Truthful Matching with Online Items and Offline Agents

Michal Feldman ✉ ⬤
Blavatnik School of Computer Science, Tel Aviv University, Israel
Microsoft Research, Herzliya, Israel

Federico Fusco ✉ ⬤
Department of Computer, Control and Management Engineering "Antonio Ruberti",
Sapienza University of Rome, Italy

Simon Mauras ✉ ⬤
Blavatnik School of Computer Science, Tel Aviv University, Israel

Rebecca Reiffenhäuser ✉ ⬤
Institute for Logic, Language and Computation, University of Amsterdam, The Netherlands

Abstract

We study truthful mechanisms for welfare maximization in online bipartite matching. In our (multi-parameter) setting, every buyer is associated with a (possibly private) desired set of items, and has a private value for being assigned an item in her desired set. Unlike most online matching settings, where agents arrive online, in our setting the items arrive online in an adversarial order while the buyers are present for the entire duration of the process. This poses a significant challenge to the design of truthful mechanisms, due to the ability of buyers to strategize over future rounds. We provide an almost full picture of the competitive ratios in different scenarios, including myopic vs. non-myopic agents, tardy vs. prompt payments, and private vs. public desired sets. Among other results, we identify the frontier up to which the celebrated $e/(e-1)$ competitive ratio for the vertex-weighted online matching of Karp, Vazirani and Vazirani extends to truthful agents and online items.

1 Introduction

Matching in bipartite graphs is a fundamental model that has gained massive importance in numerous applications with the growth of the Internet. Some examples include items and buyers in e-commerce, drivers and passengers in ride-sharing platforms, ad slots and

50th International Colloquium on Automata, Languages, and Programming (ICALP 2023).
Editors: Kousha Etessami, Uriel Feige, and Gabriele Puppis; Article No. 58; pp. 58:1–58:20
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

advertisers in online ad auctions, and jobs and workers in online labor markets. In these applications, it is common that vertices on one side are known from the outset, while vertices from the other side arrive one-by-one in an online fashion. Upon the arrival of an online vertex, its information is revealed (containing, e.g., its set of adjacent edges, and their weights), and the algorithm has to immediately and irrevocably decide either to match it with an available offline partner or leave it unmatched forever. The goal is to maximize the sum of the weights along the matched edges.

A celebrated result in online matching by Karp, Vazirani, and Vazirani [18] shows that, in the unweighted setting, a simple randomized strategy called RANKING achieves a competitive ratio of $e/(e-1)$, and this is optimal. This result extends to the setting where the vertices on the offline side are weighted and the objective is to maximize the sum of the weights of the matched vertices. Although the original algorithm for this problem, PERTURBED-GREEDY [1], was designed for non-strategic settings, online matching problems have also been studied in the presence of *strategic* agents e.g., [21, 25, 11, 7]. This is not a mere theoretical exercise: in many applications of online matching the parties involved are interested in misreporting their true valuations to obtain a better outcome, e.g., combinatorial and ad-auctions, kidney exchange, school-student matching, and house allocation. In the presence of strategic agents, an agent's value is her private information, and is not directly available to the mechanism designer. The main challenge here is to design *incentive-compatible* or *truthful* mechanisms which, besides finding a good matching, also ensure that it is in the agents' best interest to report their true values. In addition to making decisions regarding the matching itself, such mechanisms can also charge some payment from the agents in order to incentivize them to truthfully report their values. Here, each agent strives to maximize her *quasi-linear utility*, i.e. the value she obtains from her assigned item, minus the payment she has to make.

In almost all previous studies, the agents are represented by the vertices on the online side, while the items they are competing over are available offline. In many natural internet applications, e.g. selling advertising opportunities via repeated auctions, the agents are fixed and observe a stream of items arriving online. This motivates the study of a *reversed* online matching problem, where each vertex on the offline side is strategic on her value, and her set of desired items that arrive online. This variant has been considered thus far only in very restricted settings [8, 9]. This is not a coincidence: when agents are present throughout the entire matching process, many new manipulation opportunities arise, and incentivizing truthful behavior is significantly more challenging. Indeed, the online nature of the problem forces any mechanism to repeatedly make irrevocable decisions upon the arrival of goods, lacking knowledge about future opportunities that might arise to the participating agents. The agents – possibly aware of those future opportunities – may strategize to gain benefits in the future, defying standard tools applicable when agents arrive online.

Our work provides a systematic analysis of this scenario, and gives (almost) tight competitive ratios under a rich variation of natural assumptions. We study the problem along different dimensions, as follows. First, we consider two types of agents – myopic and non-myopic – that are characterized by the different information they have on the instance. Myopic agents make strategic considerations that are limited to the current time step, without looking forward into the future (see, e.g., Deng, Panigrahi and Zhang [9]), whereas non-myopic agents optimize across multiple time steps, using the up-front knowledge of the underlying (online) graph. The assumption of myopic agents clearly eradicates some of the difficulties of designing (almost tight) online mechanisms with offline strategic agents, thus allowing us to derive efficient mechanisms from known online matching algorithms, e.g., from Aggarwal et al. [1]. Second, we consider two types of private information. In the first scenario we consider, an

Table 1 Summary of our results, with $\nu = \min(m, n)$, where n is the number of agents and m the number of items.

	deterministic	randomized
prompt	2 (Theorem 12)	$e/(e-1)$ (Theorem 13)

(a) Myopic agents.

	deterministic	randomized
tardy	2 (Theorem 14)	$e/(e-1)$ (Theorem 14)
prompt	$\geq \nu$ (Theorem 3)	$\Omega(\log \nu / \log \log \nu)$ (Theorem 4)

(b) Non-myopic agents with *public* graph edges.

	deterministic	randomized
tardy	equiv. to prompt	equiv. to prompt
prompt	$\leq \nu$ (Theorem 9)	$\mathcal{O}(\log \nu)$ (Theorem 10)

(c) Non-myopic agents with *private* graph edges.

agent's private information consists only of her private value for her desired items, but the set of desired items is publicly known. In the second scenario, both the value and the set of desired items are private information. Notably, in both cases the graph structure is revealed to the mechanism step-by-step, upon the arrival of every item. Finally, we distinguish between *prompt* and *tardy* mechanisms. Both types of mechanisms make allocation decisions immediately. However, they differ in the time at which they make payment decisions. Prompt mechanisms make payment decisions immediately upon allocation, while tardy mechanisms may delay payment decisions to the end of the entire process.

1.1 Our Results and Techniques

We conduct a systematic study of online bipartite matching with online items and offline agents, in a variety of scenarios and we provide (almost) tight bounds for the settings of interest, as summarized in Table 1.

Myopic agents. The simpler setting we investigate is that of myopic agents. These agents care only about their instantaneous utility, and do not strategize over the future. As such, we only consider prompt mechanisms for this type of agents. By exploiting the myopic nature of the agents, it is not difficult to turn the best (non-truthful) algorithms into (truthful) mechanisms. In particular, we construct a deterministic prompt mechanism based on the greedy matching algorithm that is guaranteed to achieve at least a half of the optimal welfare. We also give a randomized prompt mechanism based on the algorithm for weighted online matching [1], which is $e/(e-1)$-competitive. This shows that the transition from non-strategic agents to strategic myopic agents does not lead to a deterioration in efficiency guarantees. Notably, for the special case we study, our bounds for myopic online matching improve vastly over those obtained by Deng, Panigrahi and Zhang [9] for general XOS valuations. The results for myopic agents are reported in Appendix A.

Non-myopic agents with public graph edges. In a more general setting, we consider non-myopic agents who can strategize about their values, but not about their desired items: upon the arrival of an item, the set of agents interested in it is revealed (no strategizing involved), but the agent values are reported by the agents themselves. This variant is single-parameter, for which Myerson's lemma applies [24]. We prove that, if the mechanism is allowed to wait until the end of the online phase to set prices (i.e., tardy mechanism), it is possible to

achieve the same bounds as in the myopic case – by showing that our algorithms GREEDY and PERTURBED-GREEDY maintain a certain form of *global monotonicity*. In contrast, when prices have to be fixed upon item arrival, an agent might hope to receive an item, i.e. one over which there is not as much competition, for a better price later if she waits instead of truthfully reporting her interest in the current item. To avoid this, prompt prices need to be non-decreasing throughout the mechanism, which allows us to show a sharp deterioration from tardy to prompt mechanisms: for deterministic mechanisms, we prove a $\nu = \min(m, n)$ competitive lower bound, where n and m denote the number of agents and items, respectively. For *randomized* prompt mechanisms obtaining such a lower bound is much more challenging, but as a central result we manage to establish an $\Omega(\log \nu / \log \log \nu)$ lower bound, using Yao's minimax principle. Starting from a carefully designed distribution of problem instances with exponentially increasing agent valuations, we employ a primal-dual approach together with our previous observations on the behavior of *deterministic* truthful mechanisms to bound the achievable competitive ratio. Almost matching deterministic and randomized upper bounds for prompt mechanisms are inherited from non-myopic prompt mechanisms with private graph edges. See Section 3 for our results for public graph edges.

Non-myopic agents with private graph edges. We finally consider non-myopic agents when both valuations and the set of desired items are private information. For deterministic prompt mechanisms, the ν lower bound from the case of public graph edges applies. Moreover, we show that in the case of private edges, every deterministic truthful mechanism is essentially prompt. Thus, tardy mechanisms for this case retain the ν lower bound, exhibiting a large gap between tardy mechanisms for public vs. private edges. We then provide a prompt truthful deterministic mechanism that is ν-competitive, matching the lower bound. For randomized prompt truthful mechanisms, the $\Omega(\log \nu / \log \log \nu)$ lower bound from the case of public edges applies to tardy randomized mechanisms as well, since these are probability distributions over deterministic mechanisms and, as stated above, all deterministic truthful mechanisms for private edges are prompt. On the positive side, we provide a randomized prompt truthful mechanism that gives an almost matching competitive ratio of $O(\log \nu)$. This algorithm is based on a tailored explore-exploit approach. Our results for private graph edges are reported in Section 4.

Ex-post vs. ex-ante truthfulness. Finally, we explore the notion of *ex-ante* truthfulness, as opposed to *ex-post* truthfulness, where agents' true declarations maximize their expected utility instead of their utility in *any realization* of the random choices of the mechanism. Clearly, ex-post truthfulness implies ex-ante truthfulness. In the setting with myopic buyers, we only need to consider ex-post truthfulness as we obtain tight approximation in this stronger model that closes the problem also for the ex-ante analogue. In the setting of non-myopic buyers, we show that the additional hardness introduced by truthfulness cannot be fully attributed to the fact that we require ex-post truthfulness. Specifically, we establish a lower bound of 2 for the competitive ratio of ex-ante truthful mechanisms for this setting (even with respect to randomized tardy ones), exhibiting a gap to the corresponding $e/(e-1)$ upper bound for myopic buyers. Our proof utilizes an instance for which we establish lower bounds on the expected utility of various types of agents. We then employ these to show a contradiction to the mechanism's correctness. Our results for ex-ante truthfulness are reported in Section 6.

Remark. Throughout the paper, we assume that weights are assigned to vertices (agents) rather than edges; note, it is well known that for the more general case of edge weights even the algorithmic problem is hopeless (see, e.g., Appendix G of [1]). One may also wonder why we do not study non-myopic agents with public valuations but private edges. The reason is that in the case of public valuations, it is easy to see that agents cannot benefit from misreporting their edges, implying that GREEDY and PERTURBED-GREEDY are truthful.

1.2 Further Related Work

Karp, Vazirani and Vazirani [18] introduce the online matching problem, and study it under one-sided bipartite arrivals. They observe that the trivial 1/2-competitive greedy algorithm (which matches any arriving vertex to an arbitrary unmatched neighbor, if one exists) is optimal among deterministic algorithms for this problem. They also provide a groundbreaking and elegant randomized algorithm for this problem, called RANKING, which achieves an optimal $e/(e-1)$ competitive ratio. The work of Karp, Vazirani and Vazirani [18] was extended to vertex weighted settings by Aggarwal et al. [1], who give an optimal $e/(e-1)$-competitive, randomized algorithm using *random perturbations* of weights by appropriate multiplicative factors. The same bound has been re-proven over the years [6, 10, 14, 12]. Various extensions of one sided online matching and its economic applications (e.g., display ads) have been widely studied, see e.g. the excellent survey of Mehta [22] for further reference. Online matching has also been studied under edge and general vertex arrivals, as well as in different stochastic settings (see e.g., [19, 20, 13, 17, 15, 16]).

An important generalization of assignment problems in the form of matchings are combinatorial auctions, where buyers can obtain a *subset* of the available items, instead of just one. Combinatorial auctions with offline strategic buyers and online items have been recently studied by [9] for submodular and XOS valuations in the case of myopic buyers – considered also in this work – and in the less constrained setting of items that must not be irrevocably assigned at time of arrival. Deng, Panigrahi and Zhang [9] show (for myopic buyers) a sharp separation between submodular valuations, which admit a logarithmic competitive ratio, and XOS valuations, for which a polynomial lower bound is proven. In our work, we prove tight constant bounds for myopic buyers in the important special case of a unit-demand matching.

Cole, Dobzinski and Fleischer [8] formally introduced the notions of *prompt* and *tardy* for mechanisms, after observing the severe negative aspects of many existing (tardy) methods. They study prompt truthful mechanisms for an online problem that is related to ours, but with some restrictions: while agents are still on the offline side of the graph, their items of interest are restricted to form an interval over the online steps (which corresponds to the interval buyers are present). Further, agents report their departure time (which can be public/private) once they arrive, and their arrival time is public knowledge. Babaioff, Blumrosen and Roth [4] later investigated truthful prompt mechanisms for allocating an unknown number of identical items arriving online, which can be phrased in our model as having all desired sets equal to the same prefix of the sequence of items. Both of these works [8, 4] are close to ours in spirit. They present logarithmic-competitive prompt mechanisms in restricted settings, and prove lower bounds using Yao's pinciple (≥ 2 in [8], and $\Omega(\log \log n)$ in [4]). The notions of tardy and prompt mechanisms have since been adopted in the literature, see e.g. [3, 28]. The model of offline agents and online items has been the subject of extensive investigation in economic theory in dynamic mechanism design. Despite this obvious relation to our setting, there are fundamental differences (see for example [23, 2, 5]). In dynamic mechanism design, a strategic buyer learns her valuations at time of arrival of each item. Opposed to our setting, priors on agents' valuations for each online item are usually known beforehand. Finally,

in our matching setting the agents' valuations can assume only two values, v_i and 0, and we consider unit demand buyers instead of additive valuation agents as it is customary in dynamic mechanism design.

2   Preliminaries

We are given a bipartite graph $G = (B, I; E)$, where B is a set of n vertices, corresponding to buyers, I is a set of m vertices, corresponding to items, and $E \subseteq B \times I$ is the set of edges. We denote by ν the smallest between the number of buyers n and the items m. The set of buyers is known beforehand, while the items arrive one by one in some unknown, possibly adversarial, order. Without loss of generality we assume that item j arrives at time j. Each buyer i has two pieces of private information: the set of items she is interested in, and her value v_i if she gets at least one of them (the value for other items is 0). Upon the arrival of a new item, every buyer declares if she is interested in the current item and, if yes, her value. Let $b_{i,j}$ denote the bid of buyer i for item j (with the convention that $b_{i,j} = 0$ if buyer i is not interested in item j). Without loss of generality, we may assume that buyers cannot change their declared valuation after they have declared it once[1], i.e. every nonzero bid of the same buyer is the same value b_i, and that every buyer is assigned at most one item.

A mechanism $\mathcal{M}$ is composed of an *allocation* scheme and a *payment* scheme. Upon the arrival of every item, and based on buyer bids, the mechanism decides immediately and irrevocably to either assign the new item to some buyer who has not been assigned an item yet, or leave it unassigned forever. Thus, the resulting allocation is a matching in G: every buyer receives at most one item, and every item is allocated to at most one buyer. We denote by μ the induced matching, so that μ_j denotes the buyer to whom item j is assigned (we assume that an item j can only be assigned to a buyer who declares interest in j). If j is unassigned, we write $\mu_j = \emptyset$. We also write μ_i^{-1} to denote the item assigned to buyer i, with the convention that $\mu_i^{-1} = \emptyset$ if i is left unassigned. The *allocation* is computed online; i.e., μ_j is determined using only the bids on items up to j. In addition to the allocation, the mechanism decides how much each buyer should pay. A payment scheme is denoted by p, where p_i denotes the non-negative payment of buyer i. We distinguish between two types of *payment* schemes, according to the time at which the mechanism determines the payment. A *tardy* mechanism is one where the payment vector p is computed in the end of the process. A *prompt* mechanism is one where the payment p_i of every buyer i is determined upon the assignment of buyer i (i.e., upon the arrival of item μ_i^{-1}). The mechanism's objective is to maximize the *social welfare* of the allocation μ, which is the sum of the buyer values for their assigned items. The social welfare is given by $\mathrm{SW}(\mu) = \sum_{i \in B} v_i \cdot \mathbb{1}_{\{(i, \mu_i^{-1}) \in E\}}$. Note that a mechanism can also be randomized, so that its allocation is a distribution over matchings. In case of a randomized mechanism, we measure its efficiency by the expected social welfare. We say that a mechanism gives an α approximation, or is α-competitive (where $\alpha \geq 1$), if its (expected) social welfare is at least an $1/\alpha$ fraction of the welfare of a maximum weight matching. That is, μ is α-competitive if $\mathrm{OPT} = \mathrm{SW}(\mu^\star) \leq \alpha \cdot \mathbb{E}\left[\mathrm{SW}(\mu)\right]$, where $\mu^\star$ is the maximum weight matching in G.

A *bidding strategy* $\mathcal{B}_i$ for buyer i is a sequence of bids $b_{i,j}$ that specifies, every time a new item j arrives, whether to declare interest in it and which value to report. The bid $\mathcal{B}_i$ might depend on the bids of the other agents, the actions of the mechanism, and the knowledge the buyers have on the sequence of items. Recall that once an agent declares a positive valuation

---

[1] Mechanisms can "punish" such behavior by discarding the buyer from further consideration

$b_{i,j} = b_i > 0$ for some item j, she cannot change her value thereafter; namely, all bids for future items j' can take the value of either b_i or 0. Let $\mathcal{B}$ denote the profile of buyer bidding strategies, and $\mathcal{B}_{-i}$ denote the profile of all buyer strategies excluding buyer i. We assume that every buyer has a quasi-linear utility function: $u_i(\mathcal{M}, \mathcal{B}_i, \mathcal{B}_{-i}) = v_i \cdot \mathbb{1}_{\{(i, \mu_i^{-1}) \in E\}} - p_i$.

A buyer is called *myopic* if upon the arrival of every item j, she cares only about maximizing her utility in that round, without considering its effect on future rounds. I.e., upon the arrival of item j, she maximizes the utility function $u_{i,j} = v_i \cdot \mathbb{1}_{\{\mu_j = i,\, (i,j) \in E\}} - p_i$. We consider myopic agents only in the context of prompt mechanisms, where the price p_i is determined immediately. We study the following ex-post notion of truthfulness: (i) A mechanism for myopic agents is *truthful* if it is always in the best interest of a myopic buyer to declare her value truthfully. (ii) A mechanism for non-myopic agents is *truthful* if an agent maximizes her utility for every realization of the mechanism by declaring her value truthfully. Finally, we only consider mechanisms that are *ex-post individually rational*, meaning that all agents (myopic or not) have non-negative utility, for every realization of the mechanism.

3 Prompt mechanisms with public graph edges

We start with the setting where agents are assumed to know, and strategize about, the whole sequence of items arriving. Note that this is a strong information asymmetry between agents and mechanism, as the latter only discovers the items as they are revealed online and has no information on the future. As a first step in this challenging model, in this section we study the case where agents may only lie on their valuations. Our main focus here is on establishing lower bounds, which will naturally extend to the case where the edges of the graph are private information.

3.1 Deterministic truthful mechanisms

When mechanisms are required to be prompt, the problem becomes much harder despite the fact that each agent's private information is just a single value. This is due to the online nature of the problem versus the possibly universal knowledge of the buyers, as outlined before. We first concentrate on deterministic prompt truthful mechanisms, and prove that the scope of these is quite limited. The critical item property is also used in [4] to prove a lower bound analogous to Theorem 3.

▶ **Definition 1** (critical item property). *We say that a deterministic mechanism satisfies the critical item property if and only if for every buyer i, there exists some $j \in I$ such that for any reported value b_i of i, the mechanism assigns i with item j, or none at all. Note that j may depend on the edges of the graph, and on the values of other buyers.*

▶ **Lemma 2.** *Prompt deterministic truthful mechanisms for the problem with public graph edges satisfy the critical item property.*

Proof. For the sake of contradiction, assume that there is a buyer i who gets item j_1 at price p_1 if she reports a value β_1 and gets item j_2 at price p_2 if she reports a value β_2. Without loss of generality, let $j_1 < j_2$. By truthfulness, the mechanism must give item j_1 to buyer i if she reports a value $\geq p_1$ (as far as the mechanism knows, i might not like items after j_1, and she would have incentive to lie and report β_1 if she is not given j_1). Thus, we have $p_2 \leq \beta_2 < p_1$, where the first inequality comes from individual rationality. But now, buyer i has incentive to report β_2, in order to get j_2 and pay p_2 which is less than p_1. ◀

▶ **Theorem 3.** *Any prompt deterministic truthful mechanism for the problem with public graph edges has competitive ratio of at most $\nu = \min(m, n)$.*

Proof. Consider an instance with n buyers with value 1 that are all interested in the first item. If there is a buyer i who will never get item 1 no matter what she reports, then we change the instance so that i has an arbitrary large value and is only interested in item 1, in which case i will get nothing and the mechanism does not even approximate the optimal social welfare. Conversely, if there is no such buyer, then the critical item property states that no other item can be allocated, which gives an approximation ratio of $\min(m, n)$.   ◀

3.2   Randomized truthful mechanisms

Somewhat surprisingly, the previous result has revealed a large gap between tardy and prompt deterministic mechanisms, when the topology of the graph is public knowledge: while tardy mechanisms can be implemented *for free*, i.e., maintaining the efficiency guarantees of (non-strategic) combinatorial algorithms, for prompt mechanisms the story is different. After showing that deterministic mechanisms cannot achieve anything better than ν, we turn our focus towards impossibility results for randomized mechanisms. We utilize a well-known property of randomized truthful mechanisms, which (by definition) make truthful reports utility-maximizing for any outcome of a mechanism's random decisions, even in hindsight: this implies that they are lotteries over deterministic truthful mechanisms, which satisfy the properties shown in the previous section. By Yao's minimax principle [29], it is then enough to construct a distribution over instances, such that the optimal solutions have welfare $\Omega(\log n)$, and a best-possible deterministic mechanism $\mathcal{M}$, since it satisfies the critical item property, outputs solutions with expected value $\mathcal{O}(\log \log n)$.

▶ **Theorem 4.** *Any prompt randomized truthful mechanism for the problem with public graph edges has competitive ratio of at least $\Omega(\log \nu / \log \log \nu)$.*

Proof. Fix any prompt randomized ex-post truthful mechanism for public graph edges. We argue by Yao's principle [29] that its competitive ratio is at least $\Omega(\log \nu / \log \log \nu)$. This holds due to the upcoming Lemma 5, which shows that there exists a distribution over instances, such that the optimal solutions have welfare at least $n \log(n)/2$ with high probability, and such that any deterministic mechanism (since it satisfies the critical item property) outputs solutions with expected value $\mathcal{O}(n \log \log n)$. More precisely, given a random instance r and a mechanism $\mathcal{M}_s$ with random coin flips s, recall that Yao's principle states that:

$$\min_r \left[\frac{\mathbb{E}_s[\mathcal{M}_s(r)]}{\mathrm{OPT}(r)} \right] \leq \mathbb{E}_r \left[\frac{\mathbb{E}_s[\mathcal{M}_s(r)]}{\mathrm{OPT}(r)} \right] = \mathbb{E}_s \left[\mathbb{E}_r \left[\frac{\mathcal{M}_s(r)}{\mathrm{OPT}(r)} \right] \right] \leq \max_s \left[\mathbb{E}_r \left[\frac{\mathcal{M}_s(r)}{\mathrm{OPT}(r)} \right] \right]$$
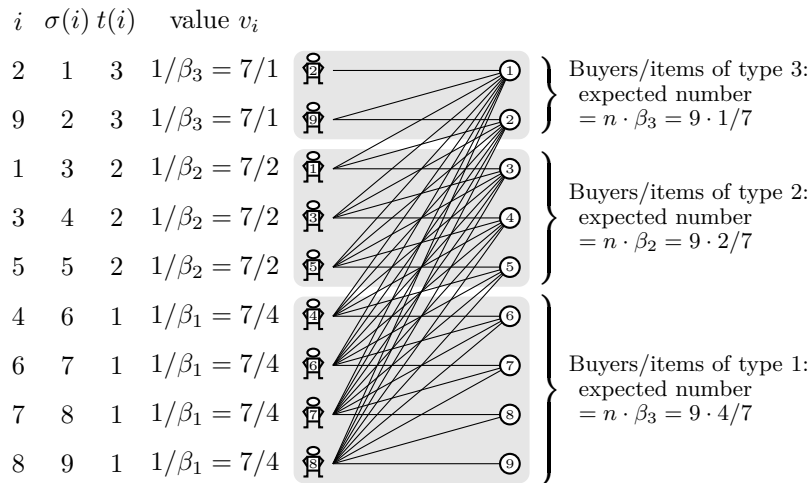
In particular, fixing the coin flips s, the mechanism $\mathcal{M}_s$ is deterministic and truthful. Hence, Lemma 5 bounds its expected approximation ratio over the random instance r, with

$$\mathbb{E}_r \left[\frac{\mathcal{M}_s(r)}{\mathrm{OPT}(r)} \right] \leq \mathbb{E}_r \left[\frac{\mathcal{M}_s(r)}{\log(n)/2} + \mathbb{1}_{\{\mathrm{OPT}(r) \leq \log(n)/2\}} \right] \leq \frac{\mathcal{O}(\log \log n)}{\log(n)/2} + \mathcal{O}(1/\log^2 n),$$

where the first inequality holds by the disjunction of whether or not $\mathrm{OPT}(r) \leq \log(n)/2$ for a given r. Combining the two inequalities concludes the proof.   ◀

▶ **Lemma 5.** *There is a distribution over instances with n buyers and n items, for which any deterministic mechanism satisfying the critical item property outputs solutions with expected value $\mathcal{O}(n \log \log n)$, and such that the optimal solution has value $\geq n \log(n)/2$ with probability at least $1 - O(1/\log^2 n)$.*

i	$\sigma(i)$	$t(i)$	value v_i
2	1	3	$1/\beta_3 = 7/1$
9	2	3	$1/\beta_3 = 7/1$
1	3	2	$1/\beta_2 = 7/2$
3	4	2	$1/\beta_2 = 7/2$
5	5	2	$1/\beta_2 = 7/2$
4	6	1	$1/\beta_1 = 7/4$
6	7	1	$1/\beta_1 = 7/4$
7	8	1	$1/\beta_1 = 7/4$
8	9	1	$1/\beta_1 = 7/4$

Buyers/items of type 3:
expected number
$= n \cdot \beta_3 = 9 \cdot 1/7$

Buyers/items of type 2:
expected number
$= n \cdot \beta_2 = 9 \cdot 2/7$

Buyers/items of type 1:
expected number
$= n \cdot \beta_3 = 9 \cdot 4/7$

■ **Figure 1** The instance from Lemma 5 with $k = 3$ and $n = 9$. Items are ordered (from top to bottom) according to their arrival times, and buyers are ordered (from top to bottom) according to σ (sort by decreasing types, breaking ties with indices). Preferences of buyers are given by the edges of the graph.

Proof. Let $k \geq 1$ be a parameter, which corresponds to the number of types of buyers, and let $\beta_1 > \cdots > \beta_k > 0$ be the probabilities of each type ($\beta_1 + \cdots + \beta_k = 1$). We choose $\beta_t = 2^{-t}/(1 - 2^{-k})$ for all t, and we set $n = 1 + 2^k$. Consider the following distribution over instances, with n buyers and n items. Each buyer i draws independently a type $t(i) \in \{1, \ldots, k\}$ with probability $\beta_{t(i)}$, and we set her value to $v_i = 1/\beta_{t(i)}$. Then, we sort buyers by decreasing $t(i)$, breaking ties using indices, and call $\sigma(i) \in \{1, \ldots, n\}$ the rank of buyer i in this ordering. We decide that buyer i is interested in all items up to the $\sigma(i)$-th one. To visualize this procedure, we refer to Figure 1. It is easy to find the optimal allocation: it consists in assigning each buyer of rank $\sigma(i)$ the $\sigma(i)$-th item, in a perfect matching. Thus the expected optimal social welfare is equal to

$$\mathbb{E}[\text{OPT}] = \sum_{i=1}^{n} \sum_{t=1}^{k} \beta_t \cdot 1/\beta_t = n \cdot k.$$

Moreover, because each type is drawn independently the variance of OPTis

$$\text{Var}(\text{OPT}) = \sum_{i=1}^{n} \text{Var}(v_i) \leq \sum_{i=1}^{n} \mathbb{E}[v_i^2] = n \cdot \sum_{t=1}^{k} \frac{1}{\beta_t} \leq 2n^2.$$

In particular, if we apply Chebyshev's inequality, we obtain

$$\mathbb{P}\left[\text{OPT} \leq \frac{nk}{2}\right] \leq \mathbb{P}\left[|\text{OPT} - nk| \geq \frac{nk}{2}\right] \leq \frac{\text{Var}(\text{OPT})}{(nk/2)^2} \leq \frac{8}{k^2}.$$

We now define the type $s(j) = t(\sigma^{-1}(j))$ of an item j as the type of the j-th buyer in the ordering σ, which corresponds to the type of its buyer in the abovementioned optimal matching. Observe that of each type, there are as many items as buyers, and that buyer i cannot be allocated an item j of type $s(j) < t(i)$. For each buyer i and for all types $t \leq s$, let $x_{s,t}^i$ be the probability (over the randomness of the types of all buyers except i) that i gets an item of type s, conditioning on the fact that i has type t. Let $x_{s,t} = \sum_i x_{s,t}^i/n$, that is, the average probability that a type t buyer will be assigned a type s item. The expected social welfare of our deterministic mechanism is equal to

$$\mathbb{E}\left[\mathrm{SW}(\mu)\right] = \sum_{i=1}^{n}\sum_{t=1}^{k}\beta_t \cdot 1/\beta_t \cdot \sum_{s=t}^{k}x_{s,t}^{i} = n\sum_{t=1}^{k}\sum_{s=t}^{k}x_{s,t}.$$

In expectation, the mechanism sells $\sum_i \sum_{t=1}^{s}\beta_t \cdot x_{s,t}^{i}$ items of type s. Because there are equally many items and buyers of each type, the expected number of items of type s is $\beta_s \cdot n$. Thus, we have the linear constraint

$$\forall 1 \leq s \leq k, \qquad \sum_{t=1}^{s}\beta_t \cdot x_{s,t} \leq \beta_s.$$

We are now going to use the critical item property. Fix a buyer i, and condition on the types of all buyers except her. We show that there exists an item $j(i) \in \{1, \ldots, n\}$, such that for every type $t(i)$, either i gets item $j(i)$, or she gets nothing. Denote as I_t the instance given by the fixed types of all buyers except i, together with buyer i who has type t. Using the critical item property with instance I_1, where i instead is of type 1 (meaning that i is interested in maximally many items), there is an item $j(i)$ such that buyer i either gets $j(i)$ or nothing. From the perspective of the mechanism, any other instance I_t (defined analogously) is identical to instance I_1 up to the point when i stops being interested in items. At this point, if buyer i has already been allocated an item, then it must be $j(i)$. Otherwise, she will not get anything.

Now that $j(i)$ is well-defined (and only depends on types of other buyers), let y_s^{i} be the probability (over the randomness of the types of all buyers except i) that there exists some type t such that if t is the type of i, then item $j(i)$ has type s. Let $y_s = \sum_i y_s^{i}/n$. Because buyer i can only get item $j(i)$, and because $j(i)$ is independent from $t(i)$, we have $x_{s,t}^{i} \leq y_s^{i}$. Thus, summing over all buyers, we have the linear constraint $x_{s,t} \leq y_s$, for all $1 \leq t \leq s \leq k$. Finally, conditioning on the types of all buyers except i, we show that there is only a small number of types that $j(i)$ can take. Recall that $s(j(i)) = t(\sigma^{-1}(j(i)))$, that is, the type of item $j(i)$ is by definition the type of the $j(i)$-th buyer in the ordering σ, where σ was obtained by sorting buyers in decreasing order of type. Consider the ordering induced by σ after excluding buyer i, and denote i_1 and i_2 the buyers of rank $j(i) - 1$ and $j(i)$. In the original ordering σ, either i comes before i_1 (in which case $s(j(i)) = t(i_1)$), or i comes after i_2 (in which case $s(j(i)) = t(i_2)$), or i comes between i_1 and i_2 (in which case $s(j(i)) = t(i)$). In any case, $t(i_1) \geq s(j(i)) \geq t(i_2)$. This shows that there are at most $2+z$ possible values for $s(j(i))$, where z denotes the number of types not seen among other buyers. By a standard computation, the expected value of z is smaller than $\sum_{t=1}^{k}(1 - \beta_t)^{n-1}$. Recall that y_s denotes the average probability over i that there exists a type for i which can make $j(i)$ have type s, where the randomness is over the instance without i. Since for every *fixed* such instance, $j(i)$ can only possibly take two of the types seen in buyers except i, for any fixed i, it holds that $\sum_{s=1}^{k}y_s^{i} \leq \alpha$, where $\alpha = 2 + \sum_{t=1}^{k}(1 - \beta_t)^{n-1}$, and therefore, the same holds also on average, i.e. for the y_s. Thus, averaging over possible types for the other buyers, and summing over i, we have the linear constraint $\sum_{s=1}^{k}y_s \leq \alpha$. If we choose $n = 1 + 2^k$ and $\beta_t = 2^{-t}/(1 - 2^{-k})$, we have

$$\sum_{t=1}^{k}(1 - \beta_t)^{n-1} \leq \sum_{t=1}^{k}e^{-2^{k-t}/(1-2^{-k})} \leq \sum_{t=0}^{+\infty}e^{-2^{t}} \leq 1,$$

and thus $\alpha \leq 3$. To conclude the proof, we use the linear constraints obtained to define a linear program (P) whose objective function is the expected value of the social welfare

$$\max \sum_{t=1}^{k} \sum_{s=t}^{k} x_{s,t} \qquad \text{(P)}$$

$$\text{s.t. } x_{s,t} \leq y_s$$
$$\sum_{t=1}^{s} \beta_t \cdot x_{s,t} \leq \beta_s$$
$$\sum_{s=1}^{k} y_s \leq \alpha$$
$$x_{s,t}, y_s \geq 0$$

$$\min \alpha \cdot w + \sum_{s=1}^{k} \beta_s \cdot v_s \qquad \text{(D)}$$

$$\text{s.t. } u_{s,t} + \beta_t \cdot v_s \geq 1$$
$$w \geq \sum_{t=1}^{s} u_{s,t}$$
$$u_{s,t}, v_s, w \geq 0$$

obtained by a deterministic truthful mechanism. We want to show that the objective function of our linear program is at most $\mathcal{O}(n \log k)$. To this end, Lemma 6 builds a solution for the dual linear program (D), whose value is an upper bound on the value of the primal linear program (for convenience, the objective function is divided by n). ◄

▶ **Lemma 6.** *Consider the linear program (P), parameterized by $\alpha > 0$ and $\beta_1 > \cdots > \beta_k > 0$. If $\beta_t = 2^{-t}/(1 - 2^{-k})$ for all $1 \leq t \leq k$, then the dual (D) has a feasible solution of value $\mathcal{O}(\alpha \log k)$.*

Proof. Set $\delta = \lceil \log_2 k \rceil$, then following solution of the dual is feasible and yields the desired objective value: $w = \delta$, $v_s = 0$ if $s < \delta$ and $2^{s-\delta}$ otherwise, while the $u_{s,t}$ are defined as:

$$\forall 1 \leq t \leq s \leq k, \quad u_{s,t} = \begin{cases} 1 & \text{if } s < \delta \\ 1 - 2^{s-\delta-t} & \text{if } 0 \leq s - \delta \leq t \\ 0 & \text{otherwise} \end{cases} \qquad ◄$$

4   Mechanisms with private graph edges

We move to the (harder) case where the graph edges are private information of the agents. The additional hardness, interestingly, severely affects the competitive guarantees for tardy truthful mechanisms. We begin by characterizing deterministic mechanisms, and then move on to results for randomized mechanisms.

4.1   Deterministic truthful mechanisms

In the previous section we assumed that the agents could not misreport their interest in items, thus reducing the problem to a single-parameter one. We now lift this assumption, and investigate the effect on the competitive ratio of deterministic truthful mechanisms. We show that deterministic truthful mechanisms can always be implemented in a prompt manner. Then, we give matching upper and lower bounds on the best approximation ratio for the social welfare.

▶ **Lemma 7.** *Tardy deterministic truthful mechanisms for the problem with private graph edges satisfy the critical item property (see Definition 1).*

Proof. For the sake of contradiction, assume that there is a buyer i who gets item j_1 at price p_1 if she reports a value β_1, and gets item j_2 at price p_2 if she reports a value β_2. Without loss of generality, we assume that $j_1 < j_2$. First, we argue that $p_1 = p_2$. Indeed, if $p_1 > p_2$ then i with value β_1 has incentive to lie and report β_2; whereas if $p_1 < p_2$ then i with value β_2 has incentive to lie and report β_1. Second, we slightly change the instance,

such that buyer i has value β_2 and is not interested in items after j_1. When allocating j_1, the mechanism has not seen any difference to the original instance, hence i has incentive to lie and report β_1 to get j_1, then lie and pretend she was interested in subsequent items to make sure she is charged p_1. ◀

▶ **Lemma 8.** *Tardy deterministic truthful mechanisms for the problem with private graph edges are prompt.*

Proof. Assume that our mechanism assigns an item j to buyer i, who reports value b_i. By Lemma 7, the mechanism satisfies the critical item property, and j is the only item which can be assigned to i. Let π be the minimum value that i could have reported and still be assigned j. By truthfulness, i must be charged exactly π. Indeed, if she is charged $p > \pi$ then i with value b_i has incentive to lie and report π; whereas if she is charged $p < \pi$ then i with value p would have incentives to lie and report b_i. Now, when the mechanism assigns j to i, it can retrospectively compute π, which proves that the mechanism is prompt. ◀

▶ **Theorem 9.** *There exists a deterministic truthful mechanism that achieves an $\nu = \min(m, n)$ approximation of the offline optimum. This result is tight in the class of deterministic truthful mechanisms, when graph edges are private.*

Proof. Consider the simple mechanism which only assigns an item to a buyer if she has the highest value seen so far (breaking ties arbitrarily), charging her the second highest value seen so far. This is a ν-competitive deterministic truthful mechanism. For the tightness, Lemma 8 shows that deterministic tardy mechanisms are in fact prompt, thus the lower bound from Theorem 4 (public graph edges) applies to this setting. ◀

5 Randomized truthful mechanisms

Recall that randomized (ex-post) truthful mechanisms are lotteries over deterministic truthful mechanisms, which in turn satisfy the characterizing properties we obtain for the deterministic case. The proof of our lower bound in Theorem 4 was based on this fact. This same argument also applies to mechanisms for private edges, even when they are tardy. On the positive side, we construct a prompt randomized truthful mechanism, the EXPLORE-EXPLOIT MECHANISM, that yields a logarithmic approximation. The EXPLORE-EXPLOIT MECHANISM divides the buyers into two types: "explore" buyers will not receive any item but are used to set the price for the "exploit" buyers. To guarantee truthfulness, we enforce monotonicity of the prices proposed by the seller during the routine: with prices always increasing, there is no way a buyer can benefit from withholding information in previous stages of the process to get something at a cheaper price later.

▶ **Theorem 10.** *The EXPLORE-EXPLOIT MECHANISM is truthful, and computes a $O(\log n)$ approximation to the optimal social welfare. This result is nearly tight (up to $\log \log n$) in the class of randomized truthful mechanisms when the edges are private information, even for tardy mechanisms.*

Proof. Buyers of type *Explore* will not get any item, and thus have no incentive to lie. Buyers of type *Exploit* only need to say if they are interested to buy an item at a given price. Because prices are non-decreasing, they have no incentive to misreport their value or their interest in an item. For each item j, we define x_j as the maximum value seen among buyers interested in items up to j.

$$\forall j \in I, \qquad x_j = \max\{v_i \text{ with } i \in B \text{ such that } \exists j' \leq j, (i, j') \in E\}$$

◾ **Algorithm 1** EXPLORE-EXPLOIT MECHANISM.

1: **Initialization:**
2:     Set $p \leftarrow 0$ and draw $k \leftarrow \text{Unif}(\{0, 1, \dots, \lceil \log_2 n \rceil\})$
3:     For each buyer i, draw type $t_i \leftarrow \text{Unif}(\{\textit{Explore}, \textit{Exploit}\})$.
4: **When an item arrives:**
5:     Buyers report if they are interested in the item.
6:     **For** each buyer i of type $t_i = \textit{Explore}$ who is interested in the item**, do**
7:         Set $p \leftarrow \max(p, v_i/2^k)$
8:     Sell the item at price p to a buyer i of type $t_i = \textit{Exploit}$, who is interested
9:     in the item and does not yet have an item, chosen arbitrarily (e.g. lowest index).

For the sake of analysis, we look at a maximum weight matching $\mu \subseteq E$, having a total value of OPT. Each edge $(i, j) \in \mu$ from the optimal solution is assigned to a bucket $\ell_{(i,j)} = \lceil \log_2(x_j/v_i) \rceil \in \mathbb{N}$. Then for each $\ell \in \mathbb{N}$ we define OPT_ℓ as the total weight of the restriction of the optimal solution to bucket ℓ.

$$\text{OPT} = \sum_{\ell \geq 0} \text{OPT}_\ell \qquad \text{where } \forall \ell \geq 0, \quad \text{OPT}_\ell = \sum_{(i,j) \in \mu} v_i \cdot \mathbb{1}_{\{\ell_{(i,j)} = \ell\}}$$

Let V be maximum value among buyers who are interested in at least one item. By optimality of μ, the corresponding buyer must be given an item, and thus $\text{OPT}_0 \geq V$. Now observe that for each $(i, j) \in \mu$ such that $\ell_{(i,j)} > \lceil \log_2 n \rceil$, we have $v_i < x_j/n \leq V/n \leq \text{OPT}_0/n$. Thus, the sum of OPT_ℓ for $\ell > \lceil \log_2 n \rceil$ is smaller than OPT_0. Therefore, buckets $0, 1, \dots, \lceil \log_2 n \rceil$ contain at least half of OPT, that is

$$\frac{\text{OPT}}{2} \leq \sum_{\ell=0}^{\lceil \log_2 n \rceil} \text{OPT}_\ell$$

For all $\ell \in \{0, 1 \dots, \lceil \log_2 n \rceil\}$, we will now show that if $k = \ell$ then the EXPLORE-EXPLOIT MECHANISM gives a solution of expected cost at least $\Omega(\text{OPT}_\ell)$. Then we will conclude the proof using the law of total probability: summing over k shows that the EXPLORE-EXPLOIT MECHANISM computes a solution of expected cost at least $\Omega(\text{OPT}/\log n)$. First, assume that $k = 0$. For each edge $(i, j) \in \mu$ in bucket $\ell_{(i,j)} = 0$, then i is the best buyer seen by the time j arrives. With probability $1/4$, buyer i has type *Exploit* and the second best buyer has type *Explore*. In that case, the EXPLORE-EXPLOIT MECHANISM gives buyer i an item (either j or one of the previous items). Using linearity of expectation, the EXPLORE-EXPLOIT MECHANISM outputs a solution of expected value at least $\text{OPT}_0/4$. Second, assume that $k = \ell$ with $\ell \in \{1, \dots, \lceil \log_2 n \rceil\}$. This case requires an amortized analysis: for each buyer i, denote X_i the random variable equal to v_i if i gets an item and 0 otherwise; and for each item j, denote Y_j the random variable equal to the value of the buyer to whom j is assigned, and 0 if j is unassigned. Notice that the EXPLORE-EXPLOIT MECHANISM outputs a solution of value $= \sum_{i \in B} X_i = \sum_{j \in I} Y_j$. Let $(i, j) \in \mu$ be an edge from bucket $\ell_{(i,j)} = \ell$. We are going to show that

$$\mathbb{E}\left[X_i + 4Y_j \mid k = \ell \text{ and } t_i = \textit{Exploit}\right] \geq v_i.$$

We condition on the fact that $k = \ell$ and $t_i = \textit{Exploit}$. If buyer i already has an item when item j arrives, then $X_i = v_i$. Otherwise, the best buyer seen so far has type *Explore* with probability $1/2$, in which case the EXPLORE-EXPLOIT MECHANISM gives item j to a

buyer of value $\geq x_j/2^\ell \geq v_i/2$. Buyer i has type $t_i = \text{\textit{Exploit}}$ with probability $1/2$, thus $v_i \leq E[2X_i + 8Y_j \mid k = \ell]$. Summing this last inequality over edges from bucket ℓ shows that the EXPLORE-EXPLOIT MECHANISM outputs a solution of expected value at least $\text{OPT}_\ell/10$.

Let's move our attention to the lower bound . Fix all random decisions of an ex-post truthful randomized mechanism. This yields a deterministic algorithm, that together with the original mechanism's payment scheme yields a (tardy) mechanism. This mechanism is deterministic, and truthful due to the definition of truthfulness. Also, such a mechanism fulfills the critical item property (Lemma 7), and can even be made prompt (Lemma 8). With this, we can follow the original proof of the lower bound. ◀

6 Ex-ante truthfulness

One might wonder if the hardness of truthful mechanisms for our problem is mainly due to the very restrictive notion of ex-post truthfulness. We state here that also for the much looser ex-ante truthfulness, the setting of non-myopic buyers separates clearly from the myopic case. The proof is via a nontrivial construction allowing bounds on agents' expected utilities.

▶ **Theorem 11.** *There exists no randomized ex-ante truthful mechanism that yields an α-approximation to the optimal social welfare, for the problem with private edges and any $\alpha < 2$. This is true even for tardy mechanisms.*

Proof. Fix $\alpha < 2$ and assume mechanism M guarantees an expected approximation ratio of α. Consider the following problem instance: there are n' buyers and $m = n' + 1$ items. Every item j has exactly one interested buyer, i_j, and all i_j have some small value $v_{i_j} = \epsilon > 0$. There exist some additional buyers $B_1 \subseteq B$ with different values who are interested only in item 1, and one buyer, i, whom we fix for our considerations. Note that $|B| = n' + n_1$, with $n_1 = |B_1|$. For n' large enough, clearly, $n'\epsilon > \max_{i' \in B_1} v_{i'}$ and the contribution of item 1 to the optimum becomes negligible with growing n'. Therefore, for M to guarantee an α-approximation, there must exist $j \in \{2, \ldots, n' + 1\}$ such that i_j is assigned the according item with probability at least $\frac{1}{\alpha}$, or in case item 1 is worth more than ϵ, at least probability $\frac{1}{\alpha} - \Delta_1$, where Δ_1 arbitrarily small for large n'.

Now, if we choose $i = i_j$, then M will assign item j to i_j w.pr. $\geq \frac{1}{\alpha} - \Delta_1$, and charge an expected price of at most ϵ. The latter is because the price cannot depend on i's bid due to incentive compatibility, and it needs to be below i's value. Assume we replace i's valuation by some $v > \epsilon$, and call this new buyer $i^{(1)}$. Since M is ex-ante truthful, still, the exp. utility $u_{i^{(1)}}$ achieved with a truthful report must be at least as large as when reporting ϵ instead of v, i.e. at least $(v - \epsilon)(\frac{1}{\alpha} - \Delta_1) > \frac{1}{2}v$, which is at least half of v because α is < 2 and ϵ, Δ_1 can be chosen arbitrarily small. We replace $i^{(1)}$ again by a different buyer $i = i^{(2)}$. She still has valuation v, however, she is now interested in items 1 *and* j. We consider the first step of M, i.e. the assignment decision made for item 1. Assuming that v is the largest value bid on item 1, and given the fact that M has no idea if any additional value will present itself in the later steps, the probability that M assigns item 1 to $i^{(2)}$ is at least $\frac{1}{\alpha} - \Delta_2$, where Δ_2 approaches 0 since the other bids on item 1 might be, in comparison, too small to matter. Note again that the assignment decision cannot depend on v itself, but only on the fact that it is the largest value bid on item 1.

We know that $i^{(2)}$ can get utility larger than $\frac{v}{2}$ by simply reporting type $i^{(1)}$ instead. We also know that since she is assigned item 1 w.pr. $> \frac{1}{2}$, she is assigned item j w.pr. $< \frac{1}{2}$. This, intuitively, means that not all of the guaranteed utility is generated by item j, not even if the price of j is always 0 – but some must be generated because her expected price paid when

item 1 is assigned is bounded away from v, i.e. $p_{i^{(2)}}(1) = v - \Delta_3$. In fact, the exp. price M charges from $i^{(2)}$ when assigning item 1 cannot be smaller if $i^{(2)}$ later reports interest in item j, since this would give a buyer of type $i^{(1)}$ incentive to also report interest in j. Also, the price charged from $i^{(2)}$ when assigning item j cannot be less than 0, and when there is no item assigned, $i^{(2)}$ is not charged anything (see preliminaries). This implies that, for $P_k(i)$ denoting the assignment probability of item k to buyer i,

$$u_{i^{(2)}} = (v - p_{i^{(2)}}(1)) \cdot P_1(i^{(2)}) + (v - p_{i^{(2)}}(j)) \cdot P_j(i^{(2)})$$
$$= \Delta_3 \cdot P_1(i^{(2)}) + (v - p_{i^{(2)}}(j)) \cdot P_j(i^{(2)}) > \frac{v}{2}$$

Otherwise, we would have a contradiction on the utility being larger than $\frac{v}{2}$, i.e. it would be beneficial for $i^{(2)}$ to only report interest in item j. In consequence, it also holds

$$u_{i^{(2)}} = \Delta_3 \cdot P_1(i^{(2)}) + (v - p_{i^{(2)}}(j)) \cdot P_j(i^{(2)}) \geq \Delta_3 \cdot P_1(i^{(2)}) + (v - v) \cdot P_j(i^{(2)}) > 0.$$

This is true because the exp. price when receiving item j can be no more than v, and $P_j(i^{(2)}) < \frac{1}{2}$. Therefore, there exists some $v^- < v$ for which it holds that

$$u_{i^-}(1) = u_{i^{(2)}}(1) - P_1(i^{(2)})(v - v^-) = (\Delta_3 - (v - v^-))P_1(i^{(2)})$$

Here, $u_{i^-}(1)$ denotes the utility obtained from being assigned item 1 of some buyer with valuation v^- for item 1, and 0 otherwise, when she reports $i^{(2)}$ as her type. Note that if buyer i^- reports value v for item 1 and 0 for all others, she will also obtain $u_{i^-}(1)$ from being assigned the first item: the assignment decision is made before the algorithm can know the difference, and the expected price paid cannot depend on the buyer's later reports due to truthfulness.

We use this to show a contradiction to the approximation ratio of M. Assume there exists, in absence of $i^{(2)}$, such a buyer i^- with smaller value v^- and utility of $u^-(1) > 0$ when reporting to have value v, who is interested in purchasing item 1, i.e. $i^- \in B_1$. Since M is ex-ante truthful, a truthful report for her will also result in positive expected utility of at least $u^-(1)$. As a direct consequence, it holds also that the probability $P_1(i^-)$ for assigning item 1 to i^- (when she reports truthfully) is lower bounded, in order to achieve above expected utility, as follows: $P_1(i^-) \geq \frac{u_{i^-}(1)}{v^-}$. Finally, we copy buyer i^- at least $\frac{v^-}{u_{i^-}(1)} + 1$ times. If necessary for tie-breaking, we distort their values a bit. Our conclusions about $i^{(2)}$'s utility hold once $i^{(2)}$ reports the largest value for item 1, regardless of other values. This means, if either of our copied v^- should decide to deviate and report to be valued like $i^{(2)}$ instead, they can recover utility $u_{i^-}(1)$. As a result, each one of the copies, when reporting truthfully, has at least the same utility, and therefore an assignment probability of at least $P_1(i^-)$. This, in sum, results in a probability of more than 1 for assigning item 1, i.e., a contradiction. ◄

7   Conclusions

We have studied vertex-weighted bipartite online matching with offline agents in various settings, obtaining an almost-complete picture of the competitive ratios achievable by mechanisms under different truthfulness notions. Our results encompass that for myopic truthfulness, the best algorithmic results [18, 1] transfer to the online agents setting. This showcases that the very general myopic bounds of [9] are far from tight for restricted settings like ours. On the other hand, we also show that equally near-optimal approximations are impossible under the assumption of classic truthfulness, even ex-ante; and for ex-post

truthfulness our seemingly simple problem already exhibits lower bounds almost matching the myopic, logarithmic competitive ratio for submodular combinatorial auctions in [9]. We leave open to what extent this additional hardness (moving from a tight $e/(e-1)$ myopic to $\Omega(\log n / \log \log n)$ truthful) already happens when imposing ex-ante truthfulness. This is an interesting subject of investigation, also for different scenarios than the one of our ≥ 2 lower bound (private edges).

--- **References** ---

1 Gagan Aggarwal, Gagan Goel, Chinmay Karande, and Aranyak Mehta. Online vertex-weighted bipartite matching and single-bid budgeted allocations. In *SODA*, pages 1253–1264. SIAM, 2011.

2 Susan Athey and Ilya Segal. An efficient dynamic mechanism. *Econometrica*, 81(6):2463–2485, 2013.

3 Yossi Azar and Ety Khaitzin. Prompt mechanism for ad placement over time. In *International Symposium on Algorithmic Game Theory SAGT*, pages 19–30. Springer, 2011.

4 Moshe Babaioff, Liad Blumrosen, and Aaron Roth. Auctions with online supply. *Games Econ. Behav.*, 90:227–246, 2015.

5 Dirk Bergemann and Juuso Välimäki. Dynamic mechanism design: An introduction. *Journal of Economic Literature*, 57(2):235–74, June 2019.

6 Benjamin Birnbaum and Claire Mathieu. On-line bipartite matching made simple. *Acm Sigact News*, 39(1):80–87, 2008.

7 Constantine Caramanis, Paul Dütting, Matthew Faw, Federico Fusco, Philip Lazos, Stefano Leonardi, Orestis Papadigenopoulos, Emmanouil Pountourakis, and Rebecca Reiffenhäuser. Single-sample prophet inequalities via greedy-ordered selection. In *SODA*, pages 1298–1325. SIAM, 2022.

8 Richard Cole, Shahar Dobzinski, and Lisa Fleischer. Prompt mechanisms for online auctions. In *International Symposium on Algorithmic Game Theory*, pages 170–181. Springer, 2008.

9 Yuan Deng, Debmalya Panigrahi, and Hanrui Zhang. Online combinatorial auctions. In *SODA*, pages 1131–1149. SIAM, 2021.

10 Nikhil R. Devanur, Kamal Jain, and Robert D. Kleinberg. Randomized primal-dual analysis of RANKING for online bipartite matching. In *SODA*, pages 101–107. SIAM, 2013.

11 Paul Dütting, Federico Fusco, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser. Efficient two-sided markets with limited information. In *STOC*, pages 1452–1465. ACM, 2021.

12 Alon Eden, Michal Feldman, Amos Fiat, and Kineret Segal. An economics-based analysis of RANKING for online bipartite matching. In *SOSA*, pages 107–110. SIAM, 2021.

13 Tomer Ezra, Michal Feldman, Nick Gravin, and Zhihao Gavin Tang. Online stochastic max-weight matching: Prophet inequality for vertex and edge arrival models. In *EC*, pages 769–787. ACM, 2020.

14 Uriel Feige. Tighter bounds for online bipartite matching. In *Building Bridges II*, pages 235–255. Springer, 2019.

15 Buddhima Gamlath, Sagar Kale, and Ola Svensson. Beating greedy for stochastic bipartite matching. In *SODA*, pages 2841–2854. SIAM, 2019.

16 Buddhima Gamlath, Michael Kapralov, Andreas Maggiori, Ola Svensson, and David Wajc. Online matching with general arrivals. In *FOCS*, pages 26–37. IEEE Computer Society, 2019.

17 Nick Gravin, Zhihao Gavin Tang, and Kangning Wang. Online stochastic matching with edge arrivals. In *ICALP*, volume 198 of *LIPIcs*, pages 74:1–74:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021.

18 Richard M. Karp, Umesh V. Vazirani, and Vijay V. Vazirani. An optimal algorithm for on-line bipartite matching. In *STOC*, pages 352–358. ACM, 1990.

19 Thomas Kesselheim, Klaus Radke, Andreas Tönnis, and Berthold Vöcking. An optimal online algorithm for weighted bipartite matching and extensions to combinatorial auctions. In *ESA*, volume 8125 of *Lecture Notes in Computer Science*, pages 589–600. Springer, 2013.

20   Nitish Korula and Martin Pál. Algorithms for secretary problems on graphs and hypergraphs.
      In *ICALP (2)*, volume 5556 of *Lecture Notes in Computer Science*, pages 508–520. Springer,
      2009.

21   Piotr Krysta and Berthold Vöcking. Online mechanism design (randomized rounding on
      the fly). In *ICALP (2)*, volume 7392 of *Lecture Notes in Computer Science*, pages 636–647.
      Springer, 2012.

22   Aranyak Mehta. Online matching and ad allocation. *Found. Trends Theor. Comput. Sci.*,
      8(4):265–368, 2013.

23   Vahab S. Mirrokni, Renato Paes Leme, Pingzhong Tang, and Song Zuo. Non-clairvoyant
      dynamic mechanism design. In *EC*, page 169. ACM, 2018.

24   Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73,
      1981.

25   Rebecca Reiffenhäuser. An optimal truthful mechanism for the online weighted bipartite
      matching problem. In *SODA*, pages 1982–1993. SIAM, 2019.

26   Tim Roughgarden. *Twenty lectures on algorithmic game theory*. Cambridge University Press,
      2016.

27   William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal
      of finance*, 16(1):8–37, 1961.

28   Xiangzhong Xiang. Prompt mechanism for online auctions with multi-unit demands. *Journal
      of Combinatorial Optimization*, 30(2):335–346, 2015.

29   Andrew Chi-Chih Yao. Probabilistic computations: Toward a unified measure of complexity
      (extended abstract). In *FOCS*, pages 222–227. IEEE Computer Society, 1977.

A   Mechanisms for Myopic Buyers

In this section we study myopic buyers. We show that for this class of agents, one can obtain
strategy proof versions of the best (non-truthful) algorithms [18]. Formally, we construct
a deterministic prompt mechanism that achieves at least half of the welfare of the best
offline matching, and a randomized prompt mechanism that is (in expectation) $e/(e-1)$-
competitive with the best offline matching. We start describing our deterministic mechanism
HONESTGREEDY, which mimics the classical GREEDY algorithm for online weighted matching
in a way that is robust to strategic bidding. Every time a new item arrives, HONESTGREEDY
runs a second price auction [27] to allocate it between the remaining (interested) buyers.
Since the buyers are myopic, every time a new item arrives, they behave like if it was the last:
clearly there is no point in lying about being interested in an item. Moreover, the truthfulness
in each step (as well as the individual rationality) is guaranteed by the well-known properties
of the second price auction. Note that the mechanism sets the price for item i immediately, so
it is prompt. The analysis of the approximation guarantee is also quite simple: the allocation
output by HONESTGREEDY is the same one that the standard GREEDY algorithm would have
computed on the same input. It is well known that GREEDY is 2-competitive with respect to
the best offline matching (see, e.g., Appendix B of [1]), and that this approximation is tight
in the class of deterministic algorithms [18]. We summarize these observations.

▶ **Theorem 12.** *The deterministic prompt mechanism HONESTGREEDY is truthful for myopic
agents and guarantees a 2 approximation to the best offline matching. The approximation is
tight even for (non-truthful) deterministic algorithms.*

We complement this deterministic 2-competitive, simple mechanism with an optimal,
randomized $e/(e-1)-$competitive alternative, HONESTPERTURBEDGREEDY, based on
PERTURBED-GREEDY of Aggarwal et al. [1]. There, each offline vertex is associated with a
random multiplier; then, every time one of the online vertices arrives, it is matched to the

▢ **Algorithm 2** HONESTPERTURBEDGREEDY.

1: **For** each buyer i, **do**
2:      Draw x_i uniformly at random from $[0, 1]$
3:      Let $y_i = 1 - e^{x_i - 1}$
4: Reveal publicly all x_i and y_i
5: **For** item j arriving online, **do**
6:      Receive bids for j and let $N(j)$ be the set of agents interested in j
7:      Allocate j to $i^\star \in \arg\max\{b_i \cdot y_i \mid i \in N(j)\}$ ▷ Allocation Rule
8:      Charge $i^\star$ with $p_{i^\star} = \max\left\{\frac{y_i}{y_{i^\star}} b_i \mid i \in N(j) \setminus \{i^\star\}\right\}$ ▷ Payment Rule
9:      Discard for further consideration $i^\star$

---

free neighbor with largest multiplier-value product. To protect from the strategic behavior of agents, HONESTPERTURBEDGREEDY declares – before the beginning of the online phase – publicly all random multipliers, and then implements Myerson's payment rule [24] for every round. For a formal description we refer to the pseudocode of HONESTPERTURBEDGREEDY, where we maintain the convention that the max of an empty set is 0 and thus if $N(j)$ is empty in line 7, then j is discarded and the mechanism passes to the next item. The properties of HONESTPERTURBEDGREEDY are summarized in the following Theorem, whose formal proof is deferred to Appendix C.

▶ **Theorem 13.** *The randomized prompt mechanism HONESTPERTURBEDGREEDY is truthful for myopic agents and achieves (in expectation) a $e/(e-1)$ approximation to the best offline matching. The approximation is tight even for (non-truthful) randomized algorithms.*

▢ B   Tardy mechanisms with public graph edges

When the graph edges are public knowledge, we can turn once again to using the algorithmic approaches outlined in the previous Section, i.e. GREEDY and PERTURBED-GREEDY. Now that agents cannot strategically withhold or misreport the existence of edges, a tardy truthful mechanism can use the whole graph structure (but of course still not the reported value b_i) when computing the price charged from any buyer i. The prompt, round-wise payment rules we designed for myopic buyers, however, do not guarantee non-myopic truthfulness. What remains to prove therefore is that these algorithms can be augmented by a different (tardy) payment rule to be made truthful. This is formally done in two steps: first, it is established that the allocations produced are monotone, and then Myerson's Lemma is employed to enforce truthfulness. All in all, we have the following.

▶ **Theorem 14.** *There exists a deterministic, respectively randomized, tardy mechanism that is truthful for non-myopic agents with public graph edges and guarantees a 2, resp. $e/(e-1)$, approximation to the best offline matching. The approximation is tight even for (non-truthful) deterministic, resp. randomized, algorithms.*

Note that the allocation computed by the mechanisms we just described are analogous to the ones computed by HONESTGREEDY and HONESTPERTURBEDGREEDY, but the payments are different! We are still using Myerson's Lemma, but the critical prices[2] are clearly different, as they are computed considering the whole run of the algorithm. To see this, consider the following example. There are two buyers, b_1 and b_2, and two items i_1 and i_2. b_1 is interested

---

[2] The critical price is the smallest bid that would have still resulted in the item being allocated to the agent. See the Appendix C for a formal definition

in both the items and has a value of 1, while b_2 only cares about i_1, with a value of 0.9. Assume also for the sake of simplicy that the perturbations y_1 and y_2 of PERTURBED-GREEDY are both 1. Both versions of PERTURBED-GREEDY would only allocate i_1 to b_1, but at two different prices: the mechanism for myopic agents would charge 0.9, while the tardy one for non-myopic agents would wait until the end of the second round and charge 0.

C   Proofs of Theorems 13 and 14

In this section we prove the properties of HONESTPERTURBEDGREEDY and of the tardy versions of GREEDY and PERTURBED-GREEDY presented in Appendices A and B. Starting from the guarantees of their non-strategic counterparts it is immediate to see that the approximation factor claimed are indeed correct. The only property to show is incentive compatibility. A crucial ingredient to prove incentive compatibility is Myerson's Lemma, that we recall here for the sake of completeness. The Lemma has been proved in Myerson's seminal paper [24]; here we follow the more modern approach by Roughgarden [26]. Since in this paper we study unit-demand agents, we restrict to consider only this type of agents.

We start introducing the notion of single-parameter environments. In such environments, there are n agents and a set X of feasible allocations of items to agents. Each agents is characterized by a private valuation to get an item and strives to maximize her quasi-linear utility. To familiarize with this notion consider the model of non-myopic buyers with public graph edges studied in the paper: those agents are indeed single-parameters, as their valuations is their only private information. At the same time, note that the "edge compatibility" is implicitly modeled by the following set of feasible allocations of items to agents: an allocation $\mathbf{x} \in \{0,1\}^n$ is feasible if and only if it is corresponds with a matching in the underlying buyers-items bipartite graph. As already mentioned in the main body, a mechanism $\mathcal{M}$ is characterized by two features: an allocation $\mathbf{x} \in X$ and a payment rule $\mathbf{p}$. While the allocation specifies who gets what, the payment rule defines how much each agent pays. Allocation and payments are functions of the bids; in particular, we use the notation $x_i(b_i, b_{-i}) \in \{0,1\}$ to specify whether the i^{th} agent is allocated an item, given her bid b_i and the $n-1$ bids b_{-i} of the other agents. We are ready for the following crucial definitions.

▶ **Definition 15** (Monotone allocation). *An allocation rule $\mathbf{x}$ for a single-parameter environment is monotone if for every bidder i and bids b_{-i} by the other bidders, the allocation $x_i(z, b_{-i})$ to i is nondecreasing in its bid z.*

▶ **Definition 16** (critical prices). *Fix and agent i and bids b_{-i} of the other agents. Then the critical price for i is defined as the smallest bid z_i such that i is allocated an item, if any. Formally, if we use the convention that the* inf *of an empty set is* 0*, we have $z_i = \inf\{z \mid x_i(z, b_{-i}) = 1\}$*

Clearly, the critical prices enforce ex-post individual rationality. Myerson showed that they also induce (ex-post) truthfulness; we report here a version of Lemma 2 of Myerson [24] that is tailored to our problem and then show the two Theorems.

▶ **Theorem 17** (Myerson's Lemma). *Fix a single-parameter mechanism. Given any monotone allocation $\mathbf{x}$, it is possible to compute a payment scheme $\mathbf{p}$ such that the resulting mechanism is truthful and individually rational. In particular, in $\mathbf{p}$, each agents that receives an item pays its critical price and 0 otherwise.*

▶ **Theorem 13.** *The randomized prompt mechanism HONESTPERTURBEDGREEDY is truthful for myopic agents and achieves (in expectation) a $e/(e-1)$ approximation to the best offline matching. The approximation is tight even for (non-truthful) randomized algorithms.*

Proof. We start the proof by arguing that HONESTPERTURBEDGREEDY is truthful and individually rational for myopic agents. First, note that when any item j arrives, there is no point for the buyers still unallocated to lie about their interest for it: if they are not interested and they bid, they would risk to get j and lose future opportunity to get allocated to something they are interested in, while if they are interested they do not want to lose the opportunity (since they have no information on the future, and the prices charged never exceed their valuations). If we restrict to consider the buyers $N(j)$ interested in item j, we see that the problem reduces to a single-parameter auction: the agents are myopic and just want to maximize their utility by getting j at a small price. All y_i are public knowledge and non-negative, so our allocation rule (line 7 of HONESTPERTURBEDGREEDY), fixing these values, is clearly monotone (the more an agent i bids, the more likely she is to exhibit the largest $y_i \cdot b_i$). The allocation is therefore implementable using the Myerson payment rule (line 8 of HONESTPERTURBEDGREEDY). We can conclude, by Myerson's Lemma, that our mechanism is truthful for myopic buyers. Moreover, it is easy to verify that the payment rule enforces individual rationality. Once we have settled the truthfulness, we can assume that all buyers declare their true bids and thus the allocation output by HONESTPERTURBEDGREEDY is the same as PERTURBED-GREEDY *for any realization* of the perturbations x_i and inherits the same approximation: HONESTPERTURBEDGREEDY is $e/(e-1)$-competitive in expectation.   ◀

▶ **Theorem 14.** *There exists a deterministic, respectively randomized, tardy mechanism that is truthful for non-myopic agents with public graph edges and guarantees a 2, resp. $e/(e-1)$, approximation to the best offline matching. The approximation is tight even for (non-truthful) deterministic, resp. randomized, algorithms.*

Proof. It is easy to see how the two mechanisms are monotone, thus it is possible to employ directly Myerson's Lemma, as the problem is single-parameter (i.e., the only private information of buyer i is the single value v_i). Therefore, GREEDY or PERTURBED-GREEDY (with fixed perturbation factors) together with the critical payments defined in Myerson's Lemma result in a truthful mechanism. Note that the greedy algorithm clearly respects our ex-post notion of truthfulness, since no randomization is involved. For the PERTURBED-GREEDY algorithm, this is also true since we fix all random decisions (perturbation) up front, and choose the payment rule accordingly.   ◀