Columbia B9136	February 18, 2025

Lecture 4 — Online Bipartite Matching

Lecturer: Will Ma *Scribe: Mohamed Lakhnichi and Zijian Liu*

1 Online Bipartite Matching

Online bipartite matching is one of the early problems to analyze the quantity

$$\sup_{P\in\Delta(\Pi)} \inf_{\tilde{V}} \frac{P(\tilde{V})}{\mathsf{OFF}(\tilde{V})};$$

i.e. the "competitive ratio" when no information about arrivals is given in advance. We earlier considered the problem of accept/rejecting agents with different valuations in the "weight-dependent" setting. However, the key tradeoff in online bipartite matching is orthogonal. In online bipartite matching, an offline vertex, if matched, yields the same reward (weight) regardless of how it was matched. The online algorithm will be presented with a sequence of choices between offline vertices to match, and must decide how to *prioritize* between them.

1.1 Preliminaries

We formally define the online bipartite matching problem now.

Definition 1 (Online Bipartite Matching)**.** An online bipartite matching problem includes the following components and requirements:

- Items $i = 1, \dots, m$.
- Agents $t = 1, \dots, T$.
- Each agent t is compatible with a set of items $N(t) \subseteq [m]$. If $i \in N(t)$, then (i, t) is said to be an edge in the graph.
- Upon seeing $N(t)$, immediately decide an item to "match" t with (if any).
- Items can only be matched once.
- Objective: maximize the total number of matches.

We denote by $\mathcal{I} = \{N(t) : t \in [T]\}$ an instance of such a problem.

- Competitive Ratio (of a policy P): $\inf_{\mathcal{I}} \frac{P(\mathcal{I})}{\mathsf{OFF}(\mathcal{I})}$.
- Competitive Ratio (of the problem): $\sup_{P\in\Delta(\Pi)} \inf_{\mathcal{I}} \frac{P(\mathcal{I})}{\mathsf{OFF}(\mathcal{I})}$.

Maybe the simplest policy one can come up with immediately is to match an item whenever possible. We now establish its theoretical justification.

Proposition 2. *If π matches whenever one if available, then*

$$\pi(\mathcal{I}) \geq \mathsf{OFF}(\mathcal{I})/2, \forall \mathcal{I}.$$

Proof. The proof relies on two key observations:

- Any such policy π ends with a maximal matching
- Whenever the policy π chooses some edge in the graph, it can only block two other edges (that could be part of the offline policy)

Combining these two, we deduce that $\pi(\mathcal{I}) \geq \mathsf{OFF}(\mathcal{I})/2, \forall \mathcal{I}$. ☐

Next, we show that the ratio 1/2 is indeed optimal.

Proposition 3. *If π is deterministic, then there exists an instance $\mathcal{I}$ such that*

$$\pi(\mathcal{I}) = \mathsf{OFF}(\mathcal{I})/2.$$

Proof. We consider the following instance with: $m = 2, T = 2, N(1) = \{1, 2\}$ such that agent 1 arrives before agent 2. In addition, the set $N(2) = \{1, 2\} \setminus \{\pi(1)\}$.

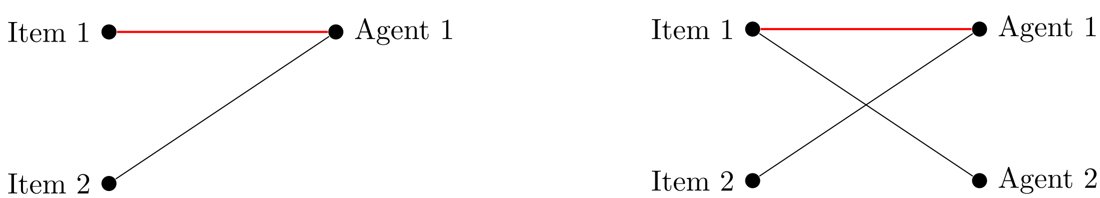

(a) Agent 1 arrives and is matched to Item 1 (b) Agent 2 arrives and remains unmatched

In other words, the instance forces the policy π to match agent 1, which is flexible, before showing agent 2 that is only compatible with the item matched to agent 1. ☐

In order to do better than the deterministic policy, the first idea that comes to our mind is randomization: we should flip a coin to decide the item to match to any arriving agent.

1.2 Policy RANDOM

Definition 4 (RANDOM)**.** The RANDOM policy matches each agent t to a uniformly random available item in $N(t)$.

Example 5. Consider the instance where $m = 2, T = 2, N(1) = \{1, 2\}$ and $N(2) = \{2\}$.

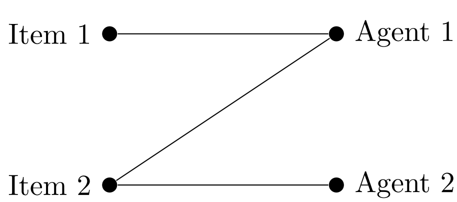

In this case:

$$\begin{cases} \mathsf{OFF}(\mathcal{I}) = 2, \\ \mathsf{RANDOM}(\mathcal{I}) = 1/2(2+1) = 1.5. \end{cases}$$

Then the competitive ratio of RANDOM for this instance is 3/4.

Proposition 6. *The Competitive Ratio (C.R.) of* RANDOM *is* $1/2$*.*

Proof. Assume m is even, $T = m$, we aim to construct an instance $\mathcal{I}_m$ satisfying $\mathsf{OFF}(\mathcal{I}_m) = m$. For the instance $\mathcal{I}_m$, let

$$N(t) = \begin{cases} \{t\} \cup \{m/2+1, \ldots, m\}, & t \leq m/2, \\ \{t\}, & t > m/2. \end{cases}$$

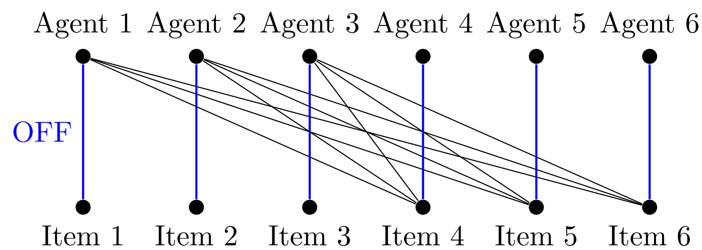

Figure 2: The hard instance for $m = T = 6$

It is clear that $\mathsf{OFF}(\mathcal{I}_m) = m$. We now analyze the policy RANDOM. We focus on items that are hard to match ($i \in [m/2]$), and derive:

$$\begin{cases} \Pr[\text{item 1 is matched}] = \frac{1}{m/2+1}, \\ \Pr[\text{item 2 is matched}] \leq \frac{1}{m/2+1-1}. \end{cases}$$

To explain, item 1 can only get matched to agent 1, who is adjacent to $m/2+1$ items in total, and hence has $\frac{1}{m/2+1}$ probability of choosing item 1. Item 2 can only get matched to agent 2, who is also adjacent to $m/2+1$ items but in the best case only $m/2+1-1$ of those items will be available (one of the items with $i > m/2$ will already be matched to agent 1). In general, for each subsequent agent, the number of available items they are adjacent to decreases by 1 in the best case. There, for any $i \leq m/2$, we have the upper bound:

$$\Pr[\text{item } i \text{ is matched}] \leq \frac{1}{m/2+2-i}.$$

Thus, we know:

$$\mathsf{RANDOM}(\mathcal{I}_m) \leq \sum_{i=1}^{m/2} \frac{1}{m/2+2-i} + \frac{m}{2} \leq \ln\left(\frac{m}{2}+1\right) + \frac{m}{2}.$$

It follows that:

$$\lim_{m\to\infty} \frac{\mathsf{RANDOM}(\mathcal{I}_m)}{\mathsf{OFF}(\mathcal{I}_m)} = \frac{1}{2}.$$

□

Even if the policy RANDOM succeeds to have a competitive ratio better than 1/2 for some instances (e.g. Example 5), Proposition 6 shows that the C.R. of the policy cannot be better than 1/2 that's the lower bound for any policy π (i.e., Proposition 2). In the following section, we consider another randomized policy that achieves a C.R. bigger than 1/2.

1.3 Policy RANKING

Definition 7 (RANKING)**.** The RANKING policy draws a random permutation ("ranking") of the items. Match each t to the highest-ranked available item in $N(t)$.

Example 8. We will highlight the difference between RANDOM and RANKING in matching agents in this instance: $m = 3, T = 2, N(1) = \{1, 2\}$ and $N(2) = \{2, 3\}$.

- When agent 1 arrives:
  - RANKING chooses the item with the highest ranking among items $\{1, 2\}$. We suppose it chooses 1.
  - RANDOM matches 1 or 2 uniformly with probability 1/2. We suppose that it also chooses 1.
- When agent 2 arrives:
  - RANKING chooses the item with the highest ranking among items $\{2, 3\}$. Knowing that item 2 has already been rejected, there is 2/3 chance of matching item 3 and 1/3 chance of matching item 2.
  - RANDOM matches 2 or 3 uniformly with probability 1/2.

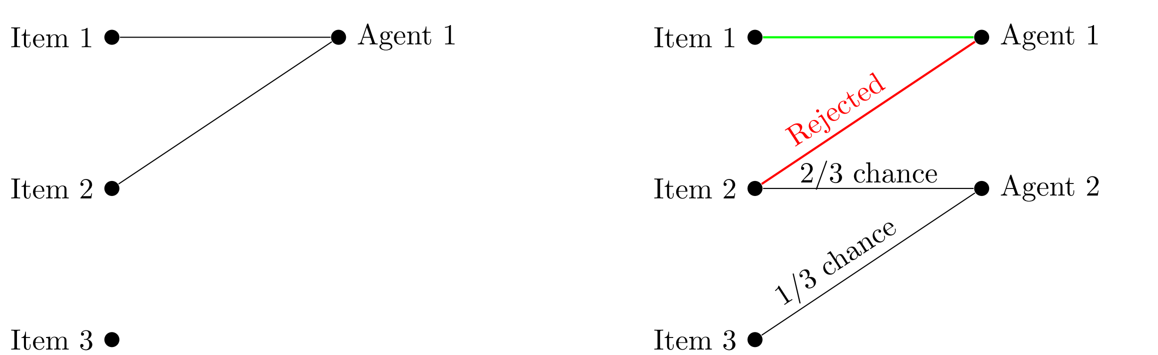

(a) Agent 1 arrives and is matched to Item 1 (b) Agent 2 arrives and is compatible with {2, 3}

In this example, we see that RANKING is more likely to reject items that have been rejected before (i.e., "correlated randomness"), which is not the case of RANDOM that keeps matching independently of whether an item has been rejected or not.

Intuition of beating the hard instances $\mathcal{I}_m$:

- Correlated randomness.
- What's been rejected will get rejected again.

Theorem 9. *The C.R. of* RANKING *is* $\geq 1 - 1/e \approx 63.2\%$.

Proof. We will model the ranking of the algorithm using random variables Y_i and a function p, which we will specify later, in the following manner:

- Items in $[m]$ draw independent seeds $Y_i \sim \text{Unif}[0, 1]$.

- Item i will have "price" $p(Y_i)$, where $p : [0,1] \to [0,1]$ is an increasing function with $p(1) = 1$ (decided later).
- Agent t has "valuation" of 1 for all $i \in N(t)$.

When we match an agent t to i:

- We collect revenue: $\mathsf{Rev}_i = p(Y_i)$.
- Agent t enjoys utility: $\mathsf{Util}_t = 1 - p(Y_i)$.

So the total number of matchings done by the policy writes: $\sum_{i=1}^m \mathsf{Rev}_i + \sum_{t=1}^T \mathsf{Util}_t$. Then:

$$\mathsf{RANKING}(\mathcal{I}) = \sum_{i=1}^m \mathbb{E}[\mathsf{Rev}_i] + \sum_{t=1}^T \mathbb{E}[\mathsf{Util}_t]$$

We use a "**Primal-Dual**" technique that upper-bounds the value of OFF using an LP. In fact:

$$\begin{aligned}
\mathsf{OFF}(\mathcal{I}) \leq \max \quad & \sum_{t, i \in N(t)} x_{t,i} \\
\text{s.t.} \quad & \sum_{t : i \in N(t)} x_{t,i} \leq 1, \quad \forall i \in [m] \\
& \sum_{i \in N(t)} x_{t,i} \leq 1, \quad \forall t \in [T] \\
& x \geq 0 \\
= \min \quad & \sum_{i \in [m]} \alpha_i + \sum_{t \in [T]} \beta_i \\
\text{s.t.} \quad & \alpha_i + \beta_t \geq 1, \quad \forall t \in [T] \quad \forall i \in N(t), \\
& \alpha, \beta \geq 0
\end{aligned}$$

It's worth mentioning that the first inequality is actually an equality in the case of Bipartite graphs: the integer LP and the relaxed LP have the same optimal objective. The key in the proof is the following lemma:

Lemma 10 (Feasibility in Expectation)**.** *$\forall t \in [T], \forall i \in N(t), \mathbb{E}[\mathsf{Rev}_i] + \mathbb{E}[\mathsf{Util}_t] \geq c$, where $c = 1 - 1/e$.*

This lemma actually suffices it implies that the variables $\alpha_i := \mathbb{E}[\mathsf{Rev}_i]/c$ and $\beta_t := \mathbb{E}[\mathsf{Util}_t]/c$ is dual feasible, which implies $\mathsf{OFF}(\mathcal{I}) \leq \frac{1}{c}(\sum_{i=1}^m \mathbb{E}[\mathsf{Rev}_i] + \sum_{t=1}^T \mathbb{E}[\mathsf{Util}_t]) = \mathsf{RANKING}(\mathcal{I})/c$.

Proof of Lemma 10. The proof of the lemma relies on two key ideas:

- **Key idea (a)**. Prove $\mathbb{E}[\mathsf{Rev}_i] + \mathbb{E}[\mathsf{Util}_t] \geq c$ holds after conditioning on any realization of $\bar{Y}_{-i} := (Y_{i'})_{i' \neq i}$.
- **Key idea (b)**. Consider a parallel universe without item i, and compare the execution of the algorithm in the real and the parallel universe.

Let u denote the utility of agent t in the parallel universe. $u = 1 - p(y)$ where y is the seed of the item purchased by t in parallel execution ($y = 1$ if t wasn't matched). Then:

Claim 1. $\mathbb{E}_{Y_i}[\mathsf{Rev}_i \mid \bar{Y}_{-i}] \geq \int_0^y p(z)dz$.

Claim 2. $\mathsf{Util}_t \geq u$ w.p. 1.

Proof of Claim 2. Consider the executions (deterministic) in the real and parallel universes over $t = 1, \dots, T$. We will prove the statement "*The set of remaining items in the real universe is a superset of the items available in the parallel universe*" by induction.
It is clear that the hypothesis is true for the base case. For the induction step, the statement won't be valid only if the item i_r chosen in the real universe is different from the item i_p chosen in the parallel universe and both items are different from i. This assumes that i_r has a higher rank than i_p in the real universe whilst it is the opposite in the parallel universe. However, this contradicts the fact that ranks are fixed from the beginning and don't depend on the universe. As a result, the induction is established and it follows that: $\mathsf{Util}_t \geq u$.

Proof of Claim 1. Conditioning on $Y_i < y$, when agent t arrives in the real universe, there are two possible cases:

1. Item i is already sold, which implies that $\mathsf{Rev}_i = p(Y_i)$.
2. Otherwise, the real and parallel universes must be identical at that point, which means that item i is necessarily matched to agent t.

It follows that:

$$\mathbb{E}_{Y_i}[\mathsf{Rev}_i \mid \bar{Y}_{-i}] \geq \int_0^y p(z)dz.$$

We have shown: $\mathbb{E}[\mathsf{Rev}_i] + \mathbb{E}[\mathsf{Util}_t] \geq 1 - p(y) + \int_0^y p(z)dz$ for any adversarially chosen y. To finish the proof, we need R.H.S., $\geq 1 - 1/e, \forall y \in [0,1]$, $p(y) = e^{y-1}$. We will do so by solving for p the differential equation:

$$\forall y \in [0,1]: \ 1 - p(y) + \int_0^y p(z)dz = c, \ p \text{ is increasing and } p(1) = 1.$$

It implies: $p'(y) - p(y) = 0$, necessarily $p(y) \propto e^y$ and the fact that $p(1) = 1$ leads to $p(y) = e^{y-1}$. □

By proving Lemma 10, we actually finished the proof of Theorem 9. □

2 Variant: k-Matching

We now consider the following variant of the online bipartite matching problem.

Definition 11 (k-Matching)**.** In the k-matching problem, each item can be matched k times.

In the online bipartite matching problem, we have seen that deterministic policies are strictly worse than randomized policies. However, when k goes large, the gap between the two kinds of policies will vanish, as indicated by the next two results.

2.1 Lower Bound

Theorem 12. *There exists a deterministic algorithm that can get C.R. $1 - 1/e$ as $k \to \infty$.*

Proof. We consider an algorithm (policy) named BALANCE defined as follows.

Definition 13 (BALANCE)**.** The BALANCE policy matches each agent t to the item $i \in N(t)$ with the most copies remaining (picks arbitrarily under a tie).

Let $c(k)$ denote the C.R. of BALANCE for a given k, then our goal is to show

$$c(1) = 1/2 \text{ and } \lim_{k\to\infty} c(k) = 1 - 1/e.$$

Before moving on, we first recall the dual LP of the k-matching problem:

$$\begin{aligned} \min \quad & \sum_{i=1}^{m} k\alpha_i + \sum_{t=1}^{T} \beta_t \\ \text{s.t.} \quad & \alpha_i + \beta_t \geq 1, \forall t \in [T], i \in N(t), \\ & \alpha_i, \beta_t \geq 0, \forall t \in [T], i \in [m]. \end{aligned} \tag{1}$$

Note that the value of the dual LP is always greater than or equal to $\mathsf{OFF}(\mathcal{I})$.

Our proof strategy is to initialize dual variables

$$\alpha_i = \beta_t = 0, \forall i \in [m], t \in [T],$$

then increase α_i over time $t = 1, \dots, T$ and set β_t at time t (when agent t arrives) such that:

- $\min_{i\in N(t)} \alpha_i \geq c(k) - \beta_t$ whenever we set up β_t at time t. Note that this ensures $\min_{i\in N(t)} \alpha_i \geq c(k) - \beta_t$ will hold at the end because we will only increase α_i over time and the value of β_t remains unchanged once it is set.
- The value of the dual LP under our α_i and β_t at time T is the same as the reward of the algorithm thus far, $\mathsf{BALANCE}(\mathcal{I})$.

To do so, we consider a sequence $p_\ell, \forall \ell \in \{0\} \cup [k]$ denoting the price of any item when ℓ units remain and satisfying $1 = p_0 > p_1 > \cdots > p_k \geq 0$ (determined later). So the BALANCE algorithm can be interpreted as picking an item with the lowest price every time. In particular, if agent t doesn't match, it is equivalent to agent t choosing an item with 0 units remaining.

At every time t, suppose agent t chooses item i, which has ℓ units remaining.

Case 1. If $\ell \geq 1$, then we know $\Delta\mathsf{BALANCE}(\mathcal{I}) = 1$, where $\Delta\mathsf{BALANCE}(\mathcal{I})$ denotes the change (in fact, the increase) of $\mathsf{BALANCE}(\mathcal{I})$. In this case, we change the dual variables as follows:

- Increase α_i by p_ℓ/k and set $\beta_t = 1 - p_\ell$.

This operation implies the increase of the dual objective (see (1)) is also 1. In other words, we have $\Delta\mathsf{DualObjective}(\mathcal{I}) = 1$.

Case 2. If $\ell = 0$ (i.e., agent t isn't matched), then we know $\Delta\mathsf{BALANCE}(\mathcal{I}) = 0$ and do the following operation:

- Keep α_i unchanged and set $\beta_t = 1 - p_\ell = 0$ (recall $p_0 = 1$).

By (1), we know $\Delta\textsf{DualObjective}(\mathcal{I}) = 0$.

Therefore, in both cases, there is always

$$\Delta\textsf{BALANCE}(\mathcal{I}) = \Delta\textsf{DualObjective}(\mathcal{I}).$$

As such, we only need to prove the dual "feasibility" at every time t, i.e.,

$$\alpha_i + \beta_t = \alpha_i + 1 - p_\ell \geq c(k), \forall i \in N(t).$$

Observe that all $i \in N(t)$ have $\leq \ell$ copies remaining now, which implies $\alpha_i \geq \frac{1}{k}\sum_{\ell' > \ell} p_{\ell'}, \forall i \in N(t)$. Hence, it is enough to show

$$1 - p_\ell + \frac{1}{k}\sum_{\ell' > \ell} p_{\ell'} \geq c(k), \forall \ell \in \{0\} \cup [k]. \tag{2}$$

To find the optimal $c(k)$ and $p_\ell, \forall \ell \in [k]$ (note that we set $p_0 = 1$ explicitly), one can imagine (2) holds as an equality. In that case, we can find:

- If $\ell = k$, we have $1 - p_k = c(k)$.
- If $\ell = k-1$, we have $1 - p_{k-1} + \frac{1}{k}p_k = c(k) \Rightarrow p_{k-1} = (1 + 1/k)p_k$.
- In general, we have for any $\ell \in [k]$,

$$1 - p_\ell + \frac{1}{k}\sum_{\ell' > \ell} p_{\ell'} - \left(1 - p_{\ell-1} + \frac{1}{k}\sum_{\ell' > \ell - 1} p_{\ell'}\right) = c(k) - c(k) = 0 \Rightarrow p_{\ell-1} = (1 + 1/k)p_\ell.$$

Recalling $p_0 = 1$, these together imply

$$p_\ell = (1 + 1/k)^{-\ell}, \forall \ell \in \{0\} \cup [k] \text{ and } c(k) = 1 - (1 + 1/k)^{-k}.$$

In particular, we have

$$c(1) = 1/2 \text{ and } \lim_{k \to \infty} c(k) = 1 - 1/e.$$

Remark 14. When k is large, it can be observed $p_\ell \approx e^{-\ell/k}$, which could be viewed as a discretized version of $e^{-U} \stackrel{\mathcal{D}}{=} e^{U-1}$ for $U = \text{Unif}[0,1]$, the same quantity shown up in the $\textsf{RANKING}$ policy.

The above argument means that we have shown for any instance $\mathcal{I}$,

$$\textsf{BALANCE}(\mathcal{I}) = \sum_{i=1}^{m} k\alpha_i + \sum_{t=1}^{T} \beta_t,$$

for some α_i's, β_t's ≥ 0 satisfying $\alpha_i + \beta_t \geq c(k), \forall t \in [T], i \in N(t)$. So $(\alpha_i/c(k))_{i\in[m]}$ and $(\beta_t/c(k))_{t\in[T]}$ are dual feasible, which concludes

$$\textsf{BALANCE}(\mathcal{I})/c(k) \geq \textsf{OFF}(\mathcal{I}) \Leftrightarrow \frac{\textsf{BALANCE}(\mathcal{I})}{\textsf{OFF}(\mathcal{I})} \geq c(k), \forall \mathcal{I}.$$

□

2.2 Upper Bound

Now, we turn our attention to the upper bound. To do so, we first recall the following result.

Lemma 15 (Yao's Minimax Lemma)**.**

$$\sup_{P\in\Delta(\Pi)} \inf_{\mathcal{I}} \frac{P(\mathcal{I})}{\mathsf{OFF}(\mathcal{I})} = \inf_{\mathcal{I}\sim\mathbf{F}} \frac{\sup_{\pi\in\Pi}\pi(\mathbf{F})}{\mathsf{OFF}(\mathbf{F})}. \tag{3}$$

By Theorem 12, we know the L.H.S. of (3) $\geq c(k)$ for any fixed k. In order to prove an upper bound, we want to use the R.H.S. of (3). In other words, we need to find a hard distribution $\mathbf{F}$. In the following, we fix k and vary m to construct a hard distribution $\mathbf{F}_m$ for every m.

Define the hard distribution $\mathbf{F}_m$: We let $T = mk$, i.e., m "groups" of size k. Moreover, agents in each group are identical (same $N(t)$).

- Group 1 arrives first, and all agents t in Group 1 have $N(t) = [m]$.
- Let σ be a uniformly random permutation of $[m]$. Any agent t in Group $g \in [m]\backslash[1]$ has

$$N(t) = [m]\backslash\{\sigma(1), \dots, \sigma(g-1)\}.$$

Each realization σ results an instance $\mathcal{I}_\sigma$. Note that $\mathsf{OFF}(\mathcal{I}_\sigma) = mk$ for any realization σ. Thus, $\mathsf{OFF}(\mathbf{F}_m) = mk$.

Let Q_{gi} denote the number of matches between group g and item $\sigma(i)$. Then there is

$$\mathbb{E}[Q_{gi}] \leq \begin{cases} 0, & \text{if } i < g, \\ \frac{k}{m-g+1}, & \text{if } i \geq g. \end{cases}$$

We know $\mathbb{E}[\sum_g Q_{gi}] \leq \min\left\{k, \sum_{g=1}^{i} \frac{k}{m-g+1}\right\}$. So there is

$$\pi(\mathbf{F}_m) \leq \sum_{i=1}^{m} k \min\left\{1, \sum_{g=1}^{i} \frac{1}{m-g+1}\right\} \leq \sum_{i=1}^{m} k \min\left\{1, \int_{m-i}^{m} \frac{1}{z}\mathrm{d}z\right\} = k\sum_{i=1}^{m} \min\left\{1, \ln\frac{m}{m-i}\right\},$$

which implies

$$\frac{\sup_{\pi\in\Pi}\pi(\mathbf{F}_m)}{\mathsf{OFF}(\mathbf{F}_m)} \leq \frac{1}{m}\sum_{i=1}^{m} \min\left\{1, \ln\frac{m}{m-i}\right\} \xrightarrow{m\to\infty} \int_0^1 \min\left\{\ln\frac{1}{1-z}, 1\right\}\mathrm{d}z$$
$$= \int_0^{1-1/e} \ln\frac{1}{1-z}\mathrm{d}z + 1/e = 1 - 1/e.$$

The above result suggests that if m can also be varied, then the ratio $1 - 1/e$ is actually unimprovable.

3 Variant: AdWords

Definition 16 (AdWords)**.** An AdWords problem includes the following components and procedures:

- advertiser $i \in [m]$ with initial budget $B_i \geq 0$, and S_i is the current spend of i.

- when impression $t \in [T]$ arrives, receive bids $b_{it} \geq 0, \forall i \in [m]$.
- "match" to advertiser i of our choosing.
- Obtain revenue $R_t = \min\{b_{it}, B_i - S_i\}$.
- $S_i \leftarrow S_i + R_t$.
- Maximize the total revenue $\sum_{t=1}^{T} R_t$.

Remark 17. By setting $B_i = k, \forall i \in [k]$ and $b_{it} \in \{0,1\}, \forall i \in [k], t \in [T]$, the AdWords problem recovers the k-matching problem studied in the previous section.

We first state a theorem without given proof (the proof uses the same ideas as above).

Theorem 18 (Proof Omitted)**.** *If* $\frac{\max_t b_{it}}{B_i} \to 0, \forall i \in [m]$ *(small bids assumption), then a generalization of* BALANCE *policy (see Definition 13) achieves C.R.* $1 - 1/e$*.*

Without the small bids assumption, whether the C.R. $1-1/e$ can be achieved remains unclear. In particular, the analysis for the RANKING policy (see Definition 7) fails. The best-known result is 1/2 until 2020 being improved to 0.5016 by Huang et al. (2020).

Theorem 19. *The C.R.* $1/2$ *can be achieved by the greedy algorithm, i.e., match advertiser* i *who can maximize revenue* R_t *at every time* t*.*

See Theorem 5.1 in Mehta (2013) for the proof of Theorem 19.

4 Discussions

Lastly, we briefly mention an open question. As one can see, the upper bound result established in Subsection 2.2 holds for any k, but where m is required to approach ∞ (and so does T). Naturally, one may wonder about the exact characterization when m is fixed with only varying T. Unfortunately, such a problem is still open even for $k = 1$, i.e., the online bipartite problem. A known lower bound of C.R. is $\geq 1 - (1 - 1/m)^m$, which is, however, smaller than the upper bound given by the upper triangular graph (the instance we used in Subsection 2.2). Once $m \geq 3$, these two bounds are not matched. People interested in this problem could refer to, for example, Feige (2020).

Bibliographical notes. The formulation of online bipartite matching and all initial results come from Karp et al. (1990). The analysis of RANKING presented here comes from Devanur et al. (2013), while the analysis of BALANCE comes from Kalyanasundaram and Pruhs (2000). The upper bound and stated results for AdWords come from Mehta et al. (2005).

References

N. R. Devanur, K. Jain, and R. D. Kleinberg. Randomized primal-dual analysis of ranking for online bipartite matching. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, pages 101–107. SIAM, 2013.

U. Feige. Tighter bounds for online bipartite matching. In *Building Bridges II: Mathematics of László Lovász*, pages 235–255. Springer, 2020.

Z. Huang, Q. Zhang, and Y. Zhang. Adwords in a panorama. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1416–1426, 2020. doi: 10.1109/FOCS46700.2020.00133.

B. Kalyanasundaram and K. R. Pruhs. An optimal deterministic algorithm for online b-matching. *Theoretical Computer Science*, 233(1-2):319–325, 2000.

R. M. Karp, U. V. Vazirani, and V. V. Vazirani. An optimal algorithm for on-line bipartite matching. In *Proceedings of the twenty-second annual ACM symposium on Theory of computing*, pages 352–358, 1990.

A. Mehta. Online matching and ad allocation. *Foundations and Trends® in Theoretical Computer Science*, 8(4):265–368, 2013. ISSN 1551-305X. doi: 10.1561/0400000057. URL `http://dx.doi.org/10.1561/0400000057`.

A. Mehta, A. Saberi, U. Vazirani, and V. Vazirani. Adwords and generalized on-line matching. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 264–273, 2005.