Secretary Markets with Local Information

Ning Chen * Martin Hoefer † Marvin Künnemann ‡ Chengyu Lin§

Peihan Miao ¶

Abstract

The secretary model is a popular framework for the analysis of online admission problems beyond the worst case. In many markets, however, decisions about admission have to be made in a decentralized fashion and under competition. We cope with this problem and design algorithms for secretary markets with limited information. In our basic model, there are m firms and each has a job to offer. n applicants arrive iteratively in random order. Upon arrival of an applicant, a value for each job is revealed. Each firm decides whether or not to offer its job to the current applicant without knowing the strategies, actions, or values of other firms. Applicants decide to accept their best offer.

We consider the social welfare of the matching and design a decentralized randomized thresholding-based algorithm with ratio $O(\log n)$ that works in a very general sampling model. It can even be used by firms hiring several applicants based on a local matroid. In contrast, even in the basic model we show a lower bound of $\Omega(\log n/(\log\log n))$ for all thresholding-based algorithms. Moreover, we provide secretary algorithms with constant competitive ratios, e.g., when values of applicants for different firms are stochastically independent. In this case, we can show a constant ratio even when each firm offers several different jobs, and even with respect to its individually optimal assignment. We also analyze several variants with stochastic correlation among applicant values.

1 Introduction

The Voice is a popular reality television singing competition to find new singing talent contested by aspiring singers. The competition employs a panel of coaches; upon the arrival of a singer, every coach critiques the artist's performance and determines in real time if he/she wants the artist to be on his/her team. Among those who express "I want you", the artist selects a favorite coach. What strategy of picking artists should coaches adopt in order to select the best candidates?

This problem is a reminiscent of the classic secretary problem [9,24]: A firm interviews a set of applicants who arrive in an online fashion. When an applicant arrives, his non-negative value is revealed, and the firm needs to make an immediate and irrevocable decision on whether to make an offer to the applicant, without knowing the values of future potential applicants. The objective is to maximize the (expected) value of the hired applicant. The secretary problem is well studied in social science and computer science. It is well known that the problem, in the worst case, does not admit an algorithm with any guaranteed competitive ratio. However, if applicants arrive in uniform random order, there is an online algorithm that hires the best applicant with optimal probability $1/e$ (see, e.g., [4]). For a more detailed discussion on the secretary problem see, e.g., [2, 12].

The scenario of The Voice is a generalization of the secretary problem from one firm to multiple firms and from one hire to multiple hires. Such a generalization yields several fundamental changes to the problem:

*Nanyang Technological University, Singapore. ningc@ntu.edu.sg

†Max-Planck-Institut für Informatik, Saarbrücken, Germany. Supported by DFG Cluster of Excellence MMCI at Saarland University. mhoefer@mpi-inf.mpg.de

‡Max-Planck-Institut für Informatik, Saarbrücken, Germany, and Saarbrücken Graduate School of Computer Science, Germany. marvin@mpi-inf.mpg.de

§Chinese University of Hong Kong, Hong Kong. cylin@cse.cuhk.edu.hk

¶University of California, Berkeley, USA. peihan@cs.berkeley.edu

Firms (i.e., coaches) are independent and compete with each other for applicants. Thus, each firm may determine on their own the strategy to adopt. Firms are decision makers; that is, there is no centralized authority and every firm can choose different strategies on its own (based on observed information). Each firm can only observe its local information, i.e., it has no knowledge about the values of other (firm, applicant)-pairs and the selected strategies of other firms. Hence, adopting a best-response strategy in a game-theoretic sense might require learning other strategies and payoffs. Given the limited feedback this can be hard or even impossible. The same issues occur in many other decentralized markets, e.g., online dating and school admission, where entities behave individually and have to make decisions based on a very limited view of the market, the preferences, and the strategies used by potential competitors.

In this paper, we design and analyze strategies for all firms in such a decentralized, competitive enviroment to enable efficient allocations. Our algorithms are evaluated both globally and individually: On the one hand, we hope the outcomes achieve good social welfare (i.e., the total value obtained by all firms). Thus, we measure the competitive ratio compared to social welfare given by the optimal centralized online algorithm. On the other hand, considering that firms are self-interested entities, we hope that our algorithms generate a nearly optimal outcome for each individual firm. That is, although given the limited feedback it can be impossible to obtain best-response strategies, we nevertheless hope that (when applied in combination) our algorithms can approximate the outcome of a best response (in hindsight with full information) of every individual firm within a small factor.

We identify several settings that admit algorithms with small constant competitive ratio both globally and individually. This implies that even in decentralized markets with very limited feedback, there are algorithms obtaining a good allocation. For the general case, we provide a strategy to approximate social welfare within a logarithmic factor, and we show almost matching lower bounds on the competitive ratio for a very natural class of algorithms. Thus, in the general case centralized control seems to be necessary in order to achieve good social welfare.

1.1 Model and Preliminaries

We first outline our *basic model*, a decentralized online scenario for hiring a single applicant per firm with random arrival. There is a complete bipartite graph $G = (U, V, w)$ with sets $U = \{u_1, u_2, \ldots, u_m\}$ and $V = \{v_1, v_2, \ldots, v_n\}$ of firms and applicants, respectively. There is a *value* or *weight function*[1] $w : U \times V \to \mathbb{R}^+$. We assume that each firm can hire at most one applicant.

The weights describe an implicit preference of each individual to the other side. Each firm $u \in U$ prefers applicants according to the decreasing order of $w(u, \cdot)$ of the edges incident to u; similarly, each applicant $v \in V$ prefers firms according to the decreasing order of $w(\cdot, v)$ of the edges incident to v.[2]

Applicants in V arrive one by one to the market and reveal their edge weights to all firms. Upon the arrival of an applicant, each firm immediately decides whether to provide an offer for the applicant or not; after collecting all job offers, the applicant then picks one that she prefers most, i.e., the one with the largest weight. Note that each firm can only see its own weights for the applicants and has no information about future applicants; in addition, all decisions cannot be revoked. A firm can make multiple offers over time until it succeeds to hire an applicant. In this paper, we mostly concentrate on the random permutation model, i.e., weights are fixed by an adversary but applicants arrive in a uniformly random order. We also consider extensions of our results to other standard models, such as the iid model (weights drawn iid from known distributions), prophet-inequality model (different known distributions, adversarial arrival), and more general models based on different mixtures of stochastic and adversarial elements.

Our goal is to design decentralized algorithms when each firm makes its decision only based on its own previously seen information and there is no centralized authority that manages different firms altogether. Due to online arrival some performance loss is unavoidable, and there are two natural objectives to quantify this loss. The standard benchmark is *social welfare*, defined to be the total weight of assigned firm-applicant

[1] To avoid technicalities, we assume that no two edges have the same weight; this assumption is without loss of generality by using small perturbations or a fixed rule to break ties.

[2] In a more general preference model there are for each pair (u, v) different values obtained by u and v; we will not consider this general case in the present paper.

pairs. For an algorithm $\mathcal{A}$, we say the algorithm has a *competitive ratio* of α if for all instances, we have $w(M^*)/\mathbb{E}\left[w(M^{\mathcal{A}})\right] \leq \alpha$. Here the expectation is taken over the random permutation, M^* is the maximum weight matching in G, and $M^{\mathcal{A}}$ is the matching returned when every firm runs algorithm $\mathcal{A}$.

In addition, we would like to approximate the *individual optimum* for each firm (i.e., the value of its best applicant) and strive to obtain a constant competitive ratio for this benchmark. This goal is obviously much more demanding than social welfare. It can be impossible, e.g., if there is a single applicant that is extremely valuable for every firm, while all others are not valuable at all. Consequently, a constant competitive ratio for individual optima can be achieved only in domains with additional structure. In this paper, we obtain them when applicant values result from stochastic processes with a sufficient amount of independence among firms.

1.2 Contribution and Techniques

As a natural first attempt, consider every firm running the classic secretary algorithm [9, 24], which samples the first $r-1$ applicants, records the best weight seen in the sample, and then makes an offer to every applicant that exceeds this threshold (see Algorithm 4 below). It turns out that such a strategy fails miserably in a decentralized market, even if each applicant has the same weight for all firms. The main idea of the proof is to construct an instance with many 'bad' applicants and only a few ($\Theta(\log n)$) 'good' applicants. With high probability, each firm will see at least one 'good' applicant over the sampling period, and compete over the 'good' applicants later on. However, at most $O(\log n)$ of the firms can hire a 'good' applicant. In contrast, to achieve optimal social welfare each firm should hire an applicant.

Proposition 1. *For any constant $\beta < 1$, when setting $r = \lfloor \beta n \rfloor + 1$, then the classic secretary algorithm has a competitive ratio of $\Omega(n/\log n)$.*

Proof. Suppose there are n applicants and $m = n$ firms. The applicants come in two types: $n_g = \gamma \ln n$ 'good' applicants with weight 2 for all edges incident to them, where $\gamma \geq 1/\beta$ is a constant, and the rest 'bad' applicants with weight 1 for all edges incident to them. (Note that to avoid ties, we can add a small perturbation $\epsilon_{u,v}$ on all pairs).

For any permutation of the applicants, we have $w(M^*) = m + n_g$. Next, we consider the matching $M^{\mathcal{A}}$ returned by the algorithm and give an upper bound on $\mathbb{E}\left[w(M^{\mathcal{A}})\right]$. In the first $r-1$ rounds, the probability that no good applicant arrives is

$$\begin{aligned} p \quad &= \quad \prod_{i=0}^{\lfloor \beta n \rfloor} \frac{n - n_g - i}{n - i} \quad \leq \quad \left(1 - \frac{n_g}{n}\right)^{\lfloor \beta n \rfloor} \quad \leq \quad \left(1 - \frac{n_g}{n}\right)^{\beta n - 1} \\ &= \quad \left(\left(1 - \frac{n_g}{n}\right)^n\right)^\beta \cdot \left(\frac{n}{n - n_g}\right) \quad \leq \quad e^{-n_g \beta} \cdot \left(\frac{n}{n - n_g}\right) \quad . \end{aligned}$$

If a firm observes a 'good' applicant, no 'bad' applicant can be hired since the threshold for the firm is set to be 2. Since a good applicant is good for all firms, a single good applicant in the first $r-1$ rounds results in all thresholds for all firms being set to 2. Hence,

$$\mathbb{E}\left[w(M^{\mathcal{A}})\right] \quad \leq \quad p(m + n_g) + (1-p) \cdot 2n_g \qquad \leq \quad \frac{n}{n - n_g} \cdot e^{-n_g \beta} \cdot (m - n_g) + 2n_g$$

Using $\gamma \geq 1/\beta$, we see

$$\frac{\mathbb{E}\left[w(M^{\mathcal{A}})\right]}{\mathbb{E}\left[w(M^*)\right]} \quad \leq \quad \frac{n}{n - n_g} \cdot e^{-n_g \beta} \cdot \frac{m - n_g}{m + n_g} + \frac{2n_g}{m + n_g} \quad = \quad \frac{n^{1 - \gamma\beta} + 2\gamma \ln n}{n + \gamma \ln n} \quad = \quad \Theta\left(\frac{\log n}{n}\right) \quad .$$

□

In contrast, we present in Section 2 a more careful approach based on sampling and thresholds that is $O(\log n)$-competitive. This algorithm can be applied beyond the basic model in a very large generality.

In fact, we prove the guarantee in a scenario where each firm u_i has a private matroid $\mathcal{S}_i$ and can accept any subset of applicants that forms an independent set in $\mathcal{S}_i$. Furthermore, our analysis extends to a general sampling model that encompasses the secretary model (random arrival, worst-case weights), prophet-inequality model (worst-case arrival, stochastic weights), as well as a variety of other mixtures of stochastic and worst-case assumptions [13]. The algorithm is essentially the approach used in [3], where it was shown that it provides a logarithmic guarantee for the matroid secretary problem. In contrast to this work, however, we must apply the algorithm based on local information, which requires a different approach for analysis. Our novel approach here to handle decentralized thresholding is to bundle all stochastic decisions and treat correlations using linearity of expectation. The effects of applicant preferences and competition can then be analyzed in a pointwise fashion.

We contrast this upper bound with an almost matching lower bound for thresholding-based algorithms in the basic model. A thresholding-based algorithm samples a number of applicants, determines a threshold, and then offers to every remaining applicant that has a weight above the threshold. Although such algorithms are nearly optimal in the centralized setting, every such algorithm must have a competitive ratio of at least $\Omega(\log n/\log\log n)$ in the decentralized setting. The lower bound carefully constructs a challenge to guess how many firms contribute to the social welfare and to avoid overly high concentration of offers on a small number of valuable applicants.

In Section 3 we show that this challenge can be overcome if there is stochastic independence between the weights of an applicant to different firms. We study this property in a generalized model for decentralized k-secretary, where each firm u_i has k_i different jobs to offer. Upon arrival, an applicant reveals k_i weights for each firm u_i, one for each position. If each firm uses a variant of the optimal e-competitive algorithm for bipartite matching [19], independence between weights of different firms allows to show a constant competitive ratio. Moreover, each firm even manages to recover a constant fraction of the individual optimum matching and therefore almost plays a best response strategy.

In Section 4 we consider two additional variants with stochastically generated weights. In both variants we can show constant competitive ratios, and in one case firms can even hire their best applicant with constant probability.

Finally, we conclude in Section 5 with a discussion of open problems.

1.3 Related Work

The secretary model is a classic approach to stopping problems and online admission [9, 24]. The classic algorithm outlined in the previous section is e-competitive, which is the best possible ratio. In the algorithmic literature, recent work has addressed secretary models for packing problems with random arrival of elements. A prominent case is the matroid secretary problem [3], for which the first general algorithm was $O(\log k)$-competitive, where k is the rank of the matroid. The ratio was very recently reduced to $O(\log\log k)$ [10, 23]. Constant-factor competitive algorithms have been obtained for numerous special cases [1, 8, 14, 17, 21, 27]. It remains a fascinating open problem whether a general constant-factor competitive algorithm exists or not.

Another popular domain is bipartite matching in the secretary model, which has many applications in online revenue maximization via ad-auctions. In Section 3 we use a variant of a recent optimal e-competitive algorithm [19], which tightened the ratio and improved it over previous algorithms [3, 7, 22]. The main idea has recently been extended to construct optimal secretary algorithms for packing linear problems [20], improving over previous approaches [6, 26]. Algorithms based on primal-dual techniques are a popular approach, especially for budgeted online matching with different stochastic input assumptions [5, 18, 25].

Our analysis of the algorithm for the general case applies in a unifying sampling model recently proposed as a framework for online maximum independent set in graphs [13]. It encompasses many stochastic adversarial models for online optimization – the secretary model, the prophet inequality model, and various other mixtures of stochastic and worst-case adversaries.

Closer to our paper are studies of a secretary problem with k queues [11], or game-theoretic approaches with complete knowledge about opponent strategies [15, 16]. These scenarios, however, have significantly different assumptions on the firms and their feedback, and they do not target markets with both decentralized control and restricted feedback that we explore in this paper.

2 General Preferences

For general weights $w : U \times V \to \mathbb{R}^+$, Proposition 1 shows that the classic secretary algorithm may perform poorly in a decentralized market. A reasonable strategy has to be more careful in adopting a threshold to avoid extensive competition over a few applicants. We overcome this obstacle with a randomized thresholding strategy, and we analyze it in a very general distributed matroid scenario. In the Appendix, we show that our bounds apply even within a general sampling model [13] that encompasses the secretary model, the prophet-inequality model, and many other approaches for stochastic online optimization.

For the combinatorial structure of the scenario, we consider that each firm u_i holds a possibly different matroid $\mathcal{S}_i$ over the set of applicants. Firm u_i can accept an applicant as long as the set of accepted applicants forms an independent set in $\mathcal{S}_i$. Special cases include hiring a single applicant or any subset of at most k_i many applicants. Each firm strives to maximize the sum of weights of hired applicants. The structure of $\mathcal{S}_i$ does not have to be known in advance. u_i only needs an oracle to test if a set of arrived applicants is an independent set in $\mathcal{S}_i$.

As a simple baseline, we can trivially obtain the following guarantee. Suppose there is an α-competitive algorithm $\mathcal{A}$' for a single firm. We assume that every firm u_i executes $\mathcal{A}$' in exactly the same way as if it was the only firm in the market. In particular, this adjusted algorithm $\mathcal{A}$ pretends that every applicant that gets an offer also accepts it.

Proposition 2. *Let algorithm $\mathcal{A}$' be any α-competitive algorithm for a single firm. Suppose every firm u_i runs a version $\mathcal{A}$ that pretends every applicant getting an offer from u_i also accepts it. Then algorithm $\mathcal{A}$ is $m\alpha$-competitive.*

Proof. For each firm u_i, consider the individual optimum M_i^* in hindsight. Clearly, there is one firm $u_{i'}$ for which this individual optimum has $w(M_{i'}^*) \geq w(M^*)/m$. Using $\mathcal{A}$', $u_{i'}$ makes offers to a set of applicants that constitute an α-approximation to the individual optimum. If an applicant decides against the offer of $u_{i'}$, it accepts a better offer from a different firm, so it secures an even larger weight in the solution $M^{\mathcal{A}}$. Hence, $\mathbb{E}[w(M^{\mathcal{A}})] \geq w(M_{i'}^*)/\alpha \geq w(M^*)/(m\alpha)$. ☐

This shows that a number of firms almost equal to the number of applicants is necessary for the lower bound in Proposition 1. Also, for general matroids it implies a competitive ratio of $O(m \log\log k_{\max})$ using the currently best algorithm [10, 23], where $k_{\max}$ is the maximum rank of any of the matroids $\mathcal{S}_i$. In the following section, we describe an algorithm that significantly improves upon this trivial guarantee when m grows large.

2.1 Logarithmic Approximation

Algorithm 1 is executed in parallel by all firms u_i. We first sample a fraction of roughly $n/2$ applicants and determine a random threshold based on the maximum weight seen by firm u_i in its sample. Firm u_i then greedily makes an offer only to those applicants whose values are above the threshold.

In line 4, the algorithm relies on an upper bound $b \geq |M^*|$. A simple example is $b = n$, which is always known and results in a $O(\log n)$-competitive algorithm. In case there is additional knowledge about the cardinality of the optimum solution, the guarantee can be improved. For example, if all firms know m and $k_{\max}$, then with $b = mk_{\max}$ the algorithm is $O(\log m + \log k_{max})$-competitive. In particular, if all firms know m in the basic model, the algorithm is $O(\log m)$-competitive.

Theorem 1. *Algorithm 1 is* $32(\lceil \log_2 b \rceil + 2)$*-competitive.*

Proof. We denote by V_i^S the set of applicants in the sample and by V_i^I the other applicants. Note that by the choice of sample and the random arrival, we have that $\mathbf{Pr}\left[v_j \in V_i^S\right] = \mathbf{Pr}\left[v_j \in V_i^I\right] = 1/2$. Thus, the sampling inequalities hold for every v_j, independently of $v_{j'} \in V_i^S$ or not for all other applicants $j' \neq j$. To see this, observe that one can simulate the algorithm by first assigning every v_j independently to V_i^S or V_i^I, then compute t_i, and finally consider applicants from V_i^I in random order[3].

[3] Such a simulation is used for the analysis of secretary algorithms in, e.g., [13, 22].

Algorithm 1: Thresholding algorithm for u_i with matroids.

Draw a random number $k \sim Binom(n, 1/2)$
Reject the first k applicants, denote this set by V_i^S
$m_i \leftarrow \arg\max_{v_j \in V_i^S} w(u_i, v_j)$
$X_i \leftarrow \text{Uniform}(0, 1, \ldots, \lceil \log_2 b \rceil + 1)$, where $b \geq |M^*|$
$t_i \leftarrow w(u_i, m_i)/2^{X_i}$, $M_i \leftarrow \emptyset$
for *all remaining* v_j *over time* **do**
  if $w(u_i, v_j) \geq t_i$ *and* $M_i \cup \{v_j\}$ *is independent set in* $\mathcal{S}_i$ **then**
    make an offer to v_j
    if v_j *accepts* **then** $M_i \leftarrow M_i \cup \{v_j\}$

Let $v_i^{\max} = \text{argmax}_j w(u_i, v_j)$ and $v_i^{\text{2nd}} = \text{argmax}_{j \neq v_i^{\max}} w(u_i, v_j)$ be the best and second best applicant for firm u_i, respectively. In addition, we denote by $w_i^{\max} = w(u_i, v_i^{\max})$ and $w_i^{\text{2nd}} = w(u_i, v_i^{\text{2nd}})$ their weights for firm u_i. For most of the analysis, we consider another weight function, the *capped weights* $\tilde{w}(u_i, v_j)$, based on thresholds t_i set by the algorithm as follows

$$\tilde{w}(u_i, v_j) = \begin{cases} w_i^{\max} & \text{if } v_j \in V_i^I, t_i = w_i^{\text{2nd}}, \text{ and } w(u_i, v_j) > w_i^{\text{2nd}}, \\ t_i & \text{if } v_j \in V_i^I \text{ and } w(u_i, v_j) \geq t_i, \\ 0 & \text{otherwise.} \end{cases}$$

Observe that the definition of $\tilde{w}$ relies on several random events, i.e., $v_j \in V_i^I$ and the choice of thresholds t_i. For any outcome of these events, however, we have that $\tilde{w}(u_i, v_j) \leq w(u_i, v_j)$ for all pairs (u_i, v_j), since if $t_i = w_i^{\text{2nd}}$ and $w(u_i, v_j) > w_i^{\text{2nd}}$, then $v_j = v_i^{\max}$ (recall that we assume no ties in w). By the following lemma, in expectation over all the correlated random events, an optimal offline solution with respect to $\tilde{w}$ still gives an approximation to the optimal offline solution with respect to w.

Lemma 1. *Denote by $w(M)$ and $\tilde{w}(M)$ the weight and capped weight of a solution M. Let $\tilde{M}^*$ and M^* be optimal solutions for $\tilde{w}$ and w, respectively. Then*

$$\mathbb{E}\left[\tilde{w}(\tilde{M}^*)\right] \quad \geq \quad \frac{1}{16(\lceil \log_2(b) \rceil + 2)} \cdot w(M^*).$$

Proof. Let $(u_i, v_j) \in M^*$ be an arbitrary pair. First, assume that v_j maximizes $w(u_i, v_j)$, i.e., $v_j = v_i^{\max}$. With probability at least $1/4$, we have $v_j \in V_i^I$ and $v_i^{\text{2nd}} \in V_i^S$. For any such outcome, we have with probability $1/(\lceil \log_2(b) \rceil + 2)$ that $t_i = w_i^{\text{2nd}}$ and $\tilde{w}(u_i, v_j) = w_i^{\max}$. This yields $\mathbb{E}[\tilde{w}(u_i, v_j)] \geq w(u_i, v_j)/(4(\lceil \log_2(b) \rceil + 2))$.

Second, for any $v_j \neq v_i^{\max}$ with $w_i^{\max}/(2b) < w(u_i, v_j) < w_i^{\max}$, we know $v_i^{\max} \in V_i^S$ is an independent event which happens with probability at least $1/2$. Then, there is some $0 \leq k' \leq \lceil \log_2(b) \rceil + 1$, with $w(u_i, v_j) > w_i^{\max}/2^{k'} \geq w(u_i, v_j)/2$. With probability $1/(\lceil \log_2(b) \rceil + 2)$, we have that $X_i = k'$ and $\tilde{w}(u_i, v_j) = t_i \geq w(u_i, v_j)/2$. This yields $\mathbb{E}[\tilde{w}(u_i, v_j)] \geq w(u_i, v_j)/(8(\lceil \log_2(b) \rceil + 2))$, since $v_j \in V_i^I$ with probability at least $1/2$.

Finally, we denote by $M^>$ the set of pairs $(u_i, v_j) \in M^*$ for which $w(u_i, v_j) > w_i^{\max}/(2b)$. The expected weight of the best assignment with respect to the threshold values is thus

$$\begin{aligned} \mathbb{E}\left[\tilde{w}(\tilde{M}^*)\right] \quad & \geq \quad \sum_{(u_i, v_j) \in M^*} \mathbb{E}[\tilde{w}(u_i, v_j)] \quad \geq \quad \sum_{(u_i, v_j) \in M^>} \frac{w(u_i, v_j)}{8(\lceil \log_2(b) \rceil + 2)} \\ & = \quad \frac{1}{8(\lceil \log_2(b) \rceil + 2)} \cdot (w(M^*) - w(M^* \setminus M^>)) \\ & \geq \quad \frac{1}{16(\lceil \log_2(b) \rceil + 2)} \cdot w(M^*), \end{aligned}$$

where the last inequality results from $\sum_{(u_i,v_j)\in M^*\setminus M^>} w_i^{\max}/(2b) \le \max_i w_i^{\max}/2 \le w(M^*)/2$. □

The previous lemma bounds the weight loss due to (i) all random choices inherent in the process of input generation and threshold selection and (ii) using the capped weights. The next lemma bounds the remaining loss due to random arrival of elements in V_i^I, exploiting that $\tilde{w}$ equalizes equal-threshold firms.

Lemma 2. *Suppose subsets V_i^I and thresholds t_i are fixed arbitrarily and consider the resulting weight function $\tilde{w}$. Let $M^{\mathcal{A}}$ be the feasible solution resulting from Algorithm 1 using the thresholds t_i, for any arbitrary arrival order of applicants in $\bigcup V_i^I$. Then $w(M^{\mathcal{A}}) \ge \tilde{w}(\tilde{M}^*)/2$.*

Proof. We will account for the weight of each edge $(u_i, v_j) \in \tilde{M}^*$ under $\tilde{w}$ by the original weight $w(e)$ of a pair $e \in M^{\mathcal{A}}$, using each pair in $M^{\mathcal{A}}$ at most twice. Let $(u_i, v_j) \in \tilde{M}^*$ be arbitrary. W.l.o.g. assume $\tilde{w}(u_i, v_j) > 0$, which implies that $w(u_i, v_j) \ge t_i$. We divide the edges of $\tilde{M}^*$ in two sets – depending on whether (1) v_j gets a better offer from another firm $u_{i'} \ne u_i$ with $w(u_{i'}, v_j) > w(u_i, v_j)$, or not.

First, consider the set of all edges $(u_i, v_j) \in \tilde{M}^*$ for firm u_i, for which (1) holds. Then, each such v_j is assigned to some other firm $u_{i'}$ that makes a better offer. Hence, $(u_{i'}, v_j) \in M^{\mathcal{A}}$ with $w(u_{i'}, v_j) \ge w(u_i, v_j) \ge \tilde{w}(u_i, v_j)$, and we can charge (u_i, v_j) to $(u_{i'}, v_j)$.

Now consider the set of all edges $(u_i, v_j) \in \tilde{M}^*$ for firm u_i, for which (1) does not hold. We denote their number by k_i. More generally, consider the superset $E_i^{t_i}$ of all edges (u_i, v_j) (in or not in M^*) with $w(u_i, v_j) \ge t_i$ such that (1) does not hold. These are all edges with $\tilde{w}(u_i, v_j) > 0$ that firm u_i is able to obtain. For the ones that u_i can obtain *simultaneously*, consider a maximum-cardinality independent set in matroid $\mathcal{S}_i$ that uses only edges from $E_i^{t_i}$. Denote this maximum cardinality by ℓ_i. Obviously, $\ell_i \ge k_i$, since $\tilde{M}^*$ is feasible for firm i. Moreover, firm u_i accepts in $M^{\mathcal{A}}$ at least ℓ_i applicants, irrespective of the arrival order. This is a simple consequence of the exchange property of matroid $\mathcal{S}_i$. Thus, for each of the k_i edges $(u_i, v_j) \in \tilde{M}^*$, for which (1) does not hold, there is some edge $(u_i, v_{j'}) \in M^{\mathcal{A}}$ with $w(u_i, v_{j'}) \ge \tilde{w}(u_i, v_{j'}) = \tilde{w}(u_i, v_j)$ to which we can charge it. This step can be done such that each $(u_i, v_j) \in \tilde{M}^*$ is charged to a different $(u_i, v_{j'}) \in M^{\mathcal{A}}$.

Finally, consider an arbitrary edge $(u_{i'}, v_{j'}) \in M^{\mathcal{A}}$. In the above accounting scheme, this edge can only be used to account for $\tilde{w}(u_i, v_j)$ with $(u_i, v_j) \in \tilde{M}^*$ if either $u_i = u_{i'}$ or $v_j = v_{j'}$. Since $\tilde{M}^*$ is a feasible assignment, $(u_{i'}, v_{j'})$ is used at most twice. This proves $\tilde{w}(\tilde{M}^*) \le 2w(M^{\mathcal{A}})$. □

Combining the insights of Lemma 1 and 2, we see that that $w(M^*) \le O(\log b) \cdot \mathbb{E}\left[w(M^{\mathcal{A}})\right]$, which proves the theorem. □

In Appendix B we extend the result of this theorem to a general sampling model that includes the secretary model, the prophet-inequality model and other models that combine stochastic and worst-case adversarial elements.

2.2 Lower Bound

Our general upper bound results from a thresholding-based algorithm. We constrast this result with a lower bound for thresholding-based algorithms in the basic model when every firm wants to hire only a single applicant. More formally, an algorithm $\mathcal{A}$ is called *thresholding-based* if during its execution $\mathcal{A}$ rejects applicants for some number of rounds, then determines a threshold T and afterwards enters an *acceptance phase*. In the acceptance phase, it makes an offer to exactly those applicants whose weight exceeds threshold T. Note that the number of rejecting rounds in the beginning and the threshold T can be chosen arbitrarily at random.

If the number of firms is unknown, the lower bound can be shown even for a class of *identical-firm* instances, in which for each applicant v_j all firms have the same weight, i.e., there is $w(v_j) \ge 0$ such that $w(u_i, v_j) = w(v_j)$ for every firm u_i. The main challenge here is to guess the right number of firms m in order to concentrate on the most profitable class of applicants. Note that in this class of instances we always have $m \le n$. Still, we manage to show a lower bound that depends on n, since m remains unknown and the best upper bound on $|M^*|$ that is known to the firms is n.

Alternatively, if the number of firms m is known, we can adjust the instance to contain an unknown number of firms that contribute only negligibly to social welfare. In fact, here we ensure that our instance satisfies $m = n$, and hence the lower bound can be expressed as $\Omega(\log m/(\log\log m))$. This almost matches our upper bound in the basic model when firms know m. More generally, one can always add more dummy applicants with weight 0 for all firms. This shows the lower bound in m continues to hold even more generally when $m \leq n$ and m known.

We show that the bound applies even in the iid model, which is a special case of the prophet-inequality model and can be seen as a special case of the secretary model. In the iid model, we draw the weight $w(v_j)$ for each v_j independently at random from a single distribution. The main difference from the secretary model is that M^* becomes a random variable.

Theorem 2. *Ssuppose every firm u_i uses some thresholding-based algorithm $\mathcal{A}_i$ in the basic model to hire a single applicant. There are instances I with $m \leq n$ on which the collection of algorithms $(\mathcal{A}_i)_{i\in m}$ has a competitive ratio of*

- $\Omega(\log n/\log\log n)$ *if firms do not know m. Here I is even an identical-firm instance.*
- $\Omega(\log m/\log\log m)$ *if firms know m.*

These lower bounds apply in the iid model and the secretary model.

Proof. We first prove the bound when firms firms do not know m. They can have IDs, but these IDs must not imply information about the number of firms m. For this case, the lower bound can be shown using an instance with identical firms. In the end of the proof, we make a simple observation how we can extend this construction to the case when m is known.

The proof follows a similar argument as in Proposition 1. Recall that in the proof of Proposition 1 there are two classes of 'good' and 'bad' applicants. Instead, here we construct an instance with logarithimcally many classes of applicants. The weights of applicants in larger classes is decreasing, but the total number of applicants in these classes is increasing. In total, the applicants in larger classes can generate more welfare – if there are enough firms to hire them. However, if the number of firms is too small, then offering to applicants from higher classes with lower weights is a bad strategy. These applicants generate small total weight, and since their numbers are significantly larger, the probability to hire an applicant from a lower class with high weight is very small. Thus, guessing the right order of magnitude of m represents the inherent difficulty on which the lower bound is based.

More formally, let $t \in \mathbb{N}, t \geq 3$ and $n = \sum_{j=2}^{t} t^{2j} = (t^{2t} - t^2) \cdot t^2/(t^2 - 1) = (1 + o(1)) \cdot t^{2t}$. Since $t^{2t} < n < t^{2t}(1+o(1))$, standard arguments for bounding t imply $t = \Theta(\log n/(\log\log n)$ with small constants hidden in the Θ-notation. We construct a distribution $\mathcal{I}$ on a family of identical-firm instances by drawing the weight $w(v_{j'})$ of each applicant $v_{j'}$ according to $\Pr[w(v_{j'}) = t^{-j}] = t^{2j}/n$ for $j = 2, \ldots, t$. In the secretary model, we may assume that each applicant draws $w(v_{j'})$ at the moment it arrives in the random order, since the order is chosen independently of the weights. Since all applicant weights are identically distributed, we may even completely disregard the random arrival order.

We define classes $C_2, \ldots, C_t$, where each class C_j consists of all applicants $v_{j'}$ with value $w(v_{j'}) = t^{-j}$. Consider how $\mathcal{A}$ performs on $\mathcal{I}$ for some firm u_i. We can assume that every algorithm $\mathcal{A}_i$ chooses a threshold among $\{t^{-2}, \ldots, t^{-t}\}$, since all other choices are equivalent concerning the set of applicants receiving an offer from u_i. Let X_{ij} be an indicator variable (resulting from randomness of $\mathcal{I}$ and the random choices of $\mathcal{A}_i$) that is 1 iff $\mathcal{A}_i$ picks threshold t^{-j} and 0 otherwise. Firms have access to shared as well as individual random bits (e.g., when sampling some of the initial applicants and picking a random individual threshold based on the observed values). As such, the random threshold choices of the firms might or might not be stochastically independent.

Let $F_j = \sum_{i=1}^{m} X_{ij}$ be the number of firms with threshold t^{-j}. Clearly, there is some $2 \leq k \leq t$ where $\mathbb{E}[F_j] \leq m/(t-1)$. In the following, we bound the expected competitive ratio of the collection of algorithms $(\mathcal{A}_i)_{i\in m}$ on $\mathcal{I}$ with $m = \sum_{j=2}^{k} t^{2j} < (9/8)t^{2k}$ firms. To obtain a competitive solution, most firms should choose a threshold of t^{-k}, hence the challenge is to guess m correctly and extract welfare from the right class of applicants.

Let us denote by T_i the threshold chosen by firm u_i. Consider the firms u_i with $T_i \geq t^{-(k-1)}$. Clearly, these firms can accept only applicants in $C_2, \ldots, C_{k-1}$. Observe that the total contribution to social welfare of all these applicants is bounded by at most $\sum_{j=2}^{k-1} |C_j| t^{-j}$. To bound the expectation of this term, consider iid random variables $Y_{j'}$ that take the value of the j'-th candidate if it falls into classes $C_2, \ldots, C_{k-1}$ and 0 otherwise. Formally, $Y_{j'} \in [0, 1]$ are iid with $\Pr[Y_{j'} = t^{-j}] = t^{2j}/n$ for $j = 2, \ldots, k-1$ and $\Pr[Y_{j'} = t^{-j}] = 0$ for $j = k, \ldots, t$. Then, since $t \geq 3$,

$$\mathbb{E}\left[\sum_{j=2}^{k-1} |C_j| t^{-j}\right] = \mathbb{E}\left[\sum_{j'=1}^{n} Y_{j'}\right] = \sum_{j=2}^{k-1} t^j < (3/2) t^{k-1} .$$

Now consider the set of firms choosing threshold $T_i = t^k$. Beyond the applicants from $C_2, \ldots, C_{k-1}$, their contribution to social welfare is at most $F_k \cdot t^{-k}$ from applicants in C_k. By the choice of k, we have

$$\mathbb{E}\left[F_k\right] t^{-k} \leq m/(t^k(t-1)) < (9/8)t^k/(t-1) \leq (27/16)t^{k-1} .$$

Finally, consider the firms with threshold $T_i = t^{-k'}$ with $k' \geq k+1$. We define $S = \sum_{j=2}^{k'} t^{2j}$. Let W_i be the expected value of an applicant matched to a firm u_i with threshold $T_i = t^{-k'}$. The expected value of W_i is bounded by

$$\mathbb{E}[W_i] \leq \sum_{j=2}^{k'} \frac{t^{2j}}{S} \cdot t^{-j} < \frac{3t^{k'}}{2t^{2k'}} \leq (3/2)t^{-(k+1)} , \tag{1}$$

since the next accepted applicant of value at least $t^{-k'}$ is distributed as $w(v_{j'})$ conditioned on containment in $\{t^{-2}, \ldots, t^{-(k+1)}\}$, except when we run out of applicants, in which case the value is zero.

Suppose that an applicant gets an offer by firm u_i, but decides to go to another firm. Since applicant weights are drawn iid, this can only have a negative influence on the expected value of accepted applicants for firm u_i. Hence, summarizing the arguments above, the expected social welfare of the assignment computed by the collection of algorithms $\mathcal{A} = (\mathcal{A}_i)$ is bounded from above by

$$\begin{aligned} \mathbb{E}_{(\mathcal{I},\mathcal{A})}[w(M^{\text{alg}})] &\leq \mathbb{E}\left[\sum_{j=2}^{k-1} |C_j| t^{-j}\right] + \mathbb{E}\left[F_k\right] t^{-k} + \mathbb{E}\left[\sum_{i=1}^{m} W_i\right] \\ &\leq (3/2)t^{k-1} + (27/16)t^{k-1} + (3/2)mt^{-(k+1)} \\ &< (3/2)t^{k-1} + (27/16)t^{k-1} + (27/16)t^{k-1} \\ &= (39/8)t^{k-1} \end{aligned}$$

For the optimum solution we bound $|C_j|$ for $j = 2, \ldots, k$. We use a standard Chernoff bound and observe that for every $j = 2, \ldots, k$ over $\mathcal{I}$

$$\Pr[|C_j| < (1/2)t^{2j}] \leq \exp(-t^{2j}/8) .$$

Using a union bound, $\sum_{j=2}^{k} |C_j| > (1/2) \sum_{j=2}^{k} t^{2j}$ with probability at least $1 - \sum j = 2^k \exp(-t^{2j}/8) \geq 1 - \exp(-t^2/8) > 0.67$, since $k \geq 2$ and $t \geq 3$. Since $m = \sum_{j=2}^{k} t^{2j}$, the optimum solution M^* has expected value at least

$$\mathbb{E}_{\mathcal{I}}[w(M^*)] \geq \sum_{j=2}^{k} |C_j| t^{-j} \cdot 0.67 > (2/3)t^k .$$

Recall that $t = \Omega(\log n / \log\log n)$. Thus, in the iid model, the ratio of expectations is at least

$$\frac{\mathbb{E}_{\mathcal{I}}[w(M^*)]}{\mathbb{E}_{\mathcal{I},\mathcal{A}}[w(M^{\mathcal{A}})]} > \frac{(2/3)t^k}{(39/8)t^{k-1}} > t/8 = \Omega\left(\frac{\log n}{\log\log n}\right).$$

Note that we did not optimize any constants, since we only intend to show that the asymptotics do not hide large constants. In fact, since the constants result from bounding exponential series and tail bounds, they can be expressed as functions of t that approach 1 quickly as t grows large. For the secretary problem, observe that $\mathbb{E}_{\mathcal{I}}[w(M^*)]$ is the average of $w(M^*)$ weighted over all possible instances $I \in \mathcal{I}$, while $\mathbb{E}_{\mathcal{I},\mathcal{A}}[w(M^{\mathcal{A}})]$ is the average over $\mathbb{E}_{\mathcal{A}}[w(M^{\mathcal{A}})]$ weighted over all possible instances $I \in \mathcal{I}$. As a simple consequence of standard calculus there must be some instance $I \in \mathcal{I}$ with competitive ratio

$$\frac{w(M^*)}{\mathbb{E}_{\mathcal{A}}[w(M^{\mathcal{A}})]} \quad \geq \quad \frac{\mathbb{E}_{\mathcal{I}}[w(M^*)]}{\mathbb{E}_{\mathcal{I},\mathcal{A}}[w(M^{\mathcal{A}})]} \quad = \quad \Omega\left(\frac{\log n}{\log\log n}\right).$$

This shows the lower bound on identical-firm instances when $\mathcal{A}$ does not know m and IDs do not imply information about m. We extend this result when each firm knows m and its (arbitrary) ID as follows. In the class of instances, we now always have $m = n$ firms. Each firm u_i knows its ID and runs any thresholding-based algorithm $\mathcal{A}_i$ (with shared and individual random bits). For the instance the adversary then picks m' as the number of *valuable* firms. For every non-valuable firm, every applicant value is multiplied with $\epsilon \ll t^{-t}$. Hence, the non-valuable firms do not contribute significantly to the value of any matching. Furthermore, the adversary picks these m' firms uniformly at random from the set of n firms. The lower bound continues to hold, since information about n and the firm ID does not help a thresholding-based algorithm to better estimate m', the number of firms that contribute significantly to $w(M^*)$. Since $m = n$, this also implies a lower bound of $\Omega(\log m/(\log\log m))$.

Finally, to see that this lower bound in m also holds for m known and $n \geq m$, we use the distribution $\Pr[w(v_{j'}) = t^{-j}] = t^{2j}/n$, and assign the remaining probability mass to $\Pr[w(v_{j'}) = 0] = 1 - \sum_{j=2}^{t} t^{2j}/n$. This has no influence on the expected values that upper bound $\mathbb{E}_{\mathcal{I},\mathcal{A}}[w(M^{\text{alg}})]$ or the Chernoff bound for lower bounding $w(M^*)$. □

The proof shows lower bounds if firms either know m and arbitrary IDs or are identical. If an instance with $m \leq n$ has identical firms that know m and their ID in $\{1, \ldots, m\}$, they can jointly implement algorithms for the k-secretary problem (such as, e.g., [1, 21]). For example, the Optimistic Algorithm of [1] can be turned into a thresholding-based algorithm that is e-competitive. All firms sample the first $\lfloor n/e \rfloor$ applicants. Then firm i sets its threshold to the weight of the i-th best candidate seen in the sample. Applicants have the same weight for all firms, and they are assumed to break ties such that u_i is preferred over $u_{i'}$ iff $i > i'$.

In Section 4.1 below we further explore this domain by using algorithms with multiple thresholds. This allows us to show the same result even when the weights are subject to small stochastic perturbations.

3 Independent Preferences

In this section, we show improved results for decentralized matching in the secretary model when preferences are independent among firms. More formally, we assume firm u_i has a set U_i of k_i positions available. An adversary specifies a separate set $\mathcal{P}_i$ of n *applicant profiles* for each firm u_i. An applicant profile $p \in \mathcal{P}_i$ is a function $p : U_i \to \mathbb{R}^+$. In round t, when a new applicant v_t arrives, we pick one remaining profile $p_{it} \in \mathcal{P}_i$ for each $u_i \in U$ independently and uniformly at random. The weight for position $u_{ij} \in U_i$ is then given by $w(u_{ij}, v_t) = p_{it}(u_{ij})$. We pick profiles from $\mathcal{P}_i$ uniformly at random without replacement. Special cases of this model are, e.g., when all weights for all positions are independently sampled from a certain distribution, or for each firm u_i the weights of all applicants are sampled independently from a different distribution for each position.

In contrast to the previous section, we assume that each applicant has k_i weight values for each firm u_i. A straightforward $O(\log n)$-competitive algorithm is to run Algorithm 1 separately for each position of each firm. In contrast, when $n \geq \sum_{i=1}^{m} k_i$ and $k_i \leq \alpha n$ for all $i \in [m]$ and some constant $\alpha \in (0, 1)$, we can achieve a constant competitive ratio using Algorithm 2. This algorithm resembles an optimal algorithm for secretary matching with a single firm [19]. Each firm rejects a number of applicants and enters an acceptance phase. In this phase, it maintains two virtual solutions: (1) an individual virtual optimum $M^*_{i,t}$ with respect to applicants arrived up to and including round t, and (2) a virtual solution M'_i where all applicants are

Algorithm 2: Matching algorithm for firm u_i for independent weights

Reject the first $r_i - 1$ applicants
$M_i, M'_i \leftarrow \emptyset$
for *applicant* v_t *arriving in round* $t = r_i, \ldots, n$ **do**
  Let $M^*_{i,t}$ be optimum matching for firm u_i and applicants $\{v_1, \ldots, v_t\}$
  if v_t *is matched to position* u_{ij} *in* M^*_i *and* u_{ij} *unmatched in* M'_i **then**
    Make an offer for position u_{ij} to v_t
    $M'_i \leftarrow M'_i \cup \{(u_{ij}, v_t)\}$
    if v_t *accepts* **then**
      $M_i \leftarrow M_i \cup \{(u_{ij}, v_t)\}$

assumed to accept the offers of u_i. If the newly arrived applicant v_t is matched in $M^*_{i,t}$, it is offered the same position unless this position is already filled in M'_i.

Note that for a single firm in the basic model, this algorithm reduces to the standard e-competitive algorithm discussed in the introduction. As such, our approach here is exactly the one we outlined above in Proposition 2. The main point here is to show that under the conditions on n and k_i, the properties of this specific algorithm combined with the independence among firms avoid the increase of m in the competitive ratio.

Theorem 3. *Algorithm 2 achieves a constant competitive ratio.*

Proof. Fix a firm u_i. The matching M'_i is constructed by assuming that u_i is the only firm in the market, i.e., every applicant accepts the offer of firm u_i. Consider the individual optimum $M^*_{i,n}$ in hindsight. Then, by repeating the analysis of [19, Section 2] and replacing the sampling size of $\lceil n/e \rceil$ by $r_i - 1$, the expected value of M'_i becomes

$$\mathbb{E}\left[w(M'_i)\right] \quad \geq \quad \sum_{\ell=r_i}^{n} \frac{r_i - 1}{\ell - 1} \cdot \frac{w(M^*_{i,n})}{n} \quad \geq \quad \frac{r_i - 1}{n} \ln\left(\frac{n}{r_i - 1}\right) \cdot w(M^*_{i,n}) \quad = \quad f(r_i) \cdot w(M^*_{i,n}) \ ,$$

where we denote the ratio by $f(r_i)$. Recall $k_i \leq \alpha n$. Set r_i in the interval $[\beta n, \gamma n]$ for some appropriate constants $\beta, \gamma \in (0, 1)$ such that $\beta > \alpha$. This ensures that $f(r_i)$ becomes a constant.

Let us now analyze the performance of the algorithm in the presence of competition. Consider applicant v_t in round t and the following events: (1) $P(u_i, v_t)$ is the event that u_i sends an offer to v_t, and (2) $A(u_i, v_t)$ is the event that u_i sends an offer to v_t and he accepts it. u_i's decision to offer depends only on $M^*_{i,t}$ and M'_i, but not on the acceptance decisions of earlier applicants. v_t for sure accepts an offer from u_i if u_i offers and no other firm offers. Offers from other firms $u_{i'}$ occur only if $u_{i'}$ is matched in $M^*_{i',t}$. More formally, $A(u_i, v_t)$ occurs (at least) if $P(u_i, v_t)$ and none of the $P(u_{i'}, v_t)$ occur. Since the profiles for different firms are combined independently

$$\mathbf{Pr}\left[A(u_i, v_t) \mid P(u_i, v_t)\right] \quad \geq \quad \prod_{i \neq i'} (1 - \mathbf{Pr}\left[P(u_{i'}, v_t)\right])$$

Consider the probability that v_t is matched in $M^*_{i',t}$. Since the order of profiles for $u_{i'}$ is independent of the order for u_i, we can imagine again choosing t profiles at random. Of those a random profile is chosen to be one of v_t. The t profiles determine $M^*_{i',t}$, which matches $\min(t, k_{i'})$ profiles. Since the last profile is determined at random, the probability that v_t is matched in $M^*_{i',t}$ is at most $\min(1, k_{i'}/t)$. As $t \geq r_{i'} \geq \beta n > \alpha n \geq k_{i'}$, we have

$$\mathbf{Pr}\left[P(u_{i'}, v_t)\right] \leq \begin{cases} 0 & \text{if } t \leq r_{i'} - 1, \\ k_{i'}/(\beta n) & \text{otherwise.} \end{cases}$$

Thus, for $t \geq r_{i'}$

$$\mathbf{Pr}\left[A(u_i,v_t) \mid P(u_i,v_t)\right] \geq \prod_{i\neq i'}(1-\mathbf{Pr}\left[P(u_{i'},v_t)\right]) \geq \exp\left(\sum_{i=1}^{m}\ln\left(1-\frac{k_i}{\beta n}\right)\right)$$
$$\geq \exp\left(-\sum_{i=1}^{m}\frac{1}{1-(\alpha/\beta)}\cdot\frac{k_i}{\beta n}\right) \geq \exp\left(-\frac{1}{\beta-\alpha}\right).$$

The third inequality follows since $k_i \leq \alpha n$, we have that $(1-k_i/(\beta n)) > 1-\alpha/\beta$. Furthermore, it holds that $\ln(1-x) \geq -\frac{x}{1-x}$ for all $x \in (0,1)$ (see Fact 1 in the Appendix). The last inequality is due to $n \geq \sum_j k_j$.

Consequently, $\mathbb{E}\left[w(M_i)\right]$ recovers at least a constant fraction of $\mathbb{E}\left[w(M_i')\right]$, which represents a constant factor approximation to the individual optimum $M_{i,n}^*$ for i in hindsight. By linearity of expectation, the algorithm achieves a constant competitive ratio for the expected weight of the optimum matching. □

4 Correlated Preferences

In this section, we treat the basic model where every firm strives to hire one applicant. In this section, we assume that $m \leq n$, i.e., there are more candidates than firms. We consider stochastic input generation which allows correlations on the weights incident to an applicant. Specifically, assume that each applicant v_i has a parameter q_i, measuring his built-in quality, and the weights of edges incident to v_i are generated independently from a distribution D_i with mean q_i and standard deviation σ. Note that the lower bound for the classical e-competitive algorithm for the secretary problem (Proposition 1) applies to this general setting. As a natural candidate, we consider in particular normal distributions and assume that $D_i \sim N(q_i, \sigma^2)$.

We analyze correlations in two regimes: When the random noise is small and the preference lists of each firm are unlikely to differ and when large variance has substantial effects on the preferences. In the two cases we apply different algorithms, and both can achieve constant competitive ratios.

4.1 Small Variance

We consider the case of highly correlated preferences of an applicant to all firms with possibly small fluctuations around an applicant's quality. Intuitively, all the firms are facing almost the same situation. If we apply the classic algorithm of letting each firm set a single threshold, as we have seen in the previous sections, there will be extensive competition over few 'good' candidates. So instead, we use an approach using multiple thresholds for each firm.

Consider the list-based approach of Algorithm 3. The algorithm knows the number m of firms. After sampling a linear number $r = \Theta(n)$ of applicants, it maintains a list of the top m candidates observed so far. The key observation we exploit is that Algorithm 3, in contrast to the classical algorithm for the secretary problem, can cope well with competition, provided that applicants have a global quality that all firms roughly agree on. In particular, each of the top m applicants will be matched to her best firm with constant probability.

Without loss of generality, let $q_1 \geq \cdots \geq q_n$. Formally, we define $\delta_{min} := \min_{i\neq j}|q_i - q_j|$, and $\psi = \frac{\delta_{min}}{\sigma}$.

Theorem 4. *Let $\psi = \omega(n)$ and $r \in [\beta n, \gamma n]$ for constants $\beta, \gamma \in (0,1)$, $\beta \leq \gamma$. Algorithm 3 achieves a constant competitive ratio with high probability, i.e., with probability approaching 1 over the choice of weights, we have $\mathbb{E}\left[w(M^*)\right] \leq c \cdot \mathbb{E}\left[w(M^{\mathcal{A}})\right]$ for some constant $c > 0$.*

Proof. Our analysis proceeds as follows. We first assume that all firms share the same uniform preference list $(v_1,\ldots,v_n)$ over the applicants, based on which they make offers. This list ranks applicants in order of q_i and not necessarily by the realized random weights $w(u_i,v_j)$. We show that under this condition, the best m applicants get matched to their best firm with constant probability. We then show that with high probability the weights induce a uniform preference list. This yields the theorem.

Algorithm 3: List-based algorithm for firm u

Initialize list $L_u = (\ell_{u,1}, \dots, \ell_{u,m})$, initialized with $(-\infty, \dots, -\infty)$
(maintain L_u to contain the top m weights u observed so far, where $\ell_{u,1} \geq \cdots \geq \ell_{u,m}$)
Reject the first $(r-1)$ applicants, denote the set by R
for *applicant v_t arriving in round $t = r, \dots, n$* **do**
  if $w(u, v_t) \geq \ell_{u,m}$ **then**
    Update L_u: Push w_{u,v_t} into L_u and pop $\ell_{u,m}$ out.
    if *popped out $\ell_{u,m} = -\infty$ or corresponds to an applicant in R* **then**
      Make an offer to v_t, stop if v_t accepts

For convenience, we allow each firm to send offers even after it has been matched, i.e., it still sends *virtual offers*, which will always be rejected. Since we assume all firms have uniform preference lists $(v_1, \dots, v_n)$, if an applicant receives an offer from a firm, every other firm also sends her an offer which might be virtual. Note that in the algorithm every firm sends out offers at most m times, thus no more than m applicants can receive offers from the same firm. It follows that when receiving *some* (potentially virtual) offer, an applicant also sees non-virtual offers and chooses the best from them.

Observe that for uniform preference lists, our algorithm is similar to the Virtual Algorithm in [1] for the k-secretary problem. Let S be the set of all applicants who receive any offer. For following lemma see [1, Lemma 1]. Note that this probability is constant for $r \in (\beta n, \gamma n)$.

Lemma 3. *For each applicant v_j with $j \in \{1, \dots, m\}$ we have*

$$\boldsymbol{Pr}[v \in S] \quad \geq \quad f(r) \quad = \quad \frac{r-1}{n} \ln\left(\frac{n}{r-1}\right) \quad .$$

Lemma 4. *Let $r \in (\beta n, \gamma n)$ and all firms have a uniform preference list $(v_1, \dots, v_n)$. Suppose for each applicant v_i, all incident weights $\{w(u, v_i) \mid u \in U\}$ are drawn independently from the same distribution D_i. Then for each $1 \leq k \leq m$, we have that v_k is matched to her best firm with constant probability.*

Proof. Denote the (random) arrival order of applicants by τ and let $s_{\tau,i}$ be the i-th applicant who receives offers. Fix an applicant $v = v_k$, $1 \leq k \leq m$, among the best m applicants.

First, for every τ where $s_{\tau,j} = v$ and $j > 1$, by swapping the position between $s_{\tau,j-1}$ and v we can obtain a new order τ'. Note that in τ', v is arriving earlier than in τ. Since the entry $\ell_{u,m}$ is monotonically increasing over time, v will also be pushed into L_u at this earlier time. Since $s_{\tau,j-1}$ was made an offer, it popped out $-\infty$ or an applicant in R, which now also happens since the state of L is the same and v is pushed into it. As such, in the new arrival order τ', v becomes the $(j-1)$-th applicant to receive offers, i.e., $s_{\tau',j-1} = v$. Clearly, for two different arrival orders τ_1 and τ_2 with $s_{\tau_1,j} = s_{\tau_2,j} = v$, the corresponding new orders τ_1' and τ_2' are also different. Thus $|\{\tau | s_{\tau,j-1} = v\}| \geq |\{\tau | s_{\tau,j} = v\}|$. Therefore

$$\Pr_{\tau}[s_{\tau,j-1} = v \mid v \in S] \geq \Pr_{\tau}[s_{\tau,j} = v \mid v \in S], \qquad \text{for all } j > 1.$$

Now given that $s_{\tau,j} = v$, among the m offers v has received, $j-1$ of them are virtual and must be rejected. If the best offer for v is among the remaining $m-j+1$ ones, then v will get her best offer. Since all the weights of edges incident to v are generated independently from the same distribution, this event occurs with probability $\frac{m-j+1}{m}$ and is decreasing in j, therefore

$$\begin{aligned}
\mathbf{Pr}\left[v_i \text{ assigned to its best firm} \mid v_i \in S\right] \quad &= \quad \sum_{j=1}^{m} \mathbf{Pr}\left[s_{\tau,j} = v_i \mid v_i \in S\right] \cdot \frac{m-j+1}{m} \\
&\geq \quad \frac{1}{m} \sum_{j=1}^{m} \frac{m-j+1}{m} \quad = \quad \frac{1}{2} + \frac{1}{2m} \quad .
\end{aligned}$$

where the inequality follows from Chebyshev's sum inequality and $\sum_{j=1}^{m} \mathbf{Pr}[s_{\tau,j} = v_i \mid v_i \in S] = 1$. Combining this with $\mathbf{Pr}[v \in S] \geq \frac{r}{n} \ln(\frac{n}{r})$, the claim follows. □

When the fluctuations in applicants' quality are small enough to have a small effect on the preference lists, it is easy to extend the result to show a constant competitive ratio for the case of small variance. We first show that in this regime, indeed the fluctuations keep the same preference lists uniform with high probability.

Lemma 5. *Let $\psi = \omega(n)$. Then for any given sequence $q_1 > \cdots > q_n$, with probability approaching 1 it holds that*

(1) for each applicant v_i and firm u, we have $|w(u, v_i) - q_i| < \delta_{min}/2$,

(2) all firms have a uniform preference list of applicants $(v_1, \ldots, v_n)$.

Proof. For some applicant v_i and firm u, note that $w(u, v_i)$ is sampled independently from D_i with mean q_i and standard deviation σ. By Chebyshev's inequality, we conclude

$$\Pr\left[|w(u, v_i) - q_i| \leq \frac{\delta_{min}}{2}\right] \quad \leq \quad \frac{4\sigma^2}{\delta_{min}^2} \quad = \quad \frac{4}{\psi^2} \ .$$

Using a union bound, this event holds for all applicants and firms with probability at least $1 - \frac{4nm}{\psi^2}$. Given that $m \leq n$ and $\psi = \omega(n)$, this probability approaches 1 when n goes to infinity. This proves that part (1) of the lemma holds with high probability.

For the part (2), assume that part (1) holds. Fix a particular firm u, and let x_i be $x_i = w(u, v_i)$. Recall that by assumption $q_1 > q_2 > ... > q_n$. Given that $q_i - q_{i-1} \geq \delta_{min}$, we conclude that

$$x_i - x_{i-1} \quad \geq \quad (q_i - q_{i-1}) - |x_i - q_i| - |x_{i-1} - q_{i-1}| \quad > \quad (q_i - q_{i-1}) - \delta_{min} \quad \geq \quad 0 \ .$$

Hence, $x_1 > \cdots > x_n$ and u has the preference list $(v_1, \ldots, v_n)$. □

Note that Lemma 5 allows us to apply Lemma 4 by letting D_i be the truncated distribution obtained by conditioning $w(u, v_i) \sim D_i$ to be contained in $(q_i - \delta_{min}/2, q_i + \delta_{min}/2)$.

Finally, for the proof of the theorem, let $T = \{v_1, \ldots, v_m\}$ be the set of the m applicants with highest mean q_i. Lemma 5 shows that with probability approaching 1, all firms have a uniform preference list. In this case, $w(M^*) \leq \sum_{v_i \in T} \max_{u \in U} w(u, v_i)$. Also, according to Lemmas 4 and 5 the algorithm guarantees that every $v_i \in T$ will be matched to her best firm with constant probability. By linearity of expectation, $\mathbb{E}\left[w(M^{\mathcal{A}})\right] \geq \sum_{v_i \in T} c \cdot \max_{u \in U} w(u, v_i)$ for some constant $c > 0$, concluding the result. □

As a corollary, we obtain the following.

Corollary 1. *Each firm has a probability of $\Omega(1/m)$ to obtain the best applicant.*

Proof. By Lemma 5, with probability approaching 1 all firms consider the same applicant as the best. Denote the best applicant by v. By the fact that $\mathbf{Pr}[v \in S] \geq \frac{r}{n} \ln(\frac{n}{r})$ is a constant, v is matched to some firm with constant probability. Since there is no difference between the firms, each firm has a probability of $\frac{1}{m}$ to be chosen. □

4.2 Large Variance

If weights are perturbed by high-variance normal distributions, then intuitively this results in a similar scenario as with independent preferences. For this scenario we proved above that the classic algorithm (more precisely, its generalization to multiple positions per firm) achieves a constant competitive ratio. For completeness, we state the classic algorithm again as Algorithm 4.

Algorithm 4: The classic secretary problem algorithm for firm u.

Reject the first $(r-1)$ applicants, denote the set by R
$T_u \leftarrow \max_{j \in R} w(u, v_j)$
for *applicant* v_t *arriving in round* $t = r, \ldots, n$ **do**
  if $w(u, v_t) \geq T_u$ **then**
    Make an offer to v_t, stop if v_t accepts

An instance is given by the distributions $(D_i)_{i=1\ldots,n}$ from which the values of applicants $i = 1, \ldots, n$ are drawn independently. We call such an instance *strong-tailed* if for any subset of r distributions $D_{i_1}, \ldots, D_{i_r}$ with $X_{i_j} \sim D_{i_j}$ we have

$$\mathbf{Pr}\left[\exists 1 \leq j \leq r-1 : X_{i_j} > X_{i_r}\right] \quad \geq \quad \left(1 - \frac{c}{r}\right)$$

for a constant $c > 0$.

Theorem 5. *The classic secretary algorithm achieves a constant competitive ratio for strong-tailed instances when $r \in [\beta n, \gamma n]$ for constants $\beta, \gamma \in (0,1)$, $\beta \leq \gamma$. Each firm hires its best applicant with constant probability.*

Proof. In the secretary model, the adversary fixes the distributions, draws the match values $w(u, v_j) \sim D_j$ independently at random for each firm u, and then applicants arrive in uniform random order. Here it seems convenient to number applicants $\{v_1, v_2, \ldots, v_n\}$ in their *arrival order* (and not based on their expected quality as in the previous section). We denote v_j's distribution by D_j. Consequently, since arrival order is random, D_j becomes a random variable itself.

Now consider $B(u, v_i), P(u, v_i)$, and $A(u, v_i)$ as the events that v_i is the best applicant for u, v_i receives an offer from u, and v_i receives and accepts an offer from u, respectively. By random arrival, for each $i \geq r$, we have $\mathbf{Pr}[B(u, v_i)] \geq 1/n$ and $\mathbf{Pr}[P(u, v_i) \mid B(u, v_i)] \geq \frac{r-1}{i-1}$ (for the latter, note that if the most valuable applicant among $v_1, \ldots, v_{i-1}$ is among $v_1, \ldots, v_{r-1}$, u cannot have made a previous offer). Hence,

$$\mathbf{Pr}[u \text{ gets its best applicant}] \quad \geq \quad \sum_{i=r}^{n} \frac{1}{n} \cdot \frac{r-1}{i-1} \cdot \mathbf{Pr}[A(u, v_i) \mid B(u, v_i), P(u, v_i)] \quad .$$

If no other firm sends an offer to v_i, $A(u, v_i)$ must be true if $P(u, v_i)$ holds. For any $u' \neq u$, $P(u', v_i)$ depends on $\{D_i\}_{i=1}^n$ (and, hence, the arrival order) and random coin flips when drawing the exact value of applicant v_j for firm u'. Note that the latter random coin flips are independent for the firms. Thus, given $\{D_i\}_{i=1}^n$ (and, hence, the arrival order) the events $\{P(u', v_i) \mid u' \neq u\}$ are independent from each other, and are all independent from $B(u, v_i)$ and $P(u, v_i)$. Therefore

$$\mathbf{Pr}[A(u, v_i) \mid B(u, v_i), P(u, v_i)] \quad \geq \quad \mathbb{E}_{\{D_i\}_{i=1}^n}\left[\prod_{u' \neq u} \mathbf{Pr}\left[\overline{P(u', v_i)} \mid \{D_i\}_{i=1}^n\right] \middle| B(u, v_i), P(u, v_i)\right] \quad .$$

Note that if firm u' saw a better applicant during the sampling phase, then $\overline{P(u', v_i)}$ is implied[4]. Since the instance is strong-tailed, we see

$$\mathbf{Pr}\left[\overline{P(u', v_i)} \mid \{D_i\}_{i=1}^n\right] \quad \geq \quad \left(1 - \frac{c}{r}\right) \quad ,$$

and thus

$$\mathbf{Pr}[A(u, v_i) \mid B(u, v_i), P(u, v_i)] \quad \geq \quad \left(1 - \frac{c}{r}\right)^{m-1} \quad .$$

[4]In fact, $\overline{P(u', v_i)}$ holds more often, e.g., if firm u' saw only worse applicants so far but hired an applicant in a previous round.

Using $m \le n$, we compute

$$\mathbf{Pr}\left[u \text{ gets its best applicant}\right] \quad \ge \quad \sum_{i=r}^{n} \frac{1}{n} \cdot \frac{r-1}{i-1} \cdot \left(1 - \frac{c}{r}\right)^{m-1} \quad \ge \quad \left(\left(1-\frac{c}{r}\right)^{r}\right)^{\frac{n}{r}} \sum_{i=r}^{n} \frac{1}{n} \cdot \frac{r-1}{i-1} \ .$$

Since $r \in [\beta n, \gamma n]$ this term can be bounded from below by a constant depending on constants c, β, γ.

This implies that the classic algorithm for the secretary problem is almost an optimal strategy for firms. It guarantees for every single firm the best outcome with constant probability. By linearity of expectation, the expected social welfare is at least a constant fraction of the optimum, completing the proof of Theorem 5. □

For an example of strong-tailed distributions, we consider normal distributions $D_i = N(q_i, \sigma^2)$. Similar to the previous section, we rely on two parameters to bound properties of these distributions. In contrast to $\delta_{\min}$ and ψ based on the minimum difference in applicants' qualifications, we here use the maximum difference. Let $\delta_{max} = \max_{i \ne j} |q_i - q_j|$ and set $\varphi = \frac{\delta_{max}}{\sigma}$.

Lemma 6. *Let $D_i = N(q_i, \sigma^2)$ for all $i = 1, \dots, n$, with $\varphi = O(1/n^2)$. Then the instance is strong-tailed.*

Proof. Due to the definition of $\delta_{\max}$, the probability $\mathbf{Pr}\left[\exists 1 \le j \le r-1 : X_{i_j} > X_{i_r}\right]$ is minimized if all distributions $D_{i_j} = N(q_{i_r} - \delta_{\max}, \sigma^2)$ for $1 \le j \le r-1$.

Thus, let $X_1, \dots, X_{r-1}, Y \sim N(q_i - \delta_{max}, \sigma^2)$, $X_i \sim N(q_i, \sigma^2)$. Also, consider $X_i' \sim N(0, \sigma^2)$ and $Y' \sim N(-\delta_{max}, \sigma^2)$. This implies

$$\begin{aligned}
\mathbf{Pr}\left[\exists 1 \le j \le r-1 : X_j > X_i\right] &= 1 - \prod_{j=1}^{r-1} \mathbf{Pr}\left[X_j \le X_i\right] \\
&= 1 - \left(\mathbf{Pr}\left[Y \le X_i\right]\right)^{r-1} \\
&= 1 - \left(\mathbf{Pr}\left[(Y - q_i) \le (X_i - q_i)\right]\right)^{r-1} \\
&= 1 - \int_{-\infty}^{\infty} f_{X_i'}(x_i) \cdot (F_{Y'}(x_i))^{r-1} \mathrm{d}x_i \\
&= 1 - \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{x_i^2}{2\sigma^2}} \left(\frac{1}{2} + \frac{1}{2}\mathrm{erf}\left(\frac{x_i + \delta_{max}}{\sqrt{2}\sigma}\right)\right)^{r-1} \mathrm{d}x_i \ ,
\end{aligned}$$

where $\mathrm{erf}(\cdot)$ is the error function

$$\mathrm{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} \mathrm{d}t \ .$$

With

$$\frac{\mathrm{d}(\frac{1}{2} + \frac{1}{2}\mathrm{erf}(\frac{x}{\sqrt{2}\sigma}))^{r-1}}{\mathrm{d}x} \quad = \quad \frac{r-1}{\sqrt{2\pi}\sigma} e^{-\frac{x^2}{2\sigma^2}} \left(\frac{1}{2} + \frac{1}{2}\mathrm{erf}\left(\frac{x}{\sqrt{2}\sigma}\right)\right)^{r-2} \quad \le \quad \frac{r-1}{\sqrt{2\pi}\sigma} \ ,$$

by the Mean Value Theorem, we have

$$\left(\frac{1}{2} + \frac{1}{2}\mathrm{erf}\left(\frac{x_i + \delta_{max}}{\sqrt{2}\sigma}\right)\right)^{r-1} \quad \le \quad \left(\frac{1}{2} + \frac{1}{2}\mathrm{erf}\left(\frac{x_i}{\sqrt{2}\sigma}\right)\right)^{r-1} + \frac{r-1}{\sqrt{2\pi}\sigma} \cdot \delta_{max} \ .$$

Therefore,

$$\begin{aligned}
&\int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{x_i^2}{2\sigma^2}} \left(\frac{1}{2} + \frac{1}{2}\mathrm{erf}\left(\frac{x_i + \delta_{max}}{\sqrt{2}\sigma}\right)\right)^{r-1} \mathrm{d}x_i \\
\le \quad &\int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{x_i^2}{2\sigma^2}} \left(\frac{1}{2} + \frac{1}{2}\mathrm{erf}\left(\frac{x_i}{\sqrt{2}\sigma}\right)\right)^{r-1} \mathrm{d}x_i + \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{x_i^2}{2\sigma^2}} \frac{(r-1)\delta_{max}}{\sqrt{2\pi}\sigma} \mathrm{d}x_i \ .
\end{aligned}$$

Note that the first term is exactly the probability that x_i is the highest value among r random variables drawn independently from the same normal distribution $N(0, \sigma^2)$, which equals $\frac{1}{r}$. For the second term

$$\int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{x_i^2}{2\sigma^2}} \frac{(r-1)\delta_{max}}{\sqrt{2\pi}\sigma} \mathrm{d}x_i \quad = \quad \frac{(r-1)\delta_{max}}{\sqrt{2\pi}\sigma} \int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{x_i^2}{2\sigma^2}} \mathrm{d}x_i \quad = \quad \frac{(r-1)\delta_{max}}{\sqrt{2\pi}\sigma} \ ,$$

thus,

$$\int_{-\infty}^{\infty} \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{x_i^2}{2\sigma^2}} \left(\frac{1}{2} + \frac{1}{2} \mathrm{erf}\Big(\frac{x_i + \delta_{max}}{\sqrt{2}\sigma} \Big) \right)^{r-1} \mathrm{d}x_i \quad \leq \quad \frac{1}{r} + \frac{r-1}{\sqrt{2\pi}} \frac{\delta_{max}}{\sigma} \quad = \quad \frac{1}{r} + \frac{r-1}{\sqrt{2\pi}} \varphi \ .$$

Note that $\varphi \leq \frac{c}{r(r-1)}$ for some constant c. Thus,

$$\mathbf{Pr}\left[\exists 1 \leq j \leq r-1 : X_j > X_i\right] \quad \geq \quad 1 - \left(\frac{1}{r} + \frac{r-1}{\sqrt{2\pi}} \varphi \right) \quad \geq \quad 1 - \frac{1 + (c/2\pi)}{r} \ .$$

This shows the lemma. □

Theorem 5 implies the following result.

Corollary 2. *Let $D_i = N(q_i, \sigma^2)$, for all $i = 1, \ldots, n$, such that $\varphi = O(1/n^2)$, then the classic secretary algorithm achieves a constant competitive ratio. Each firm hires its best applicant with constant probability.*

Finally, observe that a relatively high variance, i.e. $\varphi = O(1/n^2)$, is not sufficient for the result.

Proposition 3. *There are distributions $(D_i)_{i=1,\ldots,n}$ with the same variance σ^2 such that $\varphi = O(1/n^2)$ and the classic secretary algorithm has a competitive ratio of $\Omega(n/\log n)$.*

Proof. The proof is a simple adjustment of Proposition 1. Recall the identical-firm instance I' with $m = n$ firms, where $\gamma \ln n$ applicants are 'good' and the remaining ones are bad. For each good applicant we construct a distribution D_i by setting its value to 2 with probability $1 - (1/n^6)$ and to $(n^5 + 2)$ with probability $1/n^6$. Thus, for a good applicant we have expected value $q_i = 2 + (1/n)$. Similarly, for each bad applicant the distribution D_i sets the value to 1 with probability $1 - (1/n^6)$ and to $(n^5 + 1)$ with probability $1/n^6$, which implies an expected value of $q_i = 1 + (1/n)$. This implies $\delta_{\max} = 1$, $\sigma^2 = \Omega(n^4)$, and $\varphi = O(1/n^2)$.

In the resulting random class $\mathcal{I}$ of instances, the original instance I' analyzed in Proposition 1 appears with probability $p_0 = (1 - (1/n^6))^n$. Then the expected value obtained by the algorithm is $\mathbb{E}_{I',\mathcal{A}}[w(M^{\text{alg}})] = O(\ln n)$. Otherwise, we overestimate the performance of the algorithm by assuming it computes the optimum that matches all applicants. Let $X_i \in \{0, 1\}$ be the indicator variable with $X_i = 1$ if applicant i has a value at least $n^5 + 1$. For $X = \sum_i X_i$, note that $\mathbb{E}[X] = 1/n^5$ and $\mathbb{E}[X \mid X \geq 1] = 1/(n^5(1 - p_0))$. This implies

$$\begin{aligned} \mathbb{E}_{\mathcal{I},\mathcal{A}}[M^{\text{alg}}] &< p_0 \cdot \mathbb{E}_{I',\mathcal{A}}[w(M^{\text{alg}})] + (1 - p_0) \cdot \left(n + \gamma \ln n + (n^5 + 2) \frac{1}{n^5(1 - p_0)} \right) \\ &< \mathbb{E}_{I',\mathcal{A}}[w(M^{\text{alg}})] + \frac{n + \gamma \ln n}{2n^5} + \frac{n^5 + 2}{n^5} \\ &< \mathbb{E}_{I',\mathcal{A}}[w(M^{\text{alg}})] + 2 \ , \end{aligned}$$

where the second inequality results from $1 - p_0 = 1 - (1 - (1/n^6))^n \leq 1 - \exp(-1/n^5) \leq 1/(2n^5)$. Thus, overall the algorithm only obtains an expected value of $O(\log n)$. Obviously, the social optimum obtains a value of $n + \gamma \ln(n)$ with probability p_0, where $p_0 > (1 - 1/n^5)$ by a union bound. □

5 Conclusion

In this paper, we have studied online secretary problems with n applicants, m firms, and local information. We analyze thresholding-based algorithms and show how to obtain a competitive ratio of $O(\log n)$, even if firms can accept sets of applicants based on a local matroid. Moreover, we show a lower bound of $O(\log n/(\log \log n))$ if all firms use thresholding-based algorithms. These bounds continue to hold in terms of m if every firm knows the number of firms in the market. For more structured domains, we show that constant competitive ratios can be achieved.

For the general case, the main point is to determine the right order of magnitude of firms that contribute significantly to social welfare. In the basic model, if a firm makes no offer or the applicant accepts, then the firm does not obtain feedback that allows to learn the number of (better) firms competing in the market. Such feedback is mostly generated by rejected offers. Thresholding-based algorithms are unable to learn based on this, since the threshold is set before any offer is made and does not get adjusted based on rejected offers. As such, these algorithms are restricted to simply guessing the number of relevant firms, which results in almost tight logarithmic ratios.

It is a fascinating open problem if our bounds can be improved, in general and for meaningful special cases. Intuitively, the feedback would have to be used to guide the offering decisions. A starting point might be domains where each firm has multiple positions (say, every firm has $k = \Theta(\log n)$ positions), where initial offers can be made simply to explore the market and generate feedback. Possibly, ideas from contention resolution protocols might be helpful in designing such protocols.

References

[1] Moshe Babaioff, Nicole Immorlica, David Kempe, and Robert Kleinberg. A knapsack secretary problem with applications. In *Proc. 10th Workshop Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, pages 16–28, 2007.

[2] Moshe Babaioff, Nicole Immorlica, David Kempe, and Robert Kleinberg. Online auctions and generalized secretary problems. *SIGecom Exchanges*, 7(2), 2008.

[3] Moshe Babaioff, Nicole Immorlica, and Robert Kleinberg. Matroids, secretary problems, and online mechanisms. In *Proc. 18th Symp. Discrete Algorithms (SODA)*, pages 434–443, 2007.

[4] Niv Buchbinder, Kamal Jain, and Mohit Singh. Secretary problems via linear programming. In *Proc. 14th Intl. Conf. Integer Programming and Combinatorial Optimization (IPCO)*, pages 163–176, 2010.

[5] Nikhil Devanur and Thomas Hayes. The adwords problem: Online keyword matching with budgeted bidders under random permutations. In *Proc. 10th Conf. Electronic Commerce (EC)*, pages 71–78, 2009.

[6] Nikhil Devanur, Kamal Jain, Balasubramanian Sivan, and Christopher Wilkens. Near optimal online algorithms and fast approximation algorithms for resource allocation problems. In *Proc. 12th Conf. Electronic Commerce (EC)*, pages 29–38, 2011.

[7] Nedialko Dimitrov and Greg Plaxton. Competitive weighted matching in transversal matroids. *Algorithmica*, 62(1-2):333–348, 2012.

[8] Michael Dinitz and Guy Kortsarz. Matroid secretary for regular and decomposable matroids. *SIAM J. Comput.*, 43(5):1807–1830, 2014.

[9] Eugene Dynkin. The optimum choice of the instant for stopping a Markov process. In *Sov. Math. Dokl*, volume 4, pages 627–629, 1963.

[10] Moran Feldman, Ola Svensson, and Rico Zenklusen. A simple $O(\log\log(\text{rank}))$-competitive algorithm for the matroid secretary problem. In *Proc. 26th Symp. Discrete Algorithms (SODA)*, pages 1189–1201, 2015.

[11] Moran Feldman and Moshe Tennenholtz. Interviewing secretaries in parallel. In *Proc. 13th Conf. Electronic Commerce (EC)*, pages 550–567, 2012.

[12] Thomas Ferguson. Who solved the secretary problem? *Statistical Science*, 4(3):282–289, 1989.

[13] Oliver Göbel, Martin Hoefer, Thomas Kesselheim, Thomas Schleiden, and Berthold Vöcking. Online independent set beyond the worst-case: Secretaries, prophets and periods. In *Proc. 41st Intl. Coll. Automata, Languages and Programming (ICALP)*, volume 2, pages 508–519, 2014.

[14] Sungjin Im and Yajun Wang. Secretary problems: Laminar matroid and interval scheduling. In *Proc. 22nd Symp. Discrete Algorithms (SODA)*, pages 1265–1274, 2011.

[15] Nicole Immorlica, Adam Kalai, Brendan Lucier, Ankur Moitra, Andrew Postlewaite, and Moshe Tennenholtz. Dueling algorithms. In *Proc. 43rd Symp. Theory of Computing (STOC)*, pages 215–224, 2011.

[16] Nicole Immorlica, Robert Kleinberg, and Mohammad Mahdian. Secretary problems with competing employers. In *Proc. 2nd Workshop Internet & Network Economics (WINE)*, pages 389–400, 2006.

[17] Patrick Jaillet, José Soto, and Rico Zenklusen. Advances on matroid secretary problems: Free order model and laminar case. In *Proc. 16th Intl. Conf. Integer Programming and Combinatorial Optimization (IPCO)*, pages 254–265, 2013.

[18] Chinmay Karande, Aranyak Mehta, and Pushkar Tripathi. Online bipartite matching with unknown distributions. In *Proc. 43rd Symp. Theory of Computing (STOC)*, pages 587–596, 2011.

[19] Thomas Kesselheim, Klaus Radke, Andreas Tönnis, and Berthold Vöcking. An optimal online algorithm for weighted bipartite matching and extensions to combinatorial auctions. In *Proc. 21st European Symp. Algorithms (ESA)*, pages 589–600, 2013.

[20] Thomas Kesselheim, Klaus Radke, Andreas Tönnis, and Berthold Vöcking. Primal beats dual on online packing LPs in the random-order model. In *Proc. 46th Symp. Theory of Computing (STOC)*, pages 303–312, 2014.

[21] Robert Kleinberg. A multiple-choice secretary algorithm with applications to online auctions. In *Proc. 16th Symp. Discrete Algorithms (SODA)*, pages 630–631, 2005.

[22] Nitish Korula and Martin Pál. Algorithms for secretary problems on graphs and hypergraphs. In *Proc. 36th Intl. Coll. Automata, Languages and Programming (ICALP)*, pages 508–520, 2009.

[23] Oded Lachish. O(log log rank) competitive ratio for the matroid secretary problem. In *Proc. 55th Symp. Foundations of Computer Science (FOCS)*, pages 326–335, 2014.

[24] Dennis Lindley. Dynamic programming and decision theory. *Applied Statistics*, 10:39–51, 1961.

[25] Aranyak Mehta, Amin Saberi, Umesh Vazirani, and Vijay Vazirani. Adwords and generalized online matching. *J. ACM*, 54(5), 2007.

[26] Marco Molinaro and R. Ravi. Geometry of online packing linear programs. *Math. Oper. Res.*, 39(1):46–59, 2014.

[27] José Soto. Matroid secretary problem in the random-assignment model. *SIAM J. Comput.*, 42(1):178–211, 2013.

A Useful Facts

Fact 1. *For all $x \in [0, 1)$ it holds that*

$$\ln(1 - x) \quad \geq \quad -\frac{x}{1 - x} \ .$$

Proof. For $x = 0$ we have equality. The derivative of left- and right-hand sides are $-(1-x)^{-1}$ and $-(1-x)^{-2}$, respectively. Hence, the right-hand side drops faster when $x > 0$ grows towards 1, so the inequality holds for the entire interval. □

B Extension to the Sampling Model

In this section, we extend our logarithmic approximation to a general sampling model presented in [13]. This model extends the secretary model (adversarial values, random-order arrival), the prophet-inequality model (stochastic values from known distributions, adversarial arrival) as well as other mixtures of stochastic and worst-case aspects.

Formally, in the sampling model we have two values for each firm-applicant pair (u_i, v_j), a non-negative *sample value* $w^S(u_i, v_j)$ and a non-negative *input value* $w^I(u_i, v_j)$. The sample and input values are both drawn from possibly different, unknown distributions. For a single applicant v_j the sample and input distributions can be arbitrarily correlated among different firms and among each other. However, there is no correlation among distributions of different applicants. This defines a probability space over a class of instances $\mathcal{I}$.

The arrival process proceeds as follows. First, the adversary draws all values $w^S(u_i, v_j)$ and $w^I(u_i, v_j)$ for all pairs (u_i, v_j). It then reveals to firm u_i all drawn sample values $w^S(u_i, v_j)$, for all applicants v_j. Subsequently, depending on the drawn values w^I it chooses a worst-case arrival order of applicants. Upon arrival, an applicant v_j reveals its "real" value $w^I(u_i, v_j)$ to firm u_i. The algorithm $\mathcal{A}$ for firm u_i decides whether to make an offer to v_j, and applicant v_j accepts an offer that maximizes $w^I(u_i, v_j)$. Then the next applicant arrives. Decisions made in earlier rounds cannot be revoked. The goal for the algorithm is to maximize social welfare, i.e., to generate an assignment $M^{\mathcal{A}}$ that minimizes the competitive ratio $\mathbb{E}_{\mathcal{I}}[w^I(M^*)]/\mathbb{E}_{\mathcal{I},\mathcal{A}}[w^I(M^{\mathcal{A}})]$.

Clearly, if sample values are completely unrelated to input values, no algorithm $\mathcal{A}$ can obtain a bounded competitive ratio. Towards this end, we assume that for each value k, there is a similar probability that w^I and w^S have value k for pair (u_i, v_j). We here restrict attention to discrete distributions over integers. It is straightforward that our results hold for general distributions, but this minor extension does not justify the notational and technical overhead in presentation. More formally, we assume

- *Stochastic similarity*: Suppose $c > 1$ is a fixed constant. For every pair (u_i, v_j) and every integer $k > 0$, we assume that $\mathbf{Pr}\left[w^I(u_i, v_j) = k\right] \leq c \cdot \mathbf{Pr}\left[w^S(u_i, v_j) = k\right]$ and $\mathbf{Pr}\left[w^S(u_i, v_j) = k\right] \leq c \cdot \mathbf{Pr}\left[w^I(u_i, v_j) = k\right]$.
- *Stochastic independence*: For every pair (u_i, v_j), the weights $w^I(u_i, v_j)$ and $w^S(u_i, v_j)$ do not depend on the weights w^S and w^I of other candidates $v_{j'} \neq v_j$.

For further discussion of the sampling model and an exposition how to formulate secretary and prophet-inequality models within this framework, see [13].

Consider Algorithm 5, which is an extension of Algorithm 1 to the sampling model. It can be applied when every firm has a local matroid $\mathcal{S}_i$ that determines the set of applicants the firm can hire simultaneously. It is executed in parallel by all firms u_i. The algorithm first simplifies the structure of input and sample values by assuming that no candidate has $w^S(u_i, v_j) > 0$ and $w^I(u_i, v_j) > 0$. This loses a factor of at most 2 in the expected value of the solution. Again, we assume that every firm knows b, an upper bound on the maximum cardinality of the optimum solution, which is used to determine a random threshold based on the maximum weight seen by firm u_i in its simplified sample. Then the algorithm greedily makes an offer only to those applicants whose simplified input values are above the threshold.

Algorithm 5: Thresholding algorithm for firm u_i for general weights and matroids.

1 For each v_j flip a fair coin: if heads $w^I(u_i, v_j) \leftarrow 0$, if tails $w^S(u_i, v_j) \leftarrow 0$
2 $m_i \leftarrow \arg\max_{v_j} w^S(u_i, v_j)$
3 $X_i \leftarrow \text{Uniform}(0, 1, \ldots, \lceil \log_2 b \rceil + 1)$, where $b \geq |M^*|$
4 $t_i \leftarrow w^S(u_i, m_i)/2^{X_i}$
5 $M_i \leftarrow \emptyset$
6 **for** *all* v_j *over time* **do**
7   **if** $w^I(u_i, v_j) \geq t_i$ *and* $M_i \cup \{v_j\}$ *is independent set in* $\mathcal{S}_i$ **then**
8   make an offer to v_j
9   **if** v_j *accepts* **then**
10             $M_i \leftarrow M_i \cup \{v_j\}$

Theorem 6. *Algorithm 5 is* $16(c+1)^2(\lceil \log_2 b \rceil + 2)$*-competitive in the sampling model.*

Proof. The proof follows largely the one presented for the secretary model in Section 2 above. At first, however, we use arguments similar to [13] to capture the relation between sample and input values and to transform the scenario into a simpler domain.

The first line of our algorithm implements an adjustment of weights, so that at most one of the two weights for an applicant and a firm is positive. Let us assume w.l.o.g. that this condition holds already for the initial weights w^I and w^S. Formally, we denote

$$\hat{w}(u_i, v_j) = \max\{w^I(u_i, v_j), w^S(u_i, v_j)\}$$

and assume that $(w^I(u_i, v_j), w^S(u_i, v_j)) \in \{(0, \hat{w}(u_i, v_j)), (\hat{w}(u_i, v_j), 0)\}$. This preserves stochastic independence and similarity properties of the sampling model. Moreover, it lowers the expected value of the optimum solution by at most a factor of 2, i.e.,

$$\mathbb{E}_{\mathcal{I}}[w^I(M^*)] \quad \leq \quad 2\mathbb{E}_{\mathcal{I}}[\hat{w}(M^*)] \quad \leq \quad 2\mathbb{E}_{\mathcal{I}}[\hat{w}(\hat{M}^*)] \ ,$$

where M^* and $\hat{M}^*$ are optimal solutions for w^I and $\hat{w}$, respectively.

We condition on properties of the applicant with largest and second largest value for firm u_i. To cope with the resulting correlations, we introduce a conditional probability space. For each applicant v_j we assume that $\hat{w}(u_i, v_j)$ is fixed arbitrarily. For simplicity, we drop applicants from consideration for which $\hat{w}(u_i, v_j) = 0$. Let $V_i^I = \{v_j \mid w^I(u_i, v_j) > 0\}$ and $V_i^S = \{v_j \mid w^S(u_i, v_j) > 0\}$. Stochastic similarity implies

$$\mathbf{Pr}\left[w^I(u_i, v_j) = \hat{w}(u_i, v_j)\right] \quad \geq \quad (1/c) \cdot \mathbf{Pr}\left[w^S(u_i, v_j) = \hat{w}(u_i, v_j)\right]$$

and

$$\mathbf{Pr}\left[w^S(u_i, v_j) = \hat{w}(u_i, v_j)\right] \quad \geq \quad (1/c) \cdot \mathbf{Pr}\left[w^I(u_i, v_j) = \hat{w}(u_i, v_j)\right] \ .$$

Since $V_i^I \cap V_i^S = \emptyset$, we have

$$\mathbf{Pr}\left[v_j \in V_i^I\right] \geq \frac{1}{c+1} \qquad \text{and} \qquad \mathbf{Pr}\left[v_j \in V_i^S\right] \geq \frac{1}{c+1} \tag{2}$$

for each applicant v_j, independent of the outcome of weights of other applicants. In particular, (2) hold for every v_j, independently of $v_{j'} \in V_i^S$ or not for all other applicants $j' \neq j$.

We now execute the proof of the theorem, which proceeds very similarly to the proof of Theorem 1 above. We make two assumptions that make the analysis easier but do not hurt the overall result.

1. Based on our reformuation on a conditional probability space, we assume all $\hat{w}(u_i, v_j)$ are fixed arbitrarily. Furthermore, we assume $\hat{M}^*$ is an optimum solution when all applicants are in V_i^I for all firms u_i. As such, we assume that both $\hat{w}$ and $\hat{M}^*$ are deterministic. Our analysis is based only on the randomization expressed by the sampling inequalities (2) and randomized choice of t_i in Algorithm 5.

2. To avoid technicalities, we again assume that for each firm u_i, the values $\hat{w}(u_i, v_j)$ of all applicants are mutually disjoint.

Let $v_i^{\max} = \text{argmax}_j \hat{w}(u_i, v_j)$ and $v_i^{\text{2nd}} = \text{argmax}_{j \neq v_i^{\max}} \hat{w}(u_i, v_j)$ be the best and second best applicant for firm u_i, respectively. Here we denote by $w_i^{\max} = \hat{w}(u_i, v_i^{\max})$ and $w_i^{\text{2nd}} = \hat{w}(u_i, v_i^{\text{2nd}})$. For most of the analysis, we again work with *capped weights* $\tilde{w}(u_i, v_j)$, based on thresholds t_i set by the algorithm as follows

$$\tilde{w}(u_i, v_j) = \begin{cases} w_i^{\max} & \text{if } v_j \in V_i^I, t_i = w_i^{\text{2nd}}, \text{ and } \hat{w}(u_i, v_j) > w_i^{\text{2nd}}, \\ t_i & \text{if } v_j \in V_i^I \text{ and } \hat{w}(u_i, v_j) \geq t_i, \\ 0 & \text{otherwise.} \end{cases}$$

The definition of $\tilde{w}$ relies on random events, i.e., $v_j \in V_i^I$ and the choice of thresholds t_i. For any outcome of these events, however, $\tilde{w}(u_i, v_j) \leq \hat{w}(u_i, v_j)$ for all pairs (u_i, v_j). The following lemma adapts Lemma 1 and shows that, in expectation over all the correlated random events, an optimal offline solution with respect to $\tilde{w}$ gives a logarithmic approximation to the optimal offline solution with respect to $\hat{w}$.

Lemma 7. *Denote by $\hat{w}(M)$ and $\tilde{w}(M)$ the weight and capped weight of a solution M. Let $\tilde{M}^*$ and $\hat{M}^*$ be optimal solutions for $\tilde{w}$ and $\hat{w}$, respectively. Then*

$$\mathbb{E}\left[\tilde{w}(\tilde{M}^*)\right] \geq \frac{1}{4(c+1)^2(\lceil \log_2 b \rceil + 2)} \cdot \hat{w}(\hat{M}^*).$$

Proof. Let $(u_i, v_j) \in \hat{M}^*$ be an arbitrary pair. First, assume that v_j maximizes $\hat{w}(u_i, v_j)$, i.e., $v_j = v_i^{\max}$. By (2) with probability at least $1/(c+1)^2$, we have $v_j \in V_i^I$ and $v_i^{\text{2nd}} \in V_i^S$. For any such outcome, we have with probability $1/(\lceil \log_2 b \rceil + 2)$ that $t_i = w_i^{\text{2nd}}$ and $\tilde{w}(u_i, v_j) = w_i^{\max}$. This yields $\mathbb{E}[\tilde{w}(u_i, v_j)] \geq \hat{w}(u_i, v_j)/(2(c+1)^2(\lceil \log_2 b \rceil + 2))$.

Second, for any $v_j \neq v_i^{\max}$ with $w_i^{\max}/(2b) < \hat{w}(u_i, v_j) \leq w_i^{\max}$, by (2) we know $v_i^{\max} \in V_i^S$ is an independent event which happens with probability at least $1/(c+1)$. Then, there is some $0 \leq k' \leq \lceil \log_2 b \rceil + 1$, with $\hat{w}(u_i, v_j) > w_i^{\max}/2^{k'} \geq \hat{w}(u_i, v_j)/2$. With probability $1/(\lceil \log_2 b \rceil + 2)$, we have that $X_i = k'$ and $\tilde{w}(u_i, v_j) = t_i \geq \hat{w}(u_i, v_j)/2$. This yields $\mathbb{E}[\tilde{w}(u_i, v_j)] \geq \hat{w}(u_i, v_j)/(2(c+1)^2(\lceil \log_2 b \rceil + 2))$, since $v_j \in V_i^I$ with probability at least $1/(c+1)$ by (2).

Finally, we denote by $\hat{M}^>$ the set of pairs $(u_i, v_j) \in \hat{M}^*$ for which $\hat{w}(u_i, v_j) > w_i^{\max}/(2b)$. The expected weight of the best assignment with respect to the threshold values is thus

$$\begin{aligned} \mathbb{E}\left[\tilde{w}(\tilde{M}^*)\right] &\geq \sum_{(u_i, v_j) \in \hat{M}^*} \mathbb{E}[\tilde{w}(u_i, v_j)] \geq \sum_{(u_i, v_j) \in \hat{M}^>} \frac{\hat{w}(u_i, v_j)}{2(c+1)^2(\lceil \log_2 b \rceil + 2)} \\ &= \frac{1}{2(c+1)^2(\lceil \log_2 b \rceil + 2)} \cdot (\hat{w}(\hat{M}^*) - \hat{w}(\hat{M}^* \setminus \hat{M}^>)) \\ &\geq \frac{1}{4(c+1)^2(\lceil \log_2 b \rceil + 2)} \cdot \hat{w}(\hat{M}^*), \end{aligned}$$

since $\sum_{(u_i, v_j) \in \hat{M}^* \setminus \hat{M}^>} w_i^{\max}/(2b) \leq \max_i w_i^{\max}/2 \leq \hat{w}(\hat{M}^*)/2$. ☐

The previous lemma bounds the weight loss due to (i) all random choices inherent in the process of input generation and threshold selection and (ii) using the capped weights. The next lemma is essentially identical to Lemma 2 and bounds the remaining loss due to adversarial arrival of elements in V_i^I, exploiting that $\tilde{w}$ equalizes equal-threshold firms. Note that in Lemma 2 we already prove the result for arbitrary arrival, arbitrary weights w, and arbitrary thresholds based on w. Moreover, we define thresholds t_i based on $\hat{w}$ in exactly the same way as they we did based on w for Lemma 2. Hence, the lemma and its proof can be applied literally when using $\hat{w}$ instead of w.

Lemma 8. *Suppose subsets V_i^I and thresholds t_i are fixed arbitrarily and consider the resulting weight function $\tilde{w}$. Let $M^{\mathcal{A}}$ be the feasible solution resulting from Algorithm 5 using the thresholds t_i, for any arbitrary arrival order of applicants in $\bigcup V_i^I$. Then $\hat{w}(M^{\mathcal{A}}) \geq \tilde{w}(\tilde{M}^*)/2$.*

Combining these insights we see that that

$$
\begin{aligned}
\mathbb{E}_{\mathcal{I}}[w^I(M^*)] &\leq 2\mathbb{E}_{\mathcal{I}}[\hat{w}(\hat{M}^*)] \\
&\leq 8(c+1)^2(\lceil \log_2 b \rceil + 2)\mathbb{E}_{\mathcal{I},\mathcal{A}}[\tilde{w}(\tilde{M}^*)] \\
&\leq 16(c+1)^2(\lceil \log_2 b \rceil + 2)\mathbb{E}_{\mathcal{I},\mathcal{A}}[\hat{w}(M^{\mathcal{A}})] \\
&\leq 16(c+1)^2(\lceil \log_2 b \rceil + 2)\mathbb{E}_{\mathcal{I},\mathcal{A}}[w^I(M^{\mathcal{A}})] \ .
\end{aligned}
$$

This proves the theorem. □