
Learning Fair Division from Bandit Feedback

Hakuei Yamada

University of Electro-Communications

Junpei Komiyama

New York University

Kenshi Abe

CyberAgent

Atsushi Iwasaki

University of Electro-Communications

Abstract

This work addresses learning online fair division under uncertainty, where a central planner sequentially allocates items without precise knowledge of agents' values or utilities. Departing from conventional online algorithms, the planner here relies on noisy, estimated values obtained after allocating items. We introduce wrapper algorithms utilizing *dual averaging*, enabling gradual learning of both the type distribution of arriving items and agents' values through bandit feedback. This approach enables the algorithms to asymptotically achieve optimal Nash social welfare in linear Fisher markets with agents having additive utilities. We also empirically verify the performance of the proposed algorithms across synthetic and empirical datasets.

1 INTRODUCTION

Ensuring an equitable distribution of limited resources among individuals, or agents, is a crucial task, as it involves reconciling competing interests among the agents. In today's society, there is a growing demand for efficient and fair resource allocation. Given this context, the fair division problem has attracted substantial interest from the researchers. This problem involves the process of assigning items to individuals with idiosyncratic preferences (Brams and Taylor, 1996; Moulin, 2003). Many studies in fair division are treated as a static problem, assuming that all information about the items is known beforehand. However, in reality, it is uncommon for all items to arrive beforehand. *Online fair division* address uncertainty about what type of items can arrive dynamically, and the term "online" in this

context means that items must be irrevocably assigned to an agent (Aleksandrov and Walsh, 2020).

This paper introduces another layer of uncertainty regarding the values or utilities that agents place on the items they receive. Unlike traditional online algorithms, a central planner does not have exact knowledge of the agents' utilities. Instead, the planner relies on noisy, estimated utilities obtained after the items have been allocated, as often assumed in *multi-armed bandit* problems.

We propose wrapper algorithms employing *dual averaging* (DA) (Xiao, 2009; Gao et al., 2021), enabling the gradual learning of both the distribution of arriving item types and the values attributed to these items by the agents through bandit feedback. Gao et al. (2021) applied DA to the Eisenberg-Gale (EG) convex program (Eisenberg and Gale, 1959; Jain and Vazirani, 2010). If items are divisible, the solution aligns with *competitive equilibrium with equal income* (CEEI) (Budish, 2011). CEEI ensures both envy-freeness and Pareto optimality and coincides with an assignment that maximizes *Nash social welfare* (NSW) (Nash, 1950). Gao et al. (2021) derived a fair assignment, specifically an asymptotically envy-free and Pareto optimal online assignment, assuming types of arriving items are drawn independently and identically distributed (i.i.d.)¹ and the values for each agent must be precisely revealed before the assignment in a round. However, as demonstrated in the following examples, the precise observation of values is not always possible, necessitating the handling of noisy feedback regarding the values for the agent who has received the item.

Example 1. (Crowdsourcing) *Consider the problem of allocating tasks (items) to workers (agents), where each worker has their own area of expertise and the quality of the task may vary. In such a scenario, the importance of learning becomes evident, as the expertise of each worker is typically unknown before the task is assigned. If the aim is to maximize the quality of the task, there is a risk that only a small*

Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. Copyright 2024 by the author(s).

¹The i.i.d. assumption has recently been relaxed by Liao et al. (2022) to some extent.

group of workers will monopolize all the tasks. However, by maximizing the NSW, task assignments can be distributed more evenly among the workers.

Example 2. (Food recommendation on online food delivery) Consider recommending a restaurant (agent) to a user (item). To do so, we need to estimate the user’s taste preferences based on their past orders, such as whether they have a preference for specific types of cuisines. If the goal is solely to maximize social welfare, there is a risk that a few popular restaurants may receive a large portion of the orders, leading to delays or unfulfilled orders. Instead, distributing the workload equitably among restaurants enhances the long-term efficiency (Wang et al., 2022).

Example 3. (Humanitarian aid) Consider the problem of providing a limited number of items to disaster-stricken areas, where the supply arrives in real-time, and the prompt distribution of the items is crucial. The value of the items for each district can only be determined through actual feedback, which implies that the situation is an online problem. In this context, the value of the items needs to be learned and updated as the supplies are being distributed.

To address these challenges, we introduce two novel algorithms: DA-Explore-then-Commit (DA-EtC) and DA-Upper Confidence Bound (DA-UCB), which combine dual averaging with multi-armed bandit strategies. These algorithms gradually learn and explore the values of items through bandit feedback and the resulting allocation approaches the solution of the EG program as the number of items grows. Our algorithms handle situations where agents do not precisely observe the values of allocated items by incorporating estimated values into the DA process. First, analyzing DA-EtC are intricate due to the need to create *virtual* values fed into DA, ensuring compatibility with the i.i.d. property assumed by Xiao (2009). We demonstrate that DA-EtC achieves a regret of $\tilde{O}(T^{2/3})^2$ by meticulously dissecting the regret into the estimation error of values and the online learning error of DA, relative to the number of items T . This indicates the rate at which it converges to the true CEEI solution as the number of items grows.

Second, we introduce DA-UCB, which employs the upper confidence bound and is designed to have better empirical performance than DA-EtC. Analyzing DA-UCB is even more complicated than DA-EtC due to the non-i.i.d. nature of the virtual values, referred to here as UCB values. To make the analysis tractable, we devise a variant of DA-UCB where multiple instances of DA are executed, treating each estimator matrix as constant. This preserves the structure of DA-UCB, enabling the virtual values to be drawn in an i.i.d. manner. Consequently, we establish a regret upper bound of $\tilde{O}(\sqrt{T})$, meeting the lower bound of $\tilde{\Omega}(\sqrt{T})$ applicable to any algorithm.

²We use \tilde{O} to denote a Landau notation that disregards a polylogarithmic factor.

Let us explore the rest of the literature. Finding fair division closely relates to computing market equilibria or competitive equilibria, which has been extensively studied in algorithmic game theory, mainly for the prominent case of Fisher markets (Vazirani, 2007; Codenotti and Varadarajan, 2007). This problem is highly challenging in general. However, in a static setting, if agents have linear, additive preferences, the problem can be reduced to an optimization known as the EG convex program (Eisenberg and Gale, 1959; Jain and Vazirani, 2010). As we already noted, if the items are divisible, the solution coincides with CEEI and an assignment that maximizes NSW. In contrast, if the items are indivisible, the compatibility no longer holds. However, the solutions of the NSW maximization is considered to have several plausible properties (Caragiannis et al., 2019).³

In online settings, even if items are divisible, envy-freeness is not compatible with Pareto optimality. Some researchers have taken a relaxed approach to envy-freeness, while others have found an approximate solution through stochastic approximation schemes: Kash et al. (2014) relax envy-freeness for their online setting. Bateni et al. (2022) utilize a stochastic approximation scheme and obtain an approximate solution for the EG convex program. Sinclair et al. (2022) examine the trade-off between envy-freeness and Pareto optimality. The trade-off can be resolved if we allow the number of items to be large, or if we approximate indivisible items as divisible ones. Our work aims to achieve envy-freeness and Pareto optimality in the sense of asymptotically maximizing NSW as the number of items grows.

2 PROBLEM SETUP

This section introduces the notation for our paper. Summarized notation is found in Appendix A. Let $N = \{1, 2, \dots, n\}$ denote the set of agents, and $M = \{1, 2, \dots, m\}$ denote the set of the types of items ($|N| = n, |M| = m$). The *ex ante* (expected) value of each agent i on an item of type j is denoted as $v_{i,j} \in \mathbb{R}$, which is unknown in advance to a policymaker or an algorithm. Also, $v_{i,j}$ is non-negative. At each round $t = \{1, 2, \dots, T\}$, an (indivisible) item $j(t) \in M$ arrives and the policymaker allocates it to agent $i(t) \in N$. The type $j(t) \in M$ of item at t is drawn from a certain probability distribution $\mathcal{S} = \{s_1, s_2, \dots, s_m\}$ in an i.i.d. manner, where $\sum_{j \in M} s_j = 1$ and $s_j \geq 0$ for all $j \in M$. When the agent $i(t)$ receives the item $j(t)$, the agent observes the *ex post* (realized) utility of $u_i(t) = v_{i(t),j(t)} + \varepsilon_t$, instead of directly observing $v_{i(t),j(t)}$. The quantity ε_t is a sub-Gaussian random variable with its radius σ^2 . Namely, it satisfies $\mathbb{E}[e^{\lambda \varepsilon_t}] \leq \exp(\lambda^2 \sigma^2 / 2)$,

³First, an NSW solution is envy-free up to one item, meaning that each agent (almost) prefers her allocation better than the others. Second, it is Pareto optimal, meaning that no one can be better off without sacrificing the utility of some other agents. Third, it also has pairwise maximin share guarantee, meaning that the division aligns with the intention of each agent.

which implies ε_t to be mean zero.⁴ The cumulative utility of each agent i in T rounds, assuming the agents' values are additive, is defined as: $U_i(T) = \sum_{t:i(t)=i} u_i(t)$. The additive assumption is common in fair division (Procaccia and Wang, 2014; Caragiannis et al., 2019; Bouveret and Lemaître, 2016).

This paper aims to asymptotically find a fair allocation in this online setting. In the literature, there are several ways to define fairness: *Envy-freeness* guarantees that no agent has no envy toward another agent's allocation; *Maximin share* guarantees that the minimum value among allocated agents is maximized. In general, both cannot be guaranteed in our online setting (See also the extensive survey (Aleksandrov and Walsh, 2020)). We concentrate on asymptotically maximizing *NSW*, defined as the geometric mean of the agents' obtained values, which is known as empirically balancing allocations between envy-freeness and *Pareto efficiency*.

For our noisy, online setting, we define the *ex post* notion of NSW across time, which is equivalent to the weighted geometric mean of realized utilities: $NSW(T) = \prod_{i \in N} U_i(T)^{B_i}$. The weights $B_i > 0$ indicate the priority given to an agent, and they can be interpreted as her per-round budget rate in a Fisher market, e.g., (Gao et al., 2021). Without loss of generality, we assume $\sum_i B_i = 1$. We next consider a hindsight optimal allocation as a benchmark. The law of large numbers implies that, when we allocate an item of type j to agent i many times the mean utility converges to $v_{i,j}$, and when the number of items is sufficiently large, we can approximate the items to be divisible. By using these facts, the hindsight optimal allocation is represented as the *EG* convex program.

Let $x_{i,j} \in [0, 1]$ be the fraction of items j allocated to agent i , consider the optimization per item:

$$\begin{aligned} & \text{maximize}_{\{x_{i,j}\}} && \prod_{i \in N} \left(\sum_{j \in M} s_j v_{i,j} x_{i,j} \right)^{B_i} \\ & \text{subject to} && \forall j \in M : \sum_{i \in N} x_{i,j} \leq 1, \\ & && \forall i \in N, \forall j \in M : x_{i,j} \geq 0. \end{aligned} \quad (1)$$

We say an allocation *optimal* if it maximizes the objective for the ex ante values and say the hindsight value of the optimal allocation *Optimal NSW* (ONSW).

We adopt a metric *Regret* to measure the performance of an online allocation. The regret is the difference between the total hindsight utilities ONSW obtained from Eq. (1) and the ex post NSW with time horizon T :

$$\text{Regret}(T) = T \cdot \text{ONSW} - \text{NSW}(T). \quad (2)$$

⁴The class of sub-Gaussian distributions includes many light-tail distributions. For example, it allows $u_i(t)$ to be a discrete value such as $\{0, 1\}$ (Bernoulli distributions), or $\{1, 2, 3, 4, 5\}$ (categorical distributions).

It is not very difficult to see that the maximization of social welfare (SW) $\sum_i U_i(T)$ does not necessarily maximize the NSW. For example, if there is an agent i with a very small value $v_{i,j}$ for all types j , then the social welfare is maximized by allocating no item to the agent, which results in zero NSW. Note also that, unlike SW, NSW is free from normalization; multiplying a constant on an i -th row of a matrix $\{v_{i,j}\}$ does not affect the optimal allocation. Without loss of generality, the objective function of Eq. (1) can be replaced with the sum of weighted logarithmic utilities, i.e., $\sum_{i \in N} B_i \log \sum_{j \in M} v_{i,j} x_{i,j}$. We call the primal form P_{EG} and

its dual form D_{EG} given as

$$\begin{aligned} & \text{minimize}_{\{\beta_i, \{p_j\}\}} && \sum_{j \in M} s_j p_j - \sum_{i \in N} B_i \log \beta_i \\ & \text{subject to} && \forall i \in N, \forall j \in M : p_j \geq \beta_i v_{i,j}, \\ & && \forall j \in M : p_j \geq 0. \end{aligned}$$

The value p_j implies the *price* of item j . This program has no duality gap and it belongs to a *rational convex program* (Vazirani, 2012) where all parameters are rational numbers and which always admits a rational solution. The program, while being nonlinear, can be solved by algorithms such as the ellipsoid method, e.g., (Vishnoi, 2021).

However, solving EG once is insufficient for our aim. This is because (i) we do not know the values $\{v_{i,j}\}_{i,j}$ and need to update their estimates or empirical means for each round. (ii) Allocating items in a greedy manner based on the current estimates of $\{v_{i,j}\}_{i,j}$ results in a suboptimal NSW. In the following, we discuss ideas that address the issues above.

3 DUAL-AVERAGING FOR OUR SETTING

DA is an iterative method for solving a convex optimization problem⁵, Gao et al. (2021) have utilized DA for solving the dual problem D_{EG} . A variant of DA for our setting, described in Algorithm 1, calls a subroutine (Algorithm 2) to determine the winner at each round. Algorithm 2 uses a multiplier $\beta_i \in \mathbb{R}^+$ to balance the allocation of items among agents and prevent any one agent from monopolizing the allocation. As agent i receives more items, β_i decreases. Each iteration in DA corresponds to a round, which is an analog of the arrival of an item, in the online fair division, where an item with value $\hat{v}_{i,j(t)}^{\text{DA}}(t)$ is allocated according to a (virtual) first-price auction, where the bids of agents are weighted by the multiplier $\beta_i v_i^{\text{DA}}$. Here, the parameters $l, h, \delta_0 > 0$ restrict the range of β_i for the sake of stability. At round t , the winner $i(t)$ of the auction receives the item and obtains the utility $u_i^{\text{DA}}(t)$. If the true values $\{v_{i,j(t)}\}$ are available, they are used in $\{\hat{v}_{i,j(t)}^{\text{DA}}(t)\}$ of the subroutine.

⁵The version of DA we consider involves a regularization term and is sometimes referred to as regularized dual averaging.

Algorithm 1 DA with true values $\{v_{i,j}\}$

- 1: Initialize utility $\bar{u}_i^{\text{DA}} = 0$ for each i .
 - 2: **for** $t = 1$ to T **do**
 - 3: Observe the type $j(t)$ of an arriving item, drawn from \mathcal{S} .
 - 4: Observe the value $v_{i,j(t)}$ of the item for each i .
 - 5: Update the mean utilities $\{\bar{u}_i^{\text{DA}}(t + 1)\}_i = \text{DA-Iter}(t, \{v_{i,j(t)}\}_i, \{\bar{u}_i^{\text{DA}}(t)\}_i)$
 - 6: **end for**
-

Algorithm 2 DA-Iter

Require: $t, \{\hat{v}_{i,j(t)}^{\text{DA}}(t)\}_i, \{\bar{u}_i^{\text{DA}}(t)\}_i$

- 1: Define multiplier β_i as

$$\beta_i = \text{Proj}_{[B_i/(h(1+\delta_0)), (1+\delta_0)/l]}(B_i/\bar{u}_i^{\text{DA}}(t)) \quad (\delta_0 > 0),$$

where $\text{Proj}_{[a,b]}(x) = \max(a, \min(b, x))$.

- 2: Agents bid $\beta_i \hat{v}_{i,j(t)}^{\text{DA}}(t)$.
- 3: Winner is determined: $i(t) = \arg \max_i \beta_i \hat{v}_{i,j(t)}^{\text{DA}}(t)$, where ties are broken arbitrarily.
- 4: Winner $i(t)$ pays $\beta_{i(t)} \hat{v}_{i,j(t)}^{\text{DA}}(t)$.
- 5: Each agent receives a utility: $u_i^{\text{DA}}(t) = \mathbb{I}\{i = i(t)\} \hat{v}_{i,j(t)}^{\text{DA}}(t)$.
- 6: Update mean utility for each agent as

$$\bar{u}_i^{\text{DA}}(t + 1) = \frac{t-1}{t} \bar{u}_i^{\text{DA}}(t) + \frac{1}{t} u_i^{\text{DA}}(t)$$

- 7: **return** $\{\bar{u}_i^{\text{DA}}\}_i$ (and $i(t)$ for DA-EtC, DA-UCB)
-

This is the case of the PACE algorithm (Gao et al., 2021), which has shown to have $\tilde{O}(1/\sqrt{T})$ regret.

However, Algorithm 1 requires the true value $v_{i,j(t)}$ for agent i receiving type j item at round t , which is not available in our setting. Instead, we consider plugging-in the estimated value to the subroutine (Algorithm 2) as described in the following section.

4 PROPOSED METHOD

4.1 DA-EtC

Algorithm 3 describes the procedure of DA-EtC. It begins with a uniform exploration phase, where each agent tries every item type equally for the first T_0 rounds. At the end of this phase, the algorithm generates an estimator $\hat{v}_{i,j}^{\text{EtC}} = \hat{v}_{i,j}(T_0 + 1)$ of the expected value $v_{i,j}$ of each item type j for each agent i as in line 7. Here, $N_{i,j}(t)$ is the number of times agent i received an item of type j up

Algorithm 3 DA-EtC

- 1: **for** $t = 1$ to T_0 **do**
- 2: Observe the type $j(t)$ of an arriving item, drawn from \mathcal{S} .
- 3: Give item $j(t)$ to the agent who is chosen uniformly at random.
- 4: The agent $i(t)$ receives a utility $u_{i(t)}(t)$.
- 5: Update the cumulative utility: $U_i(t + 1) = U_i(t) + \mathbb{I}[i = i(t)]u_{i(t)}(t)$ for each i .
- 6: **end for**
- 7: Fix the estimator:

$$\hat{v}_{i,j}^{\text{EtC}} = \frac{\sum_{t:j(t)=j} u_{i(t)}(t) \mathbb{I}[i = i(t)]}{N_{i,j}(T_0 + 1)}$$

for each i, j .

- 8: Initialize the DA's utility $\bar{u}_i^{\text{DA}}(1) = 0$ for each i .
 - 9: **for** $t = T_0 + 1$ to T **do**
 - 10: The type of item t is determined: $j(t) \sim \mathcal{S}$.
 - 11: $t' = t - T_0$
 - 12: Run DA-Iter to allocate an item: $\{\bar{u}_i^{\text{DA}}(t' + 1)\}_i, i(t) = \text{DA-Iter}(t', \{\hat{v}_{i,j(t)}^{\text{EtC}}\}_i, \{\bar{u}_i^{\text{DA}}(t')\}_i)$.
 - 13: The agent $i(t)$ receives a utility $u_{i(t)}(t)$.
 - 14: Update the cumulative utility: $U_i(t + 1) = U_i(t) + u_{i(t)}(t) \mathbb{I}[i = i(t)]$ for each i .
 - 15: **end for**
-

to round $t - 1$: $\sum_{\tau < t} \mathbb{I}[i(\tau) = i \wedge j(\tau) = j]$ where $\mathbb{I}[\cdot]$ is an indicator function. After the exploration phase, the algorithm runs a DA algorithm for the remaining $T - T_0$ rounds using the estimator $\{\hat{v}_{i,j}^{\text{EtC}}\}$. In each of the rounds, the algorithm selects the winner $i(t)$ and updates the mean utilities $\bar{u}_i^{\text{DA}}(t + 1)$.

To implement DA-EtC, a certain level of attention is required. Specifically, DA-EtC uses $\hat{v}_{i,j}^{\text{EtC}}$ for $\hat{v}_{i,j}^{\text{DA}}(t)$ in DA-Iter and calculate $\bar{u}_i^{\text{DA}}(t)$. Note that the estimator $\hat{v}_{i,j}^{\text{EtC}}$ is fixed after round T_0 , primarily due to theoretical reasons to derive Lemma 1.

4.2 DA-UCB

Despite being able to control the exploration duration T_0 for minimizing regret, DA-EtC has limited adaptivity. The largest limitation stems from the uniform exploration, which may result in the unnecessary exploration of items with a significantly low value. To mitigate this, we develop DA-UCB in Algorithm 4, which employs an upper confidence bound that holds with high probability and solves the exploration and exploitation tradeoff more adaptively than DA-EtC. At each round t , DA-UCB calculates the UCB value $\hat{v}_{i,j(t)}^{\text{UCB}}(t)$ for each agent i . This value is then provided to DA-Iter to determine the winner $i(t)$. Note that the utility with respect

Algorithm 4 DA-UCB

- 1: Initialize the DA's utility $\bar{u}_i^{\text{DA}}(1) = 0$ for each i and $\hat{v}_{i,j}^{\text{UCB}}(1) = 1$ for all i, j .
- 2: **for** $t = 1$ to T **do**
- 3: Observe the type $j(t)$ of an arriving item, drawn from \mathcal{S} .
- 4: Calculate the UCB value

$$\hat{v}_{i,j(t)}^{\text{UCB}}(t) = \min\left(1, \hat{v}_{i,j(t)}(t) + \sqrt{\log t / (2N_{i,j(t)}(t))}\right)$$

where $\min(1, +\infty) = 1$.

- 5: Run DA-Iter to allocate an item: $\{\bar{u}_i^{\text{DA}}(t + 1)\}_i, i(t) = \text{DA-Iter}(t, \{\hat{v}_{i,j(t)}^{\text{UCB}}(t)\}_i, \{\bar{u}_i^{\text{DA}}(t)\}_i)$.
- 6: The agent $i(t)$ receives a utility $u_{i(t)}(t)$.
- 7: Update the number of draws $N_{i(t),j(t)}(t + 1) = N_{i(t),j(t)}(t) + 1$.
- 8: Update the cumulative utility: $U_i(t + 1) = U_i(t) + u_{i(t)}(t)\mathbb{I}[i = i(t)]$ for each i .
- 9: Update the estimator:

$$\hat{v}_{i(t),j(t)}(t + 1) = \frac{\sum_{\tau \leq t: j(\tau) = j} u_{i(t)}(\tau)\mathbb{I}[i = i(t)]}{N_{i,j}(t + 1)}$$

10: **end for**

to DA, denoted as $u_i^{\text{DA}}(t)$, differs from the actual utility $u_i(t)$ for both DA-EtC and DA-UCB. This is because the estimated value is supplied to DA.

5 ANALYSIS

In this section, we analyze the regret bounds of the algorithms. We begin by establishing a lower bound on the regret, which represents the best performance that any algorithm can achieve. Following that, we provide regret upper bounds on DA-EtC and a variant of DA-UCB. This is because directly analyzing DA-UCB poses significant challenges.

5.1 Regret lower bound

Theorem 1. (Regret lower bound) *There exists a model where the expected regret of any algorithm is lower-bounded as $\mathbb{E}[\text{Regret}(T)] = \Omega(\sqrt{mT})$.*

Proof Sketch of Theorem 1. We introduce the base model where $v_{i,j} = 1/2$ for all i, j . To maximize NSW, the algorithm must allocate (approximately) T/n of items for each agent. Let j_i indicate the type of item that agent i receives the least number of times during T . There exists a set of

n items $\{j_i\}_{i \in N}$ such that, $\sum_i N_{i,j_i} \leq T/m$. We consider an alternative model such that v_{i,j_i} is larger than $1/2$ by $\sqrt{m/T}$. To have a low regret in the alternative model, the algorithm must choose arms $\{j_i\}_{i \in N}$ frequently. However, information-theoretic results imply that the algorithm cannot differentiate the base model and the alternative model, and thus suffers a large regret in the alternative model (Kaufmann et al., 2016). \square

5.2 Convergence on Dual Averaging

We state the convergence results of DA. The following lemma is an extension of Theorem 4 in (Gao et al., 2021). To be precise, Theorem 4 in (Gao et al., 2021) requires the values to be normalized as $\sum_j s_j v_{i,j} = 1$ for each agent i , and thus it is not directly applicable to the setting where $v_{i,j}$ is unknown. Lemma 1 here generalized their results by introducing l, h so that normalization is no longer required.

Lemma 1. *Assume that: (a) we run DA for T rounds; (b) let $\{v_{i,j}^{\text{DA}}\}$ be $\hat{v}_{i,j(t)}^{\text{DA}}(t)$ for all t , and they are constant; Let $\bar{u}^{\text{DA}}(T) \in \mathbb{R}^n$ be the mean utility vector of agents at the end of round T and $u^{*,\text{DA}} \in \mathbb{R}^n$ be the solution of the corresponding EG program with $\{v_{i,j}^{\text{DA}}\}$. Then, the following inequality holds for an arbitrary $\delta_0 > 0$:*

$$\mathbb{E}[\|\bar{u}^{\text{DA}}(T) - u^{*,\text{DA}}\|^2] \leq C^{\text{DA}} \frac{6 + \log T}{T}$$

where $\|\cdot\|$ is a l2 norm vector and

$$C^{\text{DA}} = \frac{\|v^{\text{DA}}\|_\infty^2 (1 + \delta_0)^6}{l^4 (\min_{i \in N} B_i)^4} \left(\frac{h^3 \|v^{\text{DA}}\|_\infty^2}{l} \left(\frac{1}{\delta_0}\right)^2 + h^4 \right)$$

and $\|v^{\text{DA}}\|_\infty = \max_{i \in N} \|v_i^{\text{DA}}\|_\infty$.

5.3 Regret upper bound of DA-EtC

Theorem 2. (Regret upper bound of DA-EtC) *Assume $B_i = 1/n$ and $s_j = 1/m$ for all i, j . Assume that $2l \leq v_{i,j} \leq h/2$ for all i, j . Then, for T_0 that is sufficiently large such that $\sqrt{\frac{8 \max(1, \sigma^2) nm \log(nmT)}{T_0}} \leq \frac{\min(1, 2l)}{2}$ holds, the expected regret of DA-EtC is bounded as follows:*

$$\mathbb{E}[\text{Regret}(T)] = \tilde{O}\left(T_0 + \sqrt{\frac{nm}{T_0}} T + nC^{\text{DA}} \sqrt{T}\right).$$

The $\tilde{O}(nC^{\text{DA}} \sqrt{T})$ term is non-leading with respect to T . By focusing on the leading term in T and setting $T_0 = T^{2/3}(nm)^{1/3}$, we obtain $\mathbb{E}[\text{Regret}(T)] = \tilde{O}(T^{2/3}(nm)^{1/3})$.

Proof Sketch of Theorem 2. We decompose the regret into two components. Let $u^{*,\text{true}} = (u_1^{*,\text{true}}, u_2^{*,\text{true}}, u_3^{*,\text{true}}, \dots, u_n^{*,\text{true}})$ be the solution of EG with true values, i.e., Eq. (1). Let

$$\bar{u}_i^{\text{EtC}} = \frac{\sum_{t > T_0}^T u_i(t)\mathbb{I}[i(t) = i]}{T - T_0}$$

be the mean utility during the exploitation rounds. $u_i^{*,\text{DA}}$ is the solution of EG where value matrix $\{\hat{v}_{i,j}^{\text{EtC}}\}_{i,j}$. Note that \bar{u}_i^{DA} is the mean utility in view of DA-Iter, which corresponds to \bar{u}_i^{EtC} but $u_i(t)$ is replaced by $\hat{v}_{i,j}^{\text{EtC}}(t)$. Then, we demonstrate a novel decomposition for the regret-per-round that boils it down into the following terms:

$$\begin{aligned} \frac{\text{Regret}(T)}{T} &= \prod_i (u_i^{*,\text{true}})^{B_i} - \prod_i (\bar{u}_i^{\text{EtC}})^{B_i} + O(T_0) \\ &= \left(\prod_i (u_i^{*,\text{true}})^{B_i} - \prod_i (u_i^{*,\text{DA}})^{B_i} \right) \\ &\quad + \left(\prod_i (u_i^{*,\text{DA}})^{B_i} - \prod_i (\bar{u}_i^{\text{DA}})^{B_i} \right) \\ &\quad + \left(\prod_i (\bar{u}_i^{\text{DA}})^{B_i} - \prod_i (\bar{u}_i^{\text{EtC}})^{B_i} \right) + O(T_0). \quad (3) \end{aligned}$$

Intuitively, The first term of Eq. (3) is due to the estimation error of $\{v_{i,j}\}_{i,j}$, which depends on the quality of estimator built at the end of the exploration phase. The second term of Eq. (3) is the error due to the loss of DA, which we bound via Eq. (1) and by converting the error on the l2-norm into the error on the error of NSW. The third term is due to the estimation error $\{v_{i,j}\}_{i,j}$ as well as the random realization of the samples in the commitment phase. \square

5.4 Regret upper bound of a variant of DA-UCB

This section analyzes a slightly modified version of DA-UCB, which we call Repeated Dual Averaging UCB (RDA-UCB, Algorithm 5) to demonstrate that DA-UCB may be able to achieve the $\tilde{O}(\sqrt{T})$ regret. In general, the regret analysis depends on Lemma 1, which requires that inputs for DA are drawn in an i.i.d. manner. DA-EtC inevitably satisfies the requirement because it fixes a value matrix, whose elements $\hat{v}_{i,j}^{\text{EtC}}$ are the inputs for DA, at the end of the exploration period. In contrast, DA-UCB by no means satisfies the requirement, since it updates a UCB value $\hat{v}_{i,j}^{\text{UCB}}$ at each round. To overcome this, RDA-UCB executes multiple instances of DA and each of the estimator matrices becomes regarded as constant, so that it comes to satisfy the i.i.d. assumption. Note that, in practice, RDA-UCB is clearly outperformed by DA-UCB due to the complexity and thus we perform only DA-UCB in the simulation section.

Let the projected UCB value when as

$$\hat{v}_{i,j}^{\text{RUCB}}(t) = \text{Proj}_{[l,h]} \left(\hat{v}_{i,j}(t) + \sqrt{\frac{\log(T^2)}{2N_{i,j}(t)}} \right), \quad (4)$$

where $\hat{v}_{i,j}(t)$ is the empirical estimate of $v_{i,j}$ with $N_{i,j}(t)$ samples. Here, we define $\hat{v}_{i,j}^{\text{UCB}}(1) = h$ for all i, j . The RDA-UCB algorithm proceeds as follows. It begins by fixing a value matrix (Line 4) based on the current UCB value and then executes an instance of DA (Line 6). This instance continues until an entry of the count matrix $N_{i,j}(t)$ reaches a power of 2 (Line 14), at which point the current

Algorithm 5 RDA-UCB

- 1: $s = 1$.
 - 2: Initialize $\hat{v}_{i,j}^{\text{RUCB}}(1) = h$ for all i, j .
 - 3: **while** $s \leq T$ **do**
 - 4: Fix the RUCB value matrix $\hat{v}_{i,j}^{\text{RUCB}}(s)$. This RUCB value remains the same during the inner loop.
 - 5: Initialize the DA's utility $\bar{u}_i^{\text{DA}}(1) = 0$ for each i .
 - 6: **while true do**
 - 7: $t = s$.
 - 8: Observe the type $j(t)$ of an arriving item, drawn from \mathcal{S} .
 - 9: Run DA-Iter to allocate an item: $\{\bar{u}_i^{\text{DA}}(t+1)\}_i, i(t) = \text{DA-Iter}(t-s+1, \{\hat{v}_{i,j}^{\text{RUCB}}(t)\}_i, \{\bar{u}_i^{\text{DA}}(t)\}_i)$.
 - 10: The agent $i(t)$ receives a utility $u_{i(t)}(t)$.
 - 11: Update the number of draws $N_{i(t),j(t)}(t+1) = N_{i(t),j(t)}(t) + 1$.
 - 12: Update the cumulative utility: $U_i(t+1) = U_i(t) + u_{i(t)}(t)\mathbb{I}[i = i(t)]$ for each i .
 - 13: Update the estimator $\hat{v}_{i(t),j(t)}(t) = \frac{U_i(t+1)}{N_{i,j}(t+1)}$.
 - 14: **if** $N_{i(t),j(t)}(t) \in \{2, 2^2, 2^3, \dots, 2^{\lceil \log_2 T \rceil}\}$ **then**
 - 15: $s = t + 1$ and break the inner loop.
 - 16: **else if** $t \geq T$ **then**
 - 17: Terminate the algorithm.
 - 18: **end if**
 - 19: $t = t + 1$.
 - 20: **end while**
 - 21: $\hat{v}_{i,j}^{\text{RUCB}}(t) = \text{Proj}_{[l,h]} \left(\hat{v}_{i,j}(t) + \sqrt{\frac{\log(T^2)}{2N_{i,j}(t)}} \right)$.
 - 22: **end while**
-

instance of DA is terminated. Following the termination, a new value matrix is created (Line 4), and another instance of DA is run (Line 6). This process is repeated until the number of rounds reaches T . The following theorem states the optimality of the regret of RDA-UCB.

Theorem 3. (Regret bound of RDA-UCB) *Assume that l, h satisfies that $l \leq v_{i,j} \leq h$ for all $i \in N, j \in M$. Then, the following inequality holds:*

$$\mathbb{E}[\text{Regret}(T)] \leq \tilde{O} \left(\text{poly}(n, m) \sqrt{T} \right), \quad (5)$$

where $\text{poly}(n, m)$ is a polynomial of n, m that is independent of T .

5.5 Discussion on the rate of regret

In summary, our theoretical contributions are twofold: i) DA converges in our setting where the values of agent-type pairs are unobservable beforehand and ii) DA-EtC has the upper and lower bound being dependent on the numbers of agents and item types, as well as the time horizon. In what follows, we will examine and discuss what these results

imply. Theorem 2 provides DA-EtC has the regret upper bound of $\tilde{O}((nm)^{1/3}T^{2/3})$. The exponent on m here is smaller than that of lower bound of $\tilde{\Omega}((mT)^{1/2})$ in Theorem 1. This does not contradict because Theorem 2 assumes the exploration per parameter $T_0/(nm) = \Omega(1)$, or equivalently $T > nm$. Thus, for any T, n, m such that $T > m$, the upper bound is no smaller than the lower bound.

We consider giving an upper bound for DA-UCB is challenging mainly because the i.i.d. assumption is crucial for deriving the performance bound in Lemma 1. However, if we plug-in the UCB values at each round t to DA-Iter, it no longer satisfy the assumption. Started from Xiao (2009), most of existing results on DA require the data to be i.i.d. A few notable exceptions are as follows: Agarwal and Duchi (2013) used in DA for mixing process, but the application of their results to a regularized version of DA such as ours is highly non-trivial. Note also that the classes of non-stationarity that Liao et al. (2022) consider are not directly applicable to our settings. It might be possible to apply the analyses by Xiao (2009); Liao et al. (2022) to our case. However, the online regret term⁶ is very challenging to bound in our case. To avoid the nonstationarity, we invented a theoretical algorithm that we call RDA-UCB. Although this algorithm employs multiple instance of DA, the number of instances is $O(nm \log(T)) = \tilde{O}(1)$ as a function of T , and thus it does not matter to the leading term of \sqrt{T} . As a result, it achieves the optimal rate $\tilde{O}(\sqrt{T})$ with respect to the number of rounds T . We consider the dependence of RDA-UCB to n, m (available in the appendix) to be suboptimal, and a more practical algorithm is desired.

Let us finally compare our results with the rate of regret in the partial monitoring problems (Bartók et al., 2014; Komiyama et al., 2015). In particular, our results is $\tilde{\Theta}(T^{1/2})$ with respect to time horizon T . A partial monitoring problem is said to be *easy* if the chosen action itself defines how many items are assigned, e.g., multi-armed bandit problems (Lai and Robbins, 1985), which corresponds to $\tilde{\Theta}(T^{1/2})$ regret. Otherwise, the problem is *hard*, which corresponds to $\tilde{\Theta}(T^{2/3})$ regret. This categorization suggests that numerous problems involving parameters that are not directly observable from the optimal decision, such as dynamic pricing, inevitably have a regret complexity of $\tilde{\Theta}(T^{2/3})$. Our results indicate that the problem we consider belongs to the former (i.e., $\tilde{O}(\sqrt{T})$) class of partial monitoring. This is interesting because we have latent parameters (coefficients $(\beta_i)_{i \in N}$) that depend on all agents and thus are not solely determined by the values of particular user i . This instills hope that a large class of bandit optimizations that entails optimizing latent parameters can achieve $\tilde{O}(\sqrt{T})$ regret. We also note that the comparison with bandit convex optimization and our problem is in Section B in the appendix.

⁶Namely, $R_t(w)$ in Section 4.1 of (Liao et al., 2022). Note that this is different from the regret in our paper.

6 SIMULATIONS

This section evaluates DA-EtC and DA-UCB in several synthetic and real datasets, *Uniform*, *Jester*, and *Household*. The source code and the dataset is available at <https://github.com/nanahaku/DA-bandit>. For all data sets, we assume that the types of items are uniformly distributed. That is, $s_j = 1/m$ for all j . We set the objective as $B_i = 1/n$ for all i . The projection range of β_i in DA-Iter is set to $[B_i/(1 + 0.95), (1 + 0.95)]$, which is wide enough to cover all datasets. The noise $\varepsilon_t \in \{0, 1\}$ for realized utilities is determined according to a Bernoulli distribution, whose probability that the event $\varepsilon_t = 1$ occurs is specified by each associated true value $v_{i,j}$, since it is normalized in $[0, 1]$. The length of the trial rounds T_0 for DA-EtC is set to $T^{2/3}/(nm)^{1/3}$. We adjust the numbers of agents (n) and item types (m), as well as the time horizon (T), according to the dataset used. Figures 1 and 2 display the averaged amounts of regrets over 20 problem instances with different random seeds.

Benchmark algorithms: We have prepared three naive algorithms called *Random*, *UCB*, and *DA-Grdy*. Random allocates each of the arriving items to agents at uniformly random. UCB allocates each one to the agent whose UCB value is the highest for the item type. UCB values, or each estimator $\hat{v}_{i,j}$, are initialized to one and are updated according to the sampled pairs of an agent and an item type as described in Line 8 in Algorithm 4. The UCB algorithm is designed to find an allocation that maximizes social welfare. As noted in Section 2, the resulting allocation from UCB is expected to yield a lower NSW than DA-UCB since some of the agents cannot receive enough amounts of items under the SW maximizing allocation. DA-Grdy does not stop updating the estimator $\hat{v}_{i,j}(t)$ at each round, unlike DA-EtC.

Datasets: *Uniform*, *Jester*, and *Household*. First, to generate the Uniform dataset, we drew a set of values $\{v_{i,j}\} \in [0, 1]^{n \times m}$ at uniformly random, where we set $n = 10, m = 10$. We ran algorithms up to time horizon $T = 100000$.

Second, Jester dataset was built for recommender systems and collaborative filtering studies (Goldberg et al., 2001). We focus on the dataset that contains the ratings of 100 jokes by about 25000 individuals. Among them, 7200 individuals rate all 100 jokes (https://eigentaste.berkeley.edu/dataset/jester_dataset_1_1.zip (Kroer et al., 2021)). We randomly select 10 out of 7200 individuals and 50 out of 100 jokes ($n = 10$ and $m = 50$). Time horizon T is set to 300000. Since the values of ratings lie between -10 and 10 , they are normalized to $[0, 1]$ for our simulation.

Finally, Household dataset consists of data from 2876 individuals regarding their estimated willingness-to-pay for 50 household items that were selected from an online review

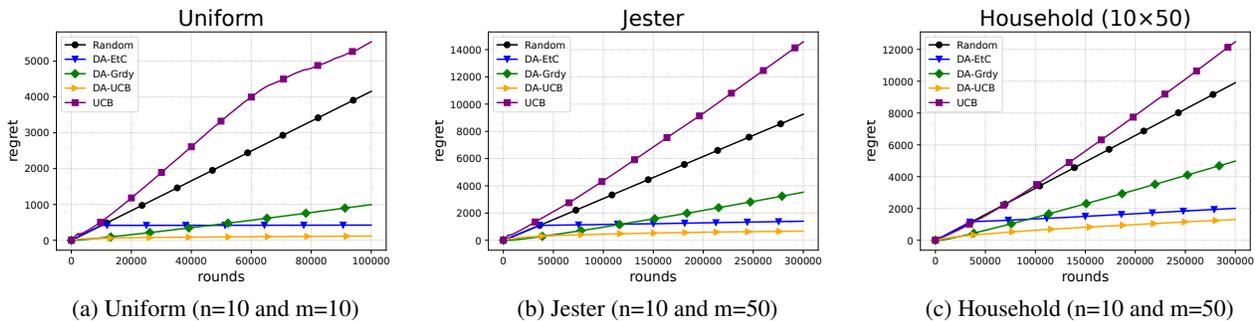


Figure 1: Regret of all algorithms.

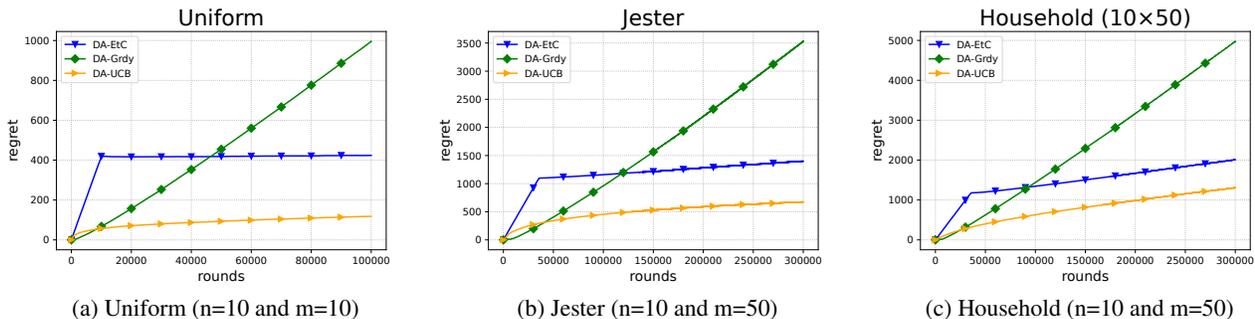


Figure 2: Regret of the three best algorithms, with the lines vertically enlarged to provide a clearer view of the data. The numbers presented are the same as Figure 1.

site (Kroer et al., 2021).⁷ We herein randomly select 10 out of 2876 individuals ($n = 10$ and $m = 50$). Time horizon T is set to 300000 as well as Jester. Since the scores of willingness-to-pay lies between 0 and 100, we normalized it to $[0, 1]$.

Results:

Figure 1 depicts Regret across the five algorithms for each of three datasets. Random and UCB are designed as a naive benchmark, and they incur significant regret compared to the DA-based algorithms. Notably, UCB performs worse than Random that is completely ignorant of the agents’ preferences. This is because UCB maximizes social welfare (the sum of agents’ utilities) instead of NSW (the product of them), and it often assigns items unequally, resulting in several agents receiving very small total utilities. Figure 2 depicts the same as Figure 1 except that the lines are vertically enlarged for a clearer view of the data. Apparently, DA-EtC and DA-UCB outperform DA-Grdy and achieve lower regret in the long run, although they do not in the short term. This is because DA-EtC and DA-UCB incur the cost for exploration before exploitation, whereas DA-Grdy attempts to exploit from the start.

DA-Grdy plug-ins empirical mean $\hat{v}_{i,j}(t)$ to DA-Iter for each round. This algorithm suffers underexploration. For

example, assume that $\hat{v}_{i,j}(t)$ is 0 for some i, j at some round t . DA is unlikely to assign future items of type j to user i , which will prevent $\hat{v}_{i,j}$ from being updated. If such underestimation occurs with a probability of $\Theta(1)$ for some (i, j) , DA-Grdy can continue to approach an incorrect optimization, which results in a regret of $\Theta(T)$.

7 CONCLUSION

This paper has considered an online fair division problem where the values of items are unknown beforehand. In this problem, there is a natural notion of regret that measures how fast we can find the optimal allocation that asymptotically maximizes NSW.

We proposed the algorithms to allocate items via dual averaging with its utility estimated from the past obtained utilities. We proposed two algorithms: DA-EtC and DA-UCB. The former is designed to feed i.i.d., data to DA. EA-EtC has the $\tilde{O}((nm)^{1/3}T^{2/3})$ regret, and the latter does not have such a regret bound but empirically performs better. Furthermore, we derived a $\tilde{O}((mT)^{1/2})$ regret lower bound irrespective of which algorithm is chosen. A version of DA-UCB called RDA-UCB achieves $\tilde{O}(\text{poly}(n, m)\sqrt{T})$ regret, which is optimal with respect to T .

Our results call for subsequent research, including but not limited to: deriving a regret bound for the DA-UCB, and

⁷The exact dataset was provided by (Kroer et al., 2021) upon our request.

investigating the efficacy of structural models, such as linear models and factorized models.

Acknowledgements

This work was partially supported by JSPS KAKENHI Grant Numbers JP21H04890 and JP23K17547. We are grateful to Alex Peysakhovich for providing the source code and dataset, which we referenced in the implementation of our experiment’s source code. We would also like to express our gratitude to the anonymous reviewers for their helpful comments.

References

- A. Agarwal and J. C. Duchi. The generalization ability of online algorithms for dependent data. *IEEE Transactions on Information Theory*, 59(1):573–587, 2013.
- M. Aleksandrov and T. Walsh. Online fair division: A survey. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34:13557–13562, 2020.
- G. Bartók, D. P. Foster, D. Pál, A. Rakhlin, and C. Szepesvári. Partial monitoring - classification, regret bounds, and algorithms. *Mathematics of Operations Research*, 39(4):967–997, 2014.
- M. H. Bateni, Y. Chen, D. F. Ciocan, and V. Mirrokni. Fair resource allocation in a volatile marketplace. *Operations Research*, 70(1):288–308, 2022.
- S. Bouveret and M. Lemaître. Efficiency and sequenceability in fair division of indivisible goods with additive preferences. In *Proceedings of the Sixth International Workshop on Computational Social Choice*, 2016.
- S. J. Brams and A. D. Taylor. *Fair division - from cake-cutting to dispute resolution*. Cambridge University Press, 1996.
- E. Budish. The combinatorial assignment problem: approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.
- I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum nash welfare. *ACM Transactions on Economics and Computation (TEAC)*, 7(3):1–32, 2019.
- B. Codenotti and K. Varadarajan. *Computation of Market Equilibria by Convex Programming*, pages 135–158. Cambridge University Press, 2007.
- E. Eisenberg and D. Gale. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics*, 30(1):165–168, 1959.
- Y. Gao, A. Peysakhovich, and C. Kroer. Online market equilibrium with application to fair division. In *Advances in Neural Information Processing Systems*, volume 34, pages 27305–27318, 2021.
- K. Goldberg, T. Roeder, D. Gupta, and C. Perkins. Eigen-taste: A constant time collaborative filtering algorithm. *Information Retrieval*, 4:133–151, 2001.
- E. Hazan and K. Levy. Bandit convex optimization: Towards tight bounds. In *Advances in Neural Information Processing Systems*, volume 27, 2014.
- K. Jain and V. V. Vazirani. Eisenberg–gale markets: Algorithms and game-theoretic properties. *Games and Economic Behavior*, 70:84–106, 2010.
- I. Kash, A. D. Procaccia, and N. Shah. No agent left behind: Dynamic fair division of multiple resources. *Journal of Artificial Intelligence Research*, 51(1):579–603, 2014.
- E. Kaufmann, O. Cappé, and A. Garivier. On the complexity of best-arm identification in multi-armed bandit models. *Journal of Machine Learning Research*, 17:1:1–1:42, 2016.
- J. Komiyama, J. Honda, and H. Nakagawa. Regret lower bound and optimal algorithm in finite stochastic partial monitoring. In *Advances in Neural Information Processing Systems*, volume 28, 2015.
- C. Kroer, A. Peysakhovich, E. Sodomka, and N. E. S. Moses. Computing large market equilibria using abstractions. *Operations Research*, 70(1):329–351, 2021.
- T. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
- L. Liao, Y. Gao, and C. Kroer. Nonstationary dual averaging and online fair allocation. In *Advances in Neural Information Processing Systems*, 2022.
- H. Moulin. *Fair Division and Collective Welfare*. The MIT Press, 2003.
- J. F. Nash. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.
- A. D. Procaccia and J. Wang. Fair enough: Guaranteeing approximate maximin shares. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, pages 675–692, 2014.
- S. R. Sinclair, S. Banerjee, and C. L. Yu. Sequential fair allocation: Achieving the optimal envy-efficiency tradeoff curve. *Operations Research*, 2022. To appear.
- V. V. Vazirani. *Combinatorial Algorithms for Market Equilibria*, chapter 5, pages 103–134. Cambridge University Press, 2007.
- V. V. Vazirani. The notion of a rational convex program, and an algorithm for the arrow-debreu nash bargaining game. *Journal of the ACM*, 59(2), 2012.
- N. K. Vishnoi. *Algorithms for Convex Optimization*. Cambridge University Press, 2021.
- Y. Wang, M. Sharma, C. Xu, S. Badam, Q. Sun, L. Richardson, L. Chung, E. H. Chi, and M. Chen. Surrogate for

long-term user experience in recommender systems. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, page 4100–4109, 2022.

L. Xiao. Dual averaging method for regularized stochastic learning and online optimization. In *Advances in Neural Information Processing Systems*, volume 22, 2009.

(e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]

5. If you used crowdsourcing or conducted research with human subjects, check if you include:

(a) The full text of instructions given to participants and screenshots. [Not Applicable]

(b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]

(c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

Checklist

1. For all models and algorithms presented, check if you include:

(a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]

(b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]

(c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [No]

2. For any theoretical claim, check if you include:

(a) Statements of the full set of assumptions of all theoretical results. [Yes]

(b) Complete proofs of all theoretical results. [Yes]

(c) Clear explanations of any assumptions. [Yes]

3. For all figures and tables that present empirical results, check if you include:

(a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes]

(b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Yes]

(c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]

(d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [No]

4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:

(a) Citations of the creator If your work uses existing assets. [Yes]

(b) The license information of the assets, if applicable. [Not Applicable]

(c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]

(d) Information about consent from data providers/curators. [No]

A NOTATION TABLE

Table 1 summarizes our notation.

Table 1: Major notation

symbol	definition
N	set of agents
M	set of the types of items
$v_{i,j}$	ex ante (expected) value item of type j for agent i
T	time horizon (the number of rounds)
$t \in T$	each round in T
$i(t) \in N$	winner, or agent who is allocated the item at round t
$j(t) \in M$	item type which arrives at round t
$s_j \in \mathcal{S}$	probability distribution where item type j is drawn
$u_i(t)$	ex post (realized) utility of agent i at round t
ϵ_t	sub-Gaussian random variable with its radius σ^2
$U_i(T)$	cumulative utility of agent i in T rounds
B_i	priority or per-period budget rate given to agent i
$x_{i,j}$	fraction of item type j allocated to agent i in the EG program (Eq. (1))
p_j	price of item type j in D_{EG}
β_i	multiplier in DA-Iter (Algorithm 2)
$\hat{v}_{i,j}^{\text{DA}}$	argument value $v_{i,j}$ in DA-Iter (Algorithm 2)
\bar{u}_i^{DA}	mean utility of $u_i^{\text{DA}}(t)$ (Algorithm 2)
T_0	exploration rounds in DA-EtC (Algorithm 3)
$\hat{v}_{i,j}^{\text{EtC}}$	Estimator in DA-EtC (Algorithm 3)
$\hat{v}_{i,j}^{\text{UCB}}$	UCB value in DA-UCB (Algorithm 4)
$N_{i,j}(t)$	number of times agent i received an item of type j up to round $t - 1$
$u^{*,\text{DA}}$	solution of EG with estimators $\{\hat{v}_{i,j}^{\text{EtC}}\}_{i,j}$
$u^{*,\text{true}}$	utility of EG problem with $v_{i,j}$
$u^{*,\text{RUCB}}$	Defined in Eq. (28).
$u^{*,\text{RUCB}(k)}$	utility of EG problem with corresponding UCB values of k th instance of DA
$u^{\text{DA},\text{true}}$	mean utility that agents receive when we run the algorithm (Eq. (27))
$u^{\text{DA},\text{RUCB}}$	Defined in Eq. (30)
$u^{\text{DA},\text{RUCB}(k)}$	mean utility of the k th instance of Dual Averaging.
$\hat{v}_{i,j}^{\text{EtC}}$	Estimator in DA-EtC (Algorithm 3)
$\hat{v}_{i,j}^{\text{UCB}}(t)$	UCB value of round t in DA-UCB (Algorithm 4)
$\hat{v}_{i,j}^{\text{RUCB}}(t)$	UCB value of round t in RDA-UCB (Algorithm 5)
T_k	number of rounds in which k th instance of DA was run

B Relation with Bandit Convex Optimization

This section discusses the online optimization (i.e., regret minimization) in this paper and bandit convex optimization (BCO) Hazan and Levy (2014), which is known to have an $O(\sqrt{T})$ bound under some assumptions.

- BCO is more challenging than our optimization in the sense that the loss function is given adversarially. This means that stochastic bandit algorithms, such as UCB, cannot be directly applied to BCO.
- On the contrary, our regret with respect to fair division optimizes latent parameters (β_i) that depend on the data sequence. This is challenging because we cannot directly observe the loss for each round.

In summary, the difficulty of BCO and our setting cannot be directly compared. Therefore, our bounds ($\tilde{O}(T^{2/3})$ upper bound and $\tilde{\Omega}(T^{1/2})$ lower bound) are nontrivial.

C PROOFS ON DA-ETC

C.1 Proof of Theorem 1

Proof of Theorem 1. We consider the case in which the type of items is equally distributed; $s_j = 1/m$ for all $j \in [M]$, and $B_i = 1/n$ for all $i \in [N]$. We consider the case where all feedback is binary. In this proof, we use the term ‘‘model’’ to denote a value matrix $\{v_{i,j}\}_{i,j} \in \mathbb{R}^{n \times m}$. We consider the following classes of models, $v : \forall_{i,j} v_{i,j} \in \{1/2, 1/2(1 + \sqrt{m/T})\}^{nm}$. We call $v^{(0)} : \forall_{i,j} v_{i,j}^{(0)} = 1/2$ the *base model*.

The following lemma, which is a version of Lemma 19 in Kaufmann et al. (2016), is used during the proof.

Lemma 2. (lower bound on any event) *Let $v^{(1)}, v^{(2)}$ be two models. Let $\mathbb{E}_{v^{(1)}}, \mathbb{E}_{v^{(2)}}$ be the corresponding expectations and $\mathbb{P}_{v^{(1)}}, \mathbb{P}_{v^{(2)}}$ be the corresponding probabilities, respectively. Then, the following inequality holds for any event \mathcal{E} .*

$$\sum_{i \in [n], j \in [m]} \mathbb{E}_{v^{(1)}} [N_{i,j}(T+1)] d(v_{i,j}^{(1)}, v_{i,j}^{(2)}) \geq d(\mathbb{P}_{v^{(1)}}(\mathcal{E}), \mathbb{P}_{v^{(2)}}(\mathcal{E})), \quad (6)$$

where $d(p, q) = p \log(p/q) + (1-p) \log((1-p)/(1-q))$ is the KL divergence between two Bernoulli distributions.

Let $j_i = \arg \min_{j'} \mathbb{E}_{v^{(0)}} [N_{i,j'}(T+1)]$ be the type of item that agent i receives the least number of times during T under the base model. By definition, $\sum_i N_{i,j_i}(T+1) \leq T/m$. Consider another model $v^{(a)}$ where $v_{i,j_i} = (1 + \sqrt{m/T})/2$ for each $i \in N$ and $v_{i,k} = 1/2$ for $k \neq j_i$. We have

$$\sum_{i \in [n], j \in [m]} \mathbb{E}_{v^{(0)}} [N_{i,j}(T+1)] d(v_{i,j}^{(0)}, v_{i,j}^{(a)}) \quad (7)$$

$$\leq (T/m) d(1/2, 1/2(1 + \sqrt{m/T})) \quad (8)$$

$$\leq (T/m) \times O(m/T) \quad (\text{by } d(1/2, 1/2 + \alpha) = O(\alpha^2)) \quad (9)$$

$$= O(1). \quad (10)$$

Consider the event

$$\mathcal{E} = \left\{ \sum_i N_{i,j_i}(T+1) \leq 2T/m \right\}.$$

By definition, $\mathbb{P}_{v^{(0)}}[\mathcal{E}] \geq 1/2$. Lemma (2) and Eq. (10) implies that

$$d(\mathbb{P}_{v^{(0)}}[\mathcal{E}], \mathbb{P}_{v^{(a)}}[\mathcal{E}]) = O(1),$$

which implies that

$$\mathbb{P}_{v^{(a)}}[\mathcal{E}] = \Omega(1).$$

Under event \mathcal{E} on the alternative model, the regret is at least $\Omega(\sqrt{mT})$, which completes the proof. \square

C.2 Proof of Theorem 2

Proof of Theorem 2. Let $l' = 2l, h' = h/2$. EtC uniformly explores during the first T_0 rounds, and the estimator $\hat{v}_{i,j} = \hat{v}_{i,j}(T_0 + 1)$ is based on $N_{i,j}(T_0 + 1) \approx T_0/(nm)$ samples. Some care is needed because $N_{i,j}(T_0 + 1)$ itself is a random variable. Let

$$\mathcal{A} = \bigcap_{i,j} \left\{ N_{i,j}(T_0 + 1) \geq \frac{T_0}{2nm} \right\} \quad (11)$$

$$\mathcal{B} = \bigcap_{i,j} \left\{ |\hat{v}_{i,j}(T_0 + 1) - v_{i,j}| \leq \sqrt{\frac{8\sigma^2 nm \log(nmT)}{T_0}} \right\}. \quad (12)$$

Since $N_{i,j}(T_0 + 1)$ is a sum of T_0 binary random variables with its mean $1/(nm)$ (i.e., item is of type j , and agent i won it), by the multiplicative Chernoff inequality,

$$N_{i,j}(T_0 + 1) < \frac{T_0}{2nm}$$

holds with probability at most $\exp(-T_0/(8nm)) < 1/(nmT)$, and Event \mathcal{A} holds with probability at least $1 - 1/T$ by considering the union bound of it over $i \in N, j \in M$.

Second, we show that \mathcal{B} holds with a high-probability given \mathcal{A} . Given n' i.i.d. samples with its mean $v_{i,j}$ and sub-Gaussian radius σ , we have

$$|\hat{v}_{i,j}(T_0 + 1) - v_{i,j}| \leq \sqrt{\frac{4\sigma^2 \log(nmT)}{n'}}$$

with probability at least $1 - 1/(nmT)^2$. Union bound of this over possible random value $n' = N_{i,j}(T_0 + 1) \in [T]$ yields

$$|\hat{v}_{i,j}(T_0 + 1) - v_{i,j}| \leq \sqrt{\frac{8\sigma^2 nm \log(nmT)}{T_0}} \Big| N_{i,j}(T_0 + 1) \geq \frac{T_0}{2nm}$$

and taking its union bound over $i \in N, j \in M$ yields \mathcal{B} that holds with probability at least $1 - 1/(nmT)$ given \mathcal{A} . In the following steps, we assume \mathcal{B} because the probability that \mathcal{B} does not hold is $O(1/T)$, and the regret in this is at most $O(T) \times O(1/T) = O(1)$, which is negligible.

Assuming that \mathcal{B} holds at the end of round T_0 , we bound the regret.

$$\begin{aligned} & \frac{\text{Regret}(T)}{T} \\ &= \prod_i (u_i^{*,\text{true}})^{B_i} - \frac{1}{T} \prod_{i \in N} U_i(T)^{B_i} \\ &\leq \prod_i (u_i^{*,\text{true}})^{B_i} - \frac{1}{T} \prod_{i \in N} ((T - T_0) \bar{u}_i^{\text{EtC}})^{B_i} \quad (\text{the sum of utility in the first } T_0 \text{ rounds is non-negative}) \\ &= \prod_i (u_i^{*,\text{true}})^{B_i} - \frac{T - T_0}{T} \prod_{i \in N} (\bar{u}_i^{\text{EtC}})^{B_i} \quad (\text{by } \sum_i B_i = 1) \\ &= \frac{T_0}{T} \prod_i (u_i^{*,\text{true}})^{B_i} + \frac{T - T_0}{T} \left(\prod_i (u_i^{*,\text{true}})^{B_i} - \prod_i (\bar{u}_i^{\text{EtC}})^{B_i} \right) \\ &\leq \frac{T_0}{T} \prod_i (u_i^{*,\text{true}})^{B_i} + \left| \prod_i (u_i^{*,\text{true}})^{B_i} - \prod_i (\bar{u}_i^{\text{EtC}})^{B_i} \right|. \end{aligned} \quad (13)$$

Here, the second term is bounded as

$$\prod_i (u_i^{*,\text{true}})^{B_i} - \prod_i (\bar{u}_i^{\text{EtC}})^{B_i} = \left(\prod_i (u_i^{*,\text{true}})^{B_i} - \prod_i (u_i^{*,\text{DA}})^{B_i} \right) + \left(\prod_i (u_i^{*,\text{DA}})^{B_i} - \prod_i (\bar{u}_i^{\text{DA}})^{B_i} \right) \quad (14)$$

$$+ \left(\prod_i (\bar{u}_i^{\text{DA}})^{B_i} - \prod_i (\bar{u}_i^{\text{EtC}})^{B_i} \right). \quad (15)$$

From the definition of $u_i^{*,\text{DA}}$, for any $\{x_{i,j}\}$, $\prod_i (u_i^{*,\text{DA}})^{B_i} \geq \prod_i \left(\sum_{j \in M} s_j \hat{v}_{i,j}(T_0 + 1) x_{i,j} \right)^{B_i}$ holds. Hence, we have

$$\prod_i (u_i^{*,\text{true}})^{B_i} - \prod_i (u_i^{*,\text{DA}})^{B_i} \leq \prod_i (u_i^{*,\text{true}})^{B_i} - \prod_i (\hat{u}_i^{*,\text{DA}})^{B_i},$$

where $\hat{u}_i^{*,\text{DA}} = \sum_{j \in M} s_j \hat{v}_{i,j}(T_0 + 1) x_{i,j}^*$ and $\{x_{i,j}^*\}$ be the solution of the optimization (Eq. (1)) with true $\{v_{i,j}\}$. On the other hand, under \mathcal{B} ,

$$|u_i^* - \hat{u}_i^{*,\text{DA}}| = \left| \sum_{j \in M} s_j x_{i,j}^* (v_{i,j} - \hat{v}_{i,j}(T_0 + 1)) \right| \leq \sum_{j \in M} s_j x_{i,j}^* |v_{i,j} - \hat{v}_{i,j}(T_0 + 1)| \leq \tilde{O} \left(\sqrt{\frac{nm}{T_0}} \right)$$

Combining these inequalities and Lemma 3, under the assumption that $u_i^{*,\text{true}} = \Theta(1)$,

$$\prod_i (u_i^{*,\text{true}})^{B_i} - \prod_i (u_i^{*,\text{DA}})^{B_i} \leq \tilde{O} \left(\sqrt{\frac{nm}{T_0}} \right). \quad (16)$$

Moreover, Event \mathcal{B} and assumptions imply that $l'/2 \leq \hat{v}_{i,j}(T_0 + 1) \leq h' + l'/2 \leq 2h'$, and thus Lemma 4 with $l = l'/2, h = 2h'$ states that

$$\mathbb{E} \left[\prod_i (u_i^{*,\text{DA}})^{B_i} - \prod_i (\bar{u}_i^{\text{DA}})^{B_i} \right] \leq \tilde{O} \left(nC^{\text{DA}} \sqrt{\frac{1}{T}} \right), \quad (17)$$

where we consider $l', h' > 0$ to be constants.

DA gives at least $\Omega(T/(\sum_i (4h'/l')^2)) = \Omega((T/n) \times 1) = \Omega(T/n)$ items to each agent because if agent i_2 receives $(2h'/(l'/2))^2$ times more items than agent i_1 , then the ratio β_{i_2}/β_{i_1} is at least $(2h'/(l'/2))$, and agent i_1 is prioritized in the next allocation to agent i_2 no matter what type of item is. Using this fact, we have $\bar{u}_i^{\text{DA}} = \Theta(1)$. Applying a concentration inequality to the samples during the exploitation phase, with probability at least $1 - 1/T$, we have $\bar{u}_i^{\text{EtC}} = \Theta(1), |\bar{u}_i^{\text{EtC}} - \bar{u}_i^{\text{DA}}| = \tilde{O}(\sqrt{nm/T_0})$. Using this and Lemma 3, we have

$$\left(\prod_i (\bar{u}_i^{\text{DA}})^{B_i} \right) - \prod_i (\bar{u}_i^{\text{EtC}})^{B_i} = \tilde{O} \left(\sqrt{\frac{nm}{T_0}} \right). \quad (18)$$

In summary,

$$\mathbb{E} \left[\frac{\text{Regret}(T)}{T} \right] = \frac{T_0}{T} + \tilde{O} \left(\sqrt{\frac{nm}{T_0}} \right) + nC^{\text{DA}} \sqrt{\frac{1}{T}}. \quad (19)$$

□

C.3 Additional Lemmas on DA-EtC

Lemma 3. Let two vectors $u^{(1)} = (u_1^{(1)}, \dots, u_n^{(1)})$ and $u^{(2)} = (u_1^{(2)}, \dots, u_n^{(2)})$ be such that $u_i^{(1)} = \Theta(1)$ and $u_i^{(2)} \geq u_i^{(1)}(1 - r_i)$, and $r = \max_i r_i$. Then,

$$\prod_i (u_i^{(1)})^{B_i} - \prod_i (u_i^{(2)})^{B_i} \leq O(r). \quad (20)$$

Proof of Lemma 3.

$$\prod_i (u_i^{(1)})^{B_i} - \prod_i (u_i^{(2)})^{B_i} = \prod_i (u_i^{(1)})^{B_i} - \prod_i \left\{ u_i^{(1)}(1 - r_i) \right\}^{B_i} \quad (21)$$

$$\leq \prod_i (u_i^{(1)})^{B_i} - \prod_i \left\{ u_i^{(1)}(1 - \max_{i'} r_{i'}) \right\}^{B_i} \quad (22)$$

$$= \max_{i'} r_{i'} \prod_i (u_i^{(1)})^{B_i}. \quad (23)$$

□

Lemma 4. Assume that $T_0 = o(T)$. Then, the following inequality holds:

$$\mathbb{E} \left[\prod_i (u_i^{*,\text{DA}})^{B_i} - \prod_i (\bar{u}_i^{\text{DA}})^{B_i} \right] \leq \tilde{O} \left(nC^{\text{DA}} \sqrt{\frac{1}{T}} \right). \quad (24)$$

Proof of Lemma 4. Remember that \bar{u}_i^{DA} is the regret-per-round during the exploitation rounds. Let $T' = T - T_0 = \Theta(T)$. In view of DA, it is an online learning with T' rounds where the value of each item is $\{\hat{v}_{i,j}(t)\}_i$.

Lemma 1 implies that

$$\mathbb{E} \left[(u_i^{*,\text{DA}} - \bar{u}_i^{\text{DA}})^2 \right] \leq \frac{C^{\text{DA}} \log T'}{T'},$$

and Markov's inequality implies that

$$\Pr[|u_i^{*,\text{DA}} - \bar{u}_i^{\text{DA}}| \geq \epsilon'] \leq \frac{1}{\epsilon'^2} \frac{C^{\text{DA}} \log T'}{T'},$$

and by letting $\epsilon' = \epsilon u_i^{*,\text{DA}}$ we have

$$\Pr \left[\frac{|u_i^{*,\text{DA}} - \bar{u}_i^{\text{DA}}|}{u_i^{*,\text{DA}}} \geq \epsilon \right] \leq \frac{1}{(u_i^{*,\text{DA}})^2 \epsilon^2} \frac{C^{\text{DA}} \log T'}{T'}, \quad (25)$$

Taking union bound of Eq. (25) over $i \in N$, we have

$$\Pr \left[\max_i \frac{|u_i^{*,\text{DA}} - \bar{u}_i^{\text{DA}}|}{u_i^{*,\text{DA}}} \geq \epsilon \right] \leq \frac{n}{\min_i (u_i^{*,\text{DA}})^2 \epsilon^2} \frac{C^{\text{DA}} \log T'}{T'}. \quad (26)$$

Here, letting

$$R := \max_i \frac{u_i^{*,\text{DA}} - \bar{u}_i^{\text{DA}}}{u_i^{*,\text{DA}}} \leq 1,$$

then $\bar{u}_i^{\text{DA}} \geq (1 - R)u_i^{*,\text{DA}}$ for all i and we have

$$\begin{aligned} \prod_i (u_i^{*,\text{DA}})^{B_i} - \prod_i (\bar{u}_i^{\text{DA}})^{B_i} &\leq \prod_i (u_i^{*,\text{DA}})^{B_i} - \prod_i ((1 - R)u_i^{*,\text{DA}})^{B_i} \\ &\leq R \prod_i (u_i^{*,\text{DA}})^{B_i}. \quad (\text{by } \prod_i (x)^{B_i} = x) \end{aligned}$$

Using this, we have

$$\begin{aligned} &\mathbb{E} \left[\prod_i (u_i^{*,\text{DA}})^{B_i} - \prod_i (\bar{u}_i^{\text{DA}})^{B_i} \right] \\ &\leq \mathbb{E} \left[R \prod_i (u_i^{*,\text{DA}})^{B_i} \right] \\ &\leq \int_{1/\sqrt{T'}}^1 \prod_i (u_i^{*,\text{DA}})^{B_i} \Pr[R \geq x] dx + \sqrt{\frac{1}{T'}} \prod_i (u_i^{*,\text{DA}})^{B_i} \\ &= \int_{1/\sqrt{T'}}^1 \prod_i (u_i^{*,\text{DA}})^{B_i} \Pr[R \geq x] dx + O\left(\sqrt{\frac{1}{T'}}\right) \\ &\leq \int_{1/\sqrt{T'}}^1 \prod_i (u_i^{*,\text{DA}})^{B_i} \frac{n}{\min_i (u_i^{*,\text{DA}})^2 x^2} \frac{C^{\text{DA}} \log T'}{T'} dx + O\left(\sqrt{\frac{1}{T'}}\right) \quad (\text{by Eq. (26)}) \\ &= (u_i^{*,\text{DA}})^{B_i} \frac{n}{\min_i (u_i^{*,\text{DA}})^2} \frac{C^{\text{DA}} \log T'}{T'} [-1/x]_{1/\sqrt{T'}}^1 + O\left(\sqrt{\frac{1}{T'}}\right) \\ &\leq (u_i^{*,\text{DA}})^{B_i} \frac{n}{\min_i (u_i^{*,\text{DA}})^2} \frac{C^{\text{DA}} \log T'}{\sqrt{T'}} + O\left(\sqrt{\frac{1}{T'}}\right) \\ &= \tilde{O}\left(n C^{\text{DA}} \sqrt{\frac{1}{T'}}\right). \end{aligned}$$

□

D PROOFS ON RDA-UCB

D.1 Proof of Theorem 3

Proof of Theorem 3. First, let $u_i^{*,\text{true}}$ be the utility of EG problem with $v_{i,j}$. Let

$$u_i^{\text{DA},\text{true}} = U_i(T)/T \quad (27)$$

be the empirical mean utility that each agent receives when we run the algorithm. Remember that RDA-UCB utilizes several instances of DA. Let y be the number of the DA instances that appear in RDA-UCB. We use index $k = 1, 2, \dots, y$ to represent each instance of DA. Let T_k be the number of rounds where each DA was run. Note that each T_k , as well as y , are the random variables. Let $u_i^{*,\text{RUCB}}, u_i^{\text{DA},\text{RUCB}}$ be:

$$u_i^{*,\text{RUCB}} = \frac{1}{T} (T_1 u_i^{*,\text{RUCB}(1)} + T_2 u_i^{*,\text{RUCB}(2)} + \dots + T_y u_i^{*,\text{RUCB}(y)}), \quad (28)$$

$$u_i^{\text{DA},\text{RUCB}} = \frac{1}{T} \sum_{t=1}^T \mathbb{I}[I(t) = i] \hat{v}_{i(t),j(t),N_{i,j(t)}(t)}^{\text{UCB}} \quad (29)$$

$$= \frac{1}{T} (T_1 u_i^{\text{DA},\text{RUCB}(1)} + T_2 u_i^{\text{DA},\text{RUCB}(2)} + \dots + T_y u_i^{\text{DA},\text{RUCB}(y)}). \quad (30)$$

Namely, The value $u_i^{*,\text{RUCB}}$ indicates the mean true utility and the random variable $u_i^{\text{DA},\text{RUCB}}$ indicates the mean utility over DA instances.

Moreover, let $\hat{v}_{i,j,n}^{\text{RUCB}}$ is the empirical estimate of $v_{i,j}$ with n samples and two events be

$$\mathcal{G}_1 = \bigcap_{i,j,n} \{ \hat{v}_{i,j,n}^{\text{RUCB}} \geq v_{i,j} \}, \quad (31)$$

$$\mathcal{G}_2 = \bigcap_i \left\{ \sum_{t=1}^T \hat{v}_{i(t),j(t),N_{i,j(t)}(t)}^{\text{RUCB}} \leq \sum_{t=1}^T (v_{i(t),j(t)} + \epsilon_t) + C_3 n m \sqrt{T} \right\}, \quad (32)$$

where $C_3 = \tilde{O}(1)$ is defined later in Eq. (45).

From lemma 5, $\mathcal{G} := \mathcal{G}_1 \cap \mathcal{G}_2$ holds with probability at least $1 - O(1/T)$.

We have

$$\begin{aligned} \mathbb{I}[\mathcal{G}] \frac{\text{Regret}(T)}{T} &= \mathbb{I}[\mathcal{G}] \left(\prod_i (u_i^{*,\text{true}})^{B_i} - \prod_i (u_i^{*,\text{RUCB}})^{B_i} \right) + \mathbb{I}[\mathcal{G}] \left(\prod_i (u_i^{*,\text{RUCB}})^{B_i} - \prod_i (u_i^{\text{DA},\text{RUCB}})^{B_i} \right) \\ &\quad + \mathbb{I}[\mathcal{G}] \left(\prod_i (u_i^{\text{DA},\text{RUCB}})^{B_i} - \prod_i (u_i^{\text{DA},\text{true}})^{B_i} \right) \\ &\leq \mathbb{I}[\mathcal{G}_1 \cap \mathcal{G}_2] \left(\prod_i (u_i^{*,\text{true}})^{B_i} - \prod_i (u_i^{*,\text{RUCB}})^{B_i} \right) + \left| \left(\prod_i (u_i^{*,\text{RUCB}})^{B_i} - \prod_i (u_i^{\text{DA},\text{RUCB}})^{B_i} \right) \right| \\ &\quad + \mathbb{I}[\mathcal{G}_1 \cap \mathcal{G}_2] \left| \left(\prod_i (u_i^{\text{DA},\text{RUCB}})^{B_i} - \prod_i (u_i^{\text{DA},\text{true}})^{B_i} \right) \right| \\ &\leq 0 + \left| \left(\prod_i (u_i^{*,\text{RUCB}})^{B_i} - \prod_i (u_i^{\text{DA},\text{RUCB}})^{B_i} \right) \right| + \tilde{O}\left(\frac{nm\sqrt{T}}{T}\right), \end{aligned} \quad (33)$$

where we have used Lemmas 6 and 21 in the last transformation.

Therefore, the regret bound is:

$$\mathbb{E} [\text{Regret}(T)] = \mathbb{E} [\mathbb{I}[\mathcal{G}] \text{Regret}(T)] + \mathbb{E} [\mathbb{I}[\mathcal{G}^c] \text{Regret}(T)] \quad (34)$$

$$\leq \mathbb{E} [\mathbb{I}[\mathcal{G}] \text{Regret}(T)] + T \Pr [\mathcal{G}^c] \quad (35)$$

$$\leq \mathbb{E} [\mathbb{I}[\mathcal{G}] \text{Regret}(T)] + T \times \tilde{O}\left(\frac{1}{T}\right) \quad (\text{by Lemma 5}) \quad (36)$$

$$\leq \mathbb{E} \left[T \left| \prod_i (u_i^{*,\text{RUCB}})^{B_i} - \prod_i (u_i^{\text{DA},\text{RUCB}})^{B_i} \right| \right] + \tilde{O}(nm\sqrt{T}) \quad (\text{by Eq. (33)}) \quad (37)$$

Lemma 20 bounds the first term on RHS of Eq. (37) bounded as

$$\mathbb{E} \left[T \left| \prod_i (u_i^{*,\text{RUCB}})^{B_i} - \prod_i (u_i^{\text{DA},\text{RUCB}})^{B_i} \right| \right] = \tilde{O}\left(n^4 m^2 \sqrt{T}\right), \quad (38)$$

which completes the proof. \square

D.2 Lemmas on the probability of event \mathcal{G}

Lemma 5. *Event $\mathcal{G}_1 \cap \mathcal{G}_2$ holds with probability at least $1 - O(1/T)$.*

Proof of Lemma 5. Subgaussian concentration inequality implies

$$\bigcap_{i,j,n} |\hat{v}_{i,j,n}^{\text{RUCB}} - v_{i,j}| \leq \sigma \sqrt{\frac{2 \log(T^2)}{n}} \quad (39)$$

holds with probability at least $1 - 2/T^2 \times nmT = 1 - O(1/T)$. Eq. (31) easily follows from the Eq. (39). Moreover,

$$\sum_{t=1}^T (\hat{v}_{i(t),j(t),N_{i(t),j(t)}(t)}^{\text{RUCB}} - (v_{i(t),j(t)} + \epsilon_t)) \quad (40)$$

$$\leq 2h + \sum_{i,j} \sum_{p=1}^{\log_2 T} 2^p (\hat{v}_{i,j,2^p}^{\text{RUCB}} - v_{i,j}) + \left| \sum_t \epsilon_t \right| \quad (41)$$

$$\text{(by the fact that RDA-UCB resets once allocation } (i, j) \text{ reaches } N_i(t) = 2, 2^2, 2^3, \dots) \quad (42)$$

$$\leq 2h + \sum_{i,j} \sum_{p=1}^{\log_2 T} 2^p (\hat{v}_{i,j,2^p}^{\text{RUCB}} - v_{i,j}) + \sigma \sqrt{2 \log(T^2)} \sqrt{T} \quad \text{(by subgaussian concentration inequality on } \sum_t \epsilon(t)) \quad (43)$$

$$\leq 2h + nm \sum_{p=1}^{\log_2 T} 2^p \times \sigma \sqrt{\frac{2 \log(T^2)}{2^p}} + \tilde{O}(\sqrt{T}) \quad \text{(by Eq. (39))} \quad (44)$$

$$=: C_3 nm \sqrt{T}. \quad (45)$$

□

D.3 Lemmas on the first term

Lemma 6. *for any agent $i \in N$, with $\mathcal{G}_1 \cap \mathcal{G}_2$ from lemma 5,*

$$\mathbb{I}[\mathcal{G}_1 \cap \mathcal{G}_2] \left(\prod_i (u_i^{*,\text{true}})^{B_i} - \prod_i (u_i^{*,\text{RUCB}})^{B_i} \right) \leq 0. \quad (46)$$

Proof of Lemma 6.

Let T_k be the number of rounds in which k th instance of DA was run and $u_i^{*,\text{RUCB}(k)}$ be the mean utility of Dual Averaging in the k th instance of DA. Let y be the index of the last instance of DA.

Recall that

$$u_i^{*,\text{RUCB}} = \frac{1}{T} (T_1 u_i^{*,\text{RUCB}(1)} + T_2 u_i^{*,\text{RUCB}(2)} + \dots + T_y u_i^{*,\text{RUCB}(y)}). \quad (47)$$

Lemma 7 implies that the region

$$\mathcal{U} = \left\{ (u_i)_{i \in [N]} \in \mathbb{R}^N \mid \left(\prod_i (u_i^{*,\text{true}})^{B_i} \leq \left(\prod_i (u_i)^{B_i} \right) \right) \right\}$$

is convex. Event \mathcal{G}_1 implies that $u_i^{*,\text{RUCB}(k)}$ for each k lies in \mathcal{U} . Since $u_i^{*,\text{RUCB}}$ is a non-negative and Affine combination of elements in \mathcal{U} , $u_i^{*,\text{RUCB}} \in \mathcal{U}$, which is Eq. (46). □

Lemma 7. (Convexity of region) *Let $v, v^1, v^2 \in (\mathbb{R}^+)^n$ be n -dimensional vectors. If $\prod_i v_i^1, \prod_i v_i^2 \geq \prod_i v_i$, then for any $p \in [0, 1]$, $\prod_i (pv_i^1 + (1-p)v_i^2) \geq \prod_i v_i$.*

Proof. Let $c = \prod_i v_i$. Weighted inequality of arithmetic and geometric means states that

$$pv_i^1 + (1-p)v_i^2 \geq (v_i^1)^p (v_i^2)^{1-p} \quad (48)$$

for each i , and thus

$$\prod_i (pv_i^1 + (1-p)v_i^2) \geq \left(\prod_i v_i^1 \right)^p \left(\prod_i v_i^2 \right)^{1-p} \geq \left(\prod_i v_i \right)^p \left(\prod_i v_i \right)^{1-p} = \prod_i v_i. \quad (49)$$

□

D.4 Lemmas on the second term

The structure of this section is as follows. Section D.4.1 shows the results for each instance of DA. Section D.4.2 uses the results for RDA-UCB, which uses multiple instances of DA. By using these lemmas, Section D.4.3 bounds the second term in the main proof.

D.4.1 Auxiliary lemmas on a single instance of DA

This section introduces lemmas that apply for each instance of DA. For ease of notation, we drop the index k of the instances in the context it is clear. For example, β^* indicates the optimal solution of the dual EG problem for the k th instance of DA. In this optimal solution, the value matrix is the corresponding UCB values $v = \hat{v}_{i,j}^{\text{RUCB}}(t_k)$, where t_k is the first round of the k th instance. Lemmas 8–12) are used to derive Lemma 13, and Lemma 13 is used in the subsequent lemmas. Lemmas 8–12 are the version of the similar results in Xiao (2009) (Theorem 1(b) therein) that are tailored for our version of DA. During these lemmas, we use the notation of Xiao (2009). Namely, for a time step τ in view of DA, let $f(\beta, j(\tau))$ be $\max_{i \in N} v_{i,j(\tau)} \beta_i$ and $\Psi(\beta)$ be $-\sum_{i \in N} B_i \log \beta_i$. Let $R_t(\beta) = \sum_{\tau=1}^t (f(\beta^\tau, j(\tau)) - \Psi(\beta^\tau) - f(\beta^*, j(\tau)) + \Psi(\beta^*))$ be the *online regret*⁸ Xiao (2009), where β^τ be the multiplier of DA at τ . Let S_t be $-R_t(\beta) + C \log T$ where $C = \frac{\|v\|^2}{\sigma^2}$, let Ψ^* be convex conjugate of Ψ and let σ be $\sigma = \frac{l^2 \min_{i \in N} B_i}{1 + \delta_0}$.

Lemma 8. *Let $\beta^\tau \in \mathbb{R}^n$ be the solution of DA and $\beta^* \in \mathbb{R}^n$ be the solution of the corresponding EG program. Then, the following inequality holds for any $t \geq 1$ and $(j(\tau))_{\tau=1}^t$:*

$$\frac{\sigma t}{2} \|\beta^{t+1} - \beta^*\|^2 \leq t\Psi^* \left(-\frac{1}{t} s^t \right) + \sum_{\tau=1}^t (\langle g^\tau, \beta^\tau \rangle + \Psi(\beta^\tau)) + \sum_{\tau=1}^t (f(\beta^*, j(\tau)) + \Psi(\beta^*) - f(\beta^\tau, j(\tau)) - \Psi(\beta^\tau)),$$

where $g^t \in \partial f(\beta^t, j(\tau)^t)$.

Proof of Lemma 8. Let us define $s^t = \sum_{\tau=1}^t g^\tau$. From the first-order optimality condition for β^{t+1} , there exist a subgradient $b^{t+1} \in \partial \Psi(\beta^{t+1})$ such that:

$$\langle s^t + t b^{t+1}, \beta - \beta^{t+1} \rangle \geq 0, \forall \beta \in \text{dom}(\Psi). \quad (50)$$

From the strong convexity of Ψ , we have:

$$\Psi(\beta) \geq \Psi(\beta^{t+1}) + \langle b^{t+1}, \beta - \beta^{t+1} \rangle + \frac{\sigma}{2} \|\beta^{t+1} - \beta\|^2, \forall \beta \in \text{dom}(\Psi). \quad (51)$$

By combining (50) and (51), we get:

$$\begin{aligned} \frac{\sigma t}{2} \|\beta^{t+1} - \beta^*\|^2 &\leq \langle s^t, \beta^* - \beta^{t+1} \rangle + t\Psi(\beta^*) - t\Psi(\beta^{t+1}) \\ &= \langle -s^t, \beta^{t+1} \rangle + t\Psi(\beta^*) - t\Psi(\beta^{t+1}) + \langle s^t, \beta^* \rangle \\ &= \langle -s^t, \beta^{t+1} \rangle + t\Psi(\beta^*) - t\Psi(\beta^{t+1}) + \sum_{\tau=1}^t \langle g^\tau, \beta^* \rangle \\ &= \langle -s^t, \beta^{t+1} \rangle + t\Psi(\beta^*) - t\Psi(\beta^{t+1}) + \sum_{\tau=1}^t \langle g^\tau, \beta^* - \beta^\tau \rangle + \sum_{\tau=1}^t \langle g^\tau, \beta^\tau \rangle. \end{aligned}$$

⁸The online regret is different from the regret in our paper

Then, from the convexity of $f(\cdot, j(\tau))$,

$$\begin{aligned} & \frac{\sigma t}{2} \|\beta^{t+1} - \beta^*\|^2 \\ & \leq \langle -s^t, \beta^{t+1} \rangle + t\Psi(\beta^*) - t\Psi(\beta^{t+1}) + \sum_{\tau=1}^t (f(\beta^*, j(\tau)) - f(\beta^\tau, j(\tau))) + \sum_{\tau=1}^t \langle g^\tau, \beta^\tau \rangle \\ & = \langle -s^t, \beta^{t+1} \rangle - t\Psi(\beta^{t+1}) + \sum_{\tau=1}^t (\langle g^\tau, \beta^\tau \rangle + \Psi(\beta^\tau)) + \sum_{\tau=1}^t (f(\beta^*, j(\tau)) + \Psi(\beta^*) - f(\beta^\tau, j(\tau)) - \Psi(\beta^\tau)). \end{aligned}$$

Since $\beta^{t+1} = \arg \max_{\beta} \{\langle -s^t, \beta \rangle - t\Psi(\beta)\}$, we have:

$$\frac{\sigma t}{2} \|\beta^{t+1} - \beta^*\|^2 \leq t\Psi^* \left(-\frac{1}{t} s^t \right) + \sum_{\tau=1}^t (\langle g^\tau, \beta^\tau \rangle + \Psi(\beta^\tau)) + \sum_{\tau=1}^t (f(\beta^*, j(\tau)) + \Psi(\beta^*) - f(\beta^\tau, j(\tau)) - \Psi(\beta^\tau)).$$

□

Lemma 9. We have for any $t \geq 1$ and $(z_\tau)_{\tau=1}^t$:

$$\sum_{\tau=1}^t (\langle \beta^\tau, g^\tau \rangle + \Psi(\beta^\tau)) \leq -t\Psi^* \left(-\frac{1}{t} s^t \right) + \Psi(\beta^2) - \Psi(\beta^1) + \frac{1}{2\sigma} \left(\|g^1\|^2 + \sum_{\tau=2}^t \frac{\|g^\tau\|^2}{\tau-1} \right).$$

Proof of Lemma 9. We introduce the following lemma:

Lemma 10. Assume that a function v is closed and convex and its convex conjugate v^* is differentiable. Then, the gradient of v^* is given by:

$$\nabla v^*(y) = \arg \max_x \{\langle y, x \rangle - v(x)\}.$$

From Lemma 10 and the fact that $\beta^\tau = \arg \max_{\beta} \left\{ \left\langle -\frac{1}{\tau-1} s^{\tau-1}, \beta \right\rangle - \Psi(\beta) \right\}$, we have:

$$\nabla \Psi^* \left(-\frac{1}{\tau-1} s^{\tau-1} \right) = \beta^\tau. \quad (52)$$

Furthermore, since Ψ is σ -strongly convex, its convex conjugate is $\frac{1}{\sigma}$ -smooth, we get:

$$\begin{aligned} \frac{1}{\tau-1} \left\langle \nabla \Psi^* \left(-\frac{1}{\tau-1} s^{\tau-1} \right), g^\tau \right\rangle & = \left\langle \nabla \Psi^* \left(-\frac{1}{\tau-1} s^{\tau-1} \right), -\frac{1}{\tau-1} s^{\tau-1} + \frac{1}{\tau-1} s^\tau \right\rangle \\ & \leq \Psi^* \left(-\frac{1}{\tau-1} s^{\tau-1} \right) - \Psi^* \left(-\frac{1}{\tau-1} s^\tau \right) + \frac{1}{2\sigma} \left\| -\frac{1}{\tau-1} s^{\tau-1} + \frac{1}{\tau-1} s^\tau \right\|^2 \\ & = \Psi^* \left(-\frac{1}{\tau-1} s^{\tau-1} \right) - \Psi^* \left(-\frac{1}{\tau-1} s^\tau \right) + \frac{\|g^\tau\|^2}{2\sigma(\tau-1)^2}. \end{aligned} \quad (53)$$

By combining (52) and (53), we have:

$$\begin{aligned} \langle \beta^\tau, g^\tau \rangle & = (\tau-1)\Psi^* \left(-\frac{1}{\tau-1} s^{\tau-1} \right) - (\tau-1)\Psi^* \left(-\frac{1}{\tau-1} s^\tau \right) + \frac{\|g^\tau\|^2}{2\sigma(\tau-1)} \\ & = (\tau-1)\Psi^* \left(-\frac{1}{\tau-1} s^{\tau-1} \right) - (\tau-1) \max_{\beta} \left\{ \left\langle -\frac{1}{\tau-1} s^\tau, \beta \right\rangle - \Psi(\beta) \right\} + \frac{\|g^\tau\|^2}{2\sigma(\tau-1)} \\ & = (\tau-1)\Psi^* \left(-\frac{1}{\tau-1} s^{\tau-1} \right) - \max_{\beta} \{ \langle -s^\tau, \beta \rangle - (\tau-1)\Psi(\beta) \} + \frac{\|g^\tau\|^2}{2\sigma(\tau-1)} \\ & \leq (\tau-1)\Psi^* \left(-\frac{1}{\tau-1} s^{\tau-1} \right) + \langle s^\tau, \beta^{\tau+1} \rangle + (\tau-1)\Psi(\beta^{\tau+1}) + \frac{\|g^\tau\|^2}{2\sigma(\tau-1)} \\ & = (\tau-1)\Psi^* \left(-\frac{1}{\tau-1} s^{\tau-1} \right) + \langle s^\tau, \beta^{\tau+1} \rangle + \tau\Psi(\beta^{\tau+1}) - \Psi(\beta^{\tau+1}) + \frac{\|g^\tau\|^2}{2\sigma(\tau-1)}. \end{aligned}$$

Since $\beta^{\tau+1} = \arg \min_{\beta} \{\langle s^{\tau}, \beta \rangle + \tau \Psi(\beta)\}$,

$$\begin{aligned} \langle \beta^{\tau}, g^{\tau} \rangle &\leq (\tau - 1) \Psi^* \left(-\frac{1}{\tau - 1} s^{\tau-1} \right) + \min_{\beta} \{\langle s^{\tau}, \beta \rangle + \tau \Psi(\beta)\} - \Psi(\beta^{\tau+1}) + \frac{\|g^{\tau}\|^2}{2\sigma(\tau - 1)} \\ &= (\tau - 1) \Psi^* \left(-\frac{1}{\tau - 1} s^{\tau-1} \right) - \max_{\beta} \{\langle -s^{\tau}, \beta \rangle - \tau \Psi(\beta)\} - \Psi(\beta^{\tau+1}) + \frac{\|g^{\tau}\|^2}{2\sigma(\tau - 1)} \\ &= (\tau - 1) \Psi^* \left(-\frac{1}{\tau - 1} s^{\tau-1} \right) - \tau \Psi^* \left(-\frac{1}{\tau} s^{\tau} \right) - \Psi(\beta^{\tau+1}) + \frac{\|g^{\tau}\|^2}{2\sigma(\tau - 1)}. \end{aligned}$$

Thus, for any $\tau \geq 2$:

$$\langle \beta^{\tau}, g^{\tau} \rangle + \Psi(\beta^{\tau+1}) \leq (\tau - 1) \Psi^* \left(-\frac{1}{\tau - 1} s^{\tau-1} \right) - \tau \Psi^* \left(-\frac{1}{\tau} s^{\tau} \right) + \frac{\|g^{\tau}\|^2}{2\sigma(\tau - 1)} \quad (54)$$

Similarly, from Lemma 10 and the assumption that $\beta^1 = \arg \min_{\beta} \Psi(\beta)$:

$$\nabla \Psi^*(0) = \beta^1. \quad (55)$$

Furthermore, from the smoothness of Ψ^* , we get:

$$\begin{aligned} \langle \nabla \Psi^*(0), g^1 \rangle &= \langle \nabla \Psi^*(-s^0), -s^0 + s^1 \rangle \\ &\leq \Psi^*(-s^0) - \Psi^*(-s^1) + \frac{1}{2\sigma} \| -s^0 + s^1 \|^2 \\ &= \Psi^*(-s^0) - \Psi^*(-s^1) + \frac{\|g^1\|^2}{2\sigma}. \end{aligned} \quad (56)$$

By combining (55) and (56), we have:

$$\langle \beta^1, g^1 \rangle \leq \Psi^*(-s^0) - \Psi^*(-s^1) + \frac{\|g^1\|^2}{2\sigma}.$$

Thus, we have for $\tau = 1$:

$$\langle \beta^1, g^1 \rangle + \Psi(\beta^2) \leq \Psi^*(-s^0) - \Psi^*(-s^1) + \frac{\|g^1\|^2}{2\sigma} + \Psi(\beta^2). \quad (57)$$

By summing (54) and (57) from $\tau = 1$ to t , we have:

$$\begin{aligned} \sum_{\tau=1}^t (\langle \beta^{\tau}, g^{\tau} \rangle + \Psi(\beta^{\tau+1})) &\leq \Psi^*(-s^0) - t \Psi^* \left(-\frac{1}{t} s^t \right) + \Psi(\beta^2) + \frac{1}{2\sigma} \left(\|g^1\|^2 + \sum_{\tau=2}^t \frac{\|g^{\tau}\|^2}{\tau - 1} \right) \\ &= -t \Psi^* \left(-\frac{1}{t} s^t \right) + \Psi(\beta^2) - \Psi(\beta^1) + \frac{1}{2\sigma} \left(\|g^1\|^2 + \sum_{\tau=2}^t \frac{\|g^{\tau}\|^2}{\tau - 1} \right). \end{aligned}$$

Since $\beta^1 = \arg \min_{\beta} \Psi(\beta)$, we have $\Psi(\beta^{t+1}) \geq \Psi(\beta^1)$. Thus, adding $\Psi(\beta^1) - \Psi(\beta^{t+1}) \leq 0$ to the above inequality yields:

$$\sum_{\tau=1}^t (\langle \beta^{\tau}, g^{\tau} \rangle + \Psi(\beta^{\tau})) \leq -t \Psi^* \left(-\frac{1}{t} s^t \right) + \Psi(\beta^2) - \Psi(\beta^1) + \frac{1}{2\sigma} \left(\|g^1\|^2 + \sum_{\tau=2}^t \frac{\|g^{\tau}\|^2}{\tau - 1} \right).$$

□

Lemma 11. *We have:*

$$\Psi(\beta^2) - \Psi(\beta^1) \leq \frac{2}{\sigma} \|g^1\|^2.$$

Proof of Lemma 11. Since $\beta^2 = \arg \min_{\beta} \{\langle s^1, \beta \rangle + \Psi(\beta)\}$, we have:

$$\langle s^1, \beta^2 \rangle + \Psi(\beta^2) \leq \langle s^1, \beta^1 \rangle + \Psi(\beta^1),$$

and then:

$$\Psi(\beta^2) - \Psi(\beta^1) \leq \langle s^1, \beta^1 - \beta^2 \rangle \leq \|s^1\| \|\beta^1 - \beta^2\| = \|g^1\| \|\beta^1 - \beta^2\|.$$

On the other hand, by strong convexity of Ψ and the first-order optimality condition for β^1 , we have:

$$\Psi(\beta^2) \geq \Psi(\beta^1) + \frac{\sigma}{2} \|\beta^1 - \beta^2\|^2.$$

By combining these inequalities, we have:

$$\frac{\sigma}{2} \|\beta^1 - \beta^2\|^2 \leq \|g^1\| \|\beta^1 - \beta^2\|,$$

and then:

$$\|\beta^1 - \beta^2\| \leq \frac{2}{\sigma} \|g^1\|.$$

Therefore, we have:

$$\Psi(\beta^2) - \Psi(\beta^1) \leq \frac{2}{\sigma} \|g^1\|^2.$$

□

Lemma 12. For any $t \geq 1$ and $(j(\tau))_{\tau=1}^t$:

$$\|\beta^{\text{DA,UCB}(k)} - \beta^{*,\text{UCB}(k)}\|^2 \leq O\left(\frac{C \log T_k - R_{T_k}(\beta)}{T_k}\right). \quad (58)$$

where $C = \frac{\|v\|^2}{\sigma^2}$

Proof of Lemma 12. From lemma 8 ,for any $t \geq 1$ and $(j(\tau))_{\tau=1}^t$,

$$\frac{\sigma t}{2} \|\beta^{t+1} - \beta^*\|^2 \leq t\Psi^* \left(-\frac{1}{t} s^t\right) + \sum_{\tau=1}^t (\langle v^\tau, \beta^\tau \rangle + \Psi(\beta^\tau)) + \sum_{\tau=1}^t (f(\beta^*, j(\tau)) + \Psi(\beta^*) - f(\beta^\tau, j(\tau)) - \Psi(\beta^\tau)) \quad (59)$$

$$= t\Psi^* \left(-\frac{1}{t} s^t\right) + \sum_{\tau=1}^t (\langle v^\tau, \beta^\tau \rangle + \Psi(\beta^\tau)) - R_t(\beta). \quad (60)$$

$$R_t(\beta) = \sum_{\tau=1}^t (f(\beta^\tau, j(\tau)) - \Psi(\beta^\tau) - f(\beta^*, j(\tau)) + \Psi(\beta^*)).$$

By combining this inequality and lemma 9,

$$\frac{\sigma t}{2} \|\beta^{t+1} - \beta^*\|^2 \leq \Psi(\beta^2) - \Psi(\beta^1) + \frac{1}{2\sigma} \left(\|g^1\|^2 + \sum_{\tau=2}^t \frac{\|g^\tau\|^2}{\tau-1} \right) - R_t(\beta). \quad (61)$$

By combining this inequality and lemma 11,

$$\frac{\sigma t}{2} \|\beta^{t+1} - \beta^*\|^2 \leq \frac{2}{\sigma} \|g^1\|^2 + \frac{1}{2\sigma} \left(\|g^1\|^2 + \sum_{\tau=2}^t \frac{\|g^\tau\|^2}{\tau-1} \right) - R_t(\beta). \quad (62)$$

Thus,

$$\|\beta^{t+1} - \beta^*\|^2 \leq \frac{2}{\sigma^2 t} \left(2\|v^1\|^2 + \frac{1}{2} \left(\|v^1\|^2 + \sum_{\tau=2}^t \frac{\|v^\tau\|^2}{\tau-1} \right) \right) - \frac{2}{\sigma t} R_t(\beta). \quad (63)$$

This inequality implies that,

$$\|\beta^{t+1} - \beta^*\|^2 \leq O\left(\frac{C \log t - R_t(\beta)}{t}\right). \quad (64)$$

where $C = \frac{\|v\|^2}{\sigma^2}$.

□

Lemma 13. For any k ,

$$\|u^{\text{DA,RUCB}(k)} - u^{*,\text{RUCB}(k)}\| \leq O\left(\frac{1}{\min_i B_i} \sqrt{\frac{C \log T_k - R_{T_k}(\beta)}{T_k}}\right).$$

Proof of Lemma 13. From lemma 12,

$$\|\beta^{\text{DA,UCB}(k)} - \beta^{*,\text{UCB}(k)}\|^2 \leq O\left(\frac{C \log T_k - R_{T_k}(\beta)}{T_k}\right), \quad (65)$$

which implies

$$\|\beta^{\text{DA,UCB}(k)} - \beta^{*,\text{UCB}(k)}\| \leq O\left(\sqrt{\frac{C \log T_k - R_{T_k}(\beta)}{T_k}}\right). \quad (66)$$

By definition of β ,

$$\beta_i^{*,\text{UCB}(k)} = \frac{B_i}{u_i^{*,\text{RUCB}(k)}} \quad (67)$$

$$\beta_i^{\text{DA,UCB}(k)} = O\left(\frac{B_i}{u_i^{\text{DA,RUCB}(k)}}\right). \quad (68)$$

By combining (66) and definition of β ,

$$\left(\min_i B_i\right) \left\| \frac{1}{u^{\text{DA,RUCB}(k)}} - \frac{1}{u^{*,\text{RUCB}(k)}} \right\| \leq O\left(\sqrt{\frac{C \log T_k - R_{T_k}(\beta)}{T_k}}\right). \quad (69)$$

Thus,

$$\left(\min_i B_i\right) \left\| \frac{u^{\text{DA,RUCB}(k)} - u^{*,\text{RUCB}(k)}}{u^{\text{DA,RUCB}(k)} u^{*,\text{RUCB}(k)}} \right\| \leq O\left(\sqrt{\frac{C \log T_k - R_{T_k}(\beta)}{T_k}}\right). \quad (70)$$

By $\|v\|_\infty \leq 1$, we get $0 \leq u_i^{\text{DA,RUCB}(k)} \leq 1$, $0 \leq u_i^{*,\text{RUCB}(k)} \leq 1$.

Thus, (70) imply that,

$$\|u^{\text{DA,RUCB}(k)} - u^{*,\text{RUCB}(k)}\| \leq O\left(\frac{1}{\min_i B_i} \sqrt{\frac{C \log T_k - R_{T_k}(\beta)}{T_k}}\right) \quad (71)$$

Moreover, combining this inequality and $0 \leq u_i^{\text{DA,RUCB}(k)} u_i^{*,\text{RUCB}(k)} \leq 1$,

$$\|u^{\text{DA,RUCB}(k)} - u^{*,\text{RUCB}(k)}\| \leq O\left(\frac{1}{\min_i B_i} \sqrt{\frac{C \log T_k - R_{T_k}(\beta)}{T_k}}\right) \quad (72)$$

□

The following lemma bounds the online regret R_τ uniformly over the rounds τ .

Lemma 14. (Martingale bound for DA) *Assume that $\Delta_t(\beta) \leq C \log T$ for all T . Then, for any $a > 0$*

$$\Pr[\sup_{\tau} (-R_\tau) + C \log T \geq a] \leq \frac{C \log T}{a}. \quad (73)$$

Proof of Lemma 14. Since $R_t \leq C \log T$ (Theorem 1 (a) in Xiao) and R_t is a submartingale, $S_t = -R_t + C \log T$ is a non-negative supermartingale. Ville's inequality implies for any supermartingale

$$\Pr[\sup_{\tau} S_\tau \geq a] \leq \frac{S_0}{a}, \quad (74)$$

which is Eq. (73). \square

Lemma 15. *Let*

$$\mathcal{X}_{ka} = \left\{ \|u^{\text{DA,RUCB}(k)} - u^{*,\text{RUCB}(k)}\| \leq O\left(\frac{1}{\min_i B_i} \sqrt{\frac{a}{T_k}}\right) \right\}. \quad (75)$$

With probability at least $1 - \frac{C \log T}{a}$, event \mathcal{X}_{ka} holds.

Proof of Lemma 15. The lemma immediately follows from Lemmas 13 and 14. \square

D.4.2 Auxiliary lemmas over the multiple instances of DA

Next, we introduce the following lemma:

Lemma 16. *For any $a > 0$,*

$$\Pr\left[\bigcup_k \mathcal{X}_{ka}\right] \leq \frac{nmC(\log T)^2}{a}.$$

Proof of Lemma 16. DA-UCB-Reset resets the DA subroutine at most $nm \log_2 T$ times. Thus, from union bound of Lemma 15 over all instances of DA yields

$$\Pr\left[\bigcup_k \mathcal{X}_{ka}\right] \leq \sum_k \Pr[\mathcal{X}_{ka}] \leq \frac{nmC(\log T)^2}{a}. \quad (76)$$

\square

Next, We introduce the following lemma:

Lemma 17. *For any $a > 0$, with $\bigcap_k \mathcal{X}_{ka}^c$,*

$$\mathbb{I}\left[\bigcap_k \mathcal{X}_{ka}^c\right] |u_i^{*,\text{RUCB}} - u_i^{\text{DA,RUCB}}| \leq \tilde{O}\left(\sqrt{\frac{anm}{T(\min_i B_i)^2}}\right).$$

Proof of Lemma 17. By definition,

$$\mathbb{I}[\mathcal{X}_{ka}^c] \|u^{\text{DA,RUCB}(k)} - u^{*,\text{RUCB}(k)}\| \leq O\left(\frac{1}{\min_i B_i} \sqrt{\frac{a}{T_k}}\right), \quad (77)$$

which implies,

$$\mathbb{I}\left[\bigcap_k \mathcal{X}_{ka}^c\right] |u_i^{\text{DA,RUCB}(k)} - u_i^{*,\text{RUCB}(k)}| \leq O\left(\frac{1}{\min_i B_i} \sqrt{\frac{a}{T_k}}\right). \quad (78)$$

By definition of $u_i^{*,\text{RUCB}}$ and $u_i^{\text{DA,RUCB}}$,

$$u_i^{*,\text{RUCB}} = \frac{1}{T}(T_1 u_i^{*,\text{RUCB}(1)} + T_2 u_i^{*,\text{RUCB}(2)} + \dots + T_y u_i^{*,\text{RUCB}(y)}), \quad (79)$$

$$u_i^{\text{DA,RUCB}} = \frac{1}{T} \sum_{t=1}^T \mathbb{I}[I(t) = i] \hat{v}_{i(t),j(t),N_{i,j(t)}(s)}^{\text{UCB}} \quad (80)$$

$$= \frac{1}{T}(T_1 u_i^{\text{DA,RUCB}(1)} + T_2 u_i^{\text{DA,RUCB}(2)} + \dots + T_y u_i^{\text{DA,RUCB}(y)}). \quad (81)$$

Thus, from the triangle inequality,

$$\mathbb{I}\left[\bigcap_k \mathcal{X}_{ka}^c\right] |u_i^{*,\text{RUCB}} - u_i^{\text{DA,RUCB}}| \leq \mathbb{I}\left[\bigcap_k \mathcal{X}_{ka}^c\right] \left(\frac{T_1}{T} |u_i^{*,\text{RUCB}(1)} - u_i^{\text{DA,RUCB}(1)}| + \frac{T_2}{T} |u_i^{*,\text{RUCB}(2)} - u_i^{\text{DA,RUCB}(2)}| \right. \quad (82)$$

$$\left. + \dots + \frac{T_y}{T} |u_i^{*,\text{RUCB}(y)} - u_i^{\text{DA,RUCB}(y)}| \right). \quad (83)$$

For all $k = 1, 2, \dots, y$, from (78),

$$\mathbb{I}\left[\bigcap_k \mathcal{X}_{ka}^c\right] \frac{T_k}{T} |u_i^{*,\text{RUCB}(k)} - u_i^{\text{DA,RUCB}(k)}| \leq O\left(\frac{T_k}{T \min_i B_i} \sqrt{\frac{a}{T_k}}\right) \quad (84)$$

$$= \tilde{O}\left(\frac{\sqrt{T_k a}}{T \min_i B_i}\right). \quad (85)$$

Thus, with $\bigcap_k \mathcal{X}_{ka}^c$,

$$\mathbb{I}\left[\bigcap_k \mathcal{X}_{ka}^c\right] |u_i^{*,\text{RUCB}} - u_i^{\text{DA,RUCB}}| \leq \frac{T_1}{T} |u_i^{*,\text{RUCB}(1)} - u_i^{\text{DA,RUCB}(1)}| + \frac{T_2}{T} |u_i^{*,\text{RUCB}(2)} - u_i^{\text{DA,RUCB}(2)}| \quad (86)$$

$$+ \dots + \frac{T_k}{T} |u_i^{*,\text{RUCB}(k)} - u_i^{\text{DA,RUCB}(k)}| \quad (87)$$

$$\leq \tilde{O}\left(\frac{a^{\frac{1}{2}}(\sqrt{T_1} + \sqrt{T_2} + \dots + \sqrt{T_k})}{T \min_i B_i}\right). \quad (88)$$

Here, DA-UCB-Reset resets the DA subroutine at most $|NM| \log_2 T$ times, and from Cauchy–Schwarz inequality,

$$\sqrt{T_1} + \sqrt{T_2} + \dots + \sqrt{T_k} \leq \sqrt{(nm \log_2 T)T} \quad (89)$$

$$(90)$$

Therefore,

$$\mathbb{I}\left[\bigcap_k \mathcal{X}_{ka}^c\right] |u_i^{*,\text{RUCB}} - u_i^{\text{DA,RUCB}}| \leq \tilde{O}\left(\frac{\sqrt{a}(\sqrt{T_1} + \sqrt{T_2} + \dots + \sqrt{T_k})}{T \min_i B_i}\right) \quad (91)$$

$$\leq \tilde{O}\left(\frac{\sqrt{anmT}}{T \min_i B_i}\right) \quad (92)$$

$$\leq \tilde{O}\left(\sqrt{\frac{anm}{T(\min_i B_i)^2}}\right). \quad (93)$$

□

D.4.3 Bound on the second term

Lemma 18. Assume that $B_i = 1/n$. Let $X = T |\prod_i u_i^{*,\text{RUCB}} - \prod_i u_i^{\text{DA,RUCB}}|$. Then, for any $x > 0$, we have

$$\Pr[X \geq x] \leq \min\left(1, \frac{n^4 m^2 C^3 (\log T)^2 T}{x^2}\right). \quad (94)$$

Proof of Lemma 18. Lemmas 16 and 17 with $\min_i B_i = 1/n$ imply there exists a constant $C > 0$ such that

$$\Pr \left[X \geq C\sqrt{an^3mT} \right] \leq \frac{nmC(\log T)^2}{a}. \quad (95)$$

Substituting $a = x^2(n^3mT)^{-1}C^{-2}$, we can see that Eq. (95) is equivalent to

$$\Pr [X \geq x] \leq \frac{n^4m^2C^3(\log T)^2T}{x^2}. \quad (96)$$

Moreover, $\Pr[X \geq x] \leq 1$ by definition of probability. Combining this with Eq. (96) completes the lemma. \square

Lemma 19. (survival function)

Let X be non-negative random variable such that

$$P[X > x] \leq C \min \left(\frac{T}{x^2}, 1 \right) \quad (97)$$

for some $C > 0$. Then,

$$\mathbb{E}[X] \leq 2C\sqrt{T}. \quad (98)$$

Proof of Lemma 19. Let $S^{ub}(x)$ be such that $\Pr[X > x] \leq S^{ub}(x)$. Then,

$$E[X] = \int_0^\infty \Pr[X > x] dx \quad (\text{by def of survival function}) \quad (99)$$

$$\leq \int_0^\infty S^{ub}(x) dx \quad (100)$$

$$\leq C \int_0^\infty \min(T/x^2, 1) dx \quad (101)$$

$$\leq C \int_{\sqrt{T}}^\infty T/x^2 dx + \int_0^{\sqrt{T}} dx \quad (102)$$

$$\leq C \left([-T/x]_{\sqrt{T}}^\infty + \sqrt{T} \right) \quad (103)$$

$$\leq C \left(\sqrt{T} + \sqrt{T} \right), \quad (104)$$

which completes the proof of Lemma 19. \square

Lemma 20. We have

$$\mathbb{E} \left[T \left| \prod_i (u_i^{*,\text{RUCB}})^{B_i} - \prod_i (u_i^{\text{DA},\text{RUCB}})^{B_i} \right| \right] = \tilde{O} \left(n^4 m^2 \sqrt{T} \right). \quad (105)$$

Proof of Lemma 20. Combining Lemmas 18 and 19 with $C = n^4 m^2 C^3 (\log T)^2$ yields Eq. (105). \square

D.5 Lemmas on the third term

Lemma 21. for any agent $i \in N$, with $\mathcal{G}_1 \cap \mathcal{G}_2$ from lemma 5,

$$\mathbb{I}[\mathcal{G}_1 \cap \mathcal{G}_2] \left| \left(\prod_i (u_i^{\text{DA},\text{RUCB}})^{B_i} - \prod_i (u_i^{\text{DA},\text{true}})^{B_i} \right) \right| = \tilde{O} \left(\frac{nm\sqrt{T}}{T} \right). \quad (106)$$

Proof of Lemma 21. The proof is immediately derived from the definition of \mathcal{G}_2 . \square

E PROOFS ON DA

Let $v_i = (v_{i,1}, v_{i,2}, \dots, v_{i,m})$ and $s = (s_1, s_2, \dots, s_m)$ be corresponding vectors of size m and $\langle v_i, s \rangle$ be an inner product of them. Let \mathbf{e}^i be the unit vector of the i -th coordinate. Let $\beta = (\beta_1, \beta_2, \dots, \beta_n)$ and its value at the beginning of t be β^t . Let $B = (B_1, B_2, \dots, B_n)$ and $\mathbf{1} = (1, 1, \dots, 1)$.

E.1 Proof of Lemma 1

Proof of Lemma 1. Let us consider the following event that no projection occurs when updating β_i^{t+1} for buyer $i \in N$ and $t \geq 1$:

$$\mathcal{V}_i^t := \left\{ \frac{B_i}{\bar{u}_i^{\text{DA}}(t)} \in \left[\frac{B_i}{h(1+\delta_0)}, \frac{1+\delta_0}{l} \right] \right\}.$$

From Lemma 22, whenever the complementary event $(\mathcal{V}_i^t)^c = \left\{ \frac{B_i}{\bar{u}_i^{\text{DA}}(t)} \notin \left[\frac{B_i}{h(1+\delta_0)}, \frac{1+\delta_0}{l} \right] \right\}$ occurs, it holds that:

$$|\beta_i^{t+1} - \beta_i^{*,\text{DA}}| > \min \left(\frac{1+\delta_0}{l} - \beta_i^{*,\text{DA}}, \beta_i^{*,\text{DA}} - \frac{B_i}{h(1+\delta_0)} \right) > 0.$$

Therefore, we have:

$$\begin{aligned} \mathbb{E} \left[(\beta_i^{t+1} - \beta_i^{*,\text{DA}})^2 \right] &= \mathbb{E} \left[(\beta_i^{t+1} - \beta_i^{*,\text{DA}})^2 \mid \mathcal{V}_i^t \right] \mathbb{P}(\mathcal{V}_i^t) + \mathbb{E} \left[(\beta_i^{t+1} - \beta_i^{*,\text{DA}})^2 \mid (\mathcal{V}_i^t)^c \right] \mathbb{P}((\mathcal{V}_i^t)^c) \\ &\geq \mathbb{E} \left[(\beta_i^{t+1} - \beta_i^{*,\text{DA}})^2 \mid (\mathcal{V}_i^t)^c \right] \mathbb{P}((\mathcal{V}_i^t)^c) \\ &\geq \left(\min \left(\frac{1+\delta_0}{l} - \beta_i^{*,\text{DA}}, \beta_i^{*,\text{DA}} - \frac{B_i}{h(1+\delta_0)} \right) \right)^2 \mathbb{P}((\mathcal{V}_i^t)^c). \end{aligned}$$

Then, we get:

$$\mathbb{P}((\mathcal{V}_i^t)^c) \leq \frac{1}{\left(\min \left(\frac{1+\delta_0}{l} - \beta_i^{*,\text{DA}}, \beta_i^{*,\text{DA}} - \frac{B_i}{h(1+\delta_0)} \right) \right)^2} \mathbb{E} \left[(\beta_i^{t+1} - \beta_i^{*,\text{DA}})^2 \right]. \quad (107)$$

On the other hand, conditioning on \mathcal{V}_i^t , we have

$$\bar{u}_i^{\text{DA}}(t) = \frac{B_i}{\beta_i^{t+1}}. \quad (108)$$

Moreover, from Lemma 22 and the assumption that $l \leq \langle v_i^{\text{DA}}, s \rangle$, we get:

$$u_i^{*,\text{DA}} \leq h = \frac{h}{l} l \leq \frac{h}{l} \langle v_i^{\text{DA}}, s \rangle \leq \frac{h}{l} \|v_i^{\text{DA}}\|_\infty \|s\|_1 = \frac{h}{l} \|v_i^{\text{DA}}\|_\infty, \quad (109)$$

where the last inequality follows from Hölder's inequality.

By combining (107), (108), and (109), we have:

$$\begin{aligned}
 & \mathbb{E} \left[(\bar{u}_i^{\text{DA}}(T) - u_i^{*,\text{DA}})^2 \right] \\
 &= \mathbb{E} \left[\mathbb{I}\{(\mathcal{V}_i^T)^c\} (\bar{u}_i^{\text{DA}}(T) - u_i^{*,\text{DA}})^2 \right] + \mathbb{E} \left[\mathbb{I}\{\mathcal{V}_i^T\} (\bar{u}_i^{\text{DA}}(T) - u_i^{*,\text{DA}})^2 \right] \\
 &= \mathbb{E} \left[\mathbb{I}\{(\mathcal{V}_i^T)^c\} (\bar{u}_i^{\text{DA}}(T) - u_i^{*,\text{DA}})^2 \right] + \mathbb{E} \left[\mathbb{I}\{\mathcal{V}_i^T\} \left(\frac{B_i}{\beta_i^{T+1}} - u_i^{*,\text{DA}} \right)^2 \right] \\
 &\leq \mathbb{E} \left[\mathbb{I}\{(\mathcal{V}_i^T)^c\} \left(\max(u_i^{*,\text{DA}}, \|v_i^{\text{DA}}\|_\infty) \right)^2 \right] + \mathbb{E} \left[\mathbb{I}\{\mathcal{V}_i^T\} \left(\frac{B_i}{\beta_i^{T+1}} - u_i^{*,\text{DA}} \right)^2 \right] \\
 &\leq \frac{h}{l} \|v_i^{\text{DA}}\|_\infty^2 \mathbb{E} \left[\mathbb{I}\{(\mathcal{V}_i^T)^c\} \right] + (u_i^{*,\text{DA}})^2 \mathbb{E} \left[\mathbb{I}\{\mathcal{V}_i^T\} \left(\frac{B_i}{\beta_i^{T+1} u_i^{*,\text{DA}}} - 1 \right)^2 \right] \\
 &= \frac{h}{l} \|v_i^{\text{DA}}\|_\infty^2 \mathbb{P} \left[(\mathcal{V}_i^T)^c \right] + (u_i^{*,\text{DA}})^2 \mathbb{E} \left[\mathbb{I}\{\mathcal{V}_i^T\} \left(\frac{\beta_i^{*,\text{DA}}}{\beta_i^{T+1}} - 1 \right)^2 \right] \\
 &\leq \frac{h \|v_i^{\text{DA}}\|_\infty^2}{l \left(\min \left(\frac{1+\delta_0}{l} - \beta_i^{*,\text{DA}}, \beta_i^{*,\text{DA}} - \frac{B_i}{h(1+\delta_0)} \right) \right)^2} \mathbb{E} \left[(\beta_i^{T+1} - \beta_i^{*,\text{DA}})^2 \right] + \left(\frac{h(1+\delta_0)u_i^{*,\text{DA}}}{B_i} \right)^2 \mathbb{E} \left[(\beta_i^{T+1} - \beta_i^{*,\text{DA}})^2 \right],
 \end{aligned}$$

where the first inequality follow from $0 \leq \bar{u}_i^{\text{DA}} \leq \|v_i^{\text{DA}}\|_\infty$, and the third equality follows from $\beta_i^{*,\text{DA}} = \frac{B_i}{u_i^{*,\text{DA}}}$ by Theorem 1 in Gao et al. (2021). Since $\frac{B_i}{h} \leq \beta_i^{*,\text{DA}} \leq \frac{1}{l}$ from Lemma 22,

$$\begin{aligned}
 \mathbb{E} \left[(\bar{u}_i^{\text{DA}}(T) - u_i^{*,\text{DA}})^2 \right] &\leq \left(\frac{h \|v_i^{\text{DA}}\|_\infty^2}{l \left(\min \left(\frac{\delta_0}{l}, \frac{B_i \delta_0}{h(1+\delta_0)} \right) \right)^2} + \left(\frac{h^2(1+\delta_0)}{B_i} \right)^2 \right) \mathbb{E} \left[(\beta_i^{T+1} - \beta_i^{*,\text{DA}})^2 \right] \\
 &\leq \left(\left(\frac{(1+\delta_0)}{\delta_0} \right)^2 \frac{h^3 \|v_i^{\text{DA}}\|_\infty^2}{l B_i^2} + \left(\frac{h^2(1+\delta_0)}{B_i} \right)^2 \right) \mathbb{E} \left[(\beta_i^{T+1} - \beta_i^{*,\text{DA}})^2 \right] \\
 &\leq \left(\left(\frac{(1+\delta_0)}{\delta_0} \right)^2 \frac{h^3 \|v_i^{\text{DA}}\|_\infty^2}{l (\min_{i \in N} B_i)^2} + \left(\frac{h^2(1+\delta_0)}{\min_{i \in N} B_i} \right)^2 \right) \mathbb{E} \left[(\beta_i^{T+1} - \beta_i^{*,\text{DA}})^2 \right],
 \end{aligned}$$

where the second inequality follows from $\frac{\delta_0}{l} \geq \frac{\delta_0}{h} \geq \frac{\delta_0}{h(1+\delta_0)} \geq \frac{B_i \delta_0}{h(1+\delta_0)}$. From Lemma 24, by summing up the above inequality for $i \in N$, we have:

$$\begin{aligned}
 \mathbb{E} \left[\|\bar{u}^{\text{DA}}(T) - u^{*,\text{DA}}\|^2 \right] &\leq \sum_{i \in N} \left(\left(\frac{(1+\delta_0)}{\delta_0} \right)^2 \frac{h^3 \|v_i^{\text{DA}}\|_\infty^2}{l (\min_{i \in N} B_i)^2} + \left(\frac{h^2(1+\delta_0)}{\min_{i \in N} B_i} \right)^2 \right) \mathbb{E} \left[(\beta_i^{T+1} - \beta_i^{*,\text{DA}})^2 \right] \\
 &\leq \left(\left(\frac{(1+\delta_0)}{\delta_0} \right)^2 \frac{h^3 \|v^{\text{DA}}\|_\infty^2}{l (\min_{i \in N} B_i)^2} + \left(\frac{h^2(1+\delta_0)}{\min_{i \in N} B_i} \right)^2 \right) \mathbb{E} \left[\|\beta^{T+1} - \beta^{*,\text{DA}}\|^2 \right] \\
 &\leq \left(\left(\frac{(1+\delta_0)}{\delta_0} \right)^2 \frac{h^3 \|v^{\text{DA}}\|_\infty^2}{l (\min_{i \in N} B_i)^2} + \left(\frac{h^2(1+\delta_0)}{\min_{i \in N} B_i} \right)^2 \right) \frac{G^2}{\mu^2 T} (6 + \log T),
 \end{aligned}$$

where $G = \|v^{\text{DA}}\|_\infty$ and $\mu = \frac{l^2 \min_{i \in N} B_i}{(1+\delta_0)^2}$. □

F ADDITIONAL LEMMAS ON DA

Lemma 22. Assume that the values of items are given by the deterministic values $\{v_{i,j}^{\text{DA}}\}_{i \in N, j \in M}$. Furthermore, assume that the parameters l, h in DA satisfies that $l \leq \sum_{j \in M} s_j v_{i,j}^{\text{DA}} \leq h$ for all $i \in N$. Then, the equilibrium utilities satisfy $B_i l \leq u_i^{*,\text{DA}} \leq h$ and hence $\frac{B_i}{h} \leq \beta_i^{*,\text{DA}} \leq \frac{1}{l}$.

Lemma 23. $\Psi(\beta) := -\sum_{i \in N} B_i \log \beta_i$ is μ -strongly convex on $[\frac{B}{h(1+\delta_0)}, \frac{1+\delta_0}{l} \mathbf{1}]$ with $\mu = \frac{l^2 \min_{i \in N} B_i}{(1+\delta_0)^2}$.

Lemma 24. Assume that the values of items $\{v_{i,j}^{\text{DA}}\}$ are deterministic. Furthermore, assume that $l, h > 0$ satisfies that $l \leq \sum_{j \in M} s_j v_{i,j}^{\text{DA}} \leq h$ for all $i \in N$, and $\beta^1 = \arg \min_{[\frac{B}{h(1+\delta_0)}, \frac{1+\delta_0}{l} \mathbf{1}]} \Psi(\beta)$. Let $\beta^{*,\text{DA}}$ be the optimal solution of (112). Then, the following inequality holds for all $t \geq 1$:

$$\mathbb{E} [\|\beta^{t+1} - \beta^{*,\text{DA}}\|^2] \leq \frac{G^2}{\mu^2 t} (6 + \log t),$$

where $G = \|v^{\text{DA}}\|_\infty := \max_{i \in N} \|v_i^{\text{DA}}\|_\infty$ and $\mu = \frac{l^2 \min_{i \in N} B_i}{(1+\delta_0)^2}$.

F.1 Proof of Lemma 22

Proof of Lemma 22. For buyer $i \in N$, the largest utility is attainable when the entire set of items is allocated to i (given by the supply s). Thus, from the assumption that $\sum_{j \in M} s_j v_{i,j}^{\text{DA}} \leq h$, we have:

$$u_i^{*,\text{DA}} \leq \langle v_i^{\text{DA}}, s \rangle \leq h. \quad (110)$$

On the other hand, from Theorem 1 in Gao et al. (2021), we have for any market equilibrium $(x^{*,\text{DA}}, p^{*,\text{DA}})$:

$$\langle p^{*,\text{DA}}, x_i^{*,\text{DA}} \rangle = \beta_i^{*,\text{DA}} \langle v_i, x_i^{*,\text{DA}} \rangle = B_i.$$

By the assumption that $\sum_{i \in N} B_i = 1$ and market clearance $\langle p^{*,\text{DA}}, s - \sum_{i \in N} x_i^{*,\text{DA}} \rangle = 0$, we get:

$$\langle p^{*,\text{DA}}, s \rangle = \sum_{i \in N} \langle p^{*,\text{DA}}, x_i^{*,\text{DA}} \rangle = \sum_{i \in N} B_i = 1.$$

Thus,

$$\langle p^{*,\text{DA}}, B_i s \rangle = B_i.$$

This means that each buyer i can afford the proportional allocation $x_i^\circ := B_i s$ under the item price $p^{*,\text{DA}}$. Therefore, from the buyer optimality of the market equilibrium and the assumption that $l \leq \sum_{j \in M} s_j v_{i,j}^{\text{DA}}$, we have:

$$u_i^{*,\text{DA}} \geq \langle v_i^{\text{DA}}, x_i^\circ \rangle = B_i \langle v_i^{\text{DA}}, s \rangle \geq B_i l. \quad (111)$$

By combining (110) and (111), we get:

$$B_i l \leq u_i^{*,\text{DA}} \leq h.$$

Moreover, since $\beta_i^{*,\text{DA}} = \frac{B_i}{u_i^{*,\text{DA}}}$ from Theorem 1 in Gao et al. (2021), we have:

$$\frac{B_i}{h} \leq \beta_i^{*,\text{DA}} \leq \frac{1}{l}.$$

□

F.2 Proof of Lemma 23

Proof of Lemma 23. The Hessian matrix of Ψ at $\beta \in [\frac{B}{h(1+\delta_0)}, \frac{1+\delta_0}{l} \mathbf{1}]$ is given by:

$$\nabla^2 \Psi(\beta) = \begin{pmatrix} \frac{B_1}{\beta_1^2} & & \\ & \ddots & \\ & & \frac{B_n}{\beta_n^2} \end{pmatrix}.$$

Thus, the minimum eigenvalue of $\nabla^2 \Psi(\beta)$ is lower bounded as:

$$\lambda_{\min}(\nabla^2 \Psi(\beta)) \geq \min_{i \in N} \frac{B_i}{\beta_i^2} \geq \frac{l^2 \min_{i \in N} B_i}{(1+\delta_0)^2}.$$

Therefore, Ψ is μ -strongly convex on $[\frac{B}{h(1+\delta_0)}, \frac{1+\delta_0}{l} \mathbf{1}]$ with $\mu = \frac{l^2 \min_{i \in N} B_i}{(1+\delta_0)^2}$. □

F.3 Proof Lemma 24

Proof of Lemma 24. Let us assume that $l \leq \langle v_i, s \rangle \leq h$ for $i \in N$. Then, D_{EG} can be written as:

$$\beta \in \left[\frac{B}{h(1+\delta_0)}, \frac{1+\delta_0}{l} \mathbf{1} \right] \left\{ \phi(\beta) := \mathbb{E}_{j \sim s} \left[\max_{i \in N} v_{i,j} \beta_i \right] - \sum_{i \in N} B_i \log \beta_i \right\}. \quad (112)$$

Defining $f(\beta, j) = \max_{i \in N} v_{i,j} \beta_i$ and $\Psi(\beta) = -\sum_{i \in N} B_i \log \beta_i$, we have $v_{i_j^*, j} \mathbf{e}^{(i_j^*)} \in \partial f(\beta, j)$, where $i_j^* \in \arg \max_{i \in N} v_{i,j} \beta_i$. The update rule of DA is given as:

$$\begin{aligned} \beta^{t+1} &= \left(\text{Proj}_{\left[\frac{B}{h(1+\delta_0)}, \frac{1+\delta_0}{l} \mathbf{1} \right]} \left(\frac{B_i}{\frac{1}{t} \sum_{\tau=1}^t v_{i, j(\tau)} \mathbb{I}\{i = i_{j(\tau)}^*\}} \right) \right)_{i \in N}, \\ &= \arg \min_{\beta \in \left[\frac{B}{h(1+\delta_0)}, \frac{1+\delta_0}{l} \mathbf{1} \right]} \left\{ \frac{1}{t} \sum_{\tau=1}^t \langle v_{i_{j(\tau)}^*, j(\tau)} \mathbf{e}^{(i_{j(\tau)}^*)}, \beta \rangle - \sum_{i \in N} B_i \log \beta_i \right\} \\ &= \arg \min_{\beta \in \left[\frac{B}{h(1+\delta_0)}, \frac{1+\delta_0}{l} \mathbf{1} \right]} \left\{ \frac{1}{t} \sum_{\tau=1}^t \langle g_\tau, \beta \rangle + \Psi(\beta) \right\}, \end{aligned} \quad (113)$$

where $g_\tau = v_{i_{j(\tau)}^*, j(\tau)} \mathbf{e}^{(i_{j(\tau)}^*)}$.

From the first-order optimality condition for (113), there exist a subgradient $b_{t+1} \in \partial \Psi(\beta^{t+1})$ such that

$$\left\langle \sum_{\tau=1}^t g_\tau + t b_{t+1}, \beta^{*, \text{DA}} - \beta^{t+1} \right\rangle \geq 0, \quad \forall \beta \in \left[\frac{B}{h(1+\delta_0)}, \frac{1+\delta_0}{l} \mathbf{1} \right]. \quad (114)$$

From Lemma 23, Ψ is μ -strongly convex with $\mu = \frac{l^2 \min_{i \in N} B_i}{(1+\delta_0)^2}$, and then we have for any $\beta \in \left[\frac{B}{h(1+\delta_0)}, \frac{1+\delta_0}{l} \mathbf{1} \right]$,

$$\Psi(\beta) \geq \Psi(\beta^{t+1}) + \langle b_{t+1}, \beta - \beta^{t+1} \rangle + \frac{\mu}{2} \|\beta^{t+1} - \beta\|^2. \quad (115)$$

By combining (114) and (115), we get:

$$\begin{aligned} & \frac{\mu t}{2} \|\beta^{t+1} - \beta^{*, \text{DA}}\|^2 \\ & \leq \left\langle \sum_{\tau=1}^t g_\tau, \beta^{*, \text{DA}} - \beta^{t+1} \right\rangle + t \Psi(\beta^{*, \text{DA}}) - t \Psi(\beta^{t+1}) \\ & = \left\langle -\sum_{\tau=1}^t g_\tau, \beta^{t+1} - \beta^1 \right\rangle - t \Psi(\beta^{t+1}) + t \Psi(\beta^{*, \text{DA}}) + \left\langle \sum_{\tau=1}^t g_\tau, \beta^{*, \text{DA}} - \beta^1 \right\rangle \\ & = \left\langle -\sum_{\tau=1}^t g_\tau, \beta^{t+1} - \beta^1 \right\rangle - t \Psi(\beta^{t+1}) + t \Psi(\beta^{*, \text{DA}}) + \sum_{\tau=1}^t \langle g_\tau, \beta^\tau - \beta^1 \rangle + \sum_{\tau=1}^t \langle g_\tau, \beta^{*, \text{DA}} - \beta^\tau \rangle \\ & = \left\langle -\sum_{\tau=1}^t g_\tau, \beta^{t+1} - \beta^1 \right\rangle - t \Psi(\beta^{t+1}) + \sum_{\tau=1}^t (\langle g_\tau, \beta^\tau - \beta^1 \rangle + \Psi(\beta^\tau)) + \sum_{\tau=1}^t \langle g_\tau, \beta^{*, \text{DA}} - \beta^\tau \rangle + t \Psi(\beta^{*, \text{DA}}) - \sum_{\tau=1}^t \Psi(\beta^\tau) \\ & \leq \left\langle -\sum_{\tau=1}^t g_\tau, \beta^{t+1} - \beta^1 \right\rangle - t \Psi(\beta^{t+1}) + \sum_{\tau=1}^t (\langle g_\tau, \beta^\tau - \beta^1 \rangle + \Psi(\beta^\tau)) \\ & + \sum_{\tau=1}^t (f(\beta^{*, \text{DA}}, j(\tau)) - f(\beta^\tau, j(\tau))) + t \Psi(\beta^{*, \text{DA}}) - \sum_{\tau=1}^t \Psi(\beta^\tau) \\ & = \left\langle -\sum_{\tau=1}^t g_\tau, \beta^{t+1} - \beta^1 \right\rangle - t \Psi(\beta^{t+1}) + \sum_{\tau=1}^t (\langle g_\tau, \beta^\tau - \beta^1 \rangle + \Psi(\beta^\tau)) \\ & + \sum_{\tau=1}^t (f(\beta^{*, \text{DA}}, j(\tau)) + \Psi(\beta^{*, \text{DA}})) - \sum_{\tau=1}^t (f(\beta^\tau, j(\tau)) + \Psi(\beta^\tau)), \end{aligned}$$

where the second inequality follows from the convexity of $f(\cdot, j(\tau))$. Here, let us define

$$V_0(s) = \max_{\beta \in \left[\frac{B}{h(1+\delta_0)}, \frac{1+\delta_0}{t} \mathbf{1} \right]} \{ \langle s, \beta - \beta^1 \rangle - \Psi(\beta) \},$$

and for $t \geq 1$,

$$V_t(s) = \max_{\beta \in \left[\frac{B}{h(1+\delta_0)}, \frac{1+\delta_0}{t} \mathbf{1} \right]} \{ \langle s, \beta - \beta^1 \rangle - t\Psi(\beta) \}.$$

Since $\beta^{t+1} = \arg \max_{\beta \in \left[\frac{B}{h(1+\delta_0)}, \frac{1+\delta_0}{t} \mathbf{1} \right]} \{ \langle -\sum_{\tau=1}^t g_\tau, \beta - \beta^1 \rangle - t\Psi(\beta) \}$,

$$\begin{aligned} & \frac{\mu t}{2} \|\beta^{t+1} - \beta^{*,\text{DA}}\|^2 \\ & \leq V_t\left(-\sum_{\tau=1}^t g_\tau\right) + \sum_{\tau=1}^t (\langle g_\tau, \beta^\tau - \beta^1 \rangle + \Psi(\beta^\tau)) + \sum_{\tau=1}^t (f(\beta^{*,\text{DA}}, j(\tau)) + \Psi(\beta^{*,\text{DA}})) - \sum_{\tau=1}^t (f(\beta^\tau, j(\tau)) + \Psi(\beta^\tau)). \end{aligned} \quad (116)$$

For $\tau \geq 2$, we have:

$$\begin{aligned} V_\tau\left(-\sum_{s=1}^{\tau} g_s\right) - V_{\tau-1}\left(-\sum_{s=1}^{\tau} g_s\right) & \leq \langle -\sum_{s=1}^{\tau} g_s, \beta^{\tau+1} - \beta^1 \rangle - \tau\Psi(\beta^{\tau+1}) - \langle -\sum_{s=1}^{\tau} g_s, \beta^{\tau+1} - \beta^1 \rangle + (\tau-1)\Psi(\beta^{\tau+1}) \\ & = -\Psi(\beta^{\tau+1}). \end{aligned}$$

Therefore,

$$V_\tau\left(-\sum_{s=1}^{\tau} g_s\right) + \Psi(\beta^{\tau+1}) \leq V_{\tau-1}\left(-\sum_{s=1}^{\tau} g_s\right) = V_{\tau-1}\left(-\sum_{s=1}^{\tau-1} g_s - g_\tau\right). \quad (117)$$

Here, since $t\Psi(\cdot)$ is μt -strongly convex, the convex conjugate of $t\Psi(\cdot)$ is $\frac{1}{\mu t}$ -smooth. Thus,

$$V_{\tau-1}\left(-\sum_{s=1}^{\tau-1} g_s - g_\tau\right) \leq V_{\tau-1}\left(-\sum_{s=1}^{\tau-1} g_s\right) - \langle \nabla V_{\tau-1}\left(-\sum_{s=1}^{\tau-1} g_s\right), g_\tau \rangle + \frac{\|g_\tau\|^2}{2\mu(\tau-1)}. \quad (118)$$

On the other hand, the gradient of V_τ is given as:

$$\begin{aligned} \nabla V_{\tau-1}\left(-\sum_{s=1}^{\tau-1} g_s\right) & = \nabla \max_{\beta \in \left[\frac{B}{h(1+\delta_0)}, \frac{1+\delta_0}{t} \mathbf{1} \right]} \left\{ \langle -\sum_{s=1}^{\tau-1} g_s, \beta \rangle - (\tau-1)\Psi(\beta) \right\} - \beta^1 \\ & = \beta^\tau - \beta^1. \end{aligned} \quad (119)$$

By combining (117), (118), and (119), we have for $\tau \geq 2$:

$$\langle \beta^\tau - \beta^1, g_\tau \rangle + \Psi(\beta^{\tau+1}) \leq V_{\tau-1}\left(-\sum_{s=1}^{\tau-1} g_s\right) - V_\tau\left(-\sum_{s=1}^{\tau} g_s\right) + \frac{\|g_\tau\|^2}{2\mu(\tau-1)}. \quad (120)$$

Similarly, for $\tau = 1$, we have:

$$V_1(-g_1) = V_0(-g_1) \leq V_0(0) - \langle \nabla V_0(0), g_1 \rangle + \frac{\|g_1\|^2}{2\mu},$$

and

$$\begin{aligned} \nabla V_0(0) & = \nabla \max_{\beta \in \left[\frac{B}{h(1+\delta_0)}, \frac{1+\delta_0}{t} \mathbf{1} \right]} \{ \langle 0, \beta \rangle - \Psi(\beta) \} - \beta^1 \\ & = \nabla \max_{\beta \in \left[\frac{B}{h(1+\delta_0)}, \frac{1+\delta_0}{t} \mathbf{1} \right]} \{ -\Psi(\beta) \} - \beta^1 \\ & = \beta^1 - \beta^1 = 0. \end{aligned}$$

Thus,

$$\langle \beta^1 - \beta^1, g_1 \rangle + \Psi(\beta^2) \leq V_0(0) - V_1(-g_1) + \frac{\|g_1\|^2}{2\mu} + \Psi(\beta^2). \quad (121)$$

By summing (120) and (121) for $\tau = 1, \dots, t$,

$$\begin{aligned} \sum_{\tau=1}^t (\langle \beta^\tau - \beta^1, g_\tau \rangle + \Psi(\beta^{\tau+1})) &\leq V_0(0) - V_t(-\sum_{\tau=1}^t g_\tau) + \Psi(\beta^2) + \frac{1}{2\mu} \left(\|g_1\|^2 + \sum_{\tau=2}^t \frac{\|g_\tau\|^2}{\tau-1} \right) \\ &= -V_t(-\sum_{\tau=1}^t g_\tau) + \Psi(\beta^2) - \Psi(\beta^1) + \frac{1}{2\mu} \left(\|g_1\|^2 + \sum_{\tau=2}^t \frac{\|g_\tau\|^2}{\tau-1} \right). \end{aligned}$$

Since $\beta^1 = \arg \min_{\beta \in \left[\frac{B}{h(1+\delta_0)}, \frac{1+\delta_0}{l} \mathbf{1} \right]} \Psi(\beta)$, $\Psi(\beta^{t+1}) \geq \Psi(\beta^1)$ holds. Thus, adding the inequality $\Psi(\beta^1) - \Psi(\beta^{t+1}) \leq 0$ to the above inequality, we get:

$$\begin{aligned} \sum_{\tau=1}^t (\langle \beta^\tau - \beta^1, g_\tau \rangle + \Psi(\beta^\tau)) &= \sum_{\tau=1}^t (\langle \beta^\tau - \beta^1, g_\tau \rangle + \Psi(\beta^{\tau+1})) + \Psi(\beta^1) - \Psi(\beta^{t+1}) \\ &\leq -V_t(-\sum_{\tau=1}^t g_\tau) + \Psi(\beta^2) - \Psi(\beta^1) + \frac{1}{2\mu} \left(\|g_1\|^2 + \sum_{\tau=2}^t \frac{\|g_\tau\|^2}{\tau-1} \right). \end{aligned} \quad (122)$$

Adding (116) and (122) yields:

$$\begin{aligned} &\frac{\mu t}{2} \|\beta^{t+1} - \beta^{*,\text{DA}}\|^2 \\ &\leq \Psi(\beta^2) - \Psi(\beta^1) + \frac{1}{2\mu} \left(\|g_1\|^2 + \sum_{\tau=2}^t \frac{\|g_\tau\|^2}{\tau-1} \right) + \sum_{\tau=1}^t (f(\beta^{*,\text{DA}}, j(\tau)) + \Psi(\beta^{*,\text{DA}})) - \sum_{\tau=1}^t (f(\beta^\tau, j(\tau)) + \Psi(\beta^\tau)). \end{aligned} \quad (123)$$

Since $\beta^2 = \arg \min_{\beta \in \left[\frac{B}{h(1+\delta_0)}, \frac{1+\delta_0}{l} \mathbf{1} \right]} \{\langle g_1, \beta \rangle + \Psi(\beta)\}$, we have:

$$\langle g_1, \beta^2 \rangle + \Psi(\beta^2) \leq \langle g_1, \beta^1 \rangle + \Psi(\beta^1).$$

Thus, from the Cauchy–Schwarz inequality,

$$\Psi(\beta^2) - \Psi(\beta^1) \leq \langle g_1, \beta^1 - \beta^2 \rangle \leq \|g_1\| \|\beta^1 - \beta^2\|.$$

On the other hand, by strong convexity of Ψ and the first-order optimality condition for β^1 , we have:

$$\Psi(\beta^2) \geq \Psi(\beta^1) + \frac{\mu}{2} \|\beta^1 - \beta^2\|^2.$$

By combining the above inequalities, we have:

$$\frac{\mu}{2} \|\beta^1 - \beta^2\|^2 \leq \|g_1\| \|\beta^1 - \beta^2\|,$$

and then:

$$\|\beta^1 - \beta^2\| \leq \frac{2}{\mu} \|g_1\|.$$

Therefore, we have:

$$\Psi(\beta^2) - \Psi(\beta^1) \leq \frac{2\|g_1\|^2}{\mu}. \quad (124)$$

By combining (123) and (124), we get:

$$\begin{aligned} \|\beta^{t+1} - \beta^{*,\text{DA}}\|^2 &\leq \frac{2}{\mu t} \left(\frac{2\|g_1\|^2}{\mu} + \frac{1}{2\mu} \left(\|g_1\|^2 + \sum_{\tau=2}^t \frac{\|g_\tau\|^2}{\tau-1} \right) \right) \\ &\quad + \frac{2}{\mu t} \left(\sum_{\tau=1}^t (f(\beta^{*,\text{DA}}, j(\tau)) + \Psi(\beta^{*,\text{DA}})) - \sum_{\tau=1}^t (f(\beta^\tau, j(\tau)) + \Psi(\beta^\tau)) \right). \end{aligned} \quad (125)$$

Since the random variables $(j(1), \dots, j(t))$ is i.i.d, the expectation of $f(\beta^\tau, j(\tau)) + \Psi(\beta^\tau)$ is given as:

$$\begin{aligned} \mathbb{E}_{(j(s))_{s=1}^t} [f(\beta^\tau, j(\tau)) + \Psi(\beta^\tau)] &= \mathbb{E}_{(j(s))_{s=1}^{\tau-1}} [\mathbb{E}_{(j(s))_{s=\tau}^t} [f(\beta^\tau, j(\tau)) + \Psi(\beta^\tau) \mid (j(s))_{s=1}^{\tau-1}]] \\ &= \mathbb{E}_{(j(s))_{s=1}^{\tau-1}} [\mathbb{E}_{j(\tau)} [f(\beta^\tau, j(\tau)) + \Psi(\beta^\tau)]] \\ &= \mathbb{E}_{(j(s))_{s=1}^{\tau-1}} [\phi(\beta^\tau)]. \end{aligned}$$

Similarly, the expectation of $f(\beta^{*,\text{DA}}, j(\tau)) + \Psi(\beta^{*,\text{DA}})$ is given by:

$$\mathbb{E}_{(j(s))_{s=1}^t} [f(\beta^{*,\text{DA}}, j(\tau)) + \Psi(\beta^{*,\text{DA}})] = \mathbb{E}_{j(\tau)} [f(\beta^{*,\text{DA}}, j(\tau)) + \Psi(\beta^{*,\text{DA}})] = \phi(\beta^{*,\text{DA}}).$$

Thus, the expected regret is upper bounded as:

$$\mathbb{E}_{(j(s))_{s=1}^t} \left[\sum_{\tau=1}^t (f(\beta^{*,\text{DA}}, j(\tau)) + \Psi(\beta^{*,\text{DA}})) - \sum_{\tau=1}^t (f(\beta^\tau, j(\tau)) + \Psi(\beta^\tau)) \right] = \sum_{\tau=1}^t \left(\phi(\beta^{*,\text{DA}}) - \mathbb{E}_{(j(s))_{s=1}^{\tau-1}} [\phi(\beta^\tau)] \right) \leq 0. \quad (126)$$

From (126), the expectation of (125) is given by:

$$\begin{aligned} &\mathbb{E}_{(j(s))_{s=1}^t} [\|\beta^{t+1} - \beta^{*,\text{DA}}\|^2] \\ &\leq \frac{2}{\mu t} \mathbb{E}_{(j(s))_{s=1}^t} \left[\frac{2\|g_1\|^2}{\mu} + \frac{1}{2\mu} \left(\|g_1\|^2 + \sum_{\tau=2}^t \frac{\|g_\tau\|^2}{\tau-1} \right) \right] \\ &= \frac{2}{\mu t} \left(\frac{2\mathbb{E}_{j(1)} [\|v_{i_{j(1)}, j(1)}^{\text{DA}} \mathbf{e}^{(i_{j(1)})^*}\|^2]}{\mu} + \frac{1}{2\mu} \left(\mathbb{E}_{j(1)} [\|v_{i_{j(1)}, j(1)}^{\text{DA}} \mathbf{e}^{(i_{j(1)})^*}\|^2] + \sum_{\tau=2}^t \frac{\mathbb{E}_{(j(s))_{s=1}^{\tau-1}} [\|v_{i_{j(\tau)}, j(\tau)}^{\text{DA}} \mathbf{e}^{(i_{j(\tau)})^*}\|^2]}{\tau-1} \right) \right) \\ &= \frac{2}{\mu t} \left(\frac{2 \max_{i \in N} \mathbb{E}_{j \sim s} [(v_{i,j}^{\text{DA}})^2]}{\mu} + \frac{1}{2\mu} \left(\max_{i \in N} \mathbb{E}_{j \sim s} [(v_{i,j}^{\text{DA}})^2] + \sum_{\tau=2}^t \frac{\max_{i \in N} \mathbb{E}_{j \sim s} [(v_{i,j}^{\text{DA}})^2]}{\tau-1} \right) \right) \\ &\leq \frac{G^2}{\mu^2 t} \left(5 + \sum_{\tau=2}^t \frac{1}{\tau-1} \right) = \frac{G^2}{\mu^2 t} \left(5 + \sum_{\tau=1}^{t-1} \frac{1}{\tau} \right) \leq \frac{G^2}{\mu^2 t} (6 + \log t). \end{aligned}$$

□

G MORE DETAILS ON THE EXPERIMENTS

In this section, we provide the result of Household ($n=50$) in Figure 3 and the l2-loss of utilities $\|\bar{u}_i - u_i^*\|$ in Table 2.

Figure 3 depicts regrets with $n = 50$ for the Household dataset. Inevitably, the same tendency retains as in Figures 1 and 2. Table 2 illustrates the l2-loss of utilities $\|\bar{u}_i - u_i^*\|$. This metric measures the disparity between the true and the resulting utilities. DA-EtC or DA-UCB outperformed the others. The inconsistency of DA-Grdy was more apparent than the regret indicated in the figures. This observation implies that the estimated values in DA-Grdy are more inaccurate than those of DA-EtC and DA-UCB.

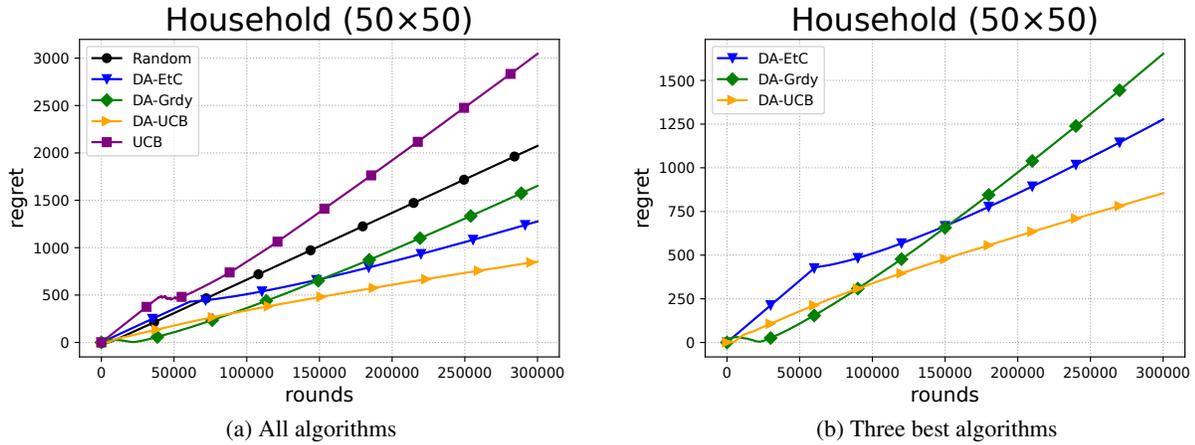


Figure 3: Regret in Household with a larger number of agents ($n = 50$).

Table 2: Loss of utilities per round. Type boldface indicates the minimum loss at the dataset.

	Uniform	Jester	Household ($n = 10$)	Household ($n = 50$)
Random	0.041	0.031	0.033	0.006
UCB	0.073	0.073	0.070	0.019
DA-Grdy	0.010	0.008	0.015	0.006
DA-EtC	0.004	0.007	0.005	0.004
DA-UCB	0.002	0.008	0.004	0.003