

Prophet Inequalities for Matching with a Single Sample ^{*}

Paul Dütting[†] Federico Fusco[‡] Philip Lazos[‡] Stefano Leonardi[‡]
Rebecca Reiffenhäuser[‡]

Abstract

We consider the prophet inequality problem for (not necessarily bipartite) matching problems with independent edge values, under both edge arrivals and vertex arrivals. We show constant-factor prophet inequalities for the case where the online algorithm has only limited access to the value distributions through samples. First, we give a 16-approximate prophet inequality for matching in general graphs under edge arrivals that uses only a single sample from each value distribution as prior information. Then, for bipartite matching and (one-sided) vertex arrivals, we show an improved bound of 8 that also uses just a single sample from each distribution. Finally, we show how to turn our 16-approximate single-sample prophet inequality into a truthful single-sample mechanism for online bipartite matching with vertex arrivals.

1 Introduction

Max-Weight Matching. We study the maximum-weight matching problem, i.e., maximizing the cumulative weight of an edge set M in a graph such that to each vertex, at most one edge from M is incident. This problem class comes up naturally in a wide range of applications, like matchmaking, two-sided markets, assigning jobs to workers, kidney exchange, housing markets, and the like. In many of these applications, the problem is *online* in that the input is not fully known from the outset. Rather it is revealed over time, and decisions have to be made immediately and irrevocably. The online matching problem comes in many different flavors, for example one can study it for restricted graph classes like bipartite graphs (e.g., jobs/workers), under different online models (e.g., a fixed set of houses for which potential buyers become known only over time), and different assumptions regarding the arrival model (random/adversarial/with expiration times etc.).

Prophet Inequalities. We concentrate here on the *prophet* model, generalizing [Krengel and Sucheston, 1977, 1978, Samuel-Cahn, 1984], in which elements arrive in adversarial order and weights are drawn independently from known distributions. In the original definition, the adversary fixes the order and distributions up front. The goal is to show a *prophet inequality*, i.e. an online algorithm such that the expected weight of the computed solution is at least a $1/\alpha$ -fraction of the expected weight of the optimal solution.

The original work of Krengel and Sucheston [1977, 1978] and Samuel-Cahn [1984] has established a tight 2-approximate prophet inequality for the problem of selecting one value out of a sequence of

^{*}This work was supported by the ERC Advanced Grant 788893 AMDROMA “Algorithmic and Mechanism Design Research in Online Markets”, the MIUR PRIN project ALGADIMAR “Algorithms, Games, and Digital Markets”.

[†]Google Research, Zürich, Switzerland. Email: duetting@google.com

[‡]Department of Computer, Control, and Management Engineering “Antonio Ruberti”, Sapienza University of Rome, Italy. Email: {fuscof, lazos, leonardi, rebeccar}@diag.uniroma1.it

n values drawn independently from not necessarily identical distributions. Motivated, in part, by a connection to posted-price mechanisms [Hajiaghayi et al., 2007, Chawla et al., 2010], recent work has established prophet inequalities for a broad range of combinatorial problems, including matroids and polymatroids [e.g., Alaei, 2011, Kleinberg and Weinberg, 2012, Dütting and Kleinberg, 2015, Feldman et al., 2016], bipartite and non-bipartite matching [e.g., Chawla et al., 2010, Kleinberg and Weinberg, 2012, Feldman et al., 2015, Gravin and Wang, 2019, Ezra et al., 2020], and combinatorial auctions [e.g., Feldman et al., 2015, Dütting et al., 2020a,b].

Specifically, for the matching problem constant-factor prophet inequalities have been established for both (one-sided) vertex arrivals in a bipartite graph [Feldman et al., 2015], and edge arrivals in bipartite and non-bipartite graphs [Chawla et al., 2010, Kleinberg and Weinberg, 2012, Gravin and Wang, 2019, Ezra et al., 2020].

The Single-Sample Paradigm. For many applications the assumption that the distributions are known exactly, as in the prophet model, is a rather strong one. It is natural to ask whether constant-factor prophet inequalities are also achievable with limited information about the distributions. A natural approach in this context (pioneered by Azar et al. [2014]) is to assume that the online algorithm has access to a limited number of *samples* from the underlying distributions, e.g., from historical data. Arguably, a minimal assumption under this approach is that the online algorithm has access to just a *single* sample from each distribution.

Surprisingly, even under this very restrictive assumption, strong positive results are possible. Importantly, for the classic single-choice prophet inequality one can obtain a factor 2-approximate prophet inequality with just a single sample from each distribution [Rubinfeld et al., 2020], matching the best possible guarantee that can be given with full knowledge of the distributions.

The same problem with identical distributions was considered by Correa et al. [2019], who showed that with one sample from each distribution it is possible to achieve a $e/(e-1)$ guarantee. This bound was subsequently improved by Kaplan et al. [2020], Correa et al. [2020a], and Correa et al. [2021] with a tight bound appearing in a recent working paper by Correa et al. [2020b]. Together these results show that there is a small constant-factor gap between the best-possible guarantees with a single sample, and the optimal bound with full knowledge of the distributions shown in Correa et al. [2017].

A natural question is whether these results for the classic single-choice prophet inequality extend to more combinatorial problems. This question was investigated first by Azar et al. [2014], who showed that for some combinatorial problems it is possible to achieve the same asymptotic guarantees that can be achieved with full knowledge of the distributions with just a single sample from each distribution.

Out of the problems considered by Azar et al. [2014], the matching problem stands out because here, the authors are only able to show a $O(d)$ -approximate prophet inequality with $O(d^2)$ samples for bipartite graphs with one-sided vertex arrivals. This leaves a significant gap between the guarantees achieved with full knowledge of the distributions, and those achieved with access to a limited number of samples.

It was unknown for a long time whether this gap between the full knowledge and single-sample model for matching is intrinsic, or whether it is in fact possible to achieve constant factor guarantees with just a single sample. Along the lines of other recent work on single-sample or limited-information models, we show for a range of problem variants that indeed, although using massively less prior information, single-sample prophet inequalities only lose a small constant factor compared to ones with full access to the distributions. With this, our research contributes to a recent collection of results Azar et al. [2014], Dhangwatnotai et al. [2015], Dütting et al. [2021] for single-sample

methods which demonstrate that also besides max-weight matching, the paradigm yields results close to those in the full information setting for combinatorially rich problems.

1.1 Results and Techniques

At the heart of our results lies a relatively simple fact. Recall that the greedy algorithm for maximum-weight matching orders all edges of a graph decreasingly by weight. Then, it traverses them in this order and whenever there is no conflict, the currently considered edge is added. It is not a new observation that this well-known and simple 2-approximation exhibits properties close to those of an online algorithm—actually, when the order of edge’s arrival is decreasing by weight, it does behave exactly like one. At the same time, the most natural idea when designing sample-based prophet inequalities is to accept a value during the online algorithm if it beats a threshold derived from the sample(s), e.g., the largest one. We combine the strategy of sample-thresholds with an analysis of the greedy algorithm on the set of *all* values involved, those from the sample as well as the actual problem instance.

Closest to our work, a similar approach of exploiting the fixed greedy order was previously applied by [Korula and Pál \[2009\]](#) in the context of secretary algorithms, i.e., in the absence of prior knowledge but with the additional assumption of random-order arrivals (see [Section 1.2](#) for a more detailed discussion).

We apply the above idea to prophet inequalities: somewhat surprisingly, this allows one to tie the run of the (online, adversarial-order) prophet algorithm closely to that of (offline, fixed-order) greedy, in the sense that the thresholds correspond to skipping certain values in the greedy algorithm.

First, we analyze the most general setting of max-weight matching, where G is an arbitrary (not necessarily bipartite) graph and the edges arrive one-by-one in adversarial order. We provide a single-sample prophet inequality that computes a matching of expected value at least a $\frac{1}{16}$ -fraction of the expected value of the maximum weight matching ([Theorem 1](#)).

Second, with considerable adjustments in the analysis and a slightly different algorithm, we derive a tighter bound of 8 for the special case of a bipartite graph and one-sided vertex arrival ([Theorem 2](#)).

Finally, since the bipartite graph can be interpreted as dividing the vertices into items and buyers, for this setting we provide also an incentive-compatible mechanism. Because our 8-approximation appears to lose its approximation guarantee when naturally adjusted to be incentive-compatible, we employ instead a generalized view on the analysis for the edge-arrival case, accommodating also the behavior of the selfish buyers. This again results in a 16-approximation ([Theorem 3](#)).

1.2 Additional Discussion

The most extensive, previous study of single-sample prophet inequalities is given by [Azar et al. \[2014\]](#). Based on the observation that the existence of an *order-oblivious* secretary algorithm implies the existence of a single-sample prophet inequality with the same ratio, they provide constant approximations for a number of matroids. They also use a variation of their reduction argument to obtain a 6.75-approximation for bipartite matching environments with edge arrivals. For this result they need to assume the vertex degree is limited by a constant d , and they require $O(d^2)$ samples, i.e., a constant number of samples per edge rather than just one.

The aforementioned results of [Korula and Pál \[2009\]](#) together with the insights from [Azar et al. \[2014\]](#) can actually be employed to derive a (truly) constant bound for max-weight matching: [Korula and Pál \[2009\]](#) present an 8-approximate secretary algorithm for bipartite weighted matching with one-sided vertex arrival. Although they do not specifically aim for it, their algorithm and a

version of their analysis is in fact order-oblivious. This, together with the reduction of Azar et al. [2014] implies the existence of a 13.5-approximate, single-sample prophet inequality for the special case of bipartite graphs where one side is known offline. This result, like the result for bipartite matching given by Azar et al. [2014], relies on the relation of single-sample prophet inequalities to order-oblivious secretary algorithms. In our work, we deviate from this traditional path. While we have the use of a greedy pricing in common with [Korula and Pál, 2009], our algorithms are directly designed to work in the prophet model, avoiding the inherently wasteful detour over order-oblivious secretary algorithms.

In concurrent and independent work, Caramanis et al. [2021] very recently provide prophet inequalities using an approach not unsimilar to ours. They, too, present a constant-factor approximation for matching in general graphs with edge arrivals, where they show a 32-approximation single-sample prophet inequality. Aiming in a different direction than ours, they then provide improved results for several special matroids, among them transversal and truncated partition matroids. We instead focus on the important case of bipartite matchings, and the problem of incentive compatibility. Together, our results demonstrate the versatility and power of a direct, greedy-order based approach for the design of single-sample prophet inequalities.

2 Preliminaries

We have a weighted graph $G = (V, E)$, where each edge e 's weight r_e is sampled independently from an unknown distribution D_e . The only information about these distributions available to the algorithm beforehand is a single sample s_e drawn from D_e for each edge. We denote with D the product distribution on E , while S and R are both drawn independently from D and denote the set of the realized weights of the samples (S) and of the actual weights (R), respectively. For the sake of simplicity, in the following we often refer to the R graph, which is simply graph G whose edge weights are given by R . The S graph is defined analogously but using as weights the values in S .

We consider two cases:

General graph and edge arrivals. All the vertices V are known in advance, but the edges arrive one-by-one in adversarial order. Before the online procedure starts, the algorithm has access to the sampled edge weights S . In the online phase, edges arrive in some (adversarial) order. Upon arrival of the generic edge e , the corresponding weight r_e , which is drawn from D_e , is revealed and the decision whether to add it the solution is made immediately and irrevocably.

Bipartite graph with vertex arrivals. The vertices are divided into two sets: buyers I (left vertex set of G), and items J (right vertices). Also in this case, before the online procedure starts, the algorithm has access to the realized S graph. The J side of the graph (the items) is fixed and available from the beginning. On the other hand, the buyers I appear in adversarial online fashion. Each time buyer i arrives, we may immediately and irrevocably match her with any item J that is still available, or let her leave forever unmatched. Contrary to our edge arrival model, the weights of *all* incident edges of buyer i are revealed at once. In this section, we also consider the issue of incentive-compatibility or truthfulness if the buyers act as selfish agents. A mechanism (assignment rule together with a pricing rule) is called *IC* or *truthful* if for every agent, her utility (value of item assigned minus price paid) is maximized by reporting her true valuation for all items on sale, i.e., for all adjacent edges in the bipartite graph.

Competitive analysis. Our objective is to maximize the weight of the output matching M , i.e. to find a one-on-one assignment that maximizes the sum of weights $w(M) = \sum_{e \in M} r_e$. Given the online nature of the problem, we follow the literature on prophet inequalities and evaluate the performance of our algorithms in terms of the *competitive ratio*, i.e. the ratio between the expected weight of an offline maximum weight matching and the expected weight of the online matching produced by the algorithm. Clearly, this notion of performance depends on the order in which the elements arrive, so we give bounds that hold for all possible (adversarial) orders. All bounds that we obtain are related to the expected value of the matching computed on the sample values. The bounds hold also for an *adaptive adversary* that knows in advance the whole sequence of values that will be realized, and it will release the next item in the sequence while knowing all the decisions previously made by the algorithm.

3 Edge Arrival in General Graphs

Our algorithm for this problem outlines very well the idea we described above: in order to bridge the gap between the fixed order employed by greedy (which is a 2-approximation) and the adversarial one of prophet inequalities (where greedily adding edges would be arbitrarily bad), we utilize the single sample $S = \{s_e | e \in E\}$ that we have from the product over all edge distributions. Recall that we assume the D_e to be independent.

The algorithm computes an (offline) greedy matching M_S on the graph with edge weights S . Then, for each vertex v , it interprets the weight of the edge incident to it in the greedy solution as a price p_v for the adversarial-order online algorithm: whenever an edge $e = (u, v)$ arrives online and both endpoints are free, it is added to the matching M if and only if $r_e \geq \max\{p_v, p_u\}$. While the algorithm is simple, the analysis is a bit more challenging. As pointed out before, the greedy method can be seen as an offline algorithm as well as an online algorithm with the (extremely simplifying) assumption of a fixed and weight-decreasing arrival order. We combine such a shift of viewpoint with a suitable reformulation of the random processes involved. More precisely, we look at our (online) algorithm as the byproduct of a run of a (fixed-order!) greedy strategy on a random part of a fictitious graph that contains each edge *twice* (once with the true weight r_e , and once with the sampled weight s_e).

To help with the analysis, together with the solution M , the algorithm builds E' , a subset of E containing all the edges such that $r_e > \max\{p_v, p_u\}$ ¹, i.e. all edges that are price-feasible. Note that E' is a superset of the actual solution M and may not be a matching.

Equivalent offline algorithm. In order to analyze the competitive ratio of [Algorithm 1](#), we state an offline, greedy-based algorithm with similar properties that is easier to work with. This offline procedure considers an equivalent stochastic process generating the edge weights in S and R . Instead of first drawing all the samples, and then all the realized weights, we consider a first stage in which two realizations of each random variable D_e are drawn. Then, they are sorted in decreasing order according to their values, breaking ties at random as explained above. Finally, the offline algorithm goes through all these values in that order; each value is then associated with S or R by a random toss of an unbiased coin. Consequently, the second, smaller weight of the same edge arriving later is associated to the remaining category with probability 1. It is not hard to see that the distributions of S and R are identical for both random processes.

¹We assume to break ties uniformly at random before running the algorithm, i.e. for the set of all drawn numbers that have value x , pick a permutation π of those uniformly at random and define a draw of the same value to be larger than another iff it comes first in π .

Algorithm 1: Prophet Matching

```
Set  $E' = \emptyset, M = \emptyset$ 
Compute the greedy matching  $M_S$  on the graph with edge weights according to sample  $S$ 
for all  $e = (u, v) \in M_S$  do
    Set  $p_u = p_v = s_e$ .
for all  $u \in V, v$  unmatched in  $M_S$  do
    Set  $p_u = 0$ 
for each arriving  $e = (u, v) \in E$  do
    if  $r_e > \max\{p_u, p_v\}$  then
        Add  $e$  to  $E'$ . // This is just for the analysis.
        if  $u, v$  are not matched in  $M$  then
            add  $e$  to  $M$  with weight  $r_e$ 
return  $M$ 
```

The pseudocode for this offline procedure is given in [Algorithm 2](#). As already mentioned, it mimics in an offline way the construction of M_S and E' of [Algorithm 1](#), as shown in the following.

Claim 1. The sets E' , M_S and M are distributed in the same way when computed by [Algorithm 1](#) as when computed by [Algorithm 2](#).

Proof. Consider any run of [Algorithm 2](#). Every edge value in A is associated either to R or to S . In fact, an edge value a is associated to the R graph if it is R used or if the other value for the same edge gets S used; the symmetric holds for S . Since the S and R values are all drawn independently from D , they are distributed in the same way no matter if we draw them first and then decide which set they belong to (like in [Algorithm 2](#)), or the other way around (like in [Algorithm 1](#)).

Once we have the equivalence of the underlying stochastic processes, we just need to show that for *any* realization of the S and R value, the corresponding sets of edges of interests are exactly the same in the offline and online settings. To avoid complication, we deterministically use the same tie breaking rules in both the online and offline case.

Notice that if we restrict the offline algorithm to consider only the edge weights in the S graph, we recover exactly the procedure to generate a greedy matching on the S graph. This is true because what happens to the R used edges, or to the second realization of S used edges (i.e. all values associated with R) has no influence on M_S . Hence, we can claim that the two versions of M_S in the different algorithms indeed follow the same distribution. Actually, this is even true step-wise: at any point in time during [Algorithm 2](#), M_S contains the greedy matching on the S graph *restricted* to the edges whose s_e is greater than the value being considered at that time.

We conclude the proof by showing that the two versions of E' coincide. Once we have that, it is enough to observe that the matchings M are extracted in the same way from E' . To this end, consider an iteration of [Algorithm 1](#) in which an edge e in the R graph arrives. We argue that e is added to E' in the offline procedure if and only if the online procedure does the same.

Fix edge $e = (u, v)$ added to E' by [Algorithm 1](#). Then, the online version of M_S contains no edges incident to u or v with larger weight than e , because this would imply a price higher than r_e . Since the M_S are the same in both algorithms, when [Algorithm 2](#) considers e , it holds $u, v \in V^S$ and [Algorithm 2](#) will also add e . The other direction is analogous. \square

Given the above result, we will be using E' , M_S and M mainly to refer to [Algorithm 2](#), but the same would apply to [Algorithm 1](#).

We are now ready for the main Theorem of the Section, in which we show that the competitive analysis of [Algorithm 1](#) is 16, i.e. for any adversarial arrival order of the edges, in expectation the

Algorithm 2: Offline Matching

Set $E' = \emptyset$, $M_S = \emptyset$, $M = \emptyset$ $V^S = V$
For each $e \in E$, draw from D_e two values $a_{e,1}$ and $a_{e,2}$
Order $A = \{a_{e,1}, a_{e,2} | e \in E\}$ in decreasing fashion
Mark each edge $e \in E$ as free
for each value $a \in A$ in the above order **do**
 if a corresponds to edge $e = (u, v)$ s.t. e is free, $u \in V^S$ and $v \in V^S$ **then**
 Flip a coin
 if Heads **then**
 Mark e as R -used and remove it from the free edges
 Add e to E' with weight a
 else
 Mark e as S -used and remove it from the free edges
 Add e to M_S with weight a
 Remove u, v from V_S
 else if a corresponds to edge $e = (u, v)$ that has been observed before, $u \in V^S$ and $v \in V^S$ **then**
 Remove u and v from V^S
 Add e to M_S with weight a
for each e in the same order as [Algorithm 1](#) **do**
 if e in E' and $\{e\} \cup M$ is a matching **then**
 Add e to M with weight r_e
return M .

online algorithm retains at least a $1/16$ fraction of the value of optimal offline matching OPT . The high level idea of the proof lies in relating the optimal offline matching and our solution by using M_S (which yields a 2 approximation of OPT) and a carefully chosen subset of E' , the *safe* edges.

Theorem 1. *For the problem of finding a maximum-weight matching in a general graph G , in the online edge arrival model, [Algorithm 1](#) is 16-competitive in expectation, i.e.*

$$16 \cdot \mathbb{E}[w(M)] \geq \mathbb{E}[OPT],$$

where OPT is the weight of an optimal matching in G .

Proof. For any run of the offline algorithm, we say that an edge is *considered* if it is added to either M_S or E' . Furthermore, let V' be the subset of V s.t. the according vertices all have at least one considered edge. Then, for each vertex v , we have the following inequality:

$$\mathbb{P}[v \in E' | v \in V'] = 1/2.$$

In fact, fix the point in the execution of the offline algorithm where the decision is made that $v \in V'$. This is the point in time the offline algorithm considers the first incident edge to v , implying that when we arrive at this value $a \in A$, both of the corresponding edges' endpoints are still free. Right after deciding to consider a value, a fair coin is flipped to determine whether the according edge is in E' or M_S .

Next, we claim the set E' has expected weight at least $w(M_S)/2$. This is true because for each considered edge, we can have one of two cases: either e is added to both M_S and E' , and in that case $r_e > s_e$, or $e \notin E' \cup M_S$. In the latter case, the decision to which set e is added was made by

flipping a fair coin. This means, M_S and E' -edges are determined in the exact same way, except that after already adding an R -version of e to E' , the S -version might still be used for M_S .

Now define $val_{\hat{E}}(v)$ for $v \in V$, and any $\hat{E} \subseteq E$ to denote the largest weight of an edge incident to vertex v in edge set \hat{E} , or 0 if there is no such edge. Then, it holds

$$\mathbb{E} \left[\sum_{v \in V} val_{E'}(v) \right] \geq \frac{1}{2} \mathbb{E} \left[\sum_{v \in V} val_{M_S}(v) \right] = \mathbb{E} [w(M_S)]$$

We define e_v as the edge (if any) that defines the value $val_{E'}(v)$, which is always the first edge incident to v that was considered by the offline algorithm. Slightly abusing the notation, we define $r_{e_v} = 0$ if the edge e_v does not exist. We rewrite (note some edges might occur twice in below sum)

$$\mathbb{E} \left[\sum_{v \in V} r_{e_v} \right] \geq \mathbb{E} [w(M_S)]$$

Fix now some edge e_v (maximal-weight among all edges incident to v in E'). We call $e_v = (u, v)$ *safe* if for v , e_v is also the only edge in E' and for u , there is no smaller incident edge in E' . The former happens iff the second edge incident to v considered by the offline algorithm is decided to be in S , which is at least with probability $1/2$ since either this is the first copy of the edge seen, or it must be in S anyway. The latter happens iff the next edge after e_v incident to u in the offline algorithm is decided to be in S , which is also the case with probability at least $1/2$ for the same reasons as before. Both events do depend on each other - one might influence the other in terms of which actual edge is the considered one or at which point in time the consideration takes place. But for the R - S - decision on the next incident edge, it holds that it is either made only at the point in time that edge is considered, or it has been made before. In the latter case, however, the second copy must be in S . This means, the probability for the next considered edge to be in S can never be manipulated to be below $1/2$. We showed that with probability at least $1/4$, any edge e_v is *safe*. Fix now any realization of the edge weights in A , and consider only the randomness of the coin tosses. Then,

$$\begin{aligned} \mathbb{E} \left[\sum_{v \in V} r_{e_v} \mathbb{1}_{\{e_v \text{ is safe}\}} \right] &= \sum_{v \in V} \sum_{e \in E} r_e \cdot \mathbb{P}[e_v \text{ is safe} | e_v = e] \cdot \mathbb{P}[e_v = e] \\ &\geq \sum_{v \in V} \sum_{e \in E} \frac{r_e}{4} \cdot \mathbb{P}[e_v = e] = \frac{1}{4} \mathbb{E} \left[\sum_{v \in V} r_{e_v} \right]. \end{aligned}$$

Finally, we need to take the set of e_v -edges and consider which subset of them will be part of the matching chosen by the prophet inequality. Note that in the *safe* set, the only edge incident to vertex v that might be chosen in the online algorithm is $e_v = (u, v)$, by definition. Now either, this indeed happens in the algorithm, or at the time e_v arrives, vertex u is already matched via some edge e'_u , which must be of larger weight than e_v due to the definition of *safe*. Now, this also implies that e_u is not safe, since apparently more than one edge is incident to u in E' . All in all, we have that for each vertex v s.t. e_v is safe, either v itself or a unique vertex u with e_u not safe will be matched via an edge of weight at least r_{e_v} . Since in case we indeed have that a safe edge e_v is added to M , e_v can possibly also be the safe edge for its other endpoint (and our sum thus counts the weight of this edge twice), the actual matching computed by the algorithm recovers at least half of the summed-up values $val_{E'}(v)$ of the safe matching:

$$\mathbb{E} [w(M)] \geq \frac{1}{2} \mathbb{E} \left[\sum_{v \in V} r_{e_v} \mathbb{1}_{\{e_v \text{ is safe}\}} \right] \geq \frac{1}{8} \mathbb{E} [w(M_S)] \geq \frac{1}{16} \mathbb{E} [OPT].$$

□

4 Vertex Arrival in Bipartite Graphs

The main difference in the bipartite setting is, as noted in the introduction, that instead of having individual edge arrivals, one side of the graph is fixed and at each step a vertex arrives with all of its incident edge weights revealed. We will refer to the fixed side as the ‘items’ and the online side as the ‘buyers’, even though in this section there will be no incentive analysis. A truthful version of these results, with slightly worse approximation guarantees, is presented in [Section 5](#).

As before, the sample S is used to calculate prices, on both sides: items (fixed side vertices) have prices to *protect* them from being sold too cheaply, while buyers (incoming vertices) have prices they need to beat in order to participate in the market, in addition to beating the item prices.

Algorithm 3: Bipartite Prophet Matching

```

Set  $E' = \emptyset$ 
Compute the greedy matching  $M_S$  on the graph with edge weights according to sample  $S$ 
for all  $e = (i, j) \in M_S$  do
    Set  $p_i = p_j = s_e$ .
for all vertices  $k$  not matched in  $M_S$  do
    Set  $p_k = 0$ .
for each arriving  $i^* \in I$  do
    Let  $\hat{e} = \arg \max\{r_e \mid e = (i^*, j) \text{ and } r_e > \max\{p_{i^*}, p_j\}\}$ 
    Add  $\hat{e}$  to  $E'$ , with weight  $r_{\hat{e}}$ . // Only used for the analysis, just as before.
    Let  $j^*$  be the item in  $\hat{e}$ 
    if  $j^*$  is still free in  $M$  then
        Add  $(i^*, j^*)$  to  $M$  with weight  $r_{\hat{e}}$ 
return  $M$ 

```

We provide some intuition on the objects of interest that appear in the pseudo-code of [Algorithm 3](#). M_S is the greedy matching on the S graph, and it does not depend in any way on the actual edge weights R . E' is a subset of the edges of the R graph: each bidder/buyer is associated with at most one item, but there may be items associated to more than one buyer. Solving these conflicts in E' in favor of the first arriving edge we obtain M . Note that M is computed in an online fashion.

In order to relate E' and M with OPT , similarly to the previous section, we consider an offline version of the algorithm, which exhibits the same distribution over the relevant allocations and can be analyzed more easily.

Equivalent offline algorithm. [Algorithm 4](#) interleaves the building of a greedy matching M_S with that of the allocation E' . All the values are drawn in advance, but they are assigned to S and R only upon arrival.

As before, we observe that the allocations E' and M are distributed in exactly the same way, whether they are computed by [Algorithm 4](#) or [Algorithm 3](#).

Claim 2. The sets E' , M_S and M are distributed the same way when computed by [Algorithm 3](#) as when computed by [Algorithm 4](#).

Proof. The proof is analogous to the one of [Claim 1](#). The R and S graphs in the offline and online setting follows the same distribution as argued in the previous Section, and M_S is still the greedy matching on the S graph, in both [Algorithm 3](#) and [Algorithm 3](#). We know that M is extracted in the same way by E' by the two algorithms, so we just need to argue that for each fixed realization of the edge weights, the online and offline versions of the edge sets E' coincide. Considering now

Algorithm 4: Bipartite Offline Matching

Set $E' = \emptyset$, $M_S = \emptyset$, $M = \emptyset$
 $I^R = I^S = I$, $J^S = J$
For each $e \in I \times J$, draw from D_e two values $a_{e,1}$ and $a_{e,2}$
Order $A = \{a_{e,1}, a_{e,2} | e \in E\}$ in a decreasing fashion
Mark each edge $e \in E$ as free
for each value a from A in above order **do**
 if a corresponds to edge e s.t. e is free **then**
 Flip a coin
 if Heads **then**
 Mark e as R -used
 else
 Mark e as S -used
 if e has just been marked R -used, or was previously S -used, $i \in I^R$ and $j \in J^S$ **then**
 Add e to E' with weight a and remove i from I^R .
 if e has just been marked S -used, or was previously R -used, $i \in I^S$ and $j \in J^S$ **then**
 Remove j from J^S , and remove i from I^S and from I^R
 Add e to M_S with weight a
for each e in the same order as [Algorithm 1](#) **do**
 if e in E' and $\{e\} \cup M$ is a matching **then**
 Add e to M with weight r_e
return M .

which edges $e = (i, j)$ from the R graph are successfully added to M' , those are exactly the ones such that for both of their endpoints, no incident edge had an S -value greater than r_e and no edge containing i with the same property but larger value exists. This is exactly the set of edges fulfilling the argmax in algorithm [Bipartite Prophet Matching](#). \square

As an intermediate step towards the solution we need to relate the weight of the edge set E' to that of the actual matching M . This is done via an intermediate matching M^+ , which is the one induced by E' if the conflicts over an item j that is allocated more than once are solved in favor of the buyer with the largest j -incident edge in E' . Clearly, $w(M) \leq w(M^+) \leq w(E')$.

Lemma 1. Consider [Algorithm 4](#) and M^+ defined above. Then the following inequality holds:

$$4 \cdot \mathbb{E} [w(M^+)] \geq \mathbb{E} [OPT]$$

Proof. We know that $w(M_S)$ is the greedy solution on the S graph, and therefore yields a 2-approximation to OPT , in expectation. This means we just need to prove that $2 \cdot \mathbb{E}[w(M^+)] \geq \mathbb{E}[w(M_S)]$.

Consider any point in time during [Algorithm 4](#) when an edge $e = (i, j)$ incident to i is considered for any fixed i , and no edge incident to i has been added to M_S yet. Then either, the according item j is blocked in the S graph (and therefore, also the R one). In this case, we cannot add e to either matching and move on. Or, j is free in both the graphs. Then, e cannot be S -used yet since both its endpoints are free. It could be R -used, but then some earlier, higher value was added to E' (and therefore, also M^+) incident to i . Finally, if e arrives undecided, we flip a coin whether to add it to M_S or E' . In case we choose M_S , we will be able to add e to it. In case we choose E' , we can either add it, or have already added a larger value for i . All in all, with probability at least $1/2$, the largest edge incident to i in E' is at least as large as the one in M_S . \square

We are now ready to show the main result of this Section.

Theorem 2. *For the problem of max-weight matching in a bipartite graph G , where vertices on one side arrive online, Algorithm [Algorithm 3](#) is 8-competitive in expectation, i.e.*

$$8 \cdot \mathbb{E}[w(M)] \geq OPT,$$

where OPT denotes the weight of an optimal matching in G .

Proof. The last step left to show the theorem is that $2 \cdot \mathbb{E}[w(M)] \geq \mathbb{E}[w(M^+)]$. In fact, recall that $4 \cdot \mathbb{E}[w(M^+)] \geq \mathbb{E}[OPT]$, by [Lemma 1](#).

Consider matching M^+ in [Algorithm 4](#). Fix any sequence of S - and R -values, and any item j . If j is not matched in M^+ , then only M can possibly generate value from this item and there is nothing to prove. If it is not the case, let $e = (i, j) \in M^+$ be the corresponding edge and consider the point in time when e is added to M^+ with weight r_e . We follow the algorithm along the sequence from there, and look for the next point in time where some edge incident to j is added to either E' or M_S . In case the latter happens, it means that in the corresponding run of [Algorithm 3](#), item j will be blocked for all the buyers but i , who has edge weight r_e when being assigned j . Because each buyer has only at most one edge in E' , then e will be selected for M and contributes the same weight to M^+ and M . In case the former happens and the next j -incident edge is added to E' instead, this is not the case. So we need a more careful analysis in order to bound the probability for this to happen.

Let us call the next j -incident edge in the sequence $e' = (i', j)$. We distinguish the following cases:

e' is free. Whenever an edge is feasible to add to E' during the algorithm, it is also feasible to add to M_S , so with probability $1/2$, we decide the current value for e' to be in S .

e' is R used. The current e' -value is in S , and we therefore add it to M_S (except of course one of its endpoints is already blocked and we just move on).

e' is S used. The current e' -value is in R . However, when the previous value for this edge arrived, we either added the edge to M_S (therefore blocking i' and j from I^R and J^S), or one of them was already blocked, i.e. either $i' \notin I^S$ (and therefore $i' \notin I^R$), or $j \notin J^S$. In both cases, we cannot add e' to E' and therefore just move on.

Concluding, the next time an edge incident to j is added to either M_S or E' , the edge will either be added to M_S with probability one (hence blocking the item j and making e the only edge in M_{off} for it), or this happens on discovery of an undecided edge e' , in which case there is always a probability of $1/2$ we try to add it to M_S (hence blocking the item j and making e the only edge in E' for it). Whenever this fails because an endpoint of e' is blocked from I^S or J^S already, it would equally have been blocked for E' . This implies there is a probability of at least $1/2$ that $e \in M$. Therefore, $\mathbb{E}[w(M)] \geq \frac{1}{2} \mathbb{E}[w(M^+)]$ and we get the claimed result. \square

5 Truthful Bipartite Matching

For the same model as above, with buyers on the left side of G arriving online and items on the right side available offline, we design an incentive-compatible version of our prophet inequality. The difficulty is that our 8-approximation in [Theorem 2](#) relies on each buyer choosing the *largest* incident edge that beats an item's price. However, if the price is close to the buyer's value, this edge might

have very low utility. The buyer might prefer to purchase an item she has lower value for instead, but where the price is sufficiently lower, weakening our analysis to the point that we can no longer show any approximation.

Clearly, a different approach is in order. We use a pricing-based algorithm that does not rely on restricting the buyer’s choice of item (among those with feasible prices that are still available). Instead, we ensure a sufficient weight for the resulting matching solely via appropriate thresholds for each buyer and item: this can be seen as doing a worst-case analysis over all threshold-feasible edges instead of choosing a good subset of them, explaining the incurred loss in the approximation. This allows us to make another small change and, among all threshold-feasible items, always assign the *utility-maximizing* one. In turn, the routine becomes truthful, and we have to mainly care about showing the approximation.

Our strategy for setting the aforementioned thresholds can be interpreted as follows. Assume we take as a starting point not our prophet inequality for the bipartite case, but the one for general graphs with edge arrivals. Obviously, arrival orders for the bipartite one-sided model we consider here are highly restricted compared to general edge-arrival. However, assuming the graph is bipartite, each arriving edge will consist of exactly one buyer i and one item j . Then, [Algorithm 1](#) assigns prices p_i and p_j , where p_i corresponds to the aforementioned buyer threshold we are planning to use.

Our mechanism works as follows: we determine greedy prices on the sample graph for all vertices just as described above. When a buyer i arrives, we present her not with the greedy price of each item, but instead charge for each item j the maximum over p_j and p_i . Then, we let the buyer choose what she likes best, i.e. we assign a utility-maximizing item.

Algorithm 5: Truthful Bipartite

```

Set  $M_T = \emptyset$ 
Compute the greedy matching  $M_S$  on the graph with edge weights according to sample  $S$ .
for all  $e = (i, j) \in M_S$  do
    Set  $p_i = p_j = s_e$ .
for all vertices  $k$  not matched in  $M_S$  do
    Set  $p_k = 0$ .
for each arriving  $i^* \in I$  do
    Set  $F(i^*) = \{e = (i^*, j) \mid r_e > \max\{p_{i^*}, p_j\} \text{ and } j \text{ is not matched in } M_T\}$ 
    Set  $\tilde{e} = \arg \max\{r_e - \max\{p_{i^*}, p_j\} \mid e \in F(i^*)\}$ 
    Add  $\tilde{e} = (i^*, j^*)$  to  $M_T$  with weight  $r_{\tilde{e}}$ 
    Charge buyer  $i^*$  a price of  $\max\{p_{i^*}, p_{j^*}\}$ 
return  $M_T$ 

```

Here, M_S is again a greedy matching on the sample graph. As before (we omit this part of the analysis because it is analogous), one can imagine an offline version of [Algorithm 5](#) that simply draws two values for each edge, and then goes through their set A in non-increasing order, deciding for each value considered whether it belongs to R or S (if this was not yet decided by the twin value for the same edge).

Also as before, the offline version will produce the same set of price-beating or *feasible* edges E' , which are the exact ones that [Algorithm 5](#) will consider (and add in case they are the buyer’s favorite and the item is still free).

Theorem 3. *For the bipartite max-weight matching problem with one-sided vertex arrival, [Algorithm 5](#) is a truthful 16-approximation, in expectation.*

Proof. Truthfulness is immediate: on arrival, the unmatched items are offered to the buyer for prices

determined independently of the buyer’s values, and she is assigned one that is maximizing her utility. By misreporting, her utility can therefore only become worse. Now for the approximation ratio, consider the edge set E' defined via the offline version of our edge-arrival algorithm [Algorithm 2](#). Note that the arrival order, be it edge-arrival or vertex-arrival, has no influence whatsoever on the computation of set E' in this algorithm. With the same arguments as in the analysis of the edge-arrival model, E' from the offline version is distributed exactly as the set of all price-feasible edges E' in [Algorithm 5](#). We base our proof on the observation that after establishing $\mathbb{E}[w(E')] \geq \frac{1}{2}\mathbb{E}[w(M_s)]$, all we do in the analysis of [Algorithm 1](#) is proving a lower bound on the weight of a maximal matching in edge set E' . Concretely, since E' is the same as for edge arrival as well as vertex arrival in the offline versions of our algorithms, all inequalities hold as before. In addition, when a bidder arrives that beats both thresholds of at least one incident edge, she also has nonnegative utility for the edge when paying the maximum over both. Therefore, if an edge in E' is incident to the buyer and the item is free at her arrival, she will match to some item (not necessarily this one, due to choosing the utility-maximizing option, but some). Therefore, no buyer $i \in I$ s.t. $\exists j \in J$ s.t. $e = (i, j) \in E'$ will ever stay unmatched unless all according items are matched already. On the other hand, fixing an item j with at least one incident edge in E' , j will either be matched or the according buyer chooses to match to a different item. Together, this implies that [Algorithm 5](#) chooses indeed a maximal matching in edge set E' . Now consider M^+ defined as before for edge arrival (and actually, computed as before, just on a subset of possible instances containing only the bipartite graphs). For each edge in $e = (i, j) \in M^+$, neither i nor j can be matched in M_T via a smaller edge than e . At the same time, since M_T is maximal in E' , either i or j must be matched, yielding $w(M^+) \leq 2w(M_T)$ analogously to the edge-arrival case. \square

6 Conclusion

The single-sample paradigm is proving more and more powerful, even for combinatorially interesting problems like max-weight matching. We have designed constant-approximation prophet inequalities for a number of settings, and our results hold without restrictive assumptions like limited degree. However, approximation ratios clearly seem to depend on the generality of arrival model, independence assumptions on the distributions, the assumed structure of the underlying graphs, and whether or not one assumes the presence of selfish agents. In the future, it will be an interesting endeavor to fully understand these relationships of model and ratio, and derive tighter approximations as well as lower bounds. Finally, a range of other problems where prophet inequalities exist is still left to explore, or derive better bounds for, in the single-sample paradigm.

References

- S. Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. In R. Ostrovsky, editor, *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, pages 512–521. IEEE Computer Society, 2011.
- P. D. Azar, R. Kleinberg, and S. M. Weinberg. Prophet inequalities with limited information. In C. Chekuri, editor, *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 1358–1377. SIAM, 2014.

- C. Caramanis, M. Faw, O. Papadigenopoulos, and E. Pountourakis. Single-sample prophet inequalities revisited. *arXiv*, 2021. URL <https://arxiv.org/abs/2103.13089>.
- S. Chawla, J. D. Hartline, D. L. Malec, and B. Sivan. Multi-parameter mechanism design and sequential posted pricing. In L. J. Schulman, editor, *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 311–320. ACM, 2010.
- J. R. Correa, P. Foncea, R. Hoeksma, T. Oosterwijk, and T. Vredeveld. Posted price mechanisms for a random stream of customers. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC '17, Cambridge, MA, USA, June 26-30, 2017*, pages 169–186. ACM, 2017.
- J. R. Correa, P. Dütting, F. A. Fischer, and K. Schewior. Prophet inequalities for I.I.D. random variables from an unknown distribution. In A. Karlin, N. Immorlica, and R. Johari, editors, *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019*, pages 3–17. ACM, 2019.
- J. R. Correa, A. Cristi, B. Epstein, and J. A. Soto. The two-sided game of googol and sample-based prophet inequalities. In S. Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 2066–2081. SIAM, 2020a.
- J. R. Correa, A. Cristi, B. Epstein, and J. A. Soto. Sample-driven optimal stopping: From the secretary problem to the i.i.d. prophet inequality. *CoRR*, abs/2011.06516, 2020b. URL <https://arxiv.org/abs/2011.06516>.
- J. R. Correa, P. Dütting, F. A. Fischer, K. Schewior, and B. Ziliotto. Unknown I.I.D. prophets: Better bounds, streaming algorithms, and a new impossibility (extended abstract). In J. R. Lee, editor, *12th Innovations in Theoretical Computer Science Conference, ITCS 2021, January 6-8, 2021, Virtual Conference*, volume 185 of *LIPICs*, pages 86:1–86:1. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.
- P. Dhangwatnotai, T. Roughgarden, and Q. Yan. Revenue maximization with a single sample. *Games and Economic Behavior*, 91:318–333, 2015.
- P. Dütting and R. Kleinberg. Polymatroid prophet inequalities. In N. Bansal and I. Finocchi, editors, *Algorithms - ESA 2015 - 23rd Annual European Symposium, Patras, Greece, September 14-16, 2015, Proceedings*, volume 9294 of *Lecture Notes in Computer Science*, pages 437–449. Springer, 2015.
- P. Dütting, M. Feldman, T. Kesselheim, and B. Lucier. Prophet inequalities made easy: Stochastic optimization by pricing nonstochastic inputs. *SIAM J. Comput.*, 49(3):540–582, 2020a.
- P. Dütting, T. Kesselheim, and B. Lucier. An $o(\log \log m)$ prophet inequality for subadditive combinatorial auctions. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 306–317. IEEE, 2020b. doi:[10.1109/FOCS46700.2020.00037](https://doi.org/10.1109/FOCS46700.2020.00037). URL <https://doi.org/10.1109/FOCS46700.2020.00037>.
- P. Dütting, F. Fusco, P. Lazos, S. Leonardi, and R. Reiffenhäuser. Efficient two-sided markets with limited information, 2021. Preprint on authors’ webpage https://www.plazos.me/pdfs/2sided_xos.pdf.

- T. Ezra, M. Feldman, N. Gravin, and Z. G. Tang. Online stochastic max-weight matching: Prophet inequality for vertex and edge arrival models. In P. Biró, J. D. Hartline, M. Ostrovsky, and A. D. Procaccia, editors, *EC '20: The 21st ACM Conference on Economics and Computation, Virtual Event, Hungary, July 13-17, 2020*, pages 769–787. ACM, 2020.
- M. Feldman, N. Gravin, and B. Lucier. Combinatorial auctions via posted prices. In P. Indyk, editor, *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015, San Diego, CA, USA, January 4-6, 2015*, pages 123–135. SIAM, 2015.
- M. Feldman, O. Svensson, and R. Zenklusen. Online contention resolution schemes. In R. Krauthgamer, editor, *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016, Arlington, VA, USA, January 10-12, 2016*, pages 1014–1033. SIAM, 2016.
- N. Gravin and H. Wang. Prophet inequality for bipartite matching: Merits of being simple and non adaptive. In A. Karlin, N. Immorlica, and R. Johari, editors, *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019*, pages 93–109. ACM, 2019.
- M. T. Hajiaghayi, R. D. Kleinberg, and T. Sandholm. Automated online mechanism design and prophet inequalities. In *Proceedings of the Twenty-Second AAAI Conference on Artificial Intelligence, July 22-26, 2007, Vancouver, British Columbia, Canada*, pages 58–65. AAAI Press, 2007.
- H. Kaplan, D. Naori, and D. Raz. Competitive analysis with a sample and the secretary problem. In S. Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 2082–2095. SIAM, 2020.
- R. Kleinberg and S. M. Weinberg. Matroid prophet inequalities. In H. J. Karloff and T. Pitassi, editors, *Proceedings of the 44th Symposium on Theory of Computing Conference, STOC 2012, New York, NY, USA, May 19 - 22, 2012*, pages 123–136, 2012.
- N. Korula and M. Pál. Algorithms for secretary problems on graphs and hypergraphs. In *Automata, Languages and Programming, 36th International Colloquium, ICALP 2009, Rhodes, Greece, July 5-12, 2009, Proceedings, Part II*, volume 5556 of *Lecture Notes in Computer Science*, pages 508–520. Springer, 2009.
- U. Krengel and L. Sucheston. Semiamarts and finite values. *Bulletin of the American Mathematical Society*, 83:745–747, 1977.
- U. Krengel and L. Sucheston. On semiamarts, amarts, and processes with finite value. *Advances in Probability and Related Topics*, 4:197—266, 1978.
- A. Rubinfeld, J. Z. Wang, and S. M. Weinberg. Optimal single-choice prophet inequalities from samples. In T. Vidick, editor, *11th Innovations in Theoretical Computer Science Conference, ITCS 2020, January 12-14, 2020, Seattle, Washington, USA*, volume 151 of *LIPICs*, pages 60:1–60:10. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.
- E. Samuel-Cahn. Comparison of threshold stop rules and maximum for independent nonnegative random variables. *Annals of Probability*, 12:1213–1216, 1984.