
Learning to Screen

Anonymous Author(s)

Affiliation

Address

email

Abstract

1 Imagine a large firm with multiple departments that plans a large recruitment.
2 Candidates arrive one-by-one, and for each candidate the firm decides, based on her
3 data (CV, skills, experience, etc), whether to summon her for an interview. The firm
4 wants to recruit the best candidates while minimizing the number of interviews.
5 We model such scenarios as an assignment problem between items (candidates)
6 and categories (departments): the items arrive one-by-one in an online manner,
7 and upon processing each item the algorithm decides, based on its value and the
8 categories it can be matched with, whether to retain or discard it (this decision is
9 irrevocable). The goal is to retain as few items as possible while guaranteeing that
10 the set of retained items contains an optimal matching.

11 We consider two variants of this problem: (i) in the first variant it is assumed that
12 the n items are drawn independently from an unknown distribution D . (ii) In the
13 second variant it is assumed that before the process starts, the algorithm has an
14 access to a training set of n items drawn independently from the same unknown
15 distribution (e.g. data of candidates from previous recruitment seasons). We give
16 tight bounds on the minimum possible number of retained items in each of these
17 variants. These results demonstrate that one can retain exponentially less items in
18 the second variant (with the training set).

19 Our algorithms and analysis utilize ideas and techniques from statistical learning
20 theory and from discrete algorithms.

21 1 Introduction

22 Matching is the bread-and-butter of many real-life problems from the fields of computer science,
23 operations research, game theory, and economics. Some examples include job scheduling where we
24 assign jobs to machines, economic markets where we allocate products to buyers, online advertising
25 where we assign advertisers to ad slots, assigning medical interns to hospitals, and many more.

26 One particular example that motivates this work is the following example from labor markets. Imagine
27 a firm that is planning a large recruitment. Candidates arrive one-by-one and the HR department
28 immediately decides whether to summon them for an interview. Moreover, the firm has multiple
29 departments, each requiring different skills and having a different target number of hires. Different
30 employees have different subsets of the required skills, and thus fit only certain departments and
31 with a certain quality. The firm's HR department, following the interviews, decides which candidates
32 to recruit and to which departments to assign them. The HR department has to maximize the total
33 quality of the hired employees such that each department gets its required number of hires with
34 the required skills. In addition, the HR uses data from the previous recruitment season in order to
35 minimize the number of interviews while not compromising the quality of the solution.

36 We study the following formulation of the problem above. We receive n items (candidates), where
37 each item has a subset of d properties (departments) denoted by P_1, \dots, P_d . We select k items out of
38 the n , subject to d constraints of the form

39
$$\text{exactly } k_i \text{ of the selected items must satisfy a property } P_i,$$

40 where $\sum_{i=1}^d k_i = k$ and we assume that $d \ll k \ll n$. Furthermore, if item c possesses property P_i , then
41 it has a value $v_i(c)$ associated with this property. Our goal is to compute a matching of maximum
42 value that associates k items to the d properties subject to the constraints above.

43 We consider matching algorithms in the following online setting. The algorithms receive n items
44 online, drawn independently from D , and either reject or retain each item. Then, the algorithm
45 utilizes the retained items and outputs an (approximately-)optimal feasible solution. We present a
46 naive greedy algorithm that returns the optimal solution with probability at least $1 - \delta$ and retains
47 $O(k \log(k/\delta))$ items in expectation. We prove that no other algorithm with the same guarantee can
48 retain less items in expectation.

49 Thus, to further reduce the number of retained items, we add an initial preprocessing phase in which
50 the algorithm learns an online policy from a *training set*. The training set is a single problem instance
51 that consists of n items drawn independently from the same unknown distribution D . We address
52 the statistical aspects of this problem and develop efficient learning algorithms. In particular, we
53 define a class of *thresholds-policies*. Each thresholds-policy is a simple rule for deciding whether
54 to retain an item. We present uniform convergence rates for both the number of items retained by a
55 thresholds policy and the value of the resulting solution. We show that these quantities deviate from
56 their expected value by order of \sqrt{k} (rather than an easier \sqrt{n} bound; recall that we assume $k \ll n$)
57 which we prove using concentration inequalities and tools from VC-theory. Using these concentration
58 inequalities, we analyze an efficient online algorithm that returns the optimal offline solution with
59 probability at least $1 - \delta$, and retains a near-optimal $O(k \log \log(1/\delta))$ number of items in expectation
60 (compare with the $O(k \log(k/\delta))$ number of retained items when no training set is given).

61 **Related work.** Our model is related to the online secretary problem in which one needs to select
62 the best secretary in an online manner (see Ferguson, 1989). Our setting differs from this classical
63 model due to the two-stage process and the complex feasibility constraints. Nonetheless, we remark
64 that there are few works on the secretary model that allow delayed selection (see Vardi, 2015, Ezra
65 et al., 2018) as well as matroid constraints [Babai et al., 2007]. These works differ from ours in
66 the way the decision is made, the feasibility constraints and the learning aspect of receiving a single
67 problem instance as a training example.

68 Another related line of work in algorithmic economics studies the statistical learnability of pricing
69 schemes (see e.g., Morgenstern and Roughgarden, 2015, 2016, Hsu et al., 2016, Balcan et al., 2018).
70 The main difference of these works from ours is that our training set consists of a single “example”
71 (namely the set of items that are used for training), and in their setting (as well as in most typical
72 statistical learning settings) the training set consists of many i.i.d examples. This difference also affects
73 the technical tools used for obtaining generalization bounds. For example, some of our bounds exploit
74 Talagrand’s concentration inequality rather than the more standard Chernoff/McDiarmid/Bernstein
75 inequalities. We note that Talagrand’s inequality and other advanced inequalities were applied in
76 machine learning in the context of learning combinatorial functions [Vondrák, 2010, Blum et al.,
77 2017]. See also the survey by Bousquet et al. [2004] or the book by Boucheron et al. [2013] for a
78 more thorough review of concentration inequalities.

79 Furthermore, there is a large body of work on online matching in which the vertices arrive in various
80 models (see Mehta et al., 2013, Gupta and Molinaro, 2016). We differ from this line of research, by
81 allowing a two-stage algorithm, and requiring to output the optimal matching is the second stage.

82 Celis et al. [2017, 2018] studies similar problems of ranking and voting with fairness constraints. In
83 fact, the optimization problem that they consider allows more general constraints and the value of a
84 candidate is determined from votes/comparisons. The main difference with our framework is that
85 they do not consider a statistical setting (i.e. there is no distribution over the items and no training set
86 for preprocessing) and focus mostly on approximation algorithms for the optimization problem.

87 2 Model and Results

88 Let X be a domain of items, where each item $c \in X$ can possess any subset of d properties denoted
89 by P_1, \dots, P_d (we view $P_i \subseteq X$ as the set of items having property P_i). Each item c has a value
90 $v_i(c) \in [0, 1]$ associated with each property P_i such that $c \in P_i$.

91 We are given a set $C \subseteq X$ of n items as well as counts k_1, \dots, k_d such that $\sum_{i=1}^d k_i = k$. Our goal is to
92 select exactly k items in total, constrained on selecting exactly k_i items with property P_i . We assume
93 that these constraints are *exclusive*, in the sense that each item in C can be used to satisfy at most
94 one of the constraints. Formally, a feasible solution is a subset $S \subseteq C$, such that $|S| = k$ and there is
95 partition S into d disjoint subsets S_1, \dots, S_d , such that $S_i \subseteq P_i$ and $|S_i| = k_i$. We aim to compute a
96 feasible subset S that maximizes $\sum_{i=1}^d \sum_{c \in S_i} v_i(c)$.

97 Furthermore, we assume that $d \ll k \ll n$. Namely, the number of constraints is much smaller than
98 the number of items that we have to select, which is much smaller than the total number of items
99 in C . In order to avoid feasibility issues we assume that there is a set C_{dummy} that contains k dummy
100 0-value items with all the d properties (we assume that the algorithm has always access to C_{dummy}
101 and do not view them as part of C).

102 **Formulation as bipartite matching.** We first discuss the offline versions of these allocation prob-
103 lems. That is, we assume that C and the capacities k_i are all given as an input before the algorithm
104 starts. We are interested in an algorithm for computing an optimal set S . That is a set of items of
105 maximum total value that satisfy the constraints. This problem is equivalent to a maximum matching
106 problem in a bipartite graph (L, R, E, w) defined as follows.

- 107 • L is the set of vertices in one side of the bipartite graph. It contains k vertices, where each
108 constraint i is represented by k_i of these vertices.
- 109 • R is the set of vertices in the other side of the bipartite graph. It contains a vertex for each
110 item $c \in C$ and for each dummy item $c' \in C_{\text{dummy}}$.
- 111 • E is the set of edges. Each vertex in R is connected to each vertex of each of the constraints
112 that it satisfies.
- 113 • The weight $w(l, r)$ of edge $(l, r) \in E$ is $v_l(r)$: the value of item r associated with property P_l .

114 There is a natural correspondence between *saturated-matchings* in this graph, that is matchings
115 in which every $l \in L$ is matched, and between *feasible solutions* (i.e., solutions that satisfy the
116 constraints) to the allocation problem. Thus, a saturated-matching of maximum value corresponds to
117 an optimal solution. It is well known that the problem of finding such a maximum weight bipartite
118 matching can be solved in polynomial time (see e.g., Lawler, 2001).

119 **Problem definition.** In this work we consider the following online learning model. We assume
120 that n items are sequentially drawn i.i.d. from an unknown distribution D over X . Upon receiving
121 each item, we decide whether to retain it, or reject it irrevocably (the first stage of the algorithm).
122 Thereafter, we select a feasible solution¹ consisting *only* of retained items (the second stage of the
123 algorithm). Most importantly, before accessing the online sequence and take irreversible online
124 decisions of which items to reject, we have access a training set C_{train} consisting of n independent
125 draws from D .

126 2.1 Results

127 2.1.1 Oblivious online screening

128 We begin by studying a greedy algorithm that does not require a training set. In the online phase,
129 this algorithm acts greedily by keeping an item if it participates in the best solution thus far. Then,
130 the algorithm computes an optimal matching among the retained items. The particular details of the
131 algorithm are given in Appendix A.1. We have the following guarantee for this greedy algorithm
132 proven in Appendix A.1.

¹In addition to the retained items, the algorithm has access to C_{dummy} , and therefore a feasible solution always exists.

133 **Theorem 1.** *Let $\delta \in (0, 1)$. The greedy algorithm outputs the optimal solution with probability at*
 134 *least $1 - \delta$ and retains $O(k \log(\min\{k/\delta, nk\}))$ items in expectation.*

135 As we shall see in the next section, learning from the training set allows one to retain exponentially
 136 less items than is implied by the theorem above. It is then natural to ask to which extent is the
 137 training phase essential in order to accommodate such an improvement. We answer this question in
 138 Appendix B.1 by proving a lower bound on the number of retained items for *any* algorithm that does
 139 not use a training phase. This lower bound already applies in the simple setting where $d = 1$: here,
 140 each item consists only of a value $v \in [0, 1]$, and the goal of the algorithm is to retain as few items as
 141 possible while guaranteeing with high probability that the top k maximal values are retained.

142 **Theorem 2.** *Let $\delta \in (0, 1)$. For every algorithm A which retains the maximal k elements with*
 143 *probability at least $1 - \delta$, there exists a distribution μ such that the expected number of retained*
 144 *elements for input sequences $v_1 \dots v_n \sim \mu^n$ is at least $\Omega(k \log(\min\{k/\delta, nk\}))$.*

145 Thus, the above theorem implies that $\Theta(k \log(n/k))$ can not be improved even if we allow failure
 146 probability $\delta = \Theta(k^2/n)$ (see Theorem 1).

147 2.1.2 Online screening with learning

148 We now design online algorithms that, before the online screening process begins, use C_{train} to learn a
 149 *thresholds-policy* $T \in \mathcal{T}$ such that with high probability: (i) the number of items that are retained in
 150 the online phase is small, and (ii) there is a feasible solution consisting of k retained items whose
 151 value is optimal (or close to optimal). Thresholds-policies are studied in Section 3 and are defined as
 152 follows.

153 **Definition 3** (Thresholds-policies). A threshold-policy is parametrized by a vector $T = (t_1, \dots, t_d)$
 154 of thresholds, where t_i corresponds to property P_i for $1 \leq i \leq d$. The semantics of T is as follows:
 155 given a sample C of n items, each item $c \in C$ is retained if and only if there exists a property P_i
 156 satisfied by c , such that its value $v_i(c)$ passes the threshold t_i . More formally, c is retained if and only
 157 if $\exists i \in \{1, \dots, d\}$ such that $c \in P_i$ and $v_i(c) \geq t_i$.

158 Having proven uniform convergence results for thresholds-policies (see Section 3.1), we show the
 159 following in Section 4.

160 **Theorem 4.** *There exists an algorithm that learns a thresholds-policy T from a single training sample*
 161 *$C_{\text{train}} \sim D^n$, such that after processing the (“real-time”) input sample $C \sim D^n$ using T :*

- 162 • *It outputs an optimal solution with probability at least $1 - \delta$.*
- 163 • *The expected number of retained items in the first phase is $O(k(\log d + \log \log(nk) +$*
 164 *$\log \log(1/\delta)))$.*

165 Thus, with the additional information given by the training set, the algorithm presented in Theo-
 166 rem 4 improves the number of retained items from $k \log(k/\delta)$ to $k \log \log(1/\delta)$. This demonstrates a
 167 significant improvement over Theorem 1.

168 Finally, in Appendix B.2 we prove that the algorithm from Theorem 4 is nearly-optimal in the sense
 169 that it is impossible to significantly improve the number of retained items even if we allow the
 170 algorithm to fully know the distribution over input items (so, in a sense, having an access to n i.i.d
 171 samples from the distribution is the same as knowing it completely).

172 **Theorem 5.** *Consider the case where $k = d$ and $k_1 = \dots k_d = 1$. There exists a universe X and a*
 173 *fixed distribution D over X such that for $C \sim D^n$ the following holds: any online learning algorithm*
 174 *(which possibly “knows” D) that retains a subset $S \subseteq C$ of items that contains an optimal solution*
 175 *with probability at least $1 - \delta$ must satisfy that $\text{Ex}[|S|] = \Omega(k \log \log(1/\delta))$.*

176 3 Thresholds-policies

177 We next discuss a framework to design algorithms that exploit the training set to learn policies that
 178 are applied in the first phase of the matching process. We would like to frame this in standard ML
 179 formalism by phrasing this problem as learning a class \mathcal{H} of policies such that:

- 180 • **\mathcal{H} is not too small:** The policies in \mathcal{H} should yield solutions with high values (optimal, or
181 near-optimal).
- 182 • **\mathcal{H} is not too large:** \mathcal{H} should satisfy some uniform convergence properties; i.e. the perfor-
183 mance of each policy in \mathcal{H} on the training set is close, with high probability, to its expected
184 real-time performance on the sampled items during the online selection process.

185 Indeed, as we now show these demands are met by the class \mathcal{T} of thresholds policies (Definition 3).
186 We first show that the class of thresholds-policies contains an optimal policy, and in the sequel we
187 show that it satisfies attractive uniform convergence properties.

188 **An assumption (values are unique).** We assume that for each constraint P_i , the marginal distribu-
189 tion over the value of $c \sim D$ conditioned on $c \in P_i$ is atomless; namely $\Pr_{c \sim D}[v(c) = v \mid c \in P_i] = 0$
190 for every $v \in [0, 1]$. This assumption can be removed by adding artificial tie-breaking rules, but
191 making it will simplify some of the technical statements.

192 **Theorem 6** (There is a thresholds policy that retains an optimal solution). *For any set of items C ,
193 there exists a thresholds vector $T \in \mathcal{T}$ that retains exactly k items that form an optimal solution for C .*

194 *Proof.* Let S denote the set of k items in an optimal solution for C , and let $S_i \subseteq S \cap P_i$ be the subset
195 of M that is assigned to the constraint P_i . Define $t_i = \min_{c \in S_i} v_i(c)$, for $i \geq 1$. Clearly, T retains all
196 the items in S . Assume towards contradiction that T retains an item $c_j \notin S$, and assume that P_j is a
197 constraint such that $c_j \in P_j$ and $v_j(c_j) \geq t_j$. Since by our assumption on D all the values $v_j(c_j)$ are
198 distinct it follows that $v_j(c_j) > t_j$. Thus, we can modify S by replacing c_j with the item of minimum
199 value in S_j and increase the total value. This contradicts the optimality of S . \square

200 We next establish generalization bounds for the class of thresholds-policies.

201 3.1 Uniform convergence of the number of retained items

202 For a sample $C \sim D^n$ and a thresholds-policy $T \in \mathcal{T}$, we denote by $R_i^T(C) = \{c : c \in P_i \text{ and } v_i(c) \geq t_i\}$
203 the set of items that are retained by the threshold t_i , and we denote its expected size by $\rho_i^T =$
204 $\mathbb{E}_{C \sim D^n}[|R_i^T(C)|]$. Similarly we denote by $R^T(C) = \cup_i R_i^T(C)$ the items retained by T , and by ρ^T its
205 expectation. We prove that the sizes of $R_i^T(C)$ and $R^T(C)$ are concentrated around their expectations
206 uniformly for all thresholds policies.

207 The following theorems establish uniform convergence results for the number of retained items.
208 Namely, with high probability we have $R_i^T \approx \rho_i^T$, $R^T \approx \rho^T$ simultaneously for all $T \in \mathcal{T}$ and $i \leq d$.

209 **Theorem 7** (Uniform convergence of the number of retained items). *With probability at least $1 - \delta$
210 over $C \sim D^n$, the following holds for all policies $T \in \mathcal{T}$ simultaneously:*

- 211 1. If $\rho^T \geq k$, then $(1 - \epsilon)\rho^T \leq |R^T(C)| \leq (1 + \epsilon)\rho^T$, and
- 212 2. if $\rho^T < k$, then $\rho^T - \epsilon k \leq |R^T(C)| \leq \rho^T + \epsilon k$,

213 where

$$\epsilon = O\left(\sqrt{\frac{d \log(d) \log(n/k) + \log(1/\delta)}{k}}\right).$$

214 **Theorem 8** (Uniform convergence of the number of retained items per constraint). *With probability
215 at least $1 - \delta$ over $C \sim D^n$, the following holds for all policies $T \in \mathcal{T}$ and all $i \leq d + 1$ simultaneously:*

- 216 1. If $\rho_i^T \geq k$, then $(1 - \epsilon)\rho_i^T \leq |R_i^T(C)| \leq (1 + \epsilon)\rho_i^T$, and
- 217 2. if $\rho_i^T < k$, then $\rho_i^T - \epsilon k \leq |R_i^T(C)| \leq \rho_i^T + \epsilon k$,

218 where

$$\epsilon = O\left(\sqrt{\frac{\log(d) \log(n/k) + \log(1/\delta)}{k}}\right).$$

219 The proofs of Theorems 7 and 8 are based on standard VC-based uniform convergence results, and
 220 technically the proof boils down to bounding the VC-dimension of the families

$$\mathcal{R} = \{R^T : T \in \mathcal{T}\} \quad \text{and} \quad \mathcal{Q} = \{R_i^T : T \in \mathcal{T}, i \leq d\}.$$

221 Indeed, in Appendix A.2 we prove the following.

222 **Lemma 9.** $VC(\mathcal{R}) = O(d \log d)$.

223 **Lemma 10.** $VC(\mathcal{Q}) = O(\log d)$.

224 Using Lemmas 9 and 10, we can now apply standard uniform convergence results from VC-theory to
 225 derive Theorems 7 and 8.

226 **Definition 11** (Relative (p, ε) -approximation; Har-Peled and Sharir, 2011). Let \mathcal{F} be a family of
 227 subsets over a domain X , and let μ be a distribution on X . $Z \subseteq X$ is a (p, ε) -approximation for \mathcal{F} if
 228 for each $f \in \mathcal{F}$ we have,

229 1. If $\mu(f) \geq p$, then $(1 - \varepsilon)\mu(f) \leq \hat{\mu}(f) \leq (1 + \varepsilon)\mu(f)$,

230 2. If $\mu(f) < p$, then $\mu(f) - \varepsilon p \leq \hat{\mu}(f) \leq \mu(f) + \varepsilon p$,

231 where $\hat{\mu}(f) = |Z \cap F|/|Z|$ is the (“empirical”) measure of f with respect to Z .

232 The proof of Theorems 7 and 8 now follows by plugging $p = k/n$ in Har-Peled and Sharir [2011,
 233 Theorem 2.11], which we state in the next proposition.

234 **Proposition 12** (Har-Peled and Sharir, 2011). *Let \mathcal{F} and μ like in Definition 11. Suppose \mathcal{F} has VC*
 235 *dimension m . Then, with provability at least $1 - \delta$, a random sample of size*

$$\Omega\left(\frac{m \log(1/p) + \log(1/\delta)}{\varepsilon^2 p}\right)$$

236 *is a relative (p, ε) -approximation for \mathcal{F} .*

237 3.2 Uniform convergence of values

238 We now prove a concentration result for the value of an optimal solution among the retained items.
 239 Unlike the number of retained items, the value of an optimal solution corresponds to a more complex
 240 random variable, and analyzing the concentration of its empirical estimate requires more advanced
 241 techniques.

242 We denote by $V^T(C)$ the value of the optimal solution among the items retained by the thresholds-
 243 policy T , and we denote its expectation by $v^T = \text{Ex}_{C \sim D^n} [V^T(C)]$. We show that $V^T(C)$ is concentrated
 244 uniformly for all thresholds policies.

245 **Theorem 13** (Uniform convergence of values). *With probability at least $1 - \delta$ over $C \sim D^n$, the*
 246 *following holds for all policies $T \in \mathcal{T}$ simultaneously:*

$$|v^T - V^T(C)| \leq \varepsilon k, \quad \text{where} \quad \varepsilon = O\left(\sqrt{\frac{d \log k + \log(1/\delta)}{k}}\right).$$

247 Note that unlike most uniform convergence results that guarantee simultaneous convergence of
 248 empirical averages to expectations, here $V^T(C)$ is not an average of the n samples, but rather a more
 249 complicated function of them. We also note that a bound of $\tilde{O}(\sqrt{n})$ (rather than $\tilde{O}(\sqrt{k})$) on the additive
 250 deviation of $V^T(C)$ from its expectation can be derived using the McDiarmid’s inequality [McDiarmid,
 251 1989]. However, this bound is meaningless when $\sqrt{n} > k$ (because k upper bounds the value of the
 252 optimal solution). We use Talagrand’s concentration inequality [Talagrand, 1995] to derive the $O(\sqrt{k})$
 253 upper bound on the additive deviation. Talagrand’s concentration inequality allows us to utilize the
 254 fact that an optimal solution uses only $k \ll n$ items, and therefore replacing an item that does not
 255 participate in the solution does not affect its value.

256 To prove the theorem we need the following concentration inequality for the value of the optimal
 257 selection in hindsight. Note that by Theorem 6 this value equals to $V^T(C)$ for some T .

258 **Lemma 14.** *Let $\text{OPT}(C)$ denote the value of the optimal solution for a sample C . We have that*

$$\Pr_{C \sim D^n} [|\text{OPT}(C) - \text{Ex}[\text{OPT}(C)]| \geq \alpha] \leq 2 \exp(-\alpha^2/2k).$$

259 So, for example, it happens that $|\text{OPT}(C) - \text{Ex}[\text{OPT}(C)]| \leq \sqrt{2k \log(2/\delta)}$ with probability at least
 260 $1 - \delta$.

261 To prove this lemma we use the following version of Talagrand's inequality (that appears for example
 262 in lecture notes by van Handel [2014]).

263 **Proposition 15** (Talagrand's Concentration Inequality). *Let $f : \mathbb{R}^n \mapsto \mathbb{R}$ be a function, and suppose
 264 that there exist $g_1, \dots, g_n : \mathbb{R}^n \mapsto \mathbb{R}$ such that for any $x, y \in \mathbb{R}^n$*

$$f(x) - f(y) \leq \sum_{i=1}^n g_i(x) 1_{[x_i \neq y_i]}. \quad (1)$$

265 Then, for independent random variables $X = (X_1, \dots, X_n)$ we have

$$\Pr [|f(X) - \text{Ex}[f(X)]| > \alpha] \leq 2 \exp \left(-\frac{\alpha^2}{2 \sup_x \sum_{i=1}^n g_i^2(x)} \right).$$

266 *Proof of Lemma 14.* We apply Talagrand's concentration inequality to the random variable $\text{OPT}(C)$.
 267 Our X_i 's are the items c_1, \dots, c_n in the order that they are given. We show that Eq. (1) holds for
 268 $g_i(C) = 1_{[c_i \in S]}$ where $S = S(C)$ is a fixed optimal solution for C (we use some arbitrary tie breaking
 269 among optimal solutions). We then have, $\sum_{i=1}^n g_i^2(C) = |S| = k$, thus completing the proof.

270 Now, let C, C' be two samples of n items. Recall that we need to show that

$$\text{OPT}(C) - \text{OPT}(C') \leq \sum_{i=1}^n g_i(C) 1_{[c_i \neq c'_i]}.$$

271 We use S to construct a solution S' for C' as follows. Let $S_j \subseteq S$ the subset of S matched to P_j . For
 272 each i , if $c_i \in S_j$ for some j , and $c_i = c'_i$, then we add i to S'_j . Otherwise, we add a dummy item from
 273 C'_{dummy} to S'_j (with value zero). Let $V(S')$ denote the value of S' . Note that the difference between
 274 the values of S and S' is the total value of all items $i \in S$ such that $c_i \neq c'_i$. Since the item values are
 275 bounded in $[0, 1]$ we get that

$$\text{OPT}(C) - V(S') = \sum_{j=1}^d \sum_{c_i \in S_j} v_j(c_i) 1_{[c_i \neq c'_i]} \leq \sum_{j=1}^d \sum_{c_i \in S_j} 1_{[c_i \neq c'_i]} = \sum_{i=1}^n g_i(C) 1_{[c_i \neq c'_i]}.$$

276 The proof is complete by noticing that $\text{OPT}(C') \geq V(S')$. □

277 We also require the following construction of a bracketing of \mathcal{T} which is formally presented in
 278 Appendix A.2.

279 **Lemma 16.** *There exists a collection of \mathcal{N} thresholds-policies such that $|\mathcal{N}| \leq k^{O(d)}$, and for every
 280 thresholds-policy $T \in \mathcal{T}$ there are $T^+, T^- \in \mathcal{N}$ such that*

281 1. $V^{T^-}(C) \leq V^T(C) \leq V^{T^+}(C)$ for every sample of items C ; note that by taking expectations
 282 this implies that $v^{T^-} \leq v^T \leq v^{T^+}$, and

283 2. $v^{T^+} - v^{T^-} \leq 10$.

284 *Proof of Theorem 13.* The items in C that are retained by T are independent samples from a distri-
 285 bution D' that is sampled as follows: (i) sample $c \sim D$, and (ii) if c is retained by T then keep it,
 286 and otherwise discard it. This means that $v^T(C)$ is in fact the optimal solution of C with respect to
 287 D' . Since Lemma 14 applies to every distribution D we can apply it to D' and get that for any fixed
 288 $T \in \mathcal{T}$

$$\Pr_{C \sim D^n} [|v^T - V^T(C)| \geq \alpha] \leq 2 \exp(-\alpha^2/2k).$$

289 Now, by the union bound for \mathcal{N} be as in Lemma 16 we get that the probability that there is $T \in \mathcal{N}$
 290 such that $|v^T - V^T(C)| \geq \alpha$ is at most $|\mathcal{N}| \cdot 2 \exp(-\alpha^2/2k)$. Thus, since $|\mathcal{N}| \leq k^{O(d)}$, it follows that
 291 with probability at least $1 - \delta$,

$$(\forall T \in \mathcal{N}) : |v^T - V^T(C)| \leq O\left(\sqrt{k(d \log k + \log(1/\delta))}\right). \quad (2)$$

292 We now show why uniform convergence for \mathcal{N} implies uniform convergence for \mathcal{T} . Combining
 293 Lemma 16 with Equation (2) we get that with probability at least $1 - \delta$, every $T \in \mathcal{T}$ satisfies:

$$\begin{aligned} |\mathbf{v}^T - V^T(C)| &\leq \max\{|\mathbf{v}^{T^+} - V^{T^+}(C)|, |\mathbf{v}^{T^-} - V^{T^-}(C)|\} && \text{(by Item 1 of Lemma 16)} \\ &\leq \max\{|\mathbf{v}^{T^+} - V^{T^+}(C)|, |\mathbf{v}^{T^-} - V^{T^-}(C)|\} + 10 && \text{(by Item 2 of Lemma 16)} \\ &\leq 10 + O\left(\sqrt{k(d \log k + \log(1/\delta))}\right). && \text{(by Eq. (2))} \end{aligned}$$

294 Here the first inequality follows from Item 1 by noticing that if $[a, b]$, $[c, d]$ are intervals on the real
 295 line and $x \in [a, b]$, $y \in [c, d]$ then $|x - y| \leq \max\{|b - c|, |d - a|\}$, and plugging in $x = \mathbf{v}^T$, $y = V^T(C)$, $a =$
 296 \mathbf{v}^{T^+} , $b = \mathbf{v}^{T^-}$, $c = V^{T^+}(C)$, $d = V^{T^-}(C)$.

297 This finishes the proof, by setting $\varepsilon \cdot k = O\left(\sqrt{k(d \log k + \log(1/\delta))}\right)$. \square

298 4 Algorithms based on learning thresholds-policies

299 We next exemplify how one can use the above properties of thresholds-policies to design algorithms.
 300 A natural algorithm would be to use the training set to learn a threshold-policy T that retains an
 301 optimal solution with k items from the training set as specified in Theorem 6, and then use this online
 302 policy to retain a subset of the n items in the first phase. Theorem 7 and Theorem 13 imply that with
 303 probability $1 - \delta$, the number of retained items is at most $m = k + O\left(\sqrt{kd \log(d) \log(n/k) + k \log(1/\delta)}\right)$
 304 and that the value of the resulting solution is at least $\text{OPT} - O\left(\sqrt{kd \log k + k \log(1/\delta)}\right)$.

305 We can improve this algorithm by combining it with the greedy algorithm of Theorem 1 described
 306 in Appendix A.1. During the first phase, we retain an item c only if (i) c is retained by T , and (ii)
 307 c participates in the optimal solution among the items that were retained thus far. Theorem 1 then
 308 implies that out of these m items greedy keeps a subset of

$$O\left(k \log \frac{m}{k}\right) = O\left(k \left(\log \log \left(\frac{n}{k}\right) + \log \log \left(\frac{1}{\delta}\right)\right)\right).$$

309 items in expectation that still contains a solution of value at least $\text{OPT} - O\left(\sqrt{kd \log k + k \log(1/\delta)}\right)$.

310 We can further improve the value of the solution and guarantee that it will be optimal (with respect
 311 to all n items) with probability $1 - \delta$. This is based on the observation that if the set of retained
 312 items contains the top k items of each property P_i then it also contains an optimal solution. Thus, we
 313 can compute a thresholds-policy T that retains the top $k + O\left(\sqrt{k \log(d) \log(n/k) + k \log(1/\delta)}\right)$ items
 314 of each property from the training set (if the training set does not have this many items with some
 315 property then set the corresponding threshold to 0). Then, it follows from Theorem 8, that with
 316 probability $1 - \delta$, T will retain the top k items of each property in the first online phase and therefore
 317 will retain an optimal solution. Now, Theorem 8 implies that with probability $1 - \delta$ the total number
 318 of items that are retained by T in real-time is at most $m = dk + O\left(d \sqrt{k \log(d) \log(n/k) + k \log(1/\delta)}\right)$.
 319 By filtering the retained elements with the greedy algorithm of Theorem 1 as before it follows that
 320 the total number of retained items is at most

$$k + k \log \left(\frac{m}{k}\right) = O\left(k \left(\log d + \log \log \left(\frac{n}{k}\right) + \log \log \left(\frac{1}{\delta}\right)\right)\right)$$

321 with probably $1 - \delta$. This proves Theorem 4.

References

- 322
- 323 M. Babaioff, N. Immorlica, and R. Kleinberg. Matroids, secretary problems, and online mechanisms.
324 In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages
325 434–443. Society for Industrial and Applied Mathematics, 2007.
- 326 M. Balcan, T. Sandholm, and E. Vitercik. A general theory of sample complexity for multi-item
327 profit maximization. In *EC*, pages 173–174. ACM, 2018.
- 328 A. Blum, I. Caragiannis, N. Haghtalab, A. D. Procaccia, E. B. Procaccia, and R. Vaish. Opting into
329 optimal matchings. In *SODA*, pages 2351–2363. SIAM, 2017.
- 330 S. Boucheron, G. Lugosi, and P. Massart. *Concentration Inequalities: A Nonasymptotic Theory of*
331 *Independence*. Oxford University Press, 2013. ISBN 9780191747106.
- 332 O. Bousquet, U. von Luxburg, and G. Rätsch, editors. *Advanced Lectures on Machine Learning, ML*
333 *Summer Schools 2003, Canberra, Australia, February 2-14, 2003, Tübingen, Germany, August*
334 *4-16, 2003, Revised Lectures*, volume 3176 of *Lecture Notes in Computer Science*, 2004. Springer.
- 335 L. E. Celis, D. Straszak, and N. K. Vishnoi. Ranking with fairness constraints. *arXiv preprint*
336 *arXiv:1704.06840*, 2017.
- 337 L. E. Celis, L. Huang, and N. K. Vishnoi. Multiwinner voting with fairness constraints. In *IJCAI*,
338 pages 144–151, 2018.
- 339 T. Ezra, M. Feldman, and I. Nehama. Prophets and secretaries with overbooking. In *Proceedings of*
340 *the 2018 ACM Conference on Economics and Computation*, pages 319–320. ACM, 2018.
- 341 T. S. Ferguson. Who solved the secretary problem? *Statistical Science*, 4(3):282–289, 1989.
- 342 S. Greenberg and M. Mohri. Tight lower bound on the probability of a binomial exceeding its
343 expectation. *CoRR*, abs/1306.1433, 2013.
- 344 A. Gupta and M. Molinaro. How the experts algorithm can help solve lps online. *Math. Oper. Res.*,
345 41(4):1404–1431, 2016.
- 346 S. Har-Peled and M. Sharir. Relative (p, ϵ) -approximations in geometry. *Discrete & Computational*
347 *Geometry*, 45(3):462–496, 2011.
- 348 J. Hsu, J. Morgenstern, R. M. Rogers, A. Roth, and R. Vohra. Do prices coordinate markets? In
349 *STOC*, pages 440–453. ACM, 2016.
- 350 E. L. Lawler. *Combinatorial optimization: networks and matroids*. Courier Corporation, 2001.
- 351 C. McDiarmid. On the method of bounded differences. In *Surveys in Combinatorics 1989*. Cambridge
352 University Press, Cambridge, 1989.
- 353 A. Mehta et al. Online matching and ad allocation. *Foundations and Trends® in Theoretical*
354 *Computer Science*, 8(4):265–368, 2013.
- 355 S. Moran, M. Snir, and U. Manber. Applications of ramsey’s theorem to decision tree complexity.
356 *Journal of the ACM (JACM)*, 32(4):938–949, 1985.
- 357 J. Morgenstern and T. Roughgarden. On the pseudo-dimension of nearly optimal auctions. In *NIPS*,
358 pages 136–144, 2015.
- 359 J. Morgenstern and T. Roughgarden. Learning simple auctions. In *COLT*, volume 49 of *JMLR*
360 *Workshop and Conference Proceedings*, pages 1298–1318. JMLR.org, 2016.
- 361 N. Sauer. On the density of families of sets. *J. Combinatorial Theory Ser. A*, 13:145–147, 1972.
- 362 M. Talagrand. Concentration of measure and isoperimetric inequalities in product spaces. *Publications*
363 *Mathématiques de l’Institut des Hautes Etudes Scientifiques*, 81(1):73–205, 1995.
- 364 R. van Handel. Probability in high dimension. Technical report, PRINCETON UNIV NJ, 2014.
- 365 S. Vardi. The returning secretary. In *32nd International Symposium on Theoretical Aspects of*
366 *Computer Science*, page 716, 2015.
- 367 J. Vondrák. A note on concentration of submodular functions. *CoRR*, abs/1005.2791, 2010.

368 **A Deferred Proofs**

369 **A.1 The Greedy Online Algorithm**

370 A simple way to collect a small set of items that contains the optimal solution is to select the k
 371 largest items of each property. This set clearly contains the optimal solution. A simple argument,
 372 as in the proof of Lemma 18, shows that this implementation of the first stage keeps $O(kd \log(n/k))$
 373 items on average. In the following we present a *greedy algorithm* that retains an average number of
 374 $O(k \log(k/\delta))$ items in the first phase (for a parameter $\delta \in (0, 1)$).

375 The greedy algorithm works as follows: it ignores the first $\delta n/k$ items² and then starts processing the
 376 items one by one. When we process the i 'th item, c_i , the algorithm computes the optimal solution M_i
 377 of the first i items (recall that we assume the algorithm has access to C_{dummy} , a large enough pool
 378 of zero valued items so there is always a feasible solution). The greedy algorithm retains c_i if and
 379 only if c_i participates in M_i . We assume that M_i is unique for every i (we can achieve this with an
 380 arbitrary consistent tie breaking rule, say among matchings of the same value we prefer the one that
 381 maximizes the sum of the indices of the matched items.). Since the optimal solutions correspond to
 382 maximum-weighted bipartite-matchings between the items and the constraints, we have the following
 383 lemma.

384 **Lemma 17.** *Suppose that the optimal solution, denoted by M , does not appear before round $\delta n/k$.*
 385 *Then it is a subset of the retained items.*

386 *Proof.* Let $i \geq \delta n/k$. Consider an item c matched by M and assume by contradiction that c is not
 387 matched in M_i . Consider $Z = M \Delta M_i$ (we take the symmetric difference of M and M_i as sets of edges).
 388 Since M and M_i do not necessarily match the same items then the edges in Z induce a collection of
 389 alternating paths and cycles where each path L has an item matched by M and not by M_i at one end,
 390 and an item matched by M_i and not by M at the other hand. Except for its two ends, an alternating
 391 path contains items that are matched by both M and M_i . From the optimality and the uniqueness of M
 392 follows that for each path the value of M is larger than the value of M_i .

393 Since c is matched by M and not by M_i there is a path L in Z that starts at c and ends at some item
 394 that is matched by M_i and not by M .

395 It follows that all the items in L are in M_i and if we match them according to M then the value that we
 396 gain from them increases. This contradicts the optimality of M_i .

397 (Note that, in fact, there are no cycles in Z , since they will imply that there are multiple optimal
 398 solutions, contradicting the uniqueness of M_i and M .) \square

399 Lemma 17 implies that, with high probability, if we collect all items that are in the optimal solution
 400 of the subset of items that precedes them then the set of items that we have at the end contains the
 401 optimal solution. Indeed, our algorithm fails if at least one of the items in the optimal solution M is
 402 among the first $\delta n/k$ items. The probability that this occurs is at most δ via a union bound and the fact
 403 that the probability that any fixed item in M is among the first $\delta n/k$ items is exactly δ/k .

404 The next question is: how large is the subset of the items which we retain? The next lemma answers
 405 this question in an average sense.

406 **Lemma 18.** *Assume that at the first stage the algorithm receives the items in a random order. Then*
 407 *the expected number of items that the first stage keeps is $O(k \log \min \{ \frac{n}{k}, \frac{k}{\delta} \})$.*

408 *Proof.* Let $i \geq \delta n/k$ and denote X_i as an indicator that is one if and only if the i 'th item belongs
 409 to M_i . Condition the probability space on the set L_i of the first i items (but not on their order).
 410 Each element of L_i is equally likely to arrive last. So since $|M_i| \leq k$, then the probability that the
 411 element arriving last in L_i is in OPT_i is at most k/i if $k < i$ or at most 1 otherwise. It follows that
 412 $E[X_i | L_i] \leq \min \{ \frac{k}{i}, 1 \}$. Since this holds for any L_i , it also holds unconditionally as well. Therefore,
 413 if $\delta n/k < k$ then by the fact that $\sum_{i=k+1}^n \frac{1}{i} \leq \log \frac{n}{k}$, the expected number of retained items is

$$k - \frac{\delta n}{k} + \sum_{i=k+1}^n \frac{k}{i} = O\left(k \log \frac{n}{k}\right).$$

²We assume $\delta n/k$ is an integer without loss of generality.

414 Similarly, if $\delta n/k \geq k$ then the expected number of retained items is

$$\sum_{i=\delta n/k+1}^n \frac{k}{i} = O\left(k \log \frac{k}{\delta}\right). \quad \square$$

415 A.2 Generalization and concentration

416 **Technical notation.** For $m \in \mathbb{N}$, the set $\{1, \dots, m\}$ is denoted by $[m]$. Given a family of sets F over
 417 a domain X , and $Y \subseteq X$, the family $\{f \cap Y : f \in F\}$ is denoted by $F|_Y$. Recall that the VC dimension
 418 of F is the maximum size of $Y \subseteq X$ such that $F|_Y$ contains all subsets of Y .

419 **Lemma** (restatement of Lemma 9). $VC(\mathcal{R}) = O(d \log d)$.

420 *Proof.* Let S be a set of items shattered by \mathcal{R} and denote its size by m ; since S is arbitrary, an upper
 421 bound on m implies an upper bound on $VC(\mathcal{R})$. To this end we upper bound the number of subsets
 422 in $\mathcal{R}|_S = \{S \cap R_T : R_T \in \mathcal{R}\}$. Now, there are m items in S with at most m different values. Therefore,
 423 we can restrict our attention to thresholds-policies where each threshold is picked from a fixed set of
 424 $m+1$ meaningful locations (one location in between values of two consecutive items when we sort the
 425 items by value). Thus $|\mathcal{R}|_S| \leq (m+1)^d$, but, as S is shattered, $|\mathcal{R}|_S| = 2^m$ and we get $m \leq d \log_2(m+1)$.
 426 This implies $m = O(d \log d)$ from which we conclude that $VC(\mathcal{R}) = O(d \log d)$. \square

427 **Lemma** (restatement of Lemma 10). $VC(\mathcal{Q}) = O(\log d)$.

428 *Proof.* For $i \leq d$, let $\mathcal{Q}_i = \{R_i^T : T \in \mathcal{T}\}$. Note that $\mathcal{Q} = \cup_i \mathcal{Q}_i$. We claim that $VC(\mathcal{Q}_i) = 1$ for all i .
 429 Indeed, let c', c'' be two items. Note that if $c' \notin P_i$ or $c'' \notin P_i$ then $\{c', c''\}$ is not contained by \mathcal{Q}_i and
 430 therefore not shattered by it. Therefore, assume that $c', c'' \in P_i$ and $v_i(c') \geq v_i(c'')$. Now, it follows
 431 that any threshold T that retains c'' must also retain c' , and so it follows that also in this case $\{c', c''\}$
 432 is not shattered.

433 The bound on the VC dimension of $\mathcal{Q} = \cup_{i \leq d} \mathcal{Q}_i$ follows from the next lemma.

434 **Lemma 19.** Let $m \geq 2$ and let F_1, \dots, F_m be classes with VC dimension at most 1. Then, the VC
 435 dimension of $\cup_i F_i$ is at most $10 \log m$.

436 *Proof.* We show that $\cup_i F_i$ does not shatter a set of size $10 \log m$. Let $Y \subseteq X$ of size $10 \log m$. Indeed,
 437 by the Sauer's Lemma [Sauer, 1972]:

$$|(\cup_i F_i)|_Y| \leq m \left(\binom{10 \log m}{0} + \binom{10 \log m}{1} \right) = m(1 + 10 \log m) < m^{10} = 2^{10 \log m},$$

438 and therefore, Y is not shattered by $\cup_i F_i$. \square

439 This finishes the proof of Lemma 10. \square

440 **Lemma** (restatement of Lemma 16). There exists a collection of \mathcal{N} thresholds-policies such that
 441 $|\mathcal{N}| \leq k^{O(d)}$, and for every thresholds-policy $T \in \mathcal{T}$ there are $T^+, T^- \in \mathcal{N}$ such that

442 1. $V^{T^-}(C) \leq V^T(C) \leq V^{T^+}(C)$ for every sample of items C . (By taking expectations this also
 443 implies that $\mathbf{v}^{T^-} \leq \mathbf{v}^T \leq \mathbf{v}^{T^+}$.)

444 2. $\mathbf{v}^{T^+} - \mathbf{v}^{T^-} \leq 10$.

445 *Proof.* For every $i \leq d$ and $j \leq dn$ define thresholds $t_i^j \in [0, 1]$ where $t_i^0 = 1$ and for $j > 0$ set t_i^j to
 446 satisfy³

$$\Pr_{c \sim D} [v(c) \geq t_i^j \text{ and } c \in P_i] = \frac{j}{dn}.$$

447 Note that $t_i^0 > t_i^1 > \dots$ (see Figure 1). Set

$$\mathcal{J}_i = \left\{ j : 0 \leq \frac{j}{dn} \leq \Pr_{c \sim D} [c \in P_i], j \in \mathbb{N} \right\},$$

³Such t_i^j 's exist due to our assumption that D is atomless (see Section 3).



Figure 1: An illustration of the thresholds in \mathcal{N}_i as defined in the proof of Lemma 16. Each t_i^j for $j \in \mathcal{J}_i$ satisfies $\Pr_{c \sim D}[v(c) \geq t_i^j \text{ and } c \in P_i] = \frac{j}{dn}$.

448 and define

$$\begin{aligned}\mathcal{N}_i &= \{t_i^j \mid j \in \mathcal{J}_i \cap \{0, 1, \dots, 10dk\}\} \cup \{0\} \\ \mathcal{N} &= \mathcal{N}_1 \times \mathcal{N}_2 \dots \times \mathcal{N}_d.\end{aligned}$$

449 Note that indeed $|\mathcal{N}| \leq (10dk + 2)^{d+1} = k^{O(d)}$.

450 We next show that \mathcal{N} satisfies items 1 and 2 in the statement of the lemma. Let $T \in \mathcal{T}$ be an arbitrary
451 thresholds-policy. The policies $T^- = (t_i^-)_{i \leq d}$, and $T^+ = (t_i^+)_{i \leq d}$ are derived by rounding t in each
452 coordinate up and down respectively, to the closest policies in \mathcal{N} (so, the thresholds in T^+ are smaller
453 than in T^- ; the “+” sign reflects that it retains more items and achieves a higher value). Formally,
454 $t_i^+ = \max\{t \in \mathcal{N}_i : t \leq t_i\}$ and $t_i^- = \min\{t \in \mathcal{N}_i : t \geq t_i\}$ where t_i is the threshold for property i in
455 T . Therefore, for every sample $C \sim D^n$, the set of items in C that are retained by T contains the set
456 retained by T^- and is contained in the set retained by T^+ . This implies item 1.

457 To derive item 2, observe that for every sample C : $V^{T^+}(C) - V^{T^-}(C) \leq |Z|$, where $Z \subseteq C$ denotes
458 the set of items which participate in some canonical optimal solution for T^+ that are not retained by
459 T^- . Thus, it suffices to show that $\text{Ex}[|Z|] \leq 10$. To this end put $p_i = \Pr_{c \sim D}[v(c) \geq t_i \text{ and } c \in P_i]$ and
460 partition Z into two disjoint sets $Z = E \cup F$, where E is the set of all items $c_j \in Z$ that are assigned by
461 the optimal solution of T^+ to a property P_i where $p_i < \frac{10k}{n}$, and $F = Z \setminus E$. We claim that

- 462 • $\text{Ex}[|E|] \leq 1$: for each P_i such that $p_i < \frac{10k}{n}$ let $G_i \subseteq P_i$ denote the set of items whose
463 value $v \in [t_i^+, t_i^-]$ (i.e. retained by T^+ and not by T^-). Note that $E \subseteq \cup_i G_i$, and that
464 $\Pr_{c \sim D}[c \in G_i] \leq \frac{1}{dn}$. Thus, it follows that

$$\text{Ex}_{C \sim D^n}[|E|] \leq \text{Ex}_{C \sim D^n}[|\cup_i G_i|] \leq \sum_i \text{Ex}_{C \sim D^n}[|G_i|] \leq d \cdot \frac{n}{dn} \leq 1.$$

- 465 • $\text{Ex}[|F|] \leq 9$: note that $\text{Ex}[|F|] \leq k \cdot \Pr[|F| > 0]$ (because $F \subseteq Z$ and $|Z| \leq k$). Thus,
466 it suffices to show that $\Pr[|F| > 0] \leq \frac{9}{k}$. Indeed, $F \neq \emptyset$ only if there is a property P_i with
467 $p_i \geq \frac{10k}{n}$ such that less than k items from P_i are retained by T^- . Fix a property P_i such that
468 $p_i \geq \frac{10k}{n}$ and let $p_i^- = \Pr_{c \sim D}[v(c) \geq t_i^- \text{ and } c \in P_i]$. Since $p_i^- \geq \frac{10k}{n}$, a multiplicative Chernoff
469 bound yields that

$$\Pr_{C \sim D^n}[\text{less than } k \text{ items from } P_i \text{ are retained by } T^-] \leq \exp\left(-\frac{(9/10)^2}{2} 10k\right) \leq \frac{9}{k^2} \leq \frac{9}{dk},$$

470 and a union bound over all such properties P_i implies that $\Pr[|F| > 0] \leq \frac{9d}{dk} \leq \frac{9}{k}$.

471 Thus, it follows that $v^{T^+} - v^{T^-} \leq 1 + k \cdot \frac{9}{k} = 10$, which finishes the proof.

472 □

473 **B Lower Bounds**

474 **B.1 Necessity of the training phase**

475 Let $n \in \mathbb{N}$ (sample size) and $\delta \in [0, 1]$ (confidence parameter). In this section we focus on the case
 476 where there is no training phase and $d = 1$. Thus, we consider algorithms which get as an input a
 477 sequence $v_1, \dots, v_n \in [0, 1]$ in an online manner (one after the other). In step m the algorithm needs to
 478 decide whether to retain v_m or to discard it (this decision may depend on the prefix $v_1 \dots, v_m$). The
 479 algorithm is not allowed to discard a sample after it has been retained.

480 The following property captures the utility of the algorithm: *for every distribution μ over $[0, 1]$, if*
 481 *v_1, \dots, v_n are sampled i.i.d from μ , then with probability at least $1 - \delta$, the algorithm retains v_{j_1}, \dots, v_{j_k}*
 482 *that are the largest k elements in v_1, \dots, v_n .* The goal is to achieve this while minimizing the number
 483 of retained items in expectation.

484 **Theorem** (Theorem 2 restatement). *Let $\delta \in (0, 1)$. For every algorithm A which retains the maximal*
 485 *k elements with probability at least $1 - \delta$, there exists a distribution μ such that the expected number*
 486 *of retained elements for input sequences $v_1 \dots v_n \sim \mu^n$ is at least $\Omega(k \log(\min\{n/k, k/\delta\}))$.*

487 We remind that the bound is tight for the greedy algorithm (Theorem 1).

488 *Proof.* Following [Moran et al., 1985, Corollary 3.4], we may assume that A accesses its input only
 489 using comparisons. More precisely: call two sequences v_1, \dots, v_m and u_1, \dots, u_m *order-equivalent*
 490 if $v_i \leq v_j \iff u_i \leq u_j$ for all $i, j \leq m$, and call the equivalence class of $v_1 \dots v_m$ its *order-type*.
 491 Note that if v_1, \dots, v_m are distinct, then their order-type is naturally identified with a permutation
 492 $\sigma \in \mathbb{S}_m$. Call an algorithm A *order-invariant* if for every $m \leq n$, the decision⁴ of A whether to
 493 retain v_m depends only on the order-type of v_1, \dots, v_m (equivalently, A accesses the input only using
 494 comparisons).

495 By Moran et al. [1985] it follows that for every algorithm A there is an infinite $W \subseteq [0, 1]$ such that A
 496 is order-invariant when restricted to input sequences $v_1, \dots, v_n \in W$. For the remainder of the proof
 497 we fix such an infinite set W and focus only on inputs from W .

498 Set μ to be a uniform distribution over a sufficiently large subset of W so that $v_1 \dots v_n$ are distinct
 499 with probability $1 - 1/n$. Let $\text{OPT}(S)$ denote the top k elements in S . Let T_m be the set of all sequences
 500 $v_1, \dots, v_m, \dots, v_n \in W^n$ such that $v_m \in \text{OPT}(\{v_1, \dots, v_m\})$, and let p_k denote the probability that
 501 A retains v_m conditioned on the input being from T_m . Let $T'_m \subseteq T_m$ denote the set of all sequences
 502 $v_1, \dots, v_m, \dots, v_n$ such that $v_m \in \text{OPT}(\{v_1, \dots, v_n\})$ (i.e., v_m is part of the optimal solution). The
 503 proof hinges on the following lemma:

504 **Lemma 20.** *Since A is order based, for every $m \leq n$, p_m is also the probability that A retains v_m*
 505 *conditioned on the input being from T'_m .*

506 *Proof.* The decision of A whether to retain v_m depends only on the order-type of v_1, \dots, v_m . For each
 507 $\sigma \in \mathbb{S}_m$, let $E(\sigma)$ denote the event that the order type of $v_1 \dots v_m$ is σ . Thus,

$$p_m = \Pr[A \text{ retains } v_m \mid T_m] = \sum_{\sigma \in \mathbb{S}_m} \Pr[E(\sigma) \mid T_m] \cdot \Pr[A \text{ retains } v_m \mid E(\sigma)],$$

508 and similarly

$$\Pr[A \text{ retains } v_m \mid T'_m] = \sum_{\sigma \in \mathbb{S}_m} \Pr[E(\sigma) \mid T'_m] \cdot \Pr[A \text{ retains } v_m \mid E(\sigma)].$$

509 Next, observe that for each order-type $\sigma \in \mathbb{S}_m$:

$$\Pr[E(\sigma) \mid T_m] = \Pr[E(\sigma) \mid T'_m] = \begin{cases} \frac{1}{m!}, & m \leq k \\ \frac{1}{k(m-1)!}, & v_m \in \text{OPT}(\{v_1, \dots, v_m\}), m > k \\ 0, & \text{otherwise.} \end{cases} \quad \square$$

⁴When A is randomized then the value of $\Pr[A \text{ retains } v_k]$ depends only on the order-type of $v_1 \dots v_m$.

510 With the above lemma in hand, we can finish the proof. For the remainder of the argument, we
 511 condition the probability space on the event that all elements in the sequence v_1, \dots, v_n are distinct
 512 and show that conditioned on this event, A retains at least $t = \Omega(\log(1/\delta))$ elements in expectation.
 513 Note that this will conclude the proof since by the choice of μ this event occurs with probability
 514 $\geq 1 - 1/n$, which implies that – unconditionally – A retains at least $t - n \cdot (1/n) = t - 1 = \Omega(\log(1/\delta))$
 515 elements in expectation.

516 For each $m \leq n$, v_m is among the top k elements with probability $\min\{1, k/m\}$, in which case it is
 517 retained with probability p_m . So A retains at least

$$\sum_{m=1}^n \min\left\{1, \frac{k}{m}\right\} \cdot p_m$$

518 elements in expectation. By the above lemma, the probability that A discards the maximum is

$$\sum_{m=1}^n \frac{k}{n} \cdot (1 - p_m),$$

519 which by assumption is smaller than δ . So we obtain that $\sum p_m \geq n(1 - \delta/k)$. Thus, to minimize
 520 $\sum_{m=1}^n \min\{1, k/m\} \cdot p_m$ subject to the constraint that $\sum p_m \geq n(1 - \delta/k)$ we make the last $n(1 - \delta/k)$
 521 p_m 's equal to 1 and the rest 0. This gives the desired lower bound. \square

522 B.2 The algorithm from Theorem 4 is optimal

523 In the previous section we have presented an algorithm that with probability at least $1 - \delta$ outputs an
 524 optimal solution while retaining at most $O(k(\log \log n + \log d + \log \log(1/\delta)))$ items in expectation
 525 during the first phase.

526 We now present a proof of Theorem 5. We start with the following lemma that shows the dependence
 527 on δ cannot be improved in general, even for $k = 1$, when there are no constraints, and the distribution
 528 over the items is known to the algorithm (so there is no need to train it on a sample from the
 529 distribution):

530 **Lemma 21.** *Let $v_1, \dots, v_n \in [0, 1]$ be drawn uniformly and independently, let $e^{-n/2} < \delta < 1/10$ and*
 531 *let A be an algorithm that retains the maximal value among the v_i 's with probability at least $1 - \delta$.*
 532 *Then,*

$$\text{Ex}[|S|] = \Omega\left(\log \log\left(\frac{1}{\delta}\right)\right),$$

533 where S is the set of values retained by the algorithm.

534 Thus, it follows that for $\delta = \text{poly}(1/n)$ and $k, d = O(1)$ the bound in Theorem 4 is tight.

535 *Proof.* Define $\alpha = \frac{\ln(1/\delta)}{2n} \in (1/n, 1/4)$. Let E_t denote the event that $v_t \geq 1 - \alpha$ and is the largest among
 536 v_1, \dots, v_t . We have that

$$\text{Ex}[|S|] \geq \sum_t \Pr[v_t \text{ is picked and } E_t] = \sum_t (\Pr[E_t] - \Pr[v_t \text{ is rejected and } E_t]) . \quad (3)$$

537 We show that since A errs with probability at most δ then $\sum_t \Pr[E_t \text{ and } v_t \text{ is rejected}]$ is small.

$$\begin{aligned} \delta &\geq \Pr[A \text{ rejects } v_{\max}] \geq \sum_t \Pr[A \text{ rejects } v_t \text{ and } E_t \text{ and } v_t = v_{\max}] \\ &= \sum_t \Pr[v_t = v_{\max} \mid A \text{ rejects } v_t \text{ and } E_t] \cdot \Pr[A \text{ rejects } v_t \text{ and } E_t] \\ &\geq \sum_t \Pr[v_i \leq 1 - \alpha \text{ for all } i > t \mid A \text{ rejects } v_t \text{ and } E_t] \cdot \Pr[A \text{ rejects } v_t \text{ and } E_t] \\ &= \sum_t \Pr[v_i \leq 1 - \alpha \text{ for all } i > t] \cdot \Pr[A \text{ rejects } v_t \text{ and } E_t] \end{aligned}$$

$$\begin{aligned} &\geq \sum_t (1 - \alpha)^{n-t} \cdot \Pr[A \text{ rejects } v_t \text{ and } E_t] \\ &\geq (1 - \alpha)^n \sum_t \Pr[A \text{ rejects } v_t \text{ and } E_t]. \end{aligned}$$

538 The crucial part of the above derivation is in third line. It replaces the event “ $v_t = v_{max}$ ” by the event
539 “ $v_i \leq 1 - \alpha$ for all $i > t$ ” (which is contained in the event “ $v_t = v_{max}$ ” under the above conditioning).
540 The gain is that the events “ $v_i \leq 1 - \alpha$ for all $i > t$ ” and “ A rejects v_t and E_t ” are independent (the
541 first depends only on v_i for $i > t$ and the latter on v_i for $i \leq t$). This justifies the “=” in the fourth line.
542 Rearranging, we have $\sum_t \Pr[A \text{ rejects } v_t \text{ and } E_t] \leq \frac{\delta}{(1 - \alpha)^n}$. Substituting this bound in Eq. (3),

$$\begin{aligned} \text{Ex}[|S|] &\geq \sum_t \Pr[v_t \text{ is picked and } E_t] \\ &= \sum_t (\Pr[E_t] - \Pr[v_t \text{ is rejected and } E_t]) \\ &= \sum_t \Pr[E_t] - \frac{\delta}{(1 - \alpha)^n} \\ &\geq \frac{1}{4} \ln(\alpha n) - \delta \cdot \exp(2\alpha n) && \text{(explained below)} \\ &= \frac{1}{4} \ln\left(\frac{\ln(1/\delta)}{2}\right) - \delta \exp(\ln(1/\delta)) && \text{(by the definition of } \alpha) \\ &= \frac{1}{4} \ln \ln(1/\delta) - \frac{1}{4} \ln 2 - 1 = \Omega(\log \log(1/\delta)), \end{aligned}$$

543 which is what we needed to prove. The last inequality follows because

544 (i) $\sum_t \Pr[E_t] \geq \frac{1}{4} \ln(\alpha n)$ (as is explained next), and

545 (ii) $1 - \alpha \geq \exp(-2\alpha)$ for every $\alpha \in [0, \frac{1}{4}]$ (which can be verified using basic analysis).

546 To see (i), note that

$$\sum_t \Pr[E_t] = \text{Ex}\left[\sum_t 1_{E_t}\right].$$

547 Let $z = |\{t : v_t \geq 1 - \alpha\}|$. Since the v_i 's are uniform in $[0, 1]$ then by the same argument as in the
548 proof of Lemma 18 we get that

$$\text{Ex}\left[\sum_t 1_{E_t} \mid z\right] = \sum_{i=1}^z \frac{1}{i} \geq \int_1^{z+1} \frac{1}{x} = \ln(z+1),$$

549 and therefore

$$\text{Ex}\left[\sum_t 1_{E_t}\right] = \text{Ex}_z \text{Ex}\left[\sum_t 1_{E_t} \mid z\right] \geq \text{Ex}_z [\ln(z+1)].$$

550 Let $Z \sim \text{Bin}(n, \alpha)$, and therefore we need to lower bound $\text{Ex}[\ln(Z+1)]$ for $Z \sim \text{Bin}(n, \alpha)$. To this end,
551 we use the assumption that $\alpha > 1/n$, and therefore $\Pr[Z \geq \alpha \cdot n] \geq 1/4$ (see Greenberg and Mohri, 2013
552 for a proof of this basic fact). In particular, this implies that $\text{Ex}[\ln(Z+1)] \geq \frac{1}{4} \ln(\alpha n + 1) > \frac{1}{4} \ln(\alpha n)$,
553 which finishes the proof. \square

554 Lemma 21 implies Theorem 5 as follows: set $k = d$, $k_1 = \dots = k_d = 1$ and $n \geq 100k \log(1/\delta)$. Pick
555 a distribution D which is uniform over items, each satisfying exactly one of d properties, and with
556 value drawn uniformly from $[0, 1]$.

557 It suffices to show that with probability of at least $1/3$, the algorithm retains an expected number
558 of $\Omega(\log \log(1/\delta))$ items from a constant fraction, say $1/4$, of the properties i . This follows from
559 Lemma 21 as we argue next. Let n_i denote the number of observed items of property i . Then, since
560 $\text{Ex}[n_i] = n/d = n/k \geq 100$, the multiplicative Chernoff bound implies that $n_i \geq n/2k \geq 2 \log(1/\delta)$

561 with high probability (probability = $1/2$ suffices). Therefore, the expected number of properties i 's for
562 which $n_i \geq 2 \log(1/\delta)$ is at least $k/2$. Now, consider the random variable Y which counts for how many
563 properties i we have $n_i \geq 2 \log(1/\delta)$. Since Y is at most k and $\text{Ex}[Y] \geq k/2$, then a simple averaging
564 argument implies that with probability of at least $1/3$ we have that $Y \geq k/4$. Conditioning on this
565 event (which happens with probability $\geq 1/3$), Lemma 21 implies⁵ that $\text{Ex}[|S_i|] = \Omega(\log \log(1/\delta))$ for
566 each of these i 's.

⁵Note that to apply Lemma 21 on S_i we need $\delta > e^{-n_i/2}$, which is equivalent to $n_i > 2 \ln(1/\delta)$.