17),

$$\begin{aligned} \mathbb{P} \left(\left| \{i \in \mathcal{S} : \hat{\rho}_i \leq \mu_{1-\frac{3\varepsilon}{4}}\} \right| \leq (1 - \frac{\varepsilon}{2}) |\mathcal{S}| \right) &\leq \exp \left(-|\mathcal{S}| \text{kl} \left(1 - \frac{\varepsilon}{2}, 1 - \frac{5\varepsilon}{8} \right) \right) \\ &\leq \exp \left(-2|\mathcal{S}| \frac{\varepsilon^2}{8^2} \right) \leq \frac{1}{T^2}. \end{aligned} \tag{19}$$

We conclude from (18) and (19) that with probability $1 - \frac{2}{T^2}$, we have

$$\mu_{1-\frac{3\varepsilon}{4}} \leq \hat{\mu}_{1-\frac{\varepsilon}{2}} \leq \mu_{1-\frac{\varepsilon}{4}}. \quad (20)$$

Further observe that using the same arguments as those used to upper bound the duration of the exploration phase of the ECT algorithm, the expected duration, and hence the expected regret, of the exploration phase in ET is $O\left(\frac{(\log T)^2}{\varepsilon^2}\right)$.

Test Phase. For convenience, let $\Delta = \log \log_2(2^e m^2)$. Then, ET runs at most $\tau = \lfloor \log_2(\frac{m}{\Delta}) \rfloor$ tests for each item. We define the distance \mathcal{D} between two Bernoulli distributions as follows:

$$\mathcal{D}(p, q) = \text{kl}(s, p) = \text{kl}(s, q) \quad \text{with} \quad s = \frac{\log \frac{1-p}{1-q}}{\log \frac{q(1-p)}{p(1-q)}} \quad \text{for} \quad p \neq q.$$

Let $m^\pi(\mu)$ be the expected number of users to whom a randomly selected item with parameter μ is recommended. Let $\rho^{(\ell)}(\mu)$ be the random value $\hat{\rho}_i$ of item i having μ after $\lfloor 2^\ell \Delta \rfloor$ observations.

Consider items having μ such that $\mu \leq \hat{\mu}_{1-\frac{\varepsilon}{2}}$ and $2^{-\ell} \leq \mathcal{D}(\mu, \hat{\mu}_{1-\frac{\varepsilon}{2}})$. Then, $\text{kl}(\bar{\rho}^{(\ell)}, \mu) \geq 2^{-\ell}$ and we have

$$\begin{aligned} m^\pi(\mu) &\leq 2^\ell \Delta + \sum_{r=\ell}^{\tau} 2^{r+1} \Delta \mathbb{P}\left(\hat{\rho}^{(r)}(\mu) \geq \bar{\rho}^{(r)}\right) \\ &\leq 2^\ell \Delta + \sum_{r=\ell}^{\tau} 2^{r+1} \Delta \mathbb{P}\left(\hat{\rho}^{(r)}(\mu) \geq \bar{\rho}^{(\ell)}\right) \\ &\leq 2^\ell \Delta + \sum_{r=\ell}^{\tau} 2^{r+1} \Delta \exp\left(-2^r \Delta \text{kl}(\bar{\rho}^{(\ell)}, \mu)\right) \\ &\leq 2^\ell \Delta + \sum_{r=\ell}^{\tau} 2^{r+1} \Delta \exp\left(-2^{r-\ell} \Delta\right) \\ &\leq 2^{\ell+2} \Delta. \end{aligned} \quad (21)$$

From (21),

$$m^\pi(\mu) \leq \begin{cases} 2^3 \Delta & \text{for } 2^{-1} \leq \mathcal{D}(\mu, \hat{\mu}_{1-\frac{\varepsilon}{2}}) \\ \frac{2^3 \Delta}{\mathcal{D}(\mu, \hat{\mu}_{1-\frac{\varepsilon}{2}})} & \text{for } 2^{-\ell} \leq \mathcal{D}(\mu, \hat{\mu}_{1-\frac{\varepsilon}{2}}) \leq 2^{-\ell+1} \\ m & \text{for } \mathcal{D}(\mu, \hat{\mu}_{1-\frac{\varepsilon}{2}}) \leq 2^{-\tau} \end{cases} \quad (22)$$

Next we study the expected regret generated by recommending a newly sampled item. From the regret definition and (21),

$$\begin{aligned} \frac{1}{8\Delta} \int_0^{\mu_{1-\varepsilon}} (\mu_{1-\varepsilon} - \mu) m^\pi(\mu) \zeta(\mu) d\mu &\leq \int_0^{\mu_{1-\varepsilon}} (\mu_{1-\varepsilon} - \mu) \left(1 + \frac{1}{\mathcal{D}(\mu, \mu_{1-\frac{3\varepsilon}{4}})}\right) \zeta(\mu) d\mu \\ &\leq \int_0^{\mu_{1-\varepsilon}} C(\mu_{1-\varepsilon} - \mu) \left(1 + \frac{2}{(\mu - \mu_{1-\frac{3\varepsilon}{4}})^2}\right) d\mu \\ &\leq \int_0^{\mu_{1-\varepsilon}} C(\mu_{1-\varepsilon} - \mu) \left(1 + \frac{2}{\left(\frac{\varepsilon}{4C} + \mu_{1-\varepsilon} - \mu\right)^2}\right) d\mu \\ &\leq \frac{C}{2} + \int_0^{\mu_{1-\varepsilon}} \frac{2C}{\frac{\varepsilon}{4C} + \mu_{1-\varepsilon} - \mu} d\mu \\ &\leq \frac{C}{2} + \log(4C/\varepsilon), \end{aligned} \quad (23)$$

where the second inequality stems from Pinsker's inequality $2(p - q)^2 \leq \text{kl}(p, q)$ and the definition of \mathcal{D} , and the third inequality uses the assumption $\zeta(\mu) \leq C$.

If an item has a parameter $\mu \geq \hat{\mu}_{1-\frac{\varepsilon}{2}}$, we do not remove it from \mathcal{V} with probability at least

$$\begin{aligned}
\mathbb{P}\left(\bigcap_{\ell=1}^{\tau} \{\hat{\rho}^{(\ell)}(\mu) > \bar{\rho}^{(\ell)}\}\right) &\geq 1 - \sum_{\ell=1}^{\tau} \mathbb{P}\left(\hat{\rho}^{(\ell)}(\mu) \leq \bar{\rho}^{(\ell)}\right) \\
&\geq 1 - \sum_{i=1}^{\tau} e^{-\Delta} \\
&= 1 - \frac{1}{\log_2(2^e m^2)} \frac{1 - e^{-\tau\Delta}}{1 - \frac{1}{\log_2(2^e m^2)}} \\
&\geq 1 - \frac{1}{\log_2(2^e m^2) - 1} \\
&\geq \frac{1}{2}.
\end{aligned}$$

To recommend items i with parameters $\mu_i \geq \mu_{1-\varepsilon}$ to the \bar{N} arrivals of every user, we then need, on average, $\frac{2\bar{N}}{\varepsilon}$ sampled items. From (23), we conclude that the satisficing regret of ET satisfies:

$$R_{\varepsilon}^{\pi}(T) = \mathcal{O}\left(\frac{\bar{N} \log(e/\varepsilon) \log \log(m)}{\varepsilon} + \frac{(\log T)^2}{\varepsilon^2}\right).$$

□

I Performance guarantees of EC-UCS: Proof of Theorem 6

These two last sections I and J of the appendix are devoted to the analysis of the regret of EC-UCS and ECB in systems with clustered items and users. The two algorithms share the same initial phase to cluster items. The next subsection is hence devoted to the analysis of this item clustering phase. Then, we present an analysis of the performance of the other phases of EC-UCS, and conclude this section with the statement and proof of lemmas used in the analysis of EC-UCS.

I.1 Clustering items in EC-UCS and ECB

The exploration phase for item clustering is of duration $10m$, and hence induces a regret upper bounded by $10m$. In what follows, we just investigate the quality of the item clusters that result from this phase.

Recall that the algorithm randomly selects a set \mathcal{S} of items to cluster. We denote by V_1, \dots, V_K the true cluster $\mathcal{S} \cap \mathcal{I}_1, \dots, \mathcal{S} \cap \mathcal{I}_K$, respectively, and assume that $m^2 \geq T(\log T)^3$ and $n = \omega(\log T)$. We let $n_0 := \min\{n, \frac{m}{(\log T)^2}\}$ be the number of sampled items. For each k , the size of V_k concentrates around $\alpha_k n_0$. Indeed, from the Chernoff-Hoeffding's inequality,

$$\mathbb{P}\left(\left|V_k\right| - \alpha_k n_0 \geq \sqrt{n_0 \log T}\right) \leq \frac{2}{T^2}. \quad (24)$$

Since $n_0 = \omega(\log T)$, we have

$$\left|V_k\right| = \alpha_k n_0 (1 + o(1)) \quad \text{for all } k \in [K]. \quad (25)$$

Then, $\left|V_k\right| \geq \bar{N}$ for all $1 \leq k \leq K$ since $n_0 = \omega(\log(m) + \frac{T}{m})$. Therefore, all users arriving after the exploration phase could be potentially recommended by items from a single cluster V_k without repetition.

Recall the procedure used by EC-UCS to cluster items in \mathcal{S} . For the $10m$ first user arrivals, it recommends items from \mathcal{S} uniformly at random. These $10m$ recommendations and the corresponding user responses are recorded in the dataset \mathcal{D} . From the dataset \mathcal{D} , the item clusters are extracted using a spectral algorithm (see Algorithm 4). This algorithm is taken from [33], and considers the *indirect edges* between items created by users. Specifically, when a user appears more than twice in \mathcal{D} , she creates an indirect edge between the items recommended to her for which she provided the same answer (1 or 0). Items with indirect edges are more likely to belong to the same cluster.

Algorithm 4 builds an adjacency matrix A from indirect edges. From Lemma 6 (presented in Appendix B), we know that at least $m/2$ users arrive twice in the first $10m$ arrivals with probability at least $1 - \frac{1}{T}$. We conclude that the construction of A is equivalent to a stochastic block model with random sampling where the number of vertices is n_0 , the sampling budget is $s \geq m/2$. We establish in the next theorem that this budget is enough to reconstruct the clusters V_1, \dots, V_K exactly using Algorithm 4. Theorem 9 is proved in Appendix J.2.

Theorem 9. *Let $\hat{I}_1, \dots, \hat{I}_K$ be the output of Algorithm 4. With probability $1 - \frac{1}{T}$, there exists permutation Γ such that*

$$\left| \bigcup_{k=1}^K (\hat{I}_{\Gamma(k)} \setminus V_k) \right| = 0.$$

I.2 Regret of EC-UCS: Proof of Theorem 6

The first component of the regret of EC-UCS is generated during the exploration phase for item clustering. This component is $\mathcal{O}(m)$. Then in view of Theorem 9, errors in item clustering cannot generate more than a $\mathcal{O}(1)$ regret. Hence, in what follows, we always assume that after the item clustering phase, we have:

$$\left| \bigcup_{k=1}^K (\hat{I}_{\Gamma(k)} \setminus V_k) \right| = 0.$$

Without loss of generality, we assume that $\Gamma(k) = k$ in the remaining of this section. After the item clustering phase, there are four sources of regret referred to as: 1. Exploration for user clustering, 2. Arrival of reference users, 3. User clustering, and 4. Optimistic assignments.

1. Exploration for user clustering. The regret induced by exploration of the users in \mathcal{U}_0 until $t \leq (10 + \log T)m$ is $\frac{m \log T}{\log T} = m$. Hence, the regret due to this step is:

$$\mathcal{O} \left(m \sum_{\ell} \beta_{\ell} (p_{\sigma_{\ell}(1)\ell} - p_{\sigma_{\ell}(K)\ell}) \right).$$

2. Arrival of reference users. If the users in \mathcal{U}^* have not arrived enough times, the algorithm cannot cluster them as intended, and this generates regret. Let n_u denote the number of times user u has arrived (until a time that will always be specified).

We define the event $\mathcal{E}_{top}^{(i)} = \{\exists u \in \mathcal{U}^* \text{ such that } n_u \leq \frac{(9+2^i)}{2} \text{ at } t = (9+2^i)m\}$ for $0 \leq i \leq \log_2 \log T$. Then, by Lemma 14, the regret due to $\mathcal{E}_{top}^{(i)}$ until $t = \lfloor (10 + \log T)m \rfloor$ is,

$$\begin{aligned} & R_{ref}(\lfloor (10 + \log T)m \rfloor) \\ & \leq \sum_{\ell} \beta_{\ell} (p_{\sigma_{\ell}(1)\ell} - p_{\sigma_{\ell}(K)\ell}) \sum_{i=0}^{\lfloor \log_2 \log T \rfloor} 2^{i+1} m \left(\frac{me}{(\log T)^3} \right)^{(\log T)^2} \exp \left(-\frac{(9+2^i)m}{16 \log T} \right) \\ & \leq \sum_{\ell} \beta_{\ell} (p_{\sigma_{\ell}(1)\ell} - p_{\sigma_{\ell}(K)\ell}) m \left(\frac{me}{(\log T)^3} \right)^{(\log T)^2} \int_{-1}^{\log_2 \log T} 2^{x+1} \exp \left(-\frac{(9+2^x)m}{16 \log T} \right) dx \\ & = \sum_{\ell} \beta_{\ell} (p_{\sigma_{\ell}(1)\ell} - p_{\sigma_{\ell}(K)\ell}) m \left(\frac{me}{(\log T)^3} \right)^{(\log T)^2} \int_{1/2}^{\log T} \frac{2}{\log 2} \exp \left(-\frac{(9+y)m}{16 \log T} \right) dy \\ & = \sum_{\ell} \beta_{\ell} (p_{\sigma_{\ell}(1)\ell} - p_{\sigma_{\ell}(K)\ell}) m \left(\frac{me}{(\log T)^3} \right)^{(\log T)^2} \left[-\frac{32 \log T}{m \log 2} \exp \left(-\frac{(9+y)m}{16 \log T} \right) \right]_{1/2}^{\log T} \\ & = \sum_{\ell} \beta_{\ell} (p_{\sigma_{\ell}(1)\ell} - p_{\sigma_{\ell}(K)\ell}) \exp \left(-\Theta \left(\frac{m}{\log T} \right) \right). \end{aligned} \quad (26)$$

where we have used the assumption $m^2 \geq T(\log T)^3$.

Also, $\mathbb{P}(\mathcal{E}_{top}^{(\lfloor \log_2 \log T \rfloor)}) \leq 2 \left(\frac{me}{(\log T)^3} \right)^{(\log T)^2} \exp \left(-\frac{(9+\frac{1}{2} \log T)m}{4 \log T} \right) = \exp(-\Theta(m))$. Hence, in view of (26), the regret due to $\mathcal{E}_{top}^{(i)}$ satisfies

$$R_{ref}(T) \leq \sum_{\ell} \beta_{\ell} (p_{\sigma_{\ell}(1)\ell} - p_{\sigma_{\ell}(K)\ell}) \exp \left(-\Theta \left(\frac{m}{\log T} \right) \right). \quad (27)$$

Let $\mathcal{B}_1 = \left(\bigcup_{i=0}^{\lfloor \log_2 \log T \rfloor} \mathcal{E}_{top}^{(i)} \right)^c$, i.e., \mathcal{B}_1 correspond to the event where every $u \in \mathcal{U}^*$ has arrived $n_u > \frac{(9+2^i)}{2}$ times at $t = (9+2^i)m$ for all i . Then, from (27), the regret because of \mathcal{B}_1^c is

$$R_{ref}(T) \leq \exp \left(-\Theta \left(\frac{m}{\log T} \right) \right). \quad (28)$$

3. User clustering. The size of \mathcal{U}^* is sufficiently large, so that \mathcal{U}^* consists of users from all clusters. More precisely, for a well-chosen $\epsilon_1 > 0$, Lemma 15 states that with probability $1 - 1/T$, the size of $\tilde{\mathcal{U}}_{\ell} = \mathcal{U}^* \cap \mathcal{U}_{\ell}$ is greater than $\beta_{\ell}(1 - \epsilon_1)(\log T)^2$ for all ℓ . Let $\mathcal{B}_2 = \{\tilde{\mathcal{U}}_{\ell} \geq \beta_{\ell}(1 - \epsilon_1)(\log T)^2 \text{ for } 1 \leq \ell \leq L\}$. By Lemma 15, the expected regret due to the event \mathcal{B}_2^c is $\mathcal{O}(1)$.

We now assume that both \mathcal{B}_1 and \mathcal{B}_2 holds throughout the remaining of the proof.

Under \mathcal{B}_1 , we have numerous observations for users in \mathcal{U}^* . Hence, most of users in \mathcal{U}^* have their empirical average success rate vector concentrated around the true parameter vector when t is large. Therefore, under \mathcal{B}_2 , the clustering step can learn the hidden parameters very accurately (since there are clear user clusters). We formalize this observation below. Consider $t \geq T_0 = \lceil Cm \rceil$ where

$$C = \max \left(\frac{512K^3}{\min(y_{\ell r}, \delta)^2} \log \left(\frac{16K^{\frac{3}{2}}}{\min(y_{\ell r}, \delta)} \right), \frac{2\sqrt{K}}{\min_{\ell} \beta_{\ell}} \right).$$

Then, we have

$$(C1) \quad K \sqrt{\frac{8Km}{T_0} \log \frac{T_0}{m}} < \frac{1}{4} \min(y_{\ell r}, \delta),$$

$$(C2) \quad (1 - \epsilon_1) \left(1 - 2K \left(\frac{m}{T_0}\right)^2\right) \min_{\ell} \beta_{\ell} > \frac{m}{T_0}.$$

Recall that in EC-UCS (see the pseudo-code), we use a parameter $\epsilon > 0$ when clustering users. This parameter is fixed and equal to $\epsilon = K \sqrt{\frac{8Km}{t} \log \frac{t}{m}}$. From Lemma 16, under (C1) and (C2), we have

$$\|p_{\ell} - \hat{p}_{\ell}\| < \epsilon < \frac{1}{4} \min(y_{\ell r}, \delta) \quad \forall t \geq T_0 \quad \text{with probability} \quad 1 - \frac{2}{T}.$$

Hence after T_0 rounds, the algorithm has accurate estimates of the parameters. We include T_0 in the regret upper bound, but can then assume that $\|p_{\ell} - \hat{p}_{\ell}\| < \epsilon < \frac{1}{4} \min(y_{\ell r}, \delta)$ for all $t \geq T_0$ in the remaining of the proof. The expected regret generated by the complement of this event is $\mathcal{O}(1)$.

4. Optimistic assignments. When $\|p_{\ell} - \hat{p}_{\ell}\| < \epsilon$, the algorithm can exploit the learned parameters. Suppose $u_t \in \mathcal{U}_{\ell}$. Recall the notation: $\mathcal{L}(u_t) \leftarrow \{\ell \in [L_0] : \sum_{k=1}^K n_{ku_t} x_{k\ell}^2 < 2K \log n_{u_t}\}$ used in the pseudo-code of EC-UCS.

Since the probability of the event $\{\ell \notin \mathcal{L}(u_t)\}$ decreases rapidly with the number of arrivals of u_t , the regret induced by this event is $\mathcal{O}(1)$. Since $\ell \in \mathcal{L}(u_t)$ holds most of the time, the algorithm recommend optimal items in item cluster k_{ℓ}^* at least $\frac{n_{u_t}}{2K}$ times. If $r \notin \mathcal{L}^{\perp}(\ell)$, we can distinguish \mathcal{U}_{ℓ} from \mathcal{U}_r with the constant number of recommendations of optimal item k_{ℓ}^* , since $p_{k_{\ell}^* \ell} \neq p_{k_{\ell}^* r}$. On the other hand, if $r \in \mathcal{L}^{\perp}(\ell)$, the algorithm cannot distinguish them unless it plays suboptimal items. Actually, suboptimal items should be played at most $\mathcal{O}(\log \bar{N})$ times in expectation. We make the above observations precise in Lemma 17, from which we conclude that the regret generated in this phase is:

$$\mathcal{O} \left(m \sum_{\ell} \beta_{\ell} (p_{\sigma_{\ell}(1)\ell} - p_{\sigma_{\ell}(K)\ell}) \left(\sum_{r \in \mathcal{R}_{\ell} \setminus \mathcal{L}^{\perp}(\ell)} \frac{K^2 \log K}{\phi(|p_{k_{\ell}^* r} - p_{k_{\ell}^* \ell}|^2)} + \sum_{k \in \mathcal{S}_{\ell r}} \sum_{r \in \mathcal{L}^{\perp}(\ell)} \frac{K \log \bar{N}}{|\mathcal{S}_{\ell r}| |p_{k\ell} - p_{kr}|^2} \right) \right)$$

Overall, accounting for all the regret sources, we have established that the regret of EC-UCS is:

$$\begin{aligned} & \mathcal{O} \left(m \sum_{\ell} \beta_{\ell} (p_{\sigma_{\ell}(1)\ell} - p_{\sigma_{\ell}(K)\ell}) \left(\max \left(\frac{K^3 \log K}{\phi(\min(y_{\ell r}, \delta)^2)}, \frac{\sqrt{K}}{\min_{\ell} \beta_{\ell}} \right) \right. \right. \\ & \quad \left. \left. + \sum_{r \in \mathcal{R}_{\ell} \setminus \mathcal{L}^{\perp}(\ell)} \frac{K^2 \log K}{\phi(|p_{k_{\ell}^* r} - p_{k_{\ell}^* \ell}|^2)} + \sum_{k \in \mathcal{S}_{\ell r}} \sum_{r \in \mathcal{L}^{\perp}(\ell)} \frac{K \log \bar{N}}{|\mathcal{S}_{\ell r}| |p_{k\ell} - p_{kr}|^2} \right) \right) \end{aligned}$$

I.3 Technical lemmas for the proof of Theorem 6

Lemma 12 (Cramer's theorem). *For any i.i.d. sequence $(X_i)_{i \geq 1}$ of real r.v. and any closed set $F \subseteq \mathbb{R}$,*

$$\mathbb{P} \left(\frac{1}{n} \sum_{i=1}^n X_i \in F \right) \leq 2 \exp \left(-n \inf_{x \in F} I(x) \right),$$

where $I(a) = \sup_{\theta \in \mathbb{R}} (\theta a - \log E[e^{\theta X}])$.

Lemma 13. *In the 'Exploitation' step (see EC-UCS pseudo-code), the expected regret generated by the exploration of $u_t \in \mathcal{U}_0$ until $t = \lfloor (10 + m) \log T \rfloor$ is $\mathcal{O}(m)$.*

Proof. Since the probability that a user from \mathcal{U}_0 arrives for each time t is $\frac{1}{\log T}$, the regret induced when exploring for users in \mathcal{U}_0 is $\mathcal{O}(\frac{m \log T}{\log T}) = \mathcal{O}(m)$. \square

Lemma 14. *In the 'Exploitation' step, with probability $1 - 2 \left(\frac{me}{(\log T)^3}\right)^{(\log T)^2} \exp\left(-\frac{t}{16 \log T}\right)$, at least $\lfloor (\log T)^2 \rfloor$ users in \mathcal{U}_0 arrive at least $\lfloor \frac{t}{2m} \rfloor$ times within the first t arrivals.*

Proof. We denote by $N_u(t)$ the number of times a user u has arrived in the first t arrivals. For any set $A \subset \mathcal{U}_0$, let $N_A(t)$ denote the total number of arrivals of users in A among the first t arrivals.

We write the probability that less than $(\log T)^2$ users in \mathcal{U}_0 arrive for $t/2m$ times in the t first arrivals as:

$$\begin{aligned}
& \mathbb{P}\left[\sum_{u \in \mathcal{U}_0} \mathbb{1}\{N_u(t) \geq \frac{t}{2m}\} < (\log T)^2\right] = \mathbb{P}\left[\sum_{u \in \mathcal{U}_0} \mathbb{1}\{N_u(t) < \frac{t}{2m}\} \geq |\mathcal{U}_0| - (\log T)^2\right] \\
& \leq \mathbb{P}[\exists A \subset \mathcal{U}_0 : |A| = |\mathcal{U}_0| - (\log T)^2, \forall u \in A, N_u(t) \leq \frac{t}{2m}] \\
& \leq \mathbb{P}[\exists A \subset \mathcal{U}_0 : |A| = |\mathcal{U}_0| - (\log T)^2, N_A(t) \leq \frac{t(|\mathcal{U}_0| - (\log T)^2)}{2m}] \tag{29} \\
& \stackrel{(a)}{\leq} 2 \binom{|\mathcal{U}_0|}{(\log T)^2} \exp\left(-\frac{t}{4} \log\left(1 + \frac{p}{2-2p}\right)\right) \\
& \stackrel{(b)}{\leq} 2 \left(\frac{me}{(\log T)^3}\right)^{(\log T)^2} \exp\left(-\frac{t}{16 \log T}\right),
\end{aligned}$$

where (b) follows from $\log(1+x) > \frac{x}{2}$ for $0 < x < 1$ and (a) can be proved using lemma 12. For simplicity, we define i.i.d. random variables $X_i \sim \text{Bern}(p)$ where $p = \frac{|\mathcal{U}_0| - (\log T)^2}{m}$. Then, $I(a) = a \log \frac{a(1-p)}{p(1-a)} - \log \frac{1-p}{1-a}$. Since $I(a)$ is a decreasing function in $(-\infty, p]$, $\inf_{a \leq \frac{p}{2}} I(a) = I\left(\frac{p}{2}\right)$. Therefore, (29) can be rewritten as:

$$\begin{aligned}
& \mathbb{P}\left(\frac{1}{t} \sum_{i=1}^t X_i \leq \frac{p}{2}\right) \leq 2 \exp\left(-tI\left(\frac{p}{2}\right)\right) \\
& = 2 \exp\left(-t \left(\log \frac{2-p}{2-2p} + \frac{p}{2} \log \frac{1-p}{2-p}\right)\right) \\
& \stackrel{(a)}{\leq} 2 \exp\left(-\frac{t}{4} \log \frac{2-p}{2-2p}\right).
\end{aligned}$$

where (a) holds since $\log(1+x) \geq x - \frac{x^2}{2}$ for $0 < x < 1$. \square

Lemma 15. Fix $\epsilon_1 = \sqrt{\frac{\log TK}{2(\min_{\ell} \beta_{\ell})^2 (\log T)^2}}$. Let $\tilde{\mathcal{U}}_{\ell} = \mathcal{U}^* \cap \mathcal{U}_{\ell}$ and $\mathcal{B}_2 = \{|\tilde{\mathcal{U}}_{\ell}| \geq \beta_{\ell}(1 - \epsilon_1)(\log T)^2 \text{ for } 1 \leq \ell \leq L\}$. Then, $\mathbb{P}(\mathcal{B}_2) \geq 1 - \frac{1}{T}$.

Proof. Since $\mathbb{P}(u \in \tilde{\mathcal{U}}_{\ell} | u \in \mathcal{U}^*) = \mathbb{P}(u \in \mathcal{U}_{\ell} | u \in \mathcal{U}^*) = \mathbb{P}(u \in \mathcal{U}_{\ell}) = \beta_{\ell}$,

$$\begin{aligned}
& \mathbb{P}\left(|\tilde{\mathcal{U}}_{\ell}| \leq \beta_{\ell}(1 - \epsilon_1)(\log T)^2\right) \stackrel{(a)}{\leq} \exp\left(-(\log T)^2 \text{kl}((1 - \epsilon_1)\beta_{\ell}, \beta_{\ell})\right) \\
& \stackrel{(b)}{\leq} \exp\left(-(\log T)^2 2\epsilon_1^2 \beta_{\ell}^2\right) \\
& \leq \frac{1}{TK}.
\end{aligned}$$

where (a) is from Chernoff-Hoeffding bound (b) is from Pinsker's inequality.

Hence, the event $\mathcal{B}_2 = \{|\tilde{\mathcal{U}}_{\ell}| \geq \beta_{\ell}(1 - \epsilon_1)(\log T)^2 \text{ for } 1 \leq \ell \leq L\}$ holds with probability at least $1 - \frac{1}{T}$. \square

In the remaining of this section, we fix $\epsilon_1 = \sqrt{\frac{\log TK}{2(\min_{\ell} \beta_{\ell})^2 (\log T)^2}}$ as chosen in the previous lemma.

Lemma 16. In the 'Exploitation' step, under \mathcal{B}_1 and \mathcal{B}_2 , if t is large enough to satisfy the conditions $\epsilon < \frac{1}{2} \min_{\ell \neq \ell'} \|p_{\ell} - p_{\ell'}\|$ and $(1 - \epsilon_1) \left(1 - 2K \left(\frac{m}{t}\right)^2\right) \min_{\ell} \beta_{\ell} > \frac{m}{t}$, then $\|p_{\ell} - \hat{p}_{\ell}\| < \epsilon$ with probability at least $1 - \frac{2}{T}$.

Proof. Recall that $\epsilon = K\sqrt{\frac{8Km}{t} \log \frac{t}{m}}$ and $Q_u = \{v \in \mathcal{U}^* : \|\hat{\rho}_u - \hat{\rho}_v\| \leq \epsilon\}$ for all $u \in \mathcal{U}^*$. We define a set \mathcal{C}_ℓ for $1 \leq \ell \leq L$ as: $\mathcal{C}_\ell = \{u \in \mathcal{U}^* : \|p_\ell - \hat{\rho}_u\| \leq \frac{\epsilon}{2}\}$. This set has the following properties:

- (i) $|\mathcal{C}_\ell| = \Omega((\log T)^2)$ with probability at least $1 - \frac{1}{T}$. This follows from the following argument.

$$\begin{aligned} \mathbb{P}(u \in \mathcal{C}_\ell) &\geq \mathbb{P}(u \in \mathcal{C}_\ell | u \in \mathcal{U}_\ell) \mathbb{P}(u \in \mathcal{U}_\ell) \\ &\stackrel{(a)}{\geq} \beta_\ell (1 - \epsilon_1) \mathbb{P}\left(\|p_\ell - \hat{\rho}_u\| \leq \frac{\epsilon}{2} \mid u \in \mathcal{U}_\ell\right) \\ &\stackrel{(b)}{\geq} \beta_\ell (1 - \epsilon_1) \left(1 - 2 \exp\left(-2 \frac{t}{2Km} \left(\frac{\epsilon}{2K}\right)^2\right)\right)^K \\ &\geq \beta_\ell (1 - \epsilon_1) \left(1 - 2K \exp\left(-\frac{t\epsilon^2}{4K^3m}\right)\right) \\ &\geq \beta_\ell (1 - \epsilon_1) \left(1 - 2K \left(\frac{m}{t}\right)^2\right), \end{aligned}$$

where (a) follows from the assumption that \mathcal{B}_2 holds and (b) stems from \mathcal{B}_1 and Chernoff-Hoeffding's bound.

Let $r = \beta_\ell (1 - \epsilon_1) \left(1 - 2K \left(\frac{m}{t}\right)^2\right)$. Then,

$$\begin{aligned} \mathbb{P}\left(|\mathcal{C}_\ell| < \left(r - \frac{1}{\sqrt{2 \log T}}\right) (\log T)^2\right) &\leq \exp\left(-(\log T)^2 \text{kl}\left(r - \frac{1}{\sqrt{2 \log T}}, r\right)\right) \\ &\leq \exp\left(-2(\log T)^2 \left(\frac{1}{\sqrt{2 \log T}}\right)^2\right) \\ &\leq \frac{1}{T}. \end{aligned}$$

Therefore, $|\mathcal{C}_\ell| = \Omega((\log T)^2)$ with probability at least $1 - \frac{1}{T}$.

- (ii) $|\mathcal{U}^* \setminus (\cup_{\ell=1}^L \mathcal{C}_\ell)| = \mathcal{O}\left(\frac{m(\log T)^2}{t}\right)$ with probability at least $1 - \frac{1}{T}$. To show this, we use a similar argument as in (i):

$$\begin{aligned} \mathbb{P}(u \in \mathcal{U}^* \setminus (\cup_{\ell=1}^L \mathcal{C}_\ell)) &\leq \sum_{\ell=1}^L \mathbb{P}(u \in \mathcal{U}_\ell) \mathbb{P}(u \in \mathcal{U}^* \setminus (\cup_{\ell=1}^L \mathcal{C}_\ell) | u \in \mathcal{U}_\ell) \\ &\leq \sum_{\ell=1}^L \mathbb{P}(u \in \mathcal{U}_\ell) \mathbb{P}\left(\|p_\ell - \hat{\rho}_u\| > \frac{\epsilon}{2} \mid u \in \mathcal{U}_\ell\right) \\ &\leq \sum_{\ell=1}^L \mathbb{P}(u \in \mathcal{U}_\ell) 2K \left(\frac{m}{t}\right)^2 \\ &= 2K \left(\frac{m}{t}\right)^2. \end{aligned}$$

Then, the probability that the size of $|\mathcal{U}^* \setminus (\cup_{\ell=1}^L \mathcal{C}_\ell)|$ is greater than $\frac{m(\log T)^2}{t}$ is,

$$\begin{aligned} \mathbb{P}\left(|\mathcal{U}^* \setminus (\cup_{\ell=1}^L \mathcal{C}_\ell)| \geq \frac{m(\log T)^2}{t}\right) &\leq \exp\left(-(\log T)^2 \text{kl}\left(\frac{m}{t}, 2K \left(\frac{m}{t}\right)^2\right)\right) \\ &\stackrel{(a)}{\leq} \exp\left(-(\log T)^2 \frac{m}{t} \log \frac{t}{m}\right) \\ &\stackrel{(b)}{\leq} \frac{1}{T}, \end{aligned}$$

where (a) is obtained from Lemma 4 and $t \geq 2Km$ and (b) is from $t \leq m \log T$.

- (iii) If $|\mathcal{C}_\ell \cap Q_u| \geq 1$, then $|\mathcal{C}_m \cap Q_u| = 0$ for all $\ell \neq m$. Because $\|\hat{\rho}_u - \hat{\rho}_v\| \geq \|p_\ell - p_m\| - \|p_\ell - \hat{\rho}_u\| - \|p_m - \hat{\rho}_v\| \geq \|p_\ell - p_m\| - \epsilon > \epsilon$ for $u \in \mathcal{C}_\ell$ and $j \in \mathcal{C}_m$, where the last inequality follows from $2\epsilon < \min_{\ell \neq \ell'} \|p_\ell - p_{\ell'}\|$.
- (iv) $\mathcal{C}_\ell \subset Q_u$ for all $u \in \mathcal{C}_\ell$, since $\|\hat{\rho}_u - \hat{\rho}_v\| \leq \|\hat{\rho}_u - p_\ell\| + \|\hat{\rho}_v - p_\ell\| \leq \epsilon$ for all $v \in \mathcal{C}_\ell$.

From the properties (iii) and (iv), there exists an item $u \in (\cup_{\ell=1}^L \mathcal{C}_\ell) \setminus (\cup_{r=1}^{\ell-1} Q_{i_r})$ such that $|Q_u \setminus (\cup_{r=1}^{\ell-1} Q_{i_r})| \geq m_\ell$. Here, $m_\ell = \Omega((\log T)^2)$ from property (i).

We also have $|Q_v| = \mathcal{O}(\frac{m(\log T)^2}{t})$ for v such that $|Q_v \cap (\cup_{k=1}^K \mathcal{C}_k)| = 0$ from property (ii). Since we assume $(1 - \epsilon_1) \left(1 - 2K \left(\frac{m}{t}\right)^2\right) \min_\ell \beta_\ell > \frac{m}{t}$, the item v cannot be chosen as i_k .

We can conclude that $\|p_\ell - \hat{p}_\ell\| \leq \epsilon$ with probability $1 - 2/T$, since $\|\hat{\rho}_u - p_\ell\| \leq \epsilon$ when $|Q_u \cap \mathcal{C}_\ell| \geq 1$. \square

Lemma 17. *If $\|p_\ell - \hat{p}_\ell\| < \frac{1}{4} \min(y_{\ell r}, \delta)$ for all $r \neq \ell$, the regret due to recommendations based on optimistic user assignments is,*

$$\mathcal{O} \left(m \sum_{\ell} \beta_\ell (p_{\sigma_\ell(1)\ell} - p_{\sigma_\ell(K)\ell}) \left(\sum_{r \in \mathcal{R}_\ell \setminus \mathcal{L}^\perp(\ell)} \frac{K^2 \log K}{\phi(|p_{k_\ell^* r} - p_{k_\ell^* \ell}|^2)} + \sum_{k \in \mathcal{S}_{\ell r}} \sum_{r \in \mathcal{L}^\perp(\ell)} \frac{K \log \bar{N}}{|\mathcal{S}_{\ell r}| |p_{k\ell} - p_{kr}|^2} \right) \right).$$

Proof. Recall that $x_{k\ell} = \max\{|\hat{p}_{k\ell} - \hat{\rho}_{ku_\ell}| - \epsilon, 0\}$ and $\mathcal{L}^\perp(\ell) = \{\ell' \neq \ell : k_\ell^* \neq k_{\ell'}^*, p_{k_\ell^* \ell} = p_{k_\ell^* \ell'}\}$. We take $\epsilon < \frac{1}{4} \min(y_{\ell r}, \delta)$ to satisfy the condition $\|p_\ell - \hat{p}_\ell\| < \frac{1}{4} \min(y_{\ell r}, \delta)$. When we make a recommendation to $u_t \in \mathcal{U}_\ell$ by referring to their neighbors from \mathcal{U}^* , if regret is generated, the following event holds.

$$\begin{aligned} \mathcal{E}_N &= \left\{ \sum_{k=1}^K n_{ku_t} x_{k\ell}^2 > 2K \log n_{u_t} \right\} \\ &\quad \cup \left(\bigcup_{r \neq \ell} \left\{ \sum_{k=1}^K n_{ku_t} x_{kr}^2 < 2K \log n_{u_t} \right\} \cap \left\{ \arg \max_k p_{k\ell} \neq \arg \max_k p_{kr} \right\} \right) \\ &= \left\{ \sum_{k=1}^K n_{ku_t} x_{k\ell}^2 > 2K \log n_{u_t} \right\} \cup \left(\bigcup_{r \in \mathcal{R}_\ell} \left\{ \sum_{k=1}^K n_{ku_t} x_{kr}^2 < 2K \log n_{u_t} \right\} \right) \\ &:= \mathcal{E}_{N_1} \cup \left(\bigcup_{r \in \mathcal{R}_\ell} \mathcal{E}_{N_2}^{(r)} \right) \end{aligned}$$

First, an upper bound of the probability of the event \mathcal{E}_{N_1} is

$$\begin{aligned}
\mathbb{P}(\mathcal{E}_{N_1}) &\leq \mathbb{P}\left(\bigcup_{k=1}^K \{n_{ku_t} x_{kl}^2 > 2 \log n_{u_t}\}\right) \\
&\leq \sum_{k=1}^K \mathbb{P}(n_{ku_t} x_{kl}^2 > 2 \log n_{u_t}) \\
&\leq \sum_{k=1}^K \mathbb{P}\left(|\hat{p}_{kl} - \hat{\rho}_{ku_t}| - \epsilon > \sqrt{\frac{2 \log n_{u_t}}{n_{ku_t}}}\right) \\
&\leq \sum_{k=1}^K \mathbb{P}\left(|\hat{p}_{kl} - \hat{\rho}_{ku_t}| - |p_{kl} - \hat{p}_{kl}| > \sqrt{\frac{2 \log n_{u_t}}{n_{ku_t}}}\right) \\
&\leq \sum_{k=1}^K \mathbb{P}\left(|p_{kl} - \hat{\rho}_{ku_t}| > \sqrt{\frac{2 \log n_{u_t}}{n_{ku_t}}}\right)
\end{aligned}$$

By Lemma 1, we know that u_t arrives at most \bar{N} times in expectation. Therefore, the regret induced by the event \mathcal{E}_{N_1} is

$$\begin{aligned}
R_{\mathcal{E}_{N_1}}(\bar{N}) &\leq \sum_{k=1}^K \sum_{s=1}^{\bar{N}} \mathbb{P}\left(|p_{kl} - \hat{\rho}_{ku_t}| > \sqrt{\frac{2 \log s}{s}} \mid n_{u_t} = s\right) \\
&\leq \sum_{k=1}^K \sum_{s=1}^{\bar{N}} 2 \exp(-4 \log s) \\
&\leq 2K \left(1 + \int_1^{\bar{N}} \frac{1}{x^4} dx\right) \\
&\leq 2K \left(1 + \left[-\frac{1}{3x^3}\right]_1^{\bar{N}}\right) = \frac{2K}{3} \left(4 - \frac{1}{\bar{N}^3}\right) \tag{30}
\end{aligned}$$

Next, we evaluate the probability of the event $\mathcal{E}_{N_2}^{(r)}$. First, we assume that $r \notin \mathcal{L}^\perp(\ell)$. Then,

$$\begin{aligned}
\mathbb{P}\left(\max_{\{1 \leq n_{ku_t} \leq n_{u_t}\}} n_{ku_t} x_{kl}^2 > 2 \log \left(\frac{n_{u_t}}{2}\right)\right) &\leq \sum_{s=1}^{n_{u_t}} \mathbb{P}\left(|\hat{p}_{kl} - \hat{\rho}_{ku_t}| - \epsilon > \sqrt{\frac{2 \log(\frac{n_{u_t}}{2})}{n_{ku_t}}}\right) \Big| n_{ku_t} = s \\
&\leq \sum_{s=1}^{n_{u_t}} \mathbb{P}\left(|p_{kl} - \hat{\rho}_{ku_t}| > \sqrt{\frac{2 \log(\frac{n_{u_t}}{2})}{n_{ku_t}}}\right) \Big| n_{ku_t} = s \\
&\leq \sum_{s=1}^{n_{u_t}} 2 \exp\left(-4 \log \left(\frac{n_{u_t}}{2}\right)\right) = \frac{32}{n_{u_t}^3} \tag{31}
\end{aligned}$$

If the event $\{\forall k, \max_{\{1 \leq n_{ku_t} \leq n_{u_t}\}} n_{ku_t} x_{kl}^2 < 2 \log \left(\frac{n_{u_t}}{2}\right)\}$ occurs, then the event $\{n_{k_\ell^* u_t} \geq \frac{n_{u_t}}{2K}\}$ occurs as well. Hence, we can deduce that $\mathbb{P}\left(n_{k_\ell^* u_t} < \frac{n_{u_t}}{2K}\right) \leq \frac{32K}{n_{u_t}^3}$ by (31). The probability of the event $\mathcal{E}_{N_2}^{(r)}$ satisfies:

$$\begin{aligned}
\mathbb{P}(\mathcal{E}_{N_2}^{(r)}) &\leq \mathbb{P}\left(n_{k_\ell^* u_t} x_{k_\ell^* r}^2 < 2K \log n_{u_t}\right) \\
&\leq \mathbb{P}\left(x_{k_\ell^* r} < \sqrt{\frac{2K \log n_{u_t}}{n_{k_\ell^* u_t}}}\right) \\
&\leq \mathbb{P}\left(|\hat{p}_{k_\ell^* r} - \hat{\rho}_{k_\ell^* u_t}| < \sqrt{\frac{2K \log n_{u_t}}{n_{k_\ell^* u_t}}} + \epsilon\right) \\
&\leq \mathbb{P}\left(|p_{k_\ell^* r} - p_{k_\ell^* \ell}| - |p_{k_\ell^* \ell} - \hat{\rho}_{k_\ell^* u_t}| - |p_{k_\ell^* r} - \hat{p}_{k_\ell^* r}| < \sqrt{\frac{2K \log n_{u_t}}{n_{k_\ell^* u_t}}} + \epsilon\right) \\
&\leq \mathbb{P}\left(|p_{k_\ell^* \ell} - \hat{\rho}_{k_\ell^* u_t}| > |p_{k_\ell^* r} - p_{k_\ell^* \ell}| - \sqrt{\frac{2K \log n_{u_t}}{n_{k_\ell^* u_t}}} - 2\epsilon\right) \\
&\leq \mathbb{P}\left(|p_{k_\ell^* \ell} - \hat{\rho}_{k_\ell^* u_t}| > \frac{1}{2}|p_{k_\ell^* r} - p_{k_\ell^* \ell}| - \sqrt{\frac{2K \log n_{u_t}}{n_{k_\ell^* u_t}}}\right) \\
&\leq \mathbb{P}\left(|p_{k_\ell^* \ell} - \hat{\rho}_{k_\ell^* u_t}| > \frac{1}{2}|p_{k_\ell^* r} - p_{k_\ell^* \ell}| - \sqrt{\frac{4K^2 \log n_{u_t}}{n_{u_t}}}\right) + \frac{32K}{n_{u_t}^3}.
\end{aligned}$$

Therefore, the regret induced by event $\mathcal{E}_{N_2}^{(r)}$ is

$$\begin{aligned}
&R_{\{\mathcal{E}_{N_2}^{(r)} \setminus \mathcal{E}_{N_1}\}}(\bar{N}) \\
&\leq \sum_{s=1}^{\bar{N}} \mathbb{P}\left(|p_{k_\ell^* \ell} - \hat{\rho}_{k_\ell^* u_t}| > \frac{1}{2}|p_{k_\ell^* r} - p_{k_\ell^* \ell}| - \sqrt{\frac{4K^2 \log s}{s}}\right) + \frac{32K}{s^3} \\
&\leq \frac{64K^2}{|p_{k_\ell^* r} - p_{k_\ell^* \ell}|^2} \log\left(\frac{6K}{|p_{k_\ell^* r} - p_{k_\ell^* \ell}|}\right) \\
&\quad + \sum_{s=\lceil \frac{64K^2}{|p_{k_\ell^* r} - p_{k_\ell^* \ell}|^2} \log\left(\frac{6K}{|p_{k_\ell^* r} - p_{k_\ell^* \ell}|}\right) \rceil}^{\bar{N}} 2 \exp\left(-\frac{s|p_{k_\ell^* \ell} - p_{k_\ell^* r}|^2}{8}\right) + 48K \\
&\leq \frac{64K^2}{|p_{k_\ell^* r} - p_{k_\ell^* \ell}|^2} \log\left(\frac{6K}{|p_{k_\ell^* r} - p_{k_\ell^* \ell}|}\right) + \frac{16}{|p_{k_\ell^* \ell} - p_{k_\ell^* r}|^2} + 48K \\
&\leq \frac{128K^2}{|p_{k_\ell^* r} - p_{k_\ell^* \ell}|^2} \log\left(\frac{6K}{|p_{k_\ell^* r} - p_{k_\ell^* \ell}|}\right). \tag{32}
\end{aligned}$$

Next assume that $r \in \mathcal{L}^\perp(\ell)$. Then, $k_\ell^* \neq k_r^*$ and $p_{k_\ell^* \ell} = p_{k_\ell^* r}$, and

$$\begin{aligned}
\mathbb{P}(\mathcal{E}_{N_2}^{(r)}) &\leq \mathbb{P}\left(\sum_{k \in \mathcal{S}_{\ell r}} n_{ku_t} x_{kr}^2 < 2K \log n_{u_t}\right) \\
&\leq \mathbb{P}\left(\bigcup_{k \in \mathcal{S}_{\ell r}} \left\{n_{ku_t} x_{kr}^2 < \frac{2K \log n_{u_t}}{|\mathcal{S}_{\ell r}|}\right\}\right) \\
&\leq \sum_{k \in \mathcal{S}_{\ell r}} \mathbb{P}\left(x_{kr} < \sqrt{\frac{2K \log n_{u_t}}{|\mathcal{S}_{\ell r}| n_{ku_t}}}\right) \\
&\leq \sum_{k \in \mathcal{S}_{\ell r}} \mathbb{P}\left(|p_{k\ell} - \hat{p}_{ku_t}| > \frac{1}{2}|p_{kr} - p_{k\ell}| - \sqrt{\frac{2K \log n_{u_t}}{|\mathcal{S}_{\ell r}| n_{ku_t}}}\right) \\
&\leq \sum_{k \in \mathcal{S}_{\ell r}} \mathbb{P}\left(|p_{k\ell} - \hat{p}_{ku_t}| > \frac{1}{2}|p_{kr} - p_{k\ell}| - \sqrt{\frac{2K \log \bar{N}}{|\mathcal{S}_{\ell r}| n_{ku_t}}}\right). \\
R_{\mathcal{E}_{N_2}^{(r)}}(\bar{N}) &\leq \sum_{k \in \mathcal{S}_{\ell r}} \sum_{s=1}^{\bar{N}} \mathbb{P}\left(|p_{k\ell} - \hat{p}_{ku_t}| > \frac{1}{2}|p_{kr} - p_{k\ell}| - \sqrt{\frac{2K \log \bar{N}}{|\mathcal{S}_{\ell r}| s}} |n_{ku_t} = s\right) \\
&\stackrel{(a)}{\leq} \sum_{k \in \mathcal{S}_{\ell r}} \left(\frac{32K \log \bar{N}}{|\mathcal{S}_{\ell r}| |p_{k\ell} - p_{kr}|^2} + \sum_{s=\lceil \frac{32K \log \bar{N}}{|\mathcal{S}_{\ell r}| |p_{k\ell} - p_{kr}|^2} \rceil}^{\bar{N}} 2 \exp\left(-\frac{s|p_{k\ell} - p_{kr}|^2}{8}\right) \right) \\
&\leq \sum_{k \in \mathcal{S}_{\ell r}} \frac{32}{|p_{k\ell} - p_{kr}|^2} \left(\frac{K \log \bar{N}}{|\mathcal{S}_{\ell r}|} + 1 \right). \tag{33}
\end{aligned}$$

Combining (30), (32) and (33), the expected regret due to recommendations made by referring to the nearest neighbors in \mathcal{U}^* is,

$$\begin{aligned}
R_N(T) &\leq m \sum_{\ell} \beta_{\ell} (p_{\sigma_{\ell}(1)\ell} - p_{\sigma_{\ell}(K)\ell}) \left(\frac{2K}{3} \left(4 - \frac{1}{N^3} \right) \right. \\
&\quad \left. + \sum_{r \in \mathcal{R}_{\ell} \setminus \mathcal{L}^{\perp}(\ell)} \frac{128K^2}{|p_{k_\ell^* r} - p_{k_\ell^* \ell}|^2} \log \left(\frac{6K}{|p_{k_\ell^* r} - p_{k_\ell^* \ell}|} \right) + \sum_{r \in \mathcal{L}^{\perp}(\ell)} \sum_{k \in \mathcal{S}_{\ell r}} \frac{32}{|p_{k\ell} - p_{kr}|^2} \left(\frac{K \log \bar{N}}{|\mathcal{S}_{\ell r}|} + 1 \right) \right) \\
&= \mathcal{O} \left(m \sum_{\ell} \beta_{\ell} (p_{\sigma_{\ell}(1)\ell} - p_{\sigma_{\ell}(K)\ell}) \left(\sum_{r \in \mathcal{R}_{\ell} \setminus \mathcal{L}^{\perp}(\ell)} \frac{K^2 \log K}{\phi(|p_{k_\ell^* r} - p_{k_\ell^* \ell}|^2)} + \sum_{k \in \mathcal{S}_{\ell r}} \sum_{r \in \mathcal{L}^{\perp}(\ell)} \frac{K \log \bar{N}}{|\mathcal{S}_{\ell r}| |p_{k\ell} - p_{kr}|^2} \right) \right).
\end{aligned}$$

□

J Performance guarantees of ECB and Item Clustering: Proof of Theorems 7 and 9

J.1 Performance guarantees of ECB: Proof of Theorem 7

The proof is straightforward from the results of Theorem 9. Indeed, the latter implies that we can assume that the item clusters estimated from the first phase of the algorithm are exact (the complement of this event happens with probability $1/T$, and hence generates an expected regret $\mathcal{O}(1)$).

Hence we can assume that we know the exact clusters of items in \mathcal{S} . Now observe that for each cluster V_k (a subset of \mathcal{S}), we have $|V_k| \geq \bar{N}$ (refer to Appendix I for a precise statement). As a consequence, all users can be served using items from V_1, \dots, V_K , except for a few users arriving more than \bar{N} . Actually, from Lemma 1, these exceptional arrivals induce an average regret $\mathcal{O}(1)$. ECB applies, for each user, a UCB1 algorithm [1] to select the cluster from which an item is recommended. For this exploitation period, each user in \mathcal{U}_ℓ will then induce a regret $\mathcal{O}(\sum_{k \neq k_\ell^*} \frac{\log(\bar{N})}{p_{k_\ell^*} e^{-p_{k_\ell}}})$. This completes the proof. \square

J.2 Item Clustering Phase: Proof of Theorem 9

We let $e(v, S) := \sum_{x \in S} A_{vx}$ and $e(A, B) := \sum_{v \in A} \sum_{w \in B} A_{vw}$.

Let

$$p(a, b) := (|V_b| - \mathbb{1}\{a = b\}) \frac{2}{n_0(n_0 - 1)} \sum_{\ell=1}^L \beta_\ell p_{a\ell} p_{b\ell}$$

and $p(a, 0) := 1 - \sum_{k=1}^K p(a, k)$.

We also define $e(v, V_0)$ and $e(v, S_0^{(t)})$ as follows:

$$\begin{aligned} e(v, V_0) &:= s - \sum_{k=1}^K e(v, V_k) \quad \text{and} \\ e(v, S_0^{(t)}) &:= s - \sum_{k=1}^K e(v, S_k^{(t)}), \end{aligned}$$

where s is the number of users who have received recommendations at least twice until $t = 10m$.

Proof of the theorem. The proof of Theorem 9 relies on the following random matrix concentration inequality. Specifically, from the matrix Bernstein inequality, we can bound $\|A - \mathbb{E}[A]\|$ as follows.

Lemma 18. *Assume $T > 10m$. Let A be the adjacency matrix obtained in Algorithm 4. Let $\|\cdot\|$ denote the spectral norm. Then,*

$$\mathbb{P} \left(\|A - \mathbb{E}[A]\| > 5 \sqrt{\frac{m \log m}{n_0}} \right) \leq \frac{1}{m^2}.$$

From Lemma 18, we deduce that \hat{A} (the rank- K approximation of A) is approximately the same as $\mathbb{E}[A]$. Indeed, since both \hat{A} and $\mathbb{E}[A]$ are of rank K ,

$$\begin{aligned} \|\hat{A} - \mathbb{E}[A]\|_F^2 &\leq 2K \|\hat{A} - \mathbb{E}[A]\|^2 \\ &\leq 4K (\|\hat{A} - A\|^2 + \|A - \mathbb{E}[A]\|^2) \\ &\leq 8K \|A - \mathbb{E}[A]\|^2, \end{aligned}$$

where $\|A - \mathbb{E}[A]\|$ is negligible compared to $\|\mathbb{E}[A]\| = \Omega(\frac{m}{n_0})$.

From the columns of \hat{A} , we can classify items. Here, $\|\mathbb{E}[A]_v - \mathbb{E}[A]_w\| = \Omega(\frac{m}{n_0^{3/2}})$ when v and w belong to different clusters and $\mathbb{E}[A]_v = \mathbb{E}[A]_w$ when v and w belong to the same cluster. Therefore,

the columns of \hat{A} are concentrated around the correct cluster column unless $\|\hat{A}_v - \mathbb{E}[A]_v\| = \Omega(\frac{m}{n_0^{3/2}})$. From this argument, the spectral decomposition used in the algorithms satisfies

$$|\cup_{k=1}^K (S_k^{(0)} \setminus V_k)| = O\left(\frac{n_0^2 \log m}{m}\right), \quad (34)$$

since $\sum_{v \in \mathcal{S}} \|\hat{A}_v - \mathbb{E}[A]_v\|^2 = \|A - \mathbb{E}[A]\|_F^2 = O(\frac{m \log m}{n_0})$ from Lemma 18 (cf. [26]).

In the improvement step of Algorithm 5, the algorithm refines the result of Spectral Decomposition iteratively. We denote the set of misclassified items after t -th iteration by $\mathcal{E}^{(t)}$. We also introduce $\mathcal{E}_{k\ell}^{(t)} = S_k^{(t)} \cap V_\ell$ so that $\mathcal{E}^{(t)} = \cup_{k=1}^K \cup_{\ell: \ell \neq k} \mathcal{E}_{k\ell}^{(t)}$.

Since the items move to more likely cluster with respect to $\hat{p}(i, j)$ at each step,

$$\begin{aligned} 0 &\leq \sum_{k, \ell: k \neq \ell} \sum_{v \in \mathcal{E}_{k\ell}^{(t+1)}} \sum_{j=0}^K e(v, S_j^{(t)}) \log \frac{\hat{p}(k, j)}{\hat{p}(\ell, j)} \\ &\stackrel{(a)}{\leq} \sum_{k, \ell: k \neq \ell} \sum_{v \in \mathcal{E}_{k\ell}^{(t+1)}} \sum_{j=0}^K e(v, S_j^{(t)}) \log \frac{p(k, j)}{p(\ell, j)} \\ &\quad + C_1 |\mathcal{E}^{(t+1)}| \sqrt{\frac{m \log m}{n_0}} \\ &\stackrel{(b)}{\leq} \sum_{k, \ell: k \neq \ell} \sum_{v \in \mathcal{E}_{k\ell}^{(t+1)}} \sum_{j=0}^K e(v, V_j) \log \frac{p(k, j)}{p(\ell, j)} \\ &\quad + C_1 |\mathcal{E}^{(t+1)}| \sqrt{\frac{m \log m}{n_0}} + C_2 e(\mathcal{E}^{(t)}, \mathcal{E}^{(t+1)}) \\ &\stackrel{(c)}{\leq} \sum_{k, \ell: k \neq \ell} \sum_{v \in \mathcal{E}_{k\ell}^{(t+1)}} \sum_{j=0}^K e(v, V_j) \log \frac{p(k, j)}{p(\ell, j)} \\ &\quad + C_1 |\mathcal{E}^{(t+1)}| \sqrt{\frac{m \log m}{n_0}} + C_3 \sqrt{|\mathcal{E}^{(t)}| |\mathcal{E}^{(t+1)}| \frac{m \log m}{n_0}} \\ &\stackrel{(d)}{\leq} -C_4 \frac{m}{n_0} |\mathcal{E}^{(t+1)}| + C_3 \sqrt{|\mathcal{E}^{(t)}| |\mathcal{E}^{(t+1)}| \frac{m \log m}{n_0}} \end{aligned} \quad (35)$$

where (a) is obtained from Lemma 19; (b) stems from the fact that $\frac{p(k, j)}{p(\ell, j)}$ is a positive constant for all $1 \leq j \leq K$; (c) follows from Lemma 20; and (d) is obtained from Lemma 21.

From (35), we can conclude that:

$$\begin{aligned} \frac{|\mathcal{E}^{(t+1)}|}{|\mathcal{E}^{(t)}|} &\leq \frac{C_3^2 n_0 \log m}{C_4^2 m} \\ &\stackrel{(a)}{\leq} C_5 \frac{1}{\log T}, \end{aligned}$$

where (a) is from $n_0 \leq m/(\log T)^2$ and $10m < T$.

Therefore, after $\log(n_0)$ iterations, we have recovered the perfect clusters. \square

Nest we state the lemmas used in the proof above.

Lemma 19. When $|\mathcal{E}^{(0)}| = O\left(\frac{n_0^2 \log m}{m}\right)$, with probability $1 - \frac{2}{m^2}$,

$$\sum_{i=0}^K e(v, S_i^{(t)}) \left| \log \frac{p(k, i)}{\hat{p}(k, i)} \right| = O\left(\sqrt{\frac{m \log m}{n_0}}\right).$$

Lemma 20. When $\|A - \mathbb{E}[A]\| = O\left(\sqrt{\frac{m \log m}{n_0}}\right)$, $|\mathcal{E}^{(t)}| = O\left(\frac{n_0^2 \log m}{m}\right)$, and $|\mathcal{E}^{(t+1)}| = O\left(\frac{n_0^2 \log m}{m}\right)$,

$$\sum_{v \in \mathcal{E}^{(t+1)}} e(v, \mathcal{E}^{(t)}) = O\left(\sqrt{|\mathcal{E}^{(t)}| |\mathcal{E}^{(t+1)}| \frac{m \log m}{n_0}}\right).$$

Lemma 21. With probability at least $1 - \frac{1}{m^2}$, for all k , for all $v \in V_k$,

$$\sum_{a=0}^K e(v, V_a) \log\left(\frac{p(k, a)}{p(k', a)}\right) = \Omega\left(\frac{m}{n_0}\right) \quad \text{for all } k' \neq k.$$

J.3 Proof of the lemmas

Proof of Lemma 18.

The adjacency matrix A can be considered as the sum of s samples of connected pairs for some $s \geq \frac{m}{2}$. We denote such samples as X_ℓ for $1 \leq \ell \leq s$. Then, $A = \sum_{\ell=1}^s X_\ell$.

By the matrix Bernstein inequality (cf. [30]), we have:

$$\begin{aligned} \mathbb{P}(\|A - \mathbb{E}[A]\| > t) &= \mathbb{P}\left(\left\|\sum_{\ell=1}^s (X_\ell - \mathbb{E}[X_\ell])\right\| > t\right) \\ &\leq n_0 \exp\left(-\frac{t^2}{2(\sigma^2(A) + Dt/3)}\right), \end{aligned} \quad (36)$$

where $\sigma^2(A) = \|\mathbb{E}[(A - \mathbb{E}[A])^2]\|$ and D is an upper bound of $\|X_\ell - \mathbb{E}[X_\ell]\|$ for all ℓ .

The expectation of the elements x_{ij} of matrix X_ℓ is

$$\mathbb{E}[x_{ij} | i \in V_k, j \in V_{k'}] = \frac{2}{n_0(n_0 - 1)} \sum_{\ell=1}^L \beta_\ell p_{k\ell} p_{k'\ell}.$$

Since every elements in $\mathbb{E}[X_\ell]$ is less than $\frac{2}{n_0(n_0-1)}$, we have:

$$\begin{aligned} \|X_\ell - \mathbb{E}[X_\ell]\| &\stackrel{(a)}{\leq} \sqrt{\frac{4n_0^2}{n_0^2(n_0 - 1)^2} + 2} \\ &\leq \frac{2}{n_0 - 1} + \sqrt{2}, \end{aligned} \quad (37)$$

where (a) is from the fact that Frobenius norm of the matrix is greater than its spectral norm.

From (37), we deduce that we can choose $D = \frac{3}{2}$.

Moreover, the variance of the matrix A is

$$\begin{aligned} \sigma^2(A) &= \|\mathbb{E}[(A - \mathbb{E}[A])^2]\| \\ &\stackrel{(a)}{=} \left\|\sum_{\ell=1}^s \mathbb{E}[(X_\ell - \mathbb{E}[X_\ell])^2]\right\| \\ &\leq \sum_{\ell=1}^s \|\mathbb{E}[(X_\ell - \mathbb{E}[X_\ell])^2]\| \\ &\leq \sum_{\ell=1}^s (\|\mathbb{E}[X_\ell^2]\| + \|\mathbb{E}[X_\ell]^2\|), \end{aligned} \quad (38)$$

where (a) is obtained from the independence of the X_ℓ 's.

To get an upper bound of (38), observe that the expectation of the (i, j) -th element of matrix X_ℓ^2 is

$$\mathbb{E}[(X_\ell^2)_{ij}] = \mathbb{E}\left[\sum_{k=1}^{n_0} x_{ik}x_{kj}\right] = \sum_{k=1}^{n_0} \mathbb{E}[x_{ik}x_{kj}] = 0.$$

In addition, the expectation of the (i, i) -th elements of matrix X_ℓ^2 is

$$\mathbb{E}[(X_\ell^2)_{ii}] = \sum_{k=1}^{n_0} \mathbb{E}[x_{ik}^2] = \sum_{k=1}^{n_0} \mathbb{E}[x_{ik}] \leq \frac{2}{n_0}.$$

On the other hand, the elements of $\mathbb{E}[X_\ell]^2$ are $\mathcal{O}\left(\frac{1}{n_0^2}\right)$, which implies $\|\mathbb{E}[X_\ell]^2\| = \mathcal{O}\left(\frac{1}{n_0^2}\right)$. Hence, using (38), we deduce that $\sigma^2(A) \leq \frac{2s}{n_0} + \mathcal{O}\left(\frac{s}{n_0^2}\right) \leq \frac{2m}{n_0}$.

Now, an upper bound of (36) is

$$n_0 \exp\left(-\frac{t^2}{\frac{4m}{n_0} + t}\right). \quad (39)$$

To conclude the proof of this lemma, we need to consider two cases: $n_0 = \frac{m}{(\log T)^2}$ and $n_0 = n$.

(i) When $n_0 = \frac{m}{(\log T)^2}$, $t = 5\sqrt{\log m} \log T$. So, (39) becomes:

$$\begin{aligned} \frac{m}{(\log T)^2} \exp\left(-\frac{25(\log T)^2 \log m}{4(\log T)^2 + 5\sqrt{\log m} \log T}\right) &\stackrel{(a)}{\leq} \exp\left(-3 \log m + \log\left(\frac{m}{(\log T)^2}\right)\right) \\ &\leq \frac{1}{m^2} \end{aligned}$$

where (a) is obtained from the assumption $T > 10m$.

(ii) If $n_0 = n$, $t = 5\sqrt{\frac{m \log m}{n}}$. Then, (39) becomes:

$$\begin{aligned} n \exp\left(-\frac{25m \log m}{\frac{4m}{n} + 5\sqrt{\frac{m \log m}{n}}}\right) &\leq \exp(-3 \log m + \log n) \\ &\stackrel{(a)}{\leq} \frac{1}{m^2} \end{aligned}$$

where (a) holds since $n \leq \frac{m}{(\log T)^2}$. □

Proof of Lemma 19.

Recall that

$$p(a, b) = (|V_b| - \mathbf{1}\{a = b\}) \frac{2}{n_0(n_0 - 1)} \sum_{\ell=1}^L \beta_\ell p_{a\ell} p_{b\ell}$$

and $p(a, 0) = 1 - \sum_{k=1}^K p(a, k)$. The estimations are $\hat{p}(i, j) = \frac{\sum_{v \in S_i} \sum_{v' \in S_j} A_{v, v'}}{s |S_i^{(0)}|}$ for all $1 \leq i, j \leq K$ and $\hat{p}(i, 0) = 1 - \sum_{k=1}^K \hat{p}(i, k)$.

An upper bound of $|p(i, j) - \hat{p}(i, j)|$ for $1 \leq i, j \leq K$ is

$$\begin{aligned} |\hat{p}(i, j) - p(i, j)| &\leq \frac{1}{s |S_i^{(0)}|} \left| e(S_i^{(0)}, S_j^{(0)}) - \mathbb{E}[e(S_i^{(0)}, S_j^{(0)})] \right| \\ &\quad + \frac{1}{s |S_i^{(0)}|} \left| \mathbb{E}[e(S_i^{(0)}, S_j^{(0)})] - s |S_i^{(0)}| p(i, j) \right|. \quad (40) \end{aligned}$$

Let \mathcal{A} be the set of partitions $\{S_k\}_{1 \leq k \leq K}$ of the set \mathcal{S} . Recall that \mathcal{S} is of cardinality n_0 . Then,

$$|\mathcal{A}| \leq K^{n_0}. \quad (41)$$

Now, we fix one partition $\{S_k\} \in \mathcal{A}$. Then, by Chernoff-Hoeffding bound,

$$\mathbb{P} \left(|e(S_i, S_j) - \mathbb{E}[e(S_i, S_j)]| < \sqrt{mn_0 \log m} \text{ for all } i, j \right) \geq 1 - \exp(-\Theta(n_0 \log m)). \quad (42)$$

Combining (41) and (42), we deduce that the following event holds:

$$|e(S_i, S_j) - \mathbb{E}[e(S_i, S_j)]| < \sqrt{mn_0 \log m}$$

for all i, j and $\{S_k\} \in \mathcal{A}$ with probability $1 - \exp(-\Theta(n_0 \log m))$ (just applying a union bound).

Since $\{S_k^{(0)}\} \in \mathcal{A}$, with probability $1 - \exp(-\Theta(n_0 \log m))$,

$$|e(S_i^{(0)}, S_j^{(0)}) - \mathbb{E}[e(S_i^{(0)}, S_j^{(0)})]| < \sqrt{mn_0 \log m} \quad (43)$$

for all i, j .

On the other hand, since $|\mathcal{E}^{(0)}| = \mathcal{O}\left(\frac{n_0^2 \log m}{m}\right)$ from the assumption, with probability $1 - \exp(-\Theta(n_0 \log m))$,

$$\frac{1}{s|S_i^{(0)}|} \left| \mathbb{E}[e(S_i^{(0)}, S_j^{(0)})] - s|S_i^{(0)}|p(i, j) \right| = \mathcal{O}\left(\frac{n_0 \log m}{m} p(i, j)\right), \quad (44)$$

for all i, j .

Then, conditioned on (44) for all $1 \leq i, j \leq K$, we can derive an upper bound of $|p(i, j) - \hat{p}(i, j)|$ for $1 \leq i, j \leq K$, using (40), (43) and (44):

$$|p(i, j) - \hat{p}(i, j)| = \mathcal{O}\left(\sqrt{\frac{n_0 \log m}{m}} p(i, j)\right),$$

which implies that for all $1 \leq i, j \leq K$

$$\begin{aligned} \left| \log \frac{\hat{p}(i, j)}{p(i, j)} \right| &\leq \frac{|p(i, j) - \hat{p}(i, j)|}{p(i, j)} \\ &= \mathcal{O}\left(\sqrt{\frac{n_0 \log m}{m}}\right). \end{aligned} \quad (45)$$

Furthermore, an upper bound of $\left| \log \frac{\hat{p}(i, 0)}{p(i, 0)} \right|$ is

$$\left| \log \frac{\hat{p}(i, 0)}{p(i, 0)} \right| = \mathcal{O}\left(\frac{1}{n_0} \sqrt{\frac{n_0 \log m}{m}}\right), \quad (46)$$

with probability at least $1 - \exp(-\Theta(n_0 \log m))$.

We also have $\mathbb{E}[e(v, \mathcal{S})] \leq \frac{m}{n_0}$ and from Chernoff inequality,

$$\mathbb{P}\left(|e(v, \mathcal{S}) - \mathbb{E}[e(v, \mathcal{S})]| > \sqrt{4m \log m}\right) \leq \frac{1}{m^2}. \quad (47)$$

Finally, we obtain the following from (45), (46) and (47):

$$\sum_{i=0}^K e(v, S_i^{(t)}) \left| \log \frac{p(k, i)}{\hat{p}(k, i)} \right| = \mathcal{O}\left(\sqrt{\frac{m \log m}{n_0}}\right),$$

with probability at least $1 - \frac{2}{m^2}$. This concludes the proof. \square

Proof of Lemma 20. We have:

$$\sum_{v \in \mathcal{E}^{(t+1)}} \left(e(v, \mathcal{E}^{(t)}) - \mathbb{E}[e(v, \mathcal{E}^{(t)})] \right) \leq \mathbf{1}_{\mathcal{E}^{(t)}}^T (A - \mathbb{E}[A]) \mathbf{1}_{\mathcal{E}^{(t+1)}},$$

where $\mathbf{1}_S$ is the vector whose i -th component is equal to 1 if $i \in S$ and to 0 otherwise. Since $\mathbb{E}[e(v, \mathcal{E}^{(t)})] \leq \frac{2m}{n_0} |\mathcal{E}^{(t)}| \leq 2 \log m$,

$$\begin{aligned} \sum_{v \in \mathcal{E}^{(t+1)}} e(v, \mathcal{E}^{(t)}) &\leq \sum_{v \in \mathcal{E}^{(t+1)}} \left(e(v, \mathcal{E}^{(t)}) - \mathbb{E}[e(v, \mathcal{E}^{(t)})] \right) + 2|\mathcal{E}^{(t+1)}| \log m \\ &\leq \|\mathbf{1}_{\mathcal{E}^{(t)}}^T (A - \mathbb{E}[A]) \mathbf{1}_{\mathcal{E}^{(t+1)}}\| + 2|\mathcal{E}^{(t+1)}| \log m \\ &\leq \|\mathbf{1}_{\mathcal{E}^{(t)}}^T\| \|A - \mathbb{E}[A]\| \|\mathbf{1}_{\mathcal{E}^{(t+1)}}\| + 2|\mathcal{E}^{(t+1)}| \log m \\ &\stackrel{(a)}{=} \mathcal{O} \left(\sqrt{|\mathcal{E}^{(t)}| |\mathcal{E}^{(t+1)}| \frac{m \log m}{n_0}} \right), \end{aligned}$$

where for (a), we used the assumption that $\|A - \mathbb{E}[A]\| = \mathcal{O} \left(\sqrt{\frac{m \log m}{n_0}} \right)$ and the definition of ℓ_2 norm. \square

Proof of Lemma 21.

From Chernoff-Hoeffding bound, there exists $C > 0$ such that

$$\mathbb{P} \left(|e(v, V_a) - \mathbb{E}[e(v, V_a)]| > C \sqrt{\frac{m \log m}{n_0}} \right) \leq \frac{1}{m^3}. \quad (48)$$

We also have for all $v \in V_k$ and all k ,

$$\sum_{a=0}^K \mathbb{E}[e(v, V_a)] \log \left(\frac{p(k, a)}{p(k', a)} \right) = \Omega \left(\frac{m}{n_0} \right) \quad \text{for all } k' \neq k. \quad (49)$$

Since $\log \left(\frac{p(k, a)}{p(k', a)} \right) = \Theta(1)$ and $\frac{m}{n_0} = \Omega((\log T)^2)$, from (48) and (49), we have Lemma 21. \square