



Competitive Query Minimization for Stable Matching with One-Sided Uncertainty

Evrpidis Bampis ✉ 

Sorbonne Université, CNRS, LIP6, Paris, France

Konstantinos Dogeas 

Department of Computer Science, Durham University, Durham, United Kingdom

Thomas Erlebach ✉ 


Department of Computer Science, Durham University, Durham, United Kingdom

Nicole Megow ✉ 

Faculty of Mathematics and Computer Science, University of Bremen, Bremen, Germany

Jens Schlöter ✉ 

Faculty of Mathematics and Computer Science, University of Bremen, Bremen, Germany

Amitabh Trehan ✉ 

Department of Computer Science, Durham University, Durham, United Kingdom

Abstract

We study the two-sided stable matching problem with one-sided uncertainty for two sets of agents A and B , with equal cardinality. Initially, the preference lists of the agents in A are given but the preferences of the agents in B are unknown. An algorithm can make queries to reveal information about the preferences of the agents in B . We examine three query models: comparison queries, interviews, and set queries. Using competitive analysis, our aim is to design algorithms that minimize the number of queries required to solve the problem of finding a stable matching or verifying that a given matching is stable (or stable and optimal for the agents of one side). We present various upper and lower bounds on the best possible competitive ratio as well as results regarding the complexity of the offline problem of determining the optimal query set given full information.

2012 ACM Subject Classification Theory of Computation → Design and analysis of algorithms

Keywords and phrases Matching under Preferences, Stable Marriage, Query-Competitive Algorithms, Uncertainty

Related Version An extended abstract of this paper appears in the proceedings of the International Conference on Approximation Algorithms for Combinatorial Optimization Problems (APPROX 2024).

Funding This research was supported by EPSRC grant EP/S033483/2.

Evrpidis Bampis: Partially funded by the grant ANR-19-CE48-0016 from the French National Research Agency (ANR).

Nicole Megow: Supported by DFG grant no. 547924951.

1 Introduction

In the classical two-sided stable matching problem, we are given two disjoint sets A and B of agents (often referred to as men and women) of equal cardinality n . Each agent has a complete preference list over the agents of the other set. The task is to find a *stable* matching, i.e., a one-to-one allocation in which no two agents prefer to be matched to each other rather than to their current matching partners. This problem has applications in numerous allocation markets, e.g., university admission, residency markets, distributed internet services, etc. Since its introduction by Gale and Shapley [12] this problem has been

43 widely studied in different variants from both practical and theoretical perspectives; we refer
44 to the books [14, 30, 24].

45 While the majority of the literature assumes full information about the preference lists,
46 this may not be realistic in large matching markets. It might be impractical or too costly and
47 not even necessary to gather the complete preferences. Hence, different models for uncertainty
48 in the preferences have received attention in the past decade [1, 2, 5, 6, 8, 16, 15, 28, 29]. Many
49 of these works rely on probabilistic models and guarantees. This may not be appropriate for
50 applications in which no (correct) distributional information is available, e.g. in one-time
51 markets. Further, one might ask for guaranteed properties such as stability and optimality
52 instead of probabilistic ones.

53 A different way of handling uncertainty in the preferences is to allow an algorithm to make
54 queries to learn about the unknown preferences. Various types of queries (in terms of both
55 input and output) are conceivable, with one example being *interview queries* [5, 6, 28, 29].
56 Here one asks for a query sequence where a query corresponds to an interview between two
57 potential matching partners and the outcome is the placement of the interview partners in
58 each other's preference list among all other candidates that she has interviewed so far. Hence,
59 if an agent has several such interviews then she finds out her preference order over all these
60 candidates.

61 In this paper we investigate various query models for stable matching problems with
62 *one-sided* uncertainty in the preferences. We assume that initially only the preference lists of
63 one side, A , are known but the preference lists of the other side, B , are unknown. Applications
64 include allocations between groups of different seniority or when preferences shall be kept
65 private; see also [16, 15, 17]. For illustration consider, e.g., pairing new staff with mentors or
66 new PhD students with supervisors as part of the onboarding. New staff can be asked to
67 provide a full preference list of mentors based on information about the available mentors
68 that can be made accessible with little effort, while requiring mentors to rank potential
69 mentees might be considered too burdensome for senior staff due to other significant time
70 commitments.

71 We consider three types of queries to gain information about the preferences, namely
72 (i) *comparison queries* that reveal for an agent $b \in B$ and a pair of agents from A which
73 one b prefers, (ii) *set queries* that reveal for an agent $b \in B$ and a subset $S \subseteq A$ the agent
74 in S that b prefers most, and (iii) *interview queries*.

75 We study basic problems regarding stability and optimality of matchings using these query
76 models. A stable matching is called A -optimal (resp. B -optimal) if no agent in A (resp. B)
77 prefers a different stable matching over the current one. To our knowledge, most existing
78 related work considers worst-case bounds on the absolute number of queries necessary to solve
79 the respective problem; see Further Related Work below for a discussion. For many instances,
80 however, executing such a worst-case number of queries might not be necessary. To also
81 optimize the number of queries on these instances, we analyze our algorithms using *competitive*
82 *analysis*. We say that an algorithm that makes queries until it can output a provably correct
83 answer, e.g., a stable and A -optimal matching, is ρ -competitive (or ρ -query-competitive) if it
84 makes at most ρ times as many queries as the minimum possible number of queries that also
85 output a provably correct answer for the given instance. Note that this answer may differ
86 from that of the algorithm, e.g., a different stable matching. In this paper, we design upper
87 and lower bounds on the competitive ratios for the above mentioned problems and query
88 models. Our results illustrate that worst-case instances regarding the competitive ratio are
89 very different from the worst-case instances regarding the absolute number of queries. Thus,
90 our lower bounds on the competitive ratio use different instances and our algorithms are

91 designed to optimize on different instances. Indeed, worst-case instances for the absolute
 92 number of queries turn out to be ‘easy’ for competitive analysis as the optimal solution we
 93 compare against is very large.

94 Query-competitive algorithms are often associated with the field of ‘explorable uncertainty’.
 95 Most previous work considers queries revealing an originally uncertain *value* [3, 7, 9, 18, 19,
 96 21, 25, 10], while in this work we query a *preference*.

97 **Our Contribution.** We study the stable matching problem with one-sided uncertainty in the
 98 preference lists and give the following main results. Note that we assume that the preferences
 99 of the B side are unknown, and $|A| = |B| = n$. We remark that our technically most involved
 100 main results are lower bounds on the competitive ratio and hardness results, so the results
 101 only get stronger by making these assumptions.

102 In Section 3 we focus on *comparison queries*. Firstly, we ask the question of how to verify
 103 that a given matching is stable. We show that the problem can be solved with a 1-competitive
 104 algorithm. Then we ask how to find a stable matching under one-sided uncertainty. We
 105 give a 1-competitive algorithm that finds a stable matching and, moreover, the solution is
 106 provably A -optimal. Essentially, we employ the well-known *deferred acceptance algorithm*,
 107 first analyzed by Gale and Shapley [12], and compare its number of queries carefully with
 108 the number of queries that any algorithm needs to verify a stable matching.

109 A substantially more challenging task is to find a B -optimal stable matching. Note that
 110 a trivial competitive ratio is $O(n^2 \log n)$, as it is possible to obtain the full preferences of
 111 each of the n elements in B using $O(n \log n)$ queries, and the optimum total number of
 112 queries is at least 1. One of our main contributions is a tight bound of $O(n)$. To that end,
 113 we first show that every algorithm for verifying that a given matching is B -optimal and
 114 stable requires $\Omega(n)$ queries. Then we give an $O(n)$ -competitive algorithm for the problem
 115 of finding one. This is best possible up to constant factors, which we prove with a matching
 116 lower bound that also holds for verifying that a given matching is stable and B -optimal, even
 117 for randomized algorithms.

118 We complement these results by showing that the offline problem of determining the
 119 optimal number of queries for finding the B -optimal stable matching is NP-hard, and we give
 120 an $\mathcal{O}(\log n \log \log n)$ -approximation algorithm. Here, the *offline* version of a problem is to
 121 compute, given full information about the preferences of all agents, a smallest set of queries
 122 with the property that an algorithm making exactly those queries has sufficient information
 123 to solve the problem with one-sided uncertainty.

124 Section 4 discusses interview queries. We show that the bounds on the competitive
 125 ratio and hardness results for comparison queries translate to interview queries. We remark
 126 that some of these results for interview queries, e.g., a 1-competitive algorithm for finding
 127 an A -optimal stable matching, were already proven by Rastegari et al. [29] and discuss
 128 differences to their results in the corresponding section. Interestingly, we can use essentially
 129 the same techniques as for the comparison model. This may seem surprising, especially for
 130 the lower bounds, as interview queries seem to be more powerful. For instance, n interviews
 131 are sufficient to determine the precise preference order of an agent $b \in B$, while we need
 132 $\Omega(n \log n)$ comparison queries to determine b ’s preference order. On the other hand, an
 133 instance that can be solved with a single comparison query requires two interviews. In
 134 general, we can simulate a comparison query by using two interview queries.

135 In Section 5 we discuss the *set query* model. While some bounds remain the same as in
 136 the other models, e.g., 1-competitiveness for verifying the stability of a given matching, we
 137 show that some bounds change drastically. For example, we give an $\mathcal{O}(\log n)$ -competitive

138 algorithm for verifying that a given matching is B -optimal, which is in contrast to the
 139 lower bound of $\Omega(n)$ in the other query models. It remains open whether $\mathcal{O}(1)$ -competitive
 140 algorithms exist for the problems of finding a stable matching or verifying a B -optimal
 141 matching with set queries.

142 **Further Related Work.** In classical work on stable matching with queries, the preferences on
 143 both sides can only be accessed via queries, with a query usually either asking for the i th entry
 144 in a preference list or for the rank of a specific element within a preference list (cf. e.g. [26]).
 145 Note that two rank queries are sufficient to simulate a comparison query, but up to $n - 1$
 146 comparison queries are needed to obtain the information of a single rank query. Thus, existing
 147 lower bounds on the necessary number of rank queries in these query models translate to
 148 our setting (up to a constant factor), but upper bounds do not necessarily translate. Ng
 149 and Hirschberg [26] showed that $\Theta(n^2)$ such queries are necessary to find or verify a stable
 150 matching in the worst case. The lower bound of $\Omega(n^2)$ translates to any type of queries
 151 with boolean answers, including comparison queries [13]. Further work on interview queries
 152 includes empirical results [5, 6] and complexity results [28] on several decision problems
 153 under partial uncertainty. We discuss the latter in Section 4.

154 Our setting of one-sided uncertainty and querying uncertain preferences is also related to
 155 existing work on online algorithms for *eliciting partial preferences* [20, 27, 23]. These works
 156 also consider a setting where the preferences of agents in one of the sets are uncertain but can
 157 be determined by using different types of queries. In particular, [27] also considers the set
 158 query model. The main difference to our work is that these papers assume that the elements
 159 of one set do not have any preferences at all. As a consequence, they do not consider stability
 160 at all and instead aim at computing pareto-optimal or rank-maximal matchings.

161 2 Preliminaries

162 An instance of the *two-sided stable matching problem* consists of two disjoint sets A and B of
 163 size $|A| = |B| = n$ and complete preference lists: The preference list for each agent $a \in A$
 164 is a total order \prec_a of B , the preference list of each agent $b \in B$ is a total order \prec_b of A .
 165 Here, $a_1 \prec_b a_2$ means that b prefers a_1 to a_2 . A matching is a bijection from A to B . For
 166 a matching M , we denote the element of B that is matched to $a \in A$ by $M(a)$, and the
 167 element of A that is matched to $b \in B$ by $M(b)$.

168 Given a matching M , a pair $(a, b) \in A \times B$ is a *blocking pair* in M if a is not matched
 169 to b in M , a prefers b to $M(a)$, and b prefers a to $M(b)$. A matching M is called a *stable*
 170 matching if there is no blocking pair in M .

171 In their influential paper, Gale and Shapley [12] showed that a stable matching always
 172 exists, and the *deferred acceptance algorithm* computes one in $\mathcal{O}(n^2)$ time. In this algorithm,
 173 one group (A or B) proposes matches and the other decides whether to accept or reject each
 174 proposal. The algorithm produces a stable matching that is best possible for the group X
 175 that proposes (we say *X -optimal*) and worst possible for the other group: Each element of
 176 the group that proposes gets matched to the highest-preference element to which it can be
 177 matched in any stable matching, and each element of the other group gets matched to the
 178 lowest-preference element to which it can be matched in any stable matching.

179 In this paper, we consider the setting of *one-sided uncertainty*, where initially only the
 180 preference lists of all agents in A are known, but the preference lists of $b \in B$ are unknown.
 181 An algorithm can make queries to learn about the preferences of $b \in B$. We distinguish the
 182 following types of queries:

- 183 ■ *Comparison queries:* For agents $b \in B$ and $a_1, a_2 \in A$, the query $\text{prefer}(b, a_1, a_2)$ returns
 184 a_1 if b prefers a_1 to a_2 and a_2 otherwise. These queries can also be seen as Boolean
 185 queries that return true iff b prefers a_1 to a_2 .
- 186 ■ *Set queries:* For agents $b \in B$ and any subset $S \subseteq A$, the query $\text{top}(b, S)$ returns b 's most
 187 preferred element of S .
- 188 ■ *Interview queries:* For agents $b \in B$ and $a \in A$, an interview query $\text{intq}(b, a)$ reveals the
 189 total order of the subset $\{a\} \cup P_b$ defined by \prec_b , where P_b is the set of all elements $a' \in A$
 190 for which a query $\text{intq}(b, a')$ has already been executed before the query $\text{intq}(b, a)$.

191 A *stable matching instance with one-sided uncertainty* is given by two sets A and B of
 192 size n and, for each agent $a \in A$, a total order \prec_a of the agents in B . The preferences
 193 of the agents in B are initially unknown. For a given stable matching instance with one-
 194 sided uncertainty, we consider the following problems: *finding a stable matching*, *finding*
 195 *an A -optimal stable matching*, and *finding a B -optimal stable matching*. For a given stable
 196 matching instance with one-sided uncertainty and a matching M , we consider the following
 197 problems: *verifying that M is stable*, *verifying that M is stable and A -optimal*, and *verifying*
 198 *that M is stable and B -optimal*. All problems can be considered for each query model. For
 199 the verification problems, we consider the competitive ratio only for inputs where M is
 200 indeed a stable (and A - or B -optimal) matching. If this is not the case, the algorithm must
 201 detect this, but we do not compare the number of queries it makes to the optimum. This is
 202 because any algorithm may be required to make up to $\Omega(n^2)$ comparison or interview queries
 203 to detect a blocking pair, while the optimum can prove its existence with a constant number
 204 of queries.

205 It is easy to see that for the optimum, the problem of verifying that a given matching M
 206 is stable and A -optimal (B -optimal) is the same as that of finding the A -optimal (B -optimal)
 207 stable matching. This implies that any lower bound on the number of queries required to
 208 verify that M is stable and A -optimal (B -optimal) also applies to the problem of finding the
 209 A -optimal (B -optimal) stable matching.

210 An important concept is the notion of *rotations*, which can be defined as follows (cf. [24]):
 211 Let a stable matching M be given. For an agent $a_i \in A$, let $s_A(a_i)$ denote the most-preferred
 212 element b_j on a_i 's preference list such that b_j prefers a_i to her current partner $M(b_j)$. Note
 213 that $s_A(a_i)$ must be lower than $M(a_i)$ in a_i 's preference list as otherwise $(a_i, s_A(a_i))$ would be
 214 a blocking pair. Let $\text{next}_A(a_i) = M(s_A(a_i))$. Then a rotation (exposed) in M is a sequence
 215 $(a_{i_0}, b_{j_0}), \dots, (a_{i_{r-1}}, b_{j_{r-1}})$ of pairs such that, for each k ($0 \leq k \leq r-1$), $(a_{i_k}, b_{j_k}) \in M$ and
 216 $a_{i_{k+1}} = \text{next}_A(a_{i_k})$, where addition is modulo r . The rotation can be viewed as an alternating
 217 cycle consisting of the matched edges (a_{i_k}, b_{i_k}) and the unmatched edges $(a_{i_k}, b_{i_{k+1}})$ (for
 218 $0 \leq k \leq r-1$). We refer to an edge $(a, s_A(a))$ as a *rotation edge* or *r -edge* as it can potentially
 219 be part of a rotation. Note that every vertex $a \in A$ is incident with at most one r -edge.

220 Given a rotation R in a stable matching M , we can construct a stable matching M'
 221 from M by removing all edges that are part of R and M and adding all r -edges that are
 222 part of R . We refer to this as *applying* a rotation. Observe that no agent in B is worse off in
 223 M' than in M , and some agents in B prefer M' to M . The following has been shown.

224 ► **Lemma 1** (Lemma 2.5.3 in Gusfield and Irving [14]). *If M is any stable matching other*
 225 *than the B -optimal stable matching, then there is at least one rotation exposed in M .*

226 **3 Stable Matching with Comparison Queries**

227 In this section, we consider the comparison query model with one-sided uncertainty. We first
 228 discuss our results on the problems of verifying that a given matching is stable and finding

229 an A -optimal matching, before moving on to our main results regarding the competitive ratio
 230 for finding/verifying a B -optimal matching. Finally, we briefly consider the variation with
 231 two-sided uncertainty and give tight bounds for the problem of verifying a stable matching
 232 in that model.

233 3.1 Verifying That a Given Matching Is Stable

234 In this section, we consider the *verification problem* where we are given a matching M and
 235 our task is to verify that M is indeed stable. We give a 1-competitive algorithm. As argued
 236 in the previous section, we only care about the competitive ratio if the given matching M is
 237 indeed stable. If the given matching M is not stable, the algorithm must detect this, but
 238 its number of queries can be arbitrarily much larger than the optimal number of queries for
 239 detecting that M is not stable. In the case of one-sided uncertainty, as we consider it here, a
 240 single query is sufficient for the optimum to identify a blocking pair.

241 The following auxiliary lemma shows that exploiting transitivity cannot reduce the number
 242 of comparison queries to an agent $b \in B$ if one needs to find out the preference relationship
 243 of k agents from A to one particular agent from A in b 's preference list.

244 ► **Lemma 2.** *Consider two agents $a \in A$, $b \in B$ and assume that there are k agents
 245 $a_1, \dots, a_k \in A \setminus \{a\}$ for each of which we want to know whether b prefers that agent to a or not.
 246 Then exactly k comparison queries to b are necessary and sufficient to obtain this knowledge.*

247 **Proof.** The k queries $\text{prefer}(b, a, a_i)$ for $i = 1, \dots, k$ are clearly sufficient. Assume that k'
 248 queries to b , for some $k' < k$, are sufficient to obtain the desired information. Consider the
 249 auxiliary graph H with vertex set $V_H = A$ and an edge $\{a', a''\}$ for each of those k' queries
 250 $\text{prefer}(b, a', a'')$. As the set $A' = \{a, a_1, a_2, \dots, a_k\}$ has $k+1$ vertices and H has fewer than k
 251 edges, the set A' intersects at least two different connected components of H . Let a_j , for some
 252 $1 \leq j \leq k$, be a vertex that does not lie in the same component as a . Then the k' queries do
 253 not show whether b prefers a_j to a or not, which contradicts k' queries being sufficient. ◀

254 If an algorithm obtains for agents x and y with $y \neq M(x)$ the information that $M(x) \prec_x y$
 255 (either via a direct query or via transitivity), we say that the algorithm *relates* y to $M(x)$
 256 for x . By Lemma 2, if the optimum relates k different elements to $M(x)$ for x , it needs to
 257 make k queries to x . A pair (x, y) with $y \neq M(x)$ such that the optimum relates y to $M(x)$
 258 for x is called a *relationship pair* (for x). Lemma 2 implies the following.

259 ► **Corollary 3.** *The total number of relationship pairs (for all agents x) is a lower bound on
 260 the number of comparison queries the optimum makes.*

261 ► **Theorem 4.** *Given a stable matching instance with one-sided uncertainty and a stable
 262 matching M , there is a 1-competitive algorithm that uses $\sum_{a \in A} |\{b \in B \mid b \prec_a M(a)\}|$
 263 queries for verifying that M is stable in the comparison query model.*

264 **Proof.** Since the preferences of agents on the A -side are not uncertain, for each $(a, b) \notin M$,
 265 we already know whether $M(a) \prec_a b$. If $M(a) \prec_a b$, then we do not have to execute any
 266 queries to show that $(a, b) \notin M$ is not a blocking pair. Otherwise, every feasible query set
 267 has to prove $M(b) \prec_b a$. Therefore, for each element $b \in B$, there is a uniquely determined
 268 number n_b of elements of A that any solution (including the optimum) must relate to $M(b)$
 269 for b . Let $K = \sum_{b \in B} n_b$ be the resulting number of relationship pairs.

270 Our algorithm simply queries $\text{prefer}(b, M(b), a)$ for every pair $(a, b) \notin M$ for which
 271 $b \prec_a M(a)$. These are exactly K queries. As the total number of relationship pairs is K , the
 272 optimum must also make K queries (Corollary 3). Hence, our algorithm is 1-competitive. ◀

273 The proof of Theorem 4 implies that the stable matching that maximizes the number of
 274 queries that are required to prove stability is the B -optimal matching.

275 ▶ **Corollary 5.** *The number of comparison queries needed to verify that the B -optimal*
 276 *matching is stable is $\max_{M \text{ stable}} \sum_{a \in A} |\{b \in B \mid b \prec_a M(a)\}|$.*

277 3.2 Finding an A -Optimal Stable Matching

278 We obtain the following positive result by adapting the classical deferred acceptance al-
 279 gorithm [12] with A making the proposals.

280 ▶ **Theorem 6.** *For a given stable matching instance with one-sided uncertainty, there is a*
 281 *1-competitive algorithm for finding a stable matching in the comparison query model. The*
 282 *algorithm actually finds an A -optimal stable matching.*

283 **Proof.** We utilize the classical deferred acceptance algorithm [12] where A makes the pro-
 284 posals, and we assume the reader's familiarity with it. An unmatched agent $a \in A$ makes a
 285 proposal to their preferred agent $b \in B$ by whom it has never been rejected. If b is unmatched,
 286 then b accepts the proposal and a and b get matched. If b is currently matched to some
 287 $a' \in A$, the algorithm makes a query $prefer(b, a, a')$. If the query result is that b prefers a
 288 to a' , then b accepts a 's proposal and becomes matched to a while a' becomes unmatched.
 289 Otherwise, b rejects the proposal and remains matched to a' . The algorithm terminates if all
 290 agents in A are matched or if every unmatched agent in A has been declined by all agents
 291 in B .

292 We show that this algorithm makes the minimum possible number of comparison queries.

293 We execute the deferred acceptance algorithm with A as the proposers, so it produces an
 294 A -optimal stable matching. Consider an arbitrary agent $b \in B$. Assume that b gets matched
 295 to $a \in A$ in the stable matching. Let $A_b = \{a, a_1, a_2, \dots, a_{k_b}\}$ (for some $0 \leq k_b < n$) be
 296 the set of agents of A that proposed to b during the execution of the algorithm. Note that
 297 $|A_b| = k_b + 1$ and the algorithm has executed k_b queries to b , each for two agents of A_b (the
 298 first agent of A that proposed to b did not require a query). Observe that each of a_1, \dots, a_{k_b}
 299 gets matched with an agent of B that they rank strictly lower than b in the final matching.

300 We claim that no stable matching can be identified without making at least k_b queries
 301 to b . Let M' be an arbitrary stable matching. Note that b rates $M'(b)$ at least as highly as a ,
 302 because M is the worst possible matching for B . Furthermore, for each a_i with $1 \leq i \leq k_b$,
 303 we have that a_i rates b strictly higher than $M'(a_i)$ because M is A -optimal and a_i rates
 304 $M(a_i)$ strictly lower than b . Thus, for none of the pairs (a_i, b) for $1 \leq i \leq k_b$ to be a blocking
 305 pair, the queries of any optimal query set must establish that b rates $M'(b)$ more highly than
 306 every a_i for $1 \leq i \leq k_b$. This can only be achieved with at least k_b queries.

307 The same argument applies to each $b \in B$, so we have that both the optimal number of
 308 queries and the number of queries made by the algorithm are equal to $\sum_{b \in B} k_b$. ◀

309 The proof of Theorem 6 implies that, for the A -optimal stable matching M , the optimal
 310 number of queries to prove that M is stable equals the optimal number of queries to prove
 311 that M is stable and A -optimal. Hence, proving optimality comes in this case for free.

312 3.3 Finding a B -Optimal Stable Matching

313 The problem of finding a B -optimal stable matching is substantially more challenging in
 314 general. For the special case where all A -side preference lists are equivalent, however, there
 315 exists a 1-competitive algorithm:

316 ► **Observation 7.** *If all agents of A have the same preference list, then there is a 1-competitive*
 317 *algorithm that uses $\frac{n^2-n}{2}$ queries for finding a stable matching under one-sided uncertainty.*

318 **Proof.** Let $B = \{b_1, \dots, b_n\}$ be indexed by the preference list of the elements in A , i.e., b_1 is
 319 the first choice of the elements in A and b_i is the i th choice of the elements in A .

320 Let a_1^* be the (initially unknown) first choice of b_1 and, for $i > 1$, let a_i^* denote the first
 321 choice of b_i among the elements of $A \setminus \{a_1^*, \dots, a_{i-1}^*\}$. We can inductively argue that every
 322 stable matching must match b_i to a_i^* for all $i \in [n]$.

323 Thus, every algorithm (including the optimal solution) has to find the first choice of b_i in
 324 $A \setminus \{a_1^*, \dots, a_{i-1}^*\}$ for all $i \in [n]$. For each b_i this requires a minimum of $|A \setminus \{a_1^*, \dots, a_{i-1}^*\}| -$
 325 $1 = n - i$ queries, as each query can exclude at most one element from being the first choice
 326 of b_i . This implies that every algorithm needs at least $\sum_{i=1}^n (i - 1) = \frac{n^2-n}{2}$ queries.

327 The following algorithm matches this lower bound. This is essentially the deferred
 328 acceptance algorithm used in the proof of Theorem 6 but considers the agents of B in a
 329 specific order: (i) Iterate through B in order of increasing indices (ii) For each b_i , determine
 330 the first choice among the not yet matched elements of A and match b_i to that choice.
 331 For each b_i this requires at most $n - i$ queries and, thus, a total of $\sum_{i=1}^n (i - 1) = \frac{n^2-n}{2}$
 332 queries. ◀

333 For arbitrary instances, we first describe an algorithm that is $\mathcal{O}(n)$ -competitive. Comple-
 334 menting this result, we then show that every (randomized) online algorithm has competitive
 335 ratio at least $\Omega(n)$ for finding a B -optimal stable matching. Finally, we show that the offline
 336 problem of determining the optimal number of queries for computing a B -optimal stable
 337 matching is NP-hard and give an $\mathcal{O}(\log n \log \log n)$ -approximation.

338 3.3.1 Algorithm for Computing a B -Optimal Stable Matching

339 We first consider the problem of verifying that a given stable B -optimal matching is indeed
 340 stable and B -optimal. An algorithm for this problem needs to prove that M has no blocking
 341 pair and that no alternating cycle with respect to M is a rotation. For each potential blocking
 342 pair (a, b) that cannot be ruled out because of a 's preferences, such an algorithm has to prove
 343 that it is not a blocking pair using a suitable query to b as discussed in Section 3.1.

344 The more involved part is proving that M is B -optimal. By Lemma 1, M is B -optimal
 345 if and only if it does not expose a rotation. Based on the known A -side preferences, each
 346 edge (a, b) with $M(a) \prec_a b$ could potentially be an r -edge. Thus, each cycle that alternates
 347 between such edges and edges in M could potentially be a rotation. An algorithm that proves
 348 B -optimality has to prove for each such alternating cycle that at least one non-matching
 349 edge (a, b) on that cycle is not an r -edge. By definition, there are two possible ways to prove
 350 that an edge (a, b) with $M(a) \prec_a b$ is not an r -edge:

- 351 1. Query b and find out that b prefers $M(b)$ to a . Then, b cannot be $s_A(a)$ as b does not
 352 prefer a to $M(b)$.
- 353 2. Query one b' with $M(a) \prec_a b' \prec_a b$ and find out that b' prefers a to $M(b')$. Then, b
 354 cannot be the most-preferred element in a 's list that prefers a to her current partner, as
 355 b' has that property and is preferred over b .

356 Corollary 5 gives the optimal number of queries to prove that the matching M is stable,
 357 which is a lower bound on the optimal number of queries necessary to prove that M is
 358 stable and B -optimal. Let $Q(M)$ denote this number. However, there exist instances where
 359 $Q(M) = 0$ and $Q_B(M) > 0$ for the optimal number $Q_B(M)$ of queries to prove that M
 360 is stable *and* B -optimal. Consider an instance where all elements of A have distinct first

361 choices and let M denote the matching that matches all elements of A to their respective
 362 first choice. Then, there is a realization of B -side preference lists such that the matching
 363 M is also B -optimal. For this realization we have $Q(M) = 0$ and $Q_B(M) > 0$. This implies
 364 that the lower bound of Corollary 5 is not strong enough for analyzing algorithms that verify
 365 B -optimality as we cannot prove that such an algorithm makes at most $c \cdot Q(M)$ queries.
 366 We give another lower bound on the optimal number of queries.

367 ► **Lemma 8.** *The optimal number of queries for verifying (and thus also for finding) the*
 368 *B -optimal stable matching is at least $n - 1$ for every instance of the stable matching problem*
 369 *with one-sided uncertainty.*

370 **Proof.** Let M be the B -optimal stable matching for the given instance. For $a \in A$, call a
 371 query an a -query if it reveals for some $b \in B$ with $b \neq M(a)$ whether b prefers a to her current
 372 partner or not. We claim that an optimal algorithm needs to make at least one a -query for
 373 every $a \in A$ with at most a single exception. Assume for a contradiction that the optimal
 374 algorithm makes neither an a -query nor an a' -query for two distinct elements $a, a' \in A$. If a
 375 prefers $M(a)$ over $M(a')$ and a' prefers $M(a')$ over $M(a)$, then it is impossible to exclude
 376 the possibility that $(a, M(a)), (a', M(a'))$ is a rotation exposed in M , because the only way
 377 to prove that $(a, M(a'))$ is not an r -edge is via an a -query, and similarly for $(a', M(a))$. If a
 378 prefers $M(a')$ over $M(a)$, then an a -query to $M(a')$ is necessary to exclude that $(a, M(a'))$ is
 379 a blocking pair. If a' prefers $M(a)$ over $M(a')$, then an a' -query to $M(a)$ is necessary for the
 380 analogous reason. Hence, the claim holds. We note that the $n - 1$ queries whose existence is
 381 asserted by the claim are distinct: A query to some $b \in B$ cannot be an a -query and at the
 382 same time an a' -query for some $a' \neq a$, as the query $prefer(b, a, a')$ cannot yield previously
 383 unknown information about how both a and a' compare to $M(b)$ in b 's preference list. ◀

384 Next, we give an $\mathcal{O}(n)$ -competitive algorithm for finding a B -optimal matching and
 385 analyze it by exploiting the lower bounds on the optimal number of queries of Corollary 5
 386 and Lemma 8. For pseudocode see Algorithm 1.

- 387 1. Find an A -optimal matching using the 1-competitive algorithm for A -optimal matchings.
- 388 2. Search for a rotation by asking, for every $a \in A$, the elements of B that are below $M(a)$
 389 in a 's preference list in order of \prec_a whether they prefer a to their current partner, until
 390 either an r -edge is found or we know that a has no r -edge.
- 391 3. If a rotation R is found, apply that rotation. The agents $a \in A \cap R$ then no longer have
 392 a known r -edge as their previous r -edge is now their matching edge. However, the new
 393 r -edge partner of such an agent must be further down the preference list of a than the
 394 old one. The elements $a \in A \setminus R$ that had an r -edge to an element $b \in B \cap R$ can no
 395 longer be sure that their edge to b is an r -edge since b has a new matching partner $M(b)$,
 396 so b must be asked again whether it prefers the new partner over a when searching for
 397 the new r -edge of a . The algorithm then repeats Step 2 but starts the search for the new
 398 rotation edge of an agent $a \in A$ at either the previous rotation edge (if $a \in A \setminus R$) or at
 399 the direct successor of the new $M(a)$ in \prec_a (if $a \in A \cap R$).
- 400 4. When a state is reached where it is known for every $a \in A$ what its r -edge is (or that
 401 it has no r -edge) but the r -edges do not form a rotation, the algorithm terminates and
 402 outputs M .

403 ► **Theorem 9.** *Given a stable matching instance with one-sided uncertainty, the algorithm is*
 404 *$\mathcal{O}(n)$ -competitive for finding a B -optimal stable matching using comparison queries.*

■ **Algorithm 1** Algorithm to find the B -optimal stable matching using comparison queries.

Input: Instance of the stable matching problem with one-sided uncertainty.

```

1  $M \leftarrow A$ -optimal matching computed using Theorem 6 ;
2  $N \leftarrow \{a \in A \mid M(a) \text{ is last in } \prec_a\}$  ;          /* Elements without  $r$ -edge */
3  $\forall a \in A \setminus N: p(a) \leftarrow$  first element in  $\prec_a$  after  $M(a)$ ;
4  $\forall a \in A \setminus N: r(a) \leftarrow \top$  ;  /* known  $r$ -edges or  $\top$  if  $r$ -edge still unknown */
5 foreach  $a \in A \setminus N$  do
6   repeat
7      $t \leftarrow \text{prefer}(p(a), a, M(p(a)))$  ;
8     if  $t = M(p(a))$  then
9       if  $p(a)$  is the last element of  $\prec_a$  then  $N \leftarrow N \cup \{a\}$  ;
10      else  $p(a) \leftarrow$  direct successor of  $p(a)$  in  $\prec_a$  ;
11    else
12       $r(a) \leftarrow p(a)$  ;          /*  $r(a)$  and  $a$  form an  $r$ -edge */
13    until  $r(a) \neq \top$  or  $a \in N$ ;
14 if  $M$  exposes a rotation  $R$  then
15    $M \leftarrow$  stable matching constructed from  $M$  by applying  $R$ ;
16    $N \leftarrow N \cup \{a \in A \cap R \mid M(a) \text{ is last in } \prec_a\}$ ;
17    $\forall a \in (A \cap R) \setminus N: r(a) \leftarrow \top$  and  $p(a) \leftarrow$  first element in  $\prec_a$  after  $M(a)$ ;
18    $\forall a \in (A \setminus R) \setminus N: p(a) \leftarrow r(a)$  and  $r(a) \leftarrow \top$ ;
19   Jump to Line 5;
20 return  $M$ ;


```

405 **Proof.** Let OPT denote the number of queries made by an optimal algorithm. Since finding
406 any stable matching can never require more queries than finding a B -optimal stable matching,
407 Theorem 6 implies that the algorithm makes at most OPT queries in the first step.

408 We analyze the queries executed *after* the first algorithm step. Call a query *good* if it is
409 the first query involving a specific combination of an agent $a \in A$ and an agent $p(a) \in B$, i.e.,
410 the first query of form $\text{prefer}(p(a), a, M(p(a)))$ for that specific combination of $p(a)$ and a .
411 All other queries are *bad*. By definition of good queries, the algorithm makes at most n^2 such
412 queries since this is the maximum number of good queries that can exist. Since $\text{OPT} \geq n - 1$
413 (Lemma 8), the number of good queries is $\mathcal{O}(n) \cdot \text{OPT}$.

414 Consider the bad queries and a fixed $a \in A$. In the second step of the algorithm, it
415 repeatedly executes queries of the form $\text{prefer}(p(a), a, M(p(a)))$ with $p(a) \in B$ to find out if
416 $(a, p(a))$ is an r -edge, starting with the direct successor $p(a)$ of $M(a)$ in \prec_a . If $(a, p(a))$ is not
417 an r -edge, then the next query partner $p(a)$ for a moves one spot down in the list \prec_a . This
418 is repeated until the r -edge $(a, r(a))$ of a is found or we know that a does not have an r -edge.
419 Here, $r(a)$ refers to the element that forms an r -edge with a .

420 If a does not have an r -edge, there will be no more queries for a again as all $b \in B$ that
421 are lower than $M(a)$ in the preference list of a prefer their current partner $M(b)$ over a and
422 this partner will only improve during the execution of the algorithm. Otherwise, a will be
423 considered again in the second step of the algorithm only if a rotation was found in the third
424 step. If a is part of the rotation, then $r(a) = p(a)$ will be the new matching partner of a and
425 $p(a)$ will be moved one spot down in \prec_a . Only if a is not part of the rotation, $p(a) = r(a)$
426 remains unchanged by definition of the third step. In conclusion, the next query partner
427 $p(a)$ of a moves down one spot in \prec_a after each query for a unless a rotation is found that



■ **Figure 1** Example of the lower bound construction for finding B -optimal matchings. The solid edges represent the A -optimal matching M that needs to be shown to be also B -optimal using queries. The dashed edges represent rotation edges. Each of the agents in $\{a_0, a_1, \dots, a_5\}$ also has a rotation edge to some agent in $\{b_6, b_7, b_8, b_9, b_{10}, b_{11}\}$ that is not shown. An asterisk ($*$) indicates that the remaining agents are placed in arbitrary order in the preference list. A diamond (\diamond) indicates that the adversary decides in response to the queries made by the algorithm which of the agents in $\{a_0, a_1, \dots, a_5\}$ are placed at the front of the preference list and which at the back.

428 does not contain a . This means that a bad query for a can only occur as the first query for a
 429 after a new rotation that does not involve a is found. Thus, each rotation can cause at most
 430 $|A| - 2$ bad queries (at least two members of A must be involved in the rotation). Thus, the
 431 number of bad queries is at most $(n - 2) \cdot n_r$ for the number of applied rotations n_r .

432 For each applied rotation, at least two agents of A get re-matched to agents of B that
 433 are lower down on their preference lists than their previous matching partner. This increases
 434 the lower bound on the optimal number of queries to show stability (cf. Corollary 5) by at
 435 least 2. Thus, Corollary 5 implies $\text{OPT} \geq 2 \cdot n_r$. We can conclude that the number of bad
 436 queries is at most $(n - 2) \cdot n_r \leq \mathcal{O}(n) \cdot \text{OPT}$. ◀

437 3.3.2 Lower Bound for Computing a B -Optimal Matching

438 We give a lower bound of $\Omega(n)$ on the competitive ratio for finding a B -optimal stable matching
 439 with comparison queries. This implies that the result of Theorem 9 is, asymptotically, best-
 440 possible. Further, the lower bound also holds for verifying that a given matching is B -optimal.

441 ► **Theorem 10.** *In the comparison query model, every deterministic or randomized online*
 442 *algorithm for finding a B -optimal stable matching in a stable matching instance with one-sided*
 443 *uncertainty has competitive ratio $\Omega(n)$.*

444 **Proof.** We first show the statement for deterministic algorithms. Consider the following
 445 instance (cf. Fig. 1) with two sets of agents $A = \{a_0, \dots, a_{n-1}\}$ and $B = \{b_0, \dots, b_{n-1}\}$, and
 446 assume $n/2$ to be even. If this is not the case, then the constant factor in the lower bound
 447 will be slightly worse.

448 We partition A into three subsets $A_1 = \{a_0, \dots, a_{\frac{n}{2}-1}\}$, $A_2 = \{a_{\frac{n}{2}}, \dots, a_{n-2}\}$ and
 449 $A_3 = \{a_{n-1}\}$. and B into two subsets, $B_1 = \{b_0, \dots, b_{\frac{n}{2}-1}\}$ and $B_2 = B \setminus B_1$.

450 In the following, we first define the known A-side preferences and the adversarial strategy.
 451 Then we give bounds on the optimal number of queries and the number of queries made by
 452 any deterministic algorithm.

453 **A-side preferences.** Consider the following preference lists for A . For an agent $a_i \in A_1$,
 454 the preference list consists of three parts, $P(a_i) = P_1(a_i)P_2(a_i)P_3(a_i)$. The first part of the
 455 list is the corresponding i^{th} agent of B , i.e., $P_1(a_i) = (b_i)$. The second part consists of the $\frac{n}{2}$
 456 agents of set B_2 in increasing order, i.e., $P_2(a_i) = (b_{\frac{n}{2}}, b_{\frac{n}{2}+1}, \dots, b_{n-2}, b_{n-1})$. The last part
 457 $P_3(a_i)$ consists of the agents of $B_1 \setminus \{b_i\}$ in an arbitrary order.

458 For an agent $a_i \in A_2$, the preference list starts with agent b_i , followed by the last agent
 459 b_{n-1} and finally an arbitrary order of the remaining agents in group B . For the single agent
 460 a_{n-1} in set A_3 , the preference list starts with agent b_{n-1} followed by an arbitrary order of
 461 the remaining agents of set B .

462 **Adversarial strategy.** The preference lists of the agents in set B are unknown. The
 463 instance has the A -optimal matching $M = \{(a_i, b_i) \mid 0 \leq i < n\}$. The adversary will ensure
 464 that this matching is also B -optimal. Since each a_i is matched with its top choice, proving
 465 stability does not require any queries. To prove B -optimality of M , the executed queries
 466 must prove that there is no rotation.

467 The adversary will ensure that M can be shown to be a B -optimal matching with $\mathcal{O}(n)$
 468 queries while any deterministic algorithm is forced to make $\Omega(n^2)$ queries.

469 To achieve this, the adversary sets the preferences of the agents of B_1 independent of the
 470 algorithm's actions as follows. For each odd $i \in \{1, 3, 5, \dots, (n/2) - 1\}$, we let the preference
 471 list of b_i start with a_{i-1} followed by a_i and finally all remaining agents in A in an arbitrary
 472 order. The preference list of b_{i-1} starts with a_i followed by a_{i-1} and then the remaining
 473 agents in A in an arbitrary order. Using these preferences, the sequences $(a_{i-1}, b_{i-1}), (a_i, b_i)$
 474 are potential rotations. To prove that such a sequence is not a rotation, an algorithm has to
 475 show that either (a_{i-1}, b_i) or (a_i, b_{i-1}) is not an r -edge. The only way of showing this is to
 476 prove that either a_{i-1} or a_i instead has an r -edge to some agent of B_2 .

477 Consider any deterministic algorithm. The adversary selects the preferences of the agents
 478 in B_2 in such a way that the following properties hold:

- 479 **(P1)** Agent a_{n-1} has no r -edge. Note that, by the definition of the preferences of B_1 above,
 480 a_{n-1} already cannot have a rotation edge to an agent of B_1 .
 481 **(P2)** Each agent in A_2 has an r -edge to b_{n-1} .
 482 **(P3)** Each agent a_i in A_1 has an r -edge to some agent $b_{t(i)}$ of B_2 . The choice of that agent
 483 $b_{t(i)}$ depends on the queries made by the algorithm.

484 The properties (P1)–(P3) ensure that there is no rotation, as the alternating path starting
 485 at any $a \in A \setminus \{a_{n-1}\}$ with the r -edge of that agent ends at a_{n-1} , which has no r -edge.

486 Let $t(i)$ denote the index of the agent $b_{t(i)}$ of B_2 to which $a_i \in A_1$ has an r -edge. This
 487 index is determined by the adversary in response to the queries made by the algorithm.
 488 Concretely, the adversary lets $t(i)$ be the index of the *last agent* $b_j \in B_2$ for which the
 489 algorithm makes a query of the form $\text{prefer}(b_j, a_i, *)$, where we use $\text{prefer}(b_j, a_i, *)$ as a short-
 490 hand to refer to queries $\text{prefer}(b_j, a_i, a_{i'})$ or $\text{prefer}(b_j, a_{i'}, a_i)$ for some i' . If the algorithm
 491 doesn't make queries of this form for all $b_j \in B_2$, then let $t(i)$ be an arbitrary j such that the
 492 algorithm does not make a query of this form for $b_j \in B_2$. The adversary sets the preferences
 493 of the agents in B_2 in such a way that $b_{t(i)}$ prefers a_i to her partner $a_{t(i)}$ in the A-optimal
 494 matching M while all other $b_j \in B_2$ prefer their partner in the A-optimal matching a_j to a_i .
 495 For example, if the algorithm was to make queries $\text{prefer}(b_j, a_i, a_j)$ for all $b_j \in B_2$ (which it
 496 might do in order to check whether a_i has an r -edge to one of these agents), the adversary
 497 would answer false to the first $\frac{n}{2} - 1$ such queries and true to the final one.

498 To achieve the properties (P1)–(P3), the adversary sets the preferences of each agent b_j
499 of B_2 as follows:

- 500 ■ b_j prefers $a_i \in A_1$ to a_j if and only if $j = t(i)$.
- 501 ■ If $j \neq n - 1$, b_j prefers a_j to $a_{j'}$ for all $j' \neq j$, $a_{j'} \in A_2$.
- 502 ■ If $j = n - 1$, b_j prefers $a_{j'}$ to a_j for all $j' \neq j$, $a_{j'} \in A_2$.

503 This can be done by letting the preference list of b_j contain first the agents $a_i \in A_1$ with
504 $j = t(i)$ in some order, then the agents of A_2 in some order (only ensuring for b_j that a_j
505 comes first among the agents of A_2 if $j \neq n - 1$ and that a_j comes last among the agents of
506 A_2 if $j = n - 1$), and finally the agents $a_i \in A_1$ with $j \neq t(i)$ in some order.

507 **Upper bound on the optimal query cost.** An optimal solution for the instance can
508 prove that matching M is B -optimal by verifying that the properties (P1)–(P3) indeed hold
509 by using at most $n - 1 + \frac{n}{2} - 1 + \frac{n}{2} = 2n - 2$ queries as follows:

- 510 ■ The $n - 1$ queries $\text{prefer}(b_i, a_{n-1}, a_i) = \text{false}$ for $i \leq n - 2$ show that a_{n-1} has no r -edge.
- 511 ■ Each of the $\frac{n}{2} - 1$ queries $\text{prefer}(b_{n-1}, a_{\frac{n}{2}+i}, a_{n-1}) = \text{true}$ for $0 \leq i \leq \frac{n}{2} - 2$ shows that
512 $a_{\frac{n}{2}+i}$ has an r -edge to b_{n-1} . This is because each agent of A_2 has b_{n-1} in its preference
513 list directly after its current matching partner. So if b_{n-1} prefers an agent of A_2 over its
514 current partner a_{n-1} , then this directly gives us an r -edge.
- 515 ■ Each of the $\frac{n}{2}$ queries $\text{prefer}(b_{t(i)}, a_i, a_{t(i)}) = \text{true}$ for $0 \leq i \leq \frac{n}{2} - 1$ shows that a_i has a
516 rotation edge to some agent in B_2 . Based on the result of such a query, a_i must have an
517 r -edge to either $b_{t(i)}$ or to some other agent of B_2 that is higher up in a_i 's preference list.

518 **Lower bound on the algorithm's query cost.** We provide to the algorithm the
519 information that a_{n-1} has no r -edge, that each agent of A_2 has an r -edge to b_{n-1} , and we
520 reveal the full preference lists of all agents in B_1 . Clearly, this extra information can only
521 reduce the number of queries a deterministic algorithm may need as it could simply ignore
522 the information.

523 For each agent $a_i \in A_1$, the algorithm will either make queries of the form $\text{prefer}(b_j, a_i, *)$
524 for all $b_j \in B_2$ or not. Call a_i *resolved* in the former case and *unresolved* otherwise. For any
525 resolved agent, the algorithm may have determined that it has an r -edge to an agent of B_2
526 and hence cannot be part of a rotation. For the unresolved agents, the algorithm cannot
527 know whether they have an r -edge to an agent in B_2 .

528 As argued above, for each odd $i \in \{1, 3, 5, \dots, \frac{n}{2} - 1\}$, the algorithm has to resolve either a_i
529 or a_{i-1} to prove that $(a_i, b_i), (a_{i-1}, b_{i-1})$ is not a rotation. Thus, it must resolve at least $n/4$
530 agents. For each resolved agent, the algorithm has made queries of the form $\text{prefer}(b_j, a_i, *)$
531 for each $b_j \in B_2$. This totals to at least $\frac{n}{4} \cdot \frac{n}{2} \cdot \frac{1}{2} = \frac{n^2}{16} \in \Omega(n^2)$ queries. Note that we
532 divide $\frac{n}{4} \cdot \frac{n}{2}$ by two as a single query $\text{prefer}(b_j, a_i, a_{i'})$ is of the form $\text{prefer}(b_j, a_i, *)$ and also
533 $\text{prefer}(b_j, a_{i'}, *)$.

534 Finally, we show how to extend the lower bound to work for randomized algorithms.

By Yao's principle [4, 31] we can prove the theorem by giving a randomized instance \mathcal{R}
and showing that

$$\mathbb{E}_{R \sim \mathcal{R}} \left[\frac{\text{ALG}(R)}{\text{OPT}(R)} \right] \in \Omega(n)$$

535 holds for every deterministic algorithm ALG , where $\text{ALG}(R)$ and $\text{OPT}(R)$ denote the number
536 of queries executed by the algorithm and an optimal solution, respectively, for the realization
537 R of the randomized instance \mathcal{R} .

538 To define the randomized instance \mathcal{R} , we take the instance of the deterministic lower
539 bound and introduce randomization into the uncertain preference lists. We define the
540 preference lists of A and B_1 without randomization in the same way as before. Similarly, we

541 leave the sub-list defined for the agents of A_2 in the preference list of an agent of B_2 as it is
 542 and only randomize the positions of the agents of A_1 in the preference lists of B_2 .

543 To this end, consider an odd $i \in \{1, 3, \dots, (n/2) - 1\}$. The randomized part of the
 544 instance uniformly at random picks a tuple (a_k, b_j) with $k \in \{i - 1, i\}$ and $b_j \in B_2$. For
 545 this selected tuple, we set the preferences such that b_j prefers a_k over its current matching
 546 partner a_j . For all other tuples $(a_{k'}, b_{j'})$ with $k' \in \{i - 1, i\}$, $b_{j'} \in B_2$ and either $k \neq k'$ or
 547 $j \neq j'$, we set the preference of $b_{j'}$ such that it prefers its current partner over $a_{k'}$.

548 This can be achieved by letting the preference list of b_j contain first the agents $a_k \in A_1$
 549 such that (a_k, b_j) was selected by the randomized procedure above, then the agents of A_2 in
 550 some non-randomized order (only ensuring for b_j that a_j comes first among the agents of A_2
 551 if $j \neq n - 1$ and that a_j comes last among the agents of A_2 if $j = n - 1$), and finally the
 552 agents $a_{k'} \in A_1$ such that the tuple $(a_{k'}, b_j)$ was *not* selected by the randomized procedure.
 553 Ties can be broken according to some arbitrary but fixed order.

By defining the preferences in this way, every realized instance still satisfies the properties
 (P1) and (P2) as defined in the proof of the deterministic lower bound. While the preferences
 do not satisfy property (P3), they satisfy for each odd $i \in \{1, 3, \dots, (n/2) - 1\}$ that either a_i
 or a_{i-1} has a rotation edge to some agent of B_2 . This still implies that, for every realized
 instance, the matching $M = \{(a_j, b_j) \mid j \in \{0, \dots, n - 1\}\}$ is B-optimal. Slightly adjusting
 the strategy of the deterministic proof, one can show that $\text{OPT} \leq 2n - 2$ still holds for each
 such realization. This implies

$$\mathbb{E}_{R \sim \mathcal{R}} \left[\frac{\text{ALG}(R)}{\text{OPT}(R)} \right] \geq \mathbb{E}_{R \sim \mathcal{R}} \left[\frac{\text{ALG}(R)}{2n - 2} \right] = \frac{\mathbb{E}_{R \sim \mathcal{R}}[\text{ALG}(R)]}{2n - 2}$$

554 for every deterministic algorithm ALG. So it suffices to show $\mathbb{E}_{R \sim \mathcal{R}}[\text{ALG}(R)] \in \Omega(n^2)$ to
 555 prove the theorem.

556 To that end, consider an arbitrary deterministic algorithm. As argued in the deterministic
 557 lower bound proof, the algorithm has to, for each odd $i \in \{1, 3, \dots, (n/2) - 1\}$, either prove
 558 that (a_i, b_{i-1}) or (a_{i-1}, b_i) is not an r -edge. By definition of the instance, this requires at least
 559 one query of the form *prefer* $(b_j, a_k, *)$ (as defined in the deterministic lower bound proof) for
 560 the tuple (b_j, a_k) with $b_j \in B_2$ and $k \in \{i, i - 1\}$ that was drawn by the randomized procedure
 561 above for index i . The algorithm will have to execute queries of the form *prefer* $(b_{j'}, a_{k'}, *)$
 562 with $b_{j'} \in B_2$ and $k' \in \{i, i - 1\}$ until it hits a query with $j' = j$ and $k' = k$. We call such
 563 a query *successful* if $j' = j$ and $k' = k$ and *unsuccessful* otherwise. In the same way, we
 564 call the selected tuples *successful* and all other tuples *unsuccessful*. Note that the algorithm
 565 might need further queries to prove that either (a_i, b_{i-1}) or (a_{i-1}, b_i) is not a rotation edge,
 566 but executing at least one successful query is a necessary condition.

567 Consider a fixed odd $i \in \{1, 3, \dots, (n/2) - 1\}$. We bound the expected number of queries
 568 of the form *prefer* $(b_{j'}, a_{k'}, *)$ with $b_{j'} \in B_2$ and $k' \in \{i, i - 1\}$ that the algorithm needs until
 569 one of them is successful. Let Y_i be a random variable denoting the number of queries of
 570 that form the algorithm executes. Note that the algorithm might execute different queries
 571 in-between the queries of that form, but the random variable Y_i only counts the queries of
 572 that form for the fixed i and ignores different queries that are executed in-between them.
 573 To further characterize Y_i , let $Z_{i,\ell}$ with $\ell \geq 1$ be an indicator random variable denoting
 574 whether the first ℓ queries of that form are *not* successful. Then, $Y_i = 1 + \sum_{\ell \geq 1} Z_{i,\ell}$ and
 575 $\mathbb{E}_{R \sim \mathcal{R}}[Y_i] = 1 + \sum_{\ell \geq 1} \mathbb{E}_{R \sim \mathcal{R}}[Z_{i,\ell}]$.

576 We first observe that queries that do not involve a_i and a_{i-1} do not give any information
 577 on which tuples can be successful for i . Furthermore, queries that involve a_i (or a_{i-1}) and
 578 some $b_j \in B_2$ do not admit any information on whether some tuple $(a_k, b_{j'})$ (or some tuple
 579 $(a_{i-1}, b_{j'})$) with $j' \neq j$ or $k = i - 1$ (or $k = i$) is successful or not. Thus, at any point during

580 the execution of an algorithm, all tuples (a_k, b_j) for which the algorithm did not yet execute
 581 a query of form $prefer(b_j, a_k, *)$ are equally likely to be successful (unless the algorithm
 582 already found the successful tuple).

Consider the expected value $\mathbb{E}_{R \sim \mathcal{R}}[Z_{i,\ell}] = \Pr[Z_{i,\ell} = 1]$. For $\ell = 1$, we have $\Pr[Z_{i,\ell} = 1] = \frac{n-2}{n}$ since there are n tuples $(a_{k'}, b_{j'})$ with $k' \in \{i, i-1\}$, among those only one successful tuple is drawn uniformly at random, and a query can cover at most two such tuples at the same time if it is of form $prefer(b_{j'}, a_i, a_{i-1})$. For $\ell = 2$, we have $\Pr[Z_{i,\ell} = 1] \geq \frac{n-2}{n} \cdot \frac{n-4}{n-2}$ because given that the first query is not successful there are still $n-2$ tuples that could still be successful, only one uniformly at random selected tuple is actually successful, and the second query can cover at most two of the potentially successful tuples. Continuing this argumentation, we get

$$\mathbb{E}_{R \sim \mathcal{R}}[Z_{i,\ell}] = \Pr[Z_{i,\ell} = 1] = \prod_{\ell'=1}^{\ell} \frac{n-2 \cdot \ell'}{n-2 \cdot (\ell'-1)} = 1 - \frac{2\ell}{n}$$

583 for each $1 \leq \ell \leq n/2$. This directly implies

$$\begin{aligned} \mathbb{E}_{R \sim \mathcal{R}}[Y_i] &= 1 + \sum_{\ell \geq 1} \mathbb{E}_{R \sim \mathcal{R}}[Z_{i,\ell}] \geq 1 + \sum_{\ell=1}^{n/2} \mathbb{E}_{R \sim \mathcal{R}}[Z_{i,\ell}] \\ &\geq \sum_{\ell=1}^{n/2} \left(1 - \frac{2\ell}{n}\right) = \frac{n-2}{4}. \end{aligned}$$

586 The number of queries the algorithm executes on a realization R is at least $\text{ALG}(R) \geq$
 587 $\sum_{i \in \{1,3,\dots,(n/2)-1\}} Y_i$, which implies

$$\begin{aligned} \mathbb{E}_{R \sim \mathcal{R}}[\text{ALG}(R)] &\geq \sum_{i \in \{1,3,\dots,(n/2)-1\}} \mathbb{E}_{R \sim \mathcal{R}}[Y_i] \geq \frac{n}{4} \cdot \frac{n-2}{4} \\ &= \frac{n^2 - 2n}{16} \in \Omega(n^2). \end{aligned}$$

590

◀

591 3.3.3 Offline Results for Computing B -Optimal Stable Matchings

592 We show NP-hardness for the offline problem of verifying a given matching M to be stable
 593 and B -optimal. Recall that in the offline problem we assume full knowledge of the B -side
 594 preferences but still want to compute a query set of minimum size that a third party without
 595 knowledge of the B -side preferences could use to verify the B -optimality of M .

596 ▶ **Theorem 11.** *The offline problem of computing an optimal set of comparison queries
 597 for finding (or verifying) the B -optimal stable matching in a stable matching instance with
 598 one-sided uncertainty is NP-hard.*

599 **Proof.** We give a reduction from the NP-hard *Minimum Feedback Arc Set (FAS)* problem.
 600 Given a directed graph $G = (V, E)$, a feedback arc set is a subset of edges $E' \subseteq E$ which, if
 601 removed from G , leaves the remaining graph acyclic. The FAS problem is to decide for a
 602 given directed graph and some $k \in \mathbb{Z}_+$, whether there is a feedback arc set E' with $|E'| \leq k$.

603 Given an instance of FAS with $G = (V, E)$ and some k , we construct a stable matching
 604 instance with one-sided uncertainty as follows. For each node v of G , introduce an agent
 605 v in A and an agent v' in B . Let $N^+(v)$ denote the set of out-neighbors of v in G , and

606 $d^+(v) = |N^+(v)|$. The preference list of v is such that it ends with v' followed by all u' for
 607 $u \in N^+(v)$. All other w' in B come before v' . Thus, the elements of $B \setminus \{u' \mid u \in N^+(v)\}$ are
 608 the most preferred partners of v , followed by v' and finally the elements of $\{u' \mid u \in N^+(v)\}$.
 609 Let M be the matching that matches v to v' , for all v . The preference lists of $b \in B$ are
 610 such that M is the B -optimal stable matching: Every v' has v as top preference, and the
 611 remaining agents of A follow in arbitrary order. By selecting the matching M this way, we
 612 have that, for every $v \in A$, all edges to elements of $\{u' \mid u \in N^+(v)\}$ are potential r -edges.
 613 To prove that such an edge (v, u') is not an r -edge, an algorithm has to compare u and v
 614 from the perspective of u' to prove that u' prefers $M(u') = u$ over v .

615 The number of queries $Q(M)$ needed to verify the stability of M is determined by M
 616 and is polynomial-time computable by using Theorem 4. To prove B -optimality of M , we
 617 need to show that there is no rotation (Lemma 1). Indeed, there is a query strategy with k
 618 queries for verifying that there is no rotation if and only if there is a feedback arc set in G of
 619 size k . To see this, observe that every directed cycle in G corresponds to a potential rotation
 620 in the matching instance, and every query that excludes one of the edges of the potential
 621 rotation from being an r -edge corresponds to the removal of the corresponding arc in G .

622 Note that, for the constructed instance, all queries to verify the stability of M obtain
 623 information of the form $M(b) \prec_b a$ for $a \in A$ and $b \in B$ with $b \prec_a M(a)$. On the other
 624 hand, all queries that help to verify the absence of a rotation obtain information of the form
 625 $M(b) \prec_b a$ for $a \in A$ and $b \in B$ with $M(a) \prec_a b$. As these are disjoint query sets, we can
 626 conclude that there is a query strategy that proves M to be stable and B -optimal with at
 627 most $Q(M) + k$ queries if and only if there is a feedback arc set in G of size at most k . ◀

628 We also prove the following approximation for the offline problem by exploiting an
 629 $\mathcal{O}(\log n \log \log n)$ -approximation for weighted feedback arc set by Even et al. [11].

630 ► **Theorem 12.** *The offline problem of computing an optimal set of comparison queries for*
 631 *finding the B -optimal stable matching in a stable matching instance with one-sided uncertainty*
 632 *can be approximated within ratio $\mathcal{O}(\log n \log \log n)$.*

633 **Proof.** Let M be the B -optimal matching. We give an algorithm that verifies M to be
 634 stable and B -optimal by executing at most $\mathcal{O}(\log n \log \log n) \cdot \text{OPT}$ queries, where OPT is
 635 the optimal number of queries for the same instance. First, the algorithm proves that M is
 636 stable using Theorem 4. This leads to at most OPT queries.

637 After that, the algorithm has to prove B -optimality. First, for every $a \in A$ that has an
 638 r -edge to an agent $r(a) \in B$, the algorithm queries $\text{prefer}(r(a), a, M(r(a)))$. Since $(a, r(a))$
 639 is an r -edge, this query must return that $r(a)$ prefers a over $M(r(a))$. This leads to at most
 640 $n \leq \text{OPT} + 1$ queries ($n \leq \text{OPT} + 1$ holds by Lemma 8). Note that, for an $a \in A$ with an
 641 r -edge, the query $\text{prefer}(r(a), a, M(r(a)))$ proves that a has an r -edge but is not necessarily
 642 sufficient to prove that $(a, r(a))$ is indeed the r -edge of a . If there is an agent $b \in B$ with
 643 $M(a) \prec_a b \prec_a r(a)$ for which we have not yet verified whether b prefers a over $M(b)$, then
 644 (a, b) could also still be the r -edge of a . We call such pairs (a, b) *potential r -edges* and let P
 645 denote the set of these edges.

646 It remains to consider the graph G defined by the matching edges, the r -edges R , and all
 647 potential r -edges P . If G has no cycle alternating between edges in M and edges in $P \cup R$,
 648 then we have shown that M does not expose a rotation and, thus, is B -optimal. Otherwise,
 649 the algorithm has to execute queries $\text{prefer}(b, a, M(b))$ for edges $(a, b) \in P$ to prove that they
 650 are not actually r -edges until it becomes clear that M has no rotation.

651 To select the edges $(a, b) \in P$ for which the algorithm executes such queries, we exploit
 652 the $\mathcal{O}(\log n \log \log n)$ -approximation for weighted feedback arc set by Even et al. [11]. To

653 this end, we create an instance of the weighted feedback arc set problem by considering the
 654 vertices $A \cup B$, adding the edges $M \cup R$ with weight ∞ each and adding the edges P with
 655 weight 1 each. We orient all edges in M from the B -side vertex to the A -side vertex and all
 656 edges in $R \cup P$ from the A -side vertex to the B -side vertex. The orientation ensures that
 657 all cycles in the graph alternate between M -edges and $R \cup P$ -edges. Since the matching
 658 M is B -optimal by assumption, there cannot be an alternating cycle using only edges in
 659 $M \cup R$, so there must be a feedback arc set that only uses edges in P . The choice of the edge
 660 weights ensures that every approximation algorithm for weighted feedback arc set finds such a
 661 solution. We use the $\mathcal{O}(\log n \log \log n)$ -approximation to find such a feedback arc set $F \subseteq P$.
 662 Since removing F from the instance yields an acyclic graph, querying $\text{prefer}(b, a, M(b))$
 663 for each $(a, b) \in F$ proves that M does not expose a rotation. As the minimum weight
 664 feedback arc set is the cheapest way to prove that M does not have a rotation, we have
 665 $|F| \leq \mathcal{O}(\log n \log \log n) \cdot \text{OPT}$, which implies the theorem. ◀

666 3.4 Verifying a Stable Matching with Two-Sided Uncertainty

667 We observe that the lower bound on the optimal number of queries in Corollary 3 can also be
 668 used for verifying a stable matching in a stable matching instance with uncertain preferences
 669 on both sides.

670 ▶ **Theorem 13.** *In the comparison query model, there is a 2-competitive algorithm for*
 671 *verifying that a given matching M in a stable matching instance with uncertain preferences*
 672 *on both sides is stable.*

673 **Proof.** To verify that a given matching M in a graph G is stable, we have to prove for each
 674 $(a, b) \notin M$ that (a, b) is not a blocking pair. That is, we have to prove that $M(a) \prec_a b$ or
 675 that $M(b) \prec_b a$. Note that $M(a) \prec_a b$ (or $M(b) \prec_b a$) can be verified by directly querying
 676 $\text{prefer}(a, b, M(a))$ (or $\text{prefer}(b, a, M(b))$) or indirectly via transitivity.

677 For every pair $(a, b) \notin M$, let the algorithm query $\text{prefer}(b, a, M(b))$ first and, if the
 678 answer is that $a \prec_b M(b)$, query also $\text{prefer}(a, b, M(a))$. If the answer to the latter query
 679 is that $b \prec_a M(a)$, the pair (a, b) is a blocking pair, and the algorithm outputs that M is
 680 not a stable matching. If the algorithm finds for every pair $(a, b) \notin M$ that $M(b) \prec_b a$ or
 681 $M(a) \prec_a b$, the algorithm outputs that M is a stable matching.

682 The algorithm makes at most $2(n^2 - n)$ queries, as it makes at most 2 queries for each of
 683 the $n^2 - n$ pairs $(a, b) \notin M$.

684 We now show that the optimal number of queries is at least $n^2 - n$. For each pair
 685 $(a, b) \notin M$, the optimum needs to prove $M(a) \prec_a b$ or $M(b) \prec_b a$. This means it must relate
 686 b to $M(a)$ for a , or it must relate a to $M(b)$ for b . Either way, this produces a relationship
 687 pair. As no two different pairs $(a, b) \notin M$ can produce the same relationship pair, the total
 688 number of relationship pairs is at least $n^2 - n$. By Corollary 3, this implies that the optimum
 689 makes at least $n^2 - n$ queries. As the algorithm makes at most $2(n^2 - n)$ queries, it is
 690 2-competitive. ◀

691 We can show that no deterministic algorithm can do better.

692 ▶ **Theorem 14.** *In the comparison query model, no deterministic algorithm can be better*
 693 *than 2-competitive for the problem of verifying that a given matching M in a stable matching*
 694 *instance with uncertain preferences on both sides is stable.*

695 **Proof.** Let $A = \{a_1, a_2\}$, $B = \{b_1, b_2\}$ and $M = \{(a_1, b_1), (a_2, b_2)\}$. Any algorithm has to
 696 prove that (a_1, b_2) and (a_2, b_1) are not blocking pairs. Thus, any algorithm has to either

697 prove $a_2 \prec_{b_2} a_1$ or $b_1 \prec_{a_1} b_2$ (to verify that (a_1, b_2) is not a blocking pair) and $a_1 \prec_{b_1} a_2$ or
 698 $b_2 \prec_{a_2} b_1$ (to verify that (a_2, b_1) is not a blocking pair).

699 Since the subproblems of proving that (a_1, b_2) and (a_2, b_1) are not blocking pairs are
 700 independent of each other, we can w.l.o.g. assume that the algorithm starts by proving that
 701 (a_1, b_2) is not a blocking pair. If the algorithm starts by querying $prefer(b_2, a_1, a_2)$, then
 702 the adversary reveals $a_2 \succ_{b_2} a_1$, which forces the algorithm to also query $prefer(a_1, b_1, b_2)$.
 703 We let this query reveal $b_1 \prec_{a_1} b_2$. The optimal solution only queries $prefer(a_1, b_1, b_2)$. If
 704 the algorithm starts by querying $prefer(a_1, b_1, b_2)$, we can argue symmetrically. Thus, the
 705 algorithm executes twice as many queries as the optimal solution to prove that (a_1, b_2) is not
 706 a blocking pair.

707 We can argue analogously to show that the algorithm also executes twice as many queries
 708 as the optimal solution to prove that (a_2, b_1) is not a blocking pair, which implies the
 709 result. ◀

710 4 Stable Matching with Interview Queries

711 In this section, we consider the interview query model. Most of our results and proofs are
 712 quite similar to their counterparts for comparison queries. This might be surprising as
 713 interview and comparison queries are, in a sense, incomparable: While interview queries allow
 714 us to more efficiently determine full preference lists, a comparison between two agents can
 715 be done more efficiently via a single comparison query. As we show the same (asymptotic)
 716 bounds on the competitive ratio, the latter seems to be the deciding factor.

717 4.1 Verifying and Finding a Stable Matching with Interview Queries

718 A 1-competitive algorithm for finding a stable matching and verifying a given stable matching
 719 with interview queries is implied by the results and arguments from [29] for a more general
 720 uncertainty setting and can be derived as follows.

721 Consider a given instance of stable matching with one-sided uncertainty and a given
 722 stable matching M . To verify that M is indeed stable, we have to consider all potential
 723 blocking pairs, i.e., all pairs (a, b) with $b \prec_a M(a)$. For such a pair, we have to verify that
 724 $M(b) \prec_b a$ holds to prove that (a, b) is not a blocking pair. The only way of comparing
 725 $M(b)$ and a from b 's perspective is to execute the interviews $intq(b, a)$ and $intq(b, M(b))$. For
 726 a fixed $b \in B$, this implies that the minimum number of interviews involving b necessary
 727 to prove that M is stable is $Q_b(M) = 0$ if no element of $a \in A \setminus \{M(b)\}$ prefers b over its
 728 current partner $M(a)$ and $Q_b(M) = 1 + |\{a \in A \mid b \prec_a M(a)\}|$ otherwise. We can observe
 729 the following.

730 ▶ **Observation 15.** *Consider a given instance of stable matching with one-sided uncertainty*
 731 *and a given stable matching M . The minimum number of interview queries necessary to*
 732 *verify that M is indeed stable is $Q(M) = \sum_{b \in B} Q_b(M)$ with $Q_b(M) = 0$ if no element of $a \in$*
 733 *$A \setminus \{M(b)\}$ prefers b over its current partner $M(a)$ and $Q_b(M) = 1 + |\{a \in A \mid b \prec_a M(a)\}|$*
 734 *otherwise*

735 Consider the following algorithm: For each $b \in B$ with $Q_b(M) > 0$, query $intq(b, M(b))$
 736 and $intq(b, a)$ for each $a \in A$ with $b \prec_a M(a)$. This algorithm clearly verifies the
 737 stability of M and executes exactly $\sum_{b \in B} Q_b(M)$ interview queries. Thus, the observation
 738 implies the following lemma.

739 ► **Lemma 16.** *For a given stable matching instance with one-sided uncertainty and a stable*
 740 *matching M , there is a 1-competitive algorithm for verifying that M is stable in the interview*
 741 *query model.*

742 Similar to the comparison query model, we can observe that the A -optimal matching
 743 M^* minimizes the query cost for verifying stability $Q(M) = \sum_{b \in B} Q_b(M)$ over all stable
 744 matchings M .

745 To find such an A -optimal matching, we can again just consider the deferred acceptance
 746 algorithm where A makes the proposals. Whenever an agent $a \in A$ makes a proposal to an
 747 element $b \in B$ that is currently matched to some $a' \in A$, the algorithm queries $intq(b, a)$
 748 and $intq(b, a')$. Each interview query is only executed if it has not yet been queried during
 749 the previous execution of the algorithm. If the query result is that b prefers a to a' , then b
 750 accepts a 's proposal and becomes matched to a while a' becomes unmatched. Otherwise, b
 751 rejects the proposal and remains matched to a' .

752 It is not hard to see that this algorithm executes exactly $Q(M^*) = \sum_{b \in B} Q_b(M^*)$
 753 interview queries. This implies that the deferred acceptance algorithm is 1-competitive for
 754 finding the A -optimal matching or any stable matching with interview queries.

755 4.2 Finding a B -Optimal Stable Matching with Interview Queries

756 For finding a B -optimal stable matching with interview queries, it is not hard to see that
 757 the lower bound of Theorem 10 for comparison queries nearly directly translates. We briefly
 758 sketch how to adjust that lower bound for interview queries to achieve the following theorem.

759 ► **Theorem 17.** *In the interview query model, every deterministic or randomized online*
 760 *algorithm for finding a B -optimal stable matching in a stable matching instance with one-sided*
 761 *uncertainty has competitive ratio $\Omega(n)$.*

762 **Proof sketch.** We separately sketch the deterministic and randomized lower bound.

763 **Deterministic lower bound.** Consider the same instance as in the deterministic lower
 764 bound of Theorem 10. Recall that each $a_i \in A_1$ has an r -edge to some $t(i) \in B_2$ that is
 765 selected by the adversary depending on the queries executed by the deterministic algorithm.
 766 Fix an $a_i \in A_1$. For interview queries we select $t(i)$ as the last element $b \in B_2$ for which the
 767 algorithm executes a query $intq(b, a_i)$. If the algorithm does not execute such a query for
 768 every element of B_2 , then select an arbitrary agent b of B_2 for which the query $intq(b, a_i)$
 769 has *not* been executed by the algorithm.

770 For a fixed $a_i \in A_1$, this forces any deterministic algorithm to execute at least $|B_2|$ queries
 771 $intq(b, a_i)$ with $b \in B_2$ to prove that a_i has an r -edge to some $b \in B_2$. As argued in the proof
 772 for comparison queries, any deterministic algorithm has to do this for at least $\frac{n}{4}$ members of
 773 A_1 . This leads to a total of at least $\frac{n}{2} \cdot \frac{n}{4} \in \Omega(n^2)$ interview queries for every deterministic
 774 algorithm.

775 The optimal solution on the other hand needs at most $\mathcal{O}(n)$ queries by for example
 776 executing the query strategy described in the proof for comparison queries while simulating
 777 each comparison query with at most two interviews.

778 **Randomized lower bound.** For the randomized lower bound, we again use Yao's
 779 principle, consider the same randomized instance as in the proof for comparison queries and
 780 prove that every deterministic algorithm need $\Omega(n^2)$ queries in expectation.

781 To this end, consider an arbitrary deterministic algorithm. Recall that, for each odd
 782 $i \in \{1, 3, \dots, (n/2) - 1\}$, the algorithm either has to prove that (a_i, b_{i-1}) or (a_{i-1}, b_i) is
 783 not an r -edge. By definition of the instance, this requires at least one query of the form

784 $\text{intq}(b_j, a_k)$ for the tuple (b_j, a_k) with $b_j \in B_2$ and $k \in \{i, i-1\}$ that was drawn by the
 785 randomized procedure as defined in the comparison query proof for index i . The algorithm
 786 will have to execute queries of the form $\text{intq}(b_{j'}, a_{k'})$ with $b_{j'} \in B_2$ and $k' \in \{i, i-1\}$ until it
 787 hits a query with $j' = j$ and $k' = k$. We call such a query *successful* if $j' = j$ and $k' = k$ and
 788 *unsuccessful* otherwise. In the same way, we call the selected tuples *successful* and all other
 789 tuples *unsuccessful*. Note that the algorithm might need further queries to prove that either
 790 (a_i, b_{i-1}) or (a_{i-1}, b_i) is not a rotation edge, but executing at least one successful query is a
 791 necessary condition.

792 After this slight adjustment, we can bound the expected number of queries in the same
 793 way as before (with the only difference that a single query can now cover only a single tuple
 794 and not two) to prove that every deterministic algorithm makes $\Omega(n^2)$ queries in expectation.
 795 On the other hand, the optimal solution for each realization of the randomized instance
 796 needs at most $\mathcal{O}(n)$ queries by again simulating the comparison query strategy with at most
 797 two interview queries per comparison. ◀

798 For the matching upper bound, recall that n^2 interview queries are enough to determine
 799 the full B -side preference lists. This means that we need at most n^2 interview queries to
 800 find the B -optimal matching M . We show that even the optimal solution needs at least
 801 $\Omega(n)$ interviews to find the B -optimal matching, which then implies an $\mathcal{O}(n)$ -competitive
 802 algorithm. We prove the following lemma by essentially repeating the corresponding proof
 803 for comparison queries (cf. Lemma 8)

804 ▶ **Lemma 18.** *The optimal number of queries for verifying the B -optimal stable matching*
 805 *with interview queries is at least $n - 1$ for every instance of the stable matching problem with*
 806 *one-sided uncertainty.*

807 **Proof.** Let M be the B -optimal stable matching for the given instance. Consider an arbitrary
 808 algorithm that verifies M to be B -optimal with interview queries. Assume that there are
 809 at least two distinct members a and a' of B for which the algorithm does not execute any
 810 queries. If some $b \in B$ satisfies either $b \prec_a M(a)$ or $b \prec_{a'} M(a')$, then this is a contradiction
 811 to the algorithm verifying M to be stable. Otherwise, $(a, M(a)), (a', M(a'))$ is a potential
 812 rotation so the algorithm has to prove that either $(a, M(a'))$ or $(a', M(a))$ is not an r -edge.
 813 Assume w.l.o.g. that the algorithm proves $(a, M(a'))$ to not be an r -edge. To do this, it either
 814 has to prove $a' \prec_{M(a')} a$, which is impossible without executing the interview $\text{intq}(M(a'), a)$,
 815 or it has to prove $a \prec_b M(b)$ for some b with $M(a) \prec_a b \prec_a M(a')$, which is impossible
 816 without executing the interview $\text{intq}(b, a)$. Each case leads to a contradiction. ◀

817 Together with Theorem 17, this lemma implies the following theorem.

818 ▶ **Theorem 19.** *In the interview query model, the best possible (randomized) competitive ratio*
 819 *for finding the B -optimal stable matching in an instance of stable matching with one-sided*
 820 *uncertainty is in $\Theta(n)$.*

821 4.3 NP-Hardness of the Offline Problem

822 For the offline problem of verifying a given B -optimal stable matching with interview queries,
 823 Rastegari et al. [29] show NP-hardness in a setting with partial uncertainty on both sides. As
 824 their proof exploits the possibility of giving partial information as part of the input, it does
 825 not directly translate to our setting with one-sided uncertainty. However, we can show with
 826 a similar proof as for comparison queries that the problem remains hard even in our setting.

827 ► **Theorem 20.** *The offline problem of computing an optimal set of interview queries for*
 828 *finding the B -optimal stable matching in a stable matching instance with one-sided uncertainty*
 829 *is NP-hard.*

830 **Proof.** Consider the same construction as in the proof for comparison queries (cf. Theorem 11).
 831 We add a dummy element z to A and a dummy element z' to B . Element z' has z as the
 832 top choice and afterwards all other agents of A in an arbitrary order. Agent z has z' as
 833 the last choice and before that the other elements of B in an arbitrary order. The agents
 834 of $A \setminus \{z\}$ all have z' as the top choice and afterwards the preference list as defined in the
 835 proof of Theorem 11. The elements of $B \setminus \{z'\}$ have z as the last choice and before that
 836 the preference list as defined in the proof of Theorem 11. This forces (z, z') to be part of
 837 the B -optimal matching and any algorithm has to query $\text{intq}(b, z)$ and $\text{intq}(b, M(b))$ for all
 838 $b \in B \setminus \{z'\}$ to prove stability.

839 After proving stability, each query $\text{intq}(b, a)$ for an $a \in A$ and $b \in B$ contains the
 840 information of $\text{prefer}(b, a, M(b))$ (as $\text{intq}(b, M(b))$ has already been queried to prove stability).
 841 Thus, we can now repeat the remaining part of the proof of Theorem 11 to show the
 842 theorem. ◀

843 5 Stable Matching with Set Queries

844 We consider the stable matching problem with one-sided uncertainty and set queries. Note
 845 that set queries are a natural generalization of comparison queries. For verifying any B -
 846 optimal matching, we show that the optimal number of set queries is at least $n - 1$. We also
 847 observe that there is an algorithm that makes at most n^2 queries for finding the B -optimal
 848 matching (or an A -optimal matching if we want to), as one can sort all preference lists
 849 using n^2 set queries. This implies an $\mathcal{O}(n)$ -competitive algorithm for finding the B -optimal
 850 matching. For the subproblem of verifying that a given matching is B -optimal, we give
 851 an $\mathcal{O}(\log n)$ -competitive algorithm by exploiting the additional power of set queries in an
 852 involved binary search algorithm. If we only have to verify stability for a given matching, we
 853 give a 1-competitive algorithm. Furthermore, we show that the offline problem of verifying
 854 that a given matching does not have a rotation is NP-hard.

855 5.1 Verifying That a Given Matching Is Stable

856 We start by characterizing the optimal number of queries (and query strategy) to verify that
 857 a given matching M is stable. The main difference to the comparison model is that, for a
 858 fixed $b \in B$, a single query $\text{top}(b, \{a \mid b \prec_a M(a)\} \cup \{M(b)\})$ is sufficient to prove that b is
 859 not part of any blocking pair.

860 ► **Theorem 21.** *Consider a stable matching instance with one-sided uncertainty and a*
 861 *stable matching M . The minimum number of set queries to verify that M is stable is*
 862 *$|\{b \in B \mid \exists a \in A: b \prec_a M(a)\}| \leq n$. Further, there is a 1-competitive algorithm to verify*
 863 *that M is stable.*

864 **Proof.** Consider an arbitrary $b \in B$. Let $Z(b) = \{a \in A \mid b \prec_a M(a)\}$, i.e., $Z(b)$ contains
 865 all $a \in A$ that could potentially form a blocking pair with b . Thus, M can only be stable if
 866 $M(b) \prec_b a$ holds for all $b \in B$ and $a \in Z(b)$. If $Z(b) \neq \emptyset$, at least one query to b is necessary,
 867 and the query $\text{top}(b, Z(b) \cup \{M(b)\})$ with answer $M(b)$ reveals all the required information
 868 to prove that b is not part of any blocking pair. Thus, the minimum number of queries to
 869 confirm that M is stable is $|\{b \in B \mid \exists a \in A: b \prec_a M(a)\}|$ as claimed. Furthermore, the
 870 algorithm that queries $\text{top}(b, Z(b) \cup \{M(b)\})$ for all $b \in B$ with $Z(b) \neq \emptyset$ is 1-competitive. ◀

871 5.2 Verifying That a Given Matching Is Stable and B -Optimal

872 For the problem of confirming that a given matching is B -optimal by using set queries, we
 873 show that every algorithm needs to execute at least $n - 1$ queries. This is analogous to the
 874 setting with comparison queries and uses a similar proof as Lemma 8. It implies that finding
 875 a B -optimal matching also requires at least $n - 1$ queries.

876 **► Lemma 22.** *Consider an arbitrary stable matching instance with one-sided uncertainty
 877 and the B -optimal matching M . Every algorithm needs at least $n - 1$ set queries to verify
 878 that M is indeed stable and B -optimal.*

879 **Proof.** For each $b \in B$, let $Z(b) = \{a \in A \mid b \prec_a M(a)\}$ and let $S = \{b \in B \mid Z(b) \neq \emptyset\}$. By
 880 the proof of Theorem 21, every algorithm needs to execute at least one query of the form
 881 $\text{top}(b, X)$ with $X \subseteq A$ for all $b \in S$ and this query has to return $M(b)$ as the top choice.
 882 Since verifying B -optimality includes proving stability, this leads to at least $|S|$ queries.

883 Consider an arbitrary algorithm that verifies M to be B -optimal and let $A_1 \subseteq A$ denote
 884 the agents of A that are returned as the top choice by some query of the algorithm. Then
 885 $|S| \leq |A_1|$ and $\{a \in A \mid \exists b \in S: M(b) = a\} \subseteq A_1$ by the argumentation above.

886 If $|A_1| \geq n - 1$, then the statement follows immediately, so assume $|A_1| < n - 1$ and
 887 let $A_2 = A \setminus A_1$. Since $|A_1| < n - 1$, the set A_2 has at least two distinct members a_1 and
 888 a_2 . Furthermore, we must have $M(a_1), M(a_2) \notin S$ as observed above. By definition of
 889 S , we have $M(a_1) \prec_{a_1} M(a_2)$ and $M(a_2) \prec_{a_2} M(a_1)$. This means that $(a_1, M(a_2))$ and
 890 $(a_2, M(a_1))$, based on the initially given information, could potentially be rotation edges.
 891 Thus, $(a_1, M(a_1)), (a_2, M(a_2))$ could potentially be a rotation and the algorithm has to
 892 prove that this is not the case by showing that one of $(a_1, M(a_2))$ and $(a_2, M(a_1))$ is not an
 893 r -edge. To prove that $(a_1, M(a_2))$ is not an r -edge, one has to either verify $a_2 \prec_{M(a_2)} a_1$ or
 894 $a_1 \prec_b M(b)$ for some $b \in B$ with $M(a_1) \prec_{a_1} b \prec_{a_1} M(a_2)$. However, this requires at least
 895 one query that returns either a_1 or a_2 as the top choice, and there is a symmetric argument
 896 for proving that $(a_2, M(a_1))$ is not an r -edge. Since a_1 and a_2 are never returned as the
 897 top choice by a query of the algorithm, this is a contradiction to the assumption that the
 898 algorithm verifies that M is B -optimal. ◀

899 In contrast to the comparison model, there exists an offline algorithm that asymptotically
 900 matches the lower bound of Lemma 22.

901 **► Theorem 23.** *There exists a polynomial-time offline algorithm that, given an instance of
 902 stable matching with one-sided uncertainty and the B -optimal matching M , verifies that M
 903 is indeed stable and B -optimal by executing $\mathcal{O}(n)$ set queries.*

904 **Proof.** By the proof of Theorem 21, an algorithm can prove M to be stable by executing at
 905 most n set queries, so it remains to prove that M is B -optimal by executing at most $\mathcal{O}(n)$
 906 set queries.

907 We do so by proving that M does not contain a rotation. First, for each $b \in B$, we
 908 compute the set $P(b) = \{a \in A \mid M(a) \prec_a b \text{ and } M(b) \prec_b a\}$. Each tuple (b, a) with $b \in B$
 909 and $a \in P(b)$ could be a rotation edge based on \prec_a but is not a rotation edge as $M(b) \prec_b a$.
 910 An algorithm can prove that none of these edge are actually rotation edges by executing a
 911 query $\text{top}(b, P(b) \cup \{M(b)\})$ for each $b \in B$. This leads to n additional queries.

912 If an $a \in A$ does not have a rotation edge, then the previous queries prove that this is the
 913 case. Consider an $a \in A$ that has a rotation edge. Then the second endpoint of that edge
 914 is the agent $b \in B$ of highest preference according to \prec_a among those agents that satisfy
 915 $M(a) \prec_a b$ and $a \prec_b M(b)$. Let b be that endpoint. To prove that (a, b) is indeed a rotation

916 edge, an algorithm has to verify $a \prec_b M(b)$ and $M(b') \prec_{b'} a$ for all b' with $M(a) \prec_a b' \prec_a b$.
 917 The latter has already been verified by the previous n queries and the former can be proven
 918 by an additional query $\text{top}(b, \{a, M(b)\})$. Doing this for every $a \in A$ that has a rotation edge
 919 leads to at most n further queries.

920 Executing these queries yields, for each $a \in A$, either the rotation edge of a or a proof
 921 that a does not have a rotation edge. Thus, it gives sufficient information to show that M
 922 does not have a rotation and is B -optimal. ◀

923 Next, we give an online algorithm that decides whether a given matching M is B -optimal
 924 by executing at most $\mathcal{O}(n \log n)$ set queries. In combination with Lemma 22, this yields an
 925 $\mathcal{O}(\log n)$ -competitive algorithm for verifying that a given matching is B -optimal with set
 926 queries.

927 ▶ **Theorem 24.** *There is an algorithm that decides if a given matching M in a stable matching*
 928 *instance with one-sided uncertainty is stable and B -optimal with $\mathcal{O}(n \log n)$ set queries.*

929 **Proof.** First, we can use Theorem 21 and execute $\mathcal{O}(n)$ queries to decide whether M is
 930 stable. If M turns out not to be stable, then we are done. Otherwise, we have to decide
 931 whether M is B -optimal by using at most $\mathcal{O}(n \log n)$ set queries. We do so by giving an
 932 algorithm that, for each $a \in A$, either finds the rotation edge of a or proves that a does not
 933 have a rotation edge. After executing that algorithm we clearly have sufficient information
 934 to decide whether M exposes a rotation and, thus, whether it is B -optimal.

935 For each $a \in A$, we use $R(a)$ to refer to the set of agents that could potentially form a
 936 rotation edge with a . Initially, we set $R(a) = \{b \in B \mid M(a) \prec_a b\}$ as all agents with a lower
 937 priority than $M(a)$ can potentially form a rotation edge with a based on the initially given
 938 information. During the course of our algorithm, we will update the set $R(a)$ such that it
 939 always only contains the agents of B that, based on the information obtained by all previous
 940 queries, could still form a rotation edge with a . In particular, if we obtain the information
 941 that $M(b) \prec_b a$ for some $b \in R(a)$, then (a, b) clearly cannot be a rotation edge and we
 942 can update $R(a) = R(a) \setminus \{b\}$. Similarly, if we obtain the information that $a \prec_b M(b)$ for
 943 some $b \in R(a)$, then the agents $b' \in R(a)$ with $b \prec_a b'$ cannot form a rotation edge with a
 944 anymore and we can update $R(a) = R(a) \setminus \{b' \in R(a) \mid b \prec_a b'\}$. Given the current list $R(a)$
 945 of potential rotation edge partners, we use $\bar{R}(a)$ to refer to the $\lceil \frac{|R(a)|}{2} \rceil$ agents of $R(a)$ with
 946 the highest priority in $R(a)$ according to \prec_a .

947 Our algorithm, cf. Algorithm 2, proceeds in iterations that each execute at most $\mathcal{O}(n)$
 948 set queries. Let $R_i(a)$, $a \in A$, denote the current sets of potential rotation edges at the
 949 beginning of iteration i and let $\bar{R}_i(a)$ be as defined above. We define our algorithm in a way
 950 such that each iteration i decides for each $a \in A$ whether it has a rotation edge to an agent of
 951 $\bar{R}_i(a)$ or not. Then, $|R_{i+1}(a)| \leq \frac{|R_i(a)|+1}{2}$ holds for each $a \in A$ with $|R_i(a)| > 1$ as we either
 952 get $R_{i+1}(a) \subseteq \bar{R}_i(a)$ or $R_{i+1}(a) \subseteq R_i(a) \setminus \bar{R}_i(a)$. Furthermore, if $|R_i(a)| = 1$, then iteration i
 953 either identifies the rotation edge of a or proves that it does not have one. This means that
 954 after at most $\mathcal{O}(\log n)$ such iterations, for each $a \in A$, we either found the rotation edge of a
 955 or verified that it does not have one. Since each iteration executes $\mathcal{O}(n)$ set queries, we get
 956 an algorithm that executes $\mathcal{O}(n \log n)$ set queries and decides whether M is B -optimal.

957 It remains to show that each iteration i indeed executes $\mathcal{O}(n)$ set queries and decides,
 958 for each $a \in A$, whether a has a rotation edge to some agent of $\bar{R}_i(a)$. Lines 4 to 13 of
 959 Algorithm 2 show the pseudocode for such an iteration. In each iteration i , the algorithm
 960 considers the set $U = \{a \in A \mid |R_i(a)| \geq 1\}$, i.e., the subset of A for which we do not yet
 961 know whether it has a rotation edge to some agent of $\bar{R}_i(a)$. Then, the algorithm iterates
 962 through the agents b of B and considers the set $U_b = \{a \in U \mid b \in \bar{R}_i(a)\}$. Note that, for

■ **Algorithm 2** Algorithm to decide whether a given matching is B -optimal using set queries.

Input: Stable matching instance with one-sided uncertainty and a matching M .

- 1 Decide whether M is stable using Theorem 21. If M is not stable, terminate;
- 2 $R(a) \leftarrow \{b \in B \mid M(a) \prec_a b\}$ for all $a \in A$;
- 3 **while** *We did not decide yet whether M is B -optimal* **do**
- 4 $U \leftarrow \{a \in A \mid |\bar{R}(a)| \geq 1\}$;
- 5 **for** $b \in B$ **do**
- 6 $U_b \leftarrow \{a \in U \mid b \in \bar{R}(a)\}$;
- 7 **repeat**
- 8 $t \leftarrow \text{top}(b, U_b \cup \{M(b)\})$;
- 9 **if** $t = M(b)$ **then** $R(a) \leftarrow R(a) \setminus b$ for all $a \in U_b$; $U_b \leftarrow \emptyset$;
- 10 **else**
- 11 $U \leftarrow U \setminus \{t\}$; $U_b \leftarrow U_b \setminus \{t\}$;
- 12 $R(t) \leftarrow R(t) \setminus \{b' \in R(t) \mid b \prec_t b'\}$;
- 13 **until** $U_b = \emptyset$;

963 each $a \in U_b$, it holds that if $a \prec_b M(b)$, then a has a rotation edge to some agent of $\bar{R}_i(a)$
 964 (not necessarily to b). The algorithm executes the query $\text{top}(b, U_b \cup \{M(b)\})$. If this query
 965 returns $M(b)$, then we know for sure that b does not have a rotation edge to any agent of U_b
 966 and we can discard b for the rest of the iteration and also remove b from the current $R(a)$ of
 967 all $a \in U_b$. On the other hand, if the query returns $a \neq M(b)$, then we know that a has a
 968 rotation edge to some agent of $\bar{R}_i(a)$ and we do not need to consider a for the rest of the
 969 iteration anymore. Thus, after each query within the iteration we discard an agent of either
 970 A or B , which means that the iteration terminates after at most $2n$ queries. At the end of
 971 the iteration, we know for each $a \in A$ whether it has a rotation edge to some $b \in \bar{R}_i(a)$. ◀

972 For the offline problem, we show that computing the query set of minimum size that
 973 verifies that a given matching does not have a rotation is NP-hard. However, in the instances
 974 constructed by the reduction, verifying that the given matching does not have a rotation
 975 and *is stable* is trivial as we will discuss after the proof. This means that the following result
 976 does *not* imply NP-hardness for the offline variant of finding the B -optimal matching with
 977 set queries.

978 ► **Theorem 25.** *In the set query model, the offline problem of computing an optimal set of*
 979 *queries for verifying that a given B -optimal stable matching M for a stable matching instance*
 980 *with one-sided uncertainty does not have a rotation is NP-hard.*

981 **Proof.** We show the statement by reduction from the NP-hard *feedback vertex set problem* [22].
 982 In this problem, we are given a directed graph $G = (V, E)$ and a parameter $k \in \mathbb{N}$. The goal
 983 is to decide whether there exists a subset $F \subseteq V$ with $|F| \leq k$ such that deleting F from G
 984 yields an acyclic graph.

985 We construct an instance of the stable matching problem with one-sided uncertainty and
 986 a matching M as follows:

- 987 1. For each $v \in V$, we add an agent a_v to set A and a matching partner $M(a_v)$ to set B .
- 988 2. For each $v \in V$ and $u \in V \setminus \{v\}$, we set $M(a_v) \prec_{a_v} M(a_u)$ if $(v, u) \in E$ and $M(a_u) \prec_{a_v}$
 989 $M(a_v)$ otherwise.
- 990 3. For each $v \in V$ and $u \in V \setminus \{v\}$, we set $a_v \prec_{M(v)} a_u$.

991 Based on the A -side preferences, each $(a_v, M(a_u))$ with $(v, u) \in E$ could be a rotation
 992 edge and each $(a_v, M(a_u))$ with $(v, u) \notin E$ is not a rotation edge. Consider the directed
 993 graph $G' = (A \cup B, E')$ with $E' = \{(M(a_v), a_v) \mid v \in V\} \cup \{(a_v, M(a_u)) \mid (v, u) \in E\}$. Then,
 994 based on the A -side preferences, each cycle in G' could be a rotation. Furthermore, if we
 995 contract the edges $\{(M(a_v), a_v) \mid v \in V\}$, we arrive at the given graph G .

996 Assume that there is a set $F \subseteq V$ with $|F| \leq k$ such that deleting F from G yields an
 997 acyclic graph. Consider the queries $\text{top}(M(a_v), A)$ for all $v \in F$. By the third step of the
 998 reduction, these queries prove that the agents $M(a_v)$ with $v \in F$ are not part of any rotation.
 999 This also means that the agents a_v with $v \in F$ cannot be part of a rotation. Thus, the only
 1000 edges that can still be part of a rotation are the matching edges $(M(a_v), a_v)$ with $v \notin F$ and
 1001 the edges $(a_v, M(a_u))$ with $(v, u) \in E$ but $v, u \notin F$. If we consider the graph induced by
 1002 these remaining edges and contract the matching edges, we arrive at the subgraph $G[V \setminus F]$
 1003 of the given feedback vertex set instance. Since this graph by assumption does not contain a
 1004 cycle, this implies that executing the queries proves that the constructed instance has no
 1005 rotation.

1006 Consider a query strategy that proves the constructed instance to not have a rotation
 1007 by using at most k queries. Let $A' \subseteq A$ denote the set of all agents that are returned as
 1008 the top choice by at least one of those queries. Then, by construction, the alternative query
 1009 strategy that queries $\text{top}(M(a_v), A)$ for each $a_v \in A'$ must also be feasible and uses at most
 1010 k queries. This alternative strategy proves that there exists no rotation by proving that no
 1011 $a_v \in A'$ is part of any rotation. Thus, removing all vertices a_v and $M(a_v)$ with $a_v \in A'$ from
 1012 the graph G' as defined above yields a graph without cycles. This also implies that removing
 1013 $F = \{v \in V \mid a_v \in A'\}$ from G yields a graph without cycles. Thus, F with $|F| \leq k$ is
 1014 feasible for the given feedback vertex set instance. ◀

1015 In the instances constructed within the proof, querying $\text{top}(M(a_v), A)$ for between $n - 1$
 1016 and n agents $a_v \in A$ proves that the given matching is stable and B -optimal. If $n - 1$ queries
 1017 suffice, then this is optimal by Lemma 22. Otherwise, n queries are optimal. We can decide
 1018 whether $n - 1$ queries suffice via enumerating all possible choices of the agent a_v for which
 1019 $M(a_v)$ does not receive a query $\text{top}(M(a_v), A)$.

1020 Thus, the NP-hardness for proving that no rotation exists does not directly translate to
 1021 the offline problem of proving that a given matching has no rotation *and is stable*.

1022 6 Open Problems

1023 While we understand the comparison model quite rigorously, it remains open in the set query
 1024 model what best possible competitive ratio can be achieved for finding a (A - or B -optimal)
 1025 stable matching. Further, it would be interesting to investigate the two-sided stable matching
 1026 problem with uncertainty in the preference lists on both sides further. For verifying the
 1027 stability of a given matching in this case, we have given a best possible 2-competitive
 1028 algorithm. All other questions regarding finding a stable or stable and optimal matching
 1029 remain open under two-sided uncertainty. It would also be interesting to investigate a
 1030 generalized set query model in which a query to a set $S \subseteq A$ for a $b \in B$ reveals the top- k
 1031 partners of b , that is, the k partners in S that b prefers most.

1032 References

- 1033 1 Haris Aziz, Péter Biró, Ronald de Haan, and Baharak Rastegari. Pareto optimal allocation
 1034 under uncertain preferences: uncertainty models, algorithms, and complexity. *Artif. Intell.*,
 1035 276:57–78, 2019.

- 1036 2 Haris Aziz, Péter Biró, Serge Gaspers, Ronald de Haan, Nicholas Mattei, and Baharak
1037 Rastegari. Stable matching with uncertain linear preferences. *Algorithmica*, 82(5):1410–1433,
1038 2020.
- 1039 3 Evripidis Bampis, Christoph Dürr, Thomas Erlebach, Murilo Santos de Lima, Nicole Megow,
1040 and Jens Schlöter. Orienting (hyper)graphs under explorable stochastic uncertainty. In *ESA*,
1041 volume 204 of *LIPIcs*, pages 10:1–10:18, 2021. doi:10.4230/LIPIcs.ESA.2021.10.
- 1042 4 Allan Borodin and Ran El-Yaniv. *Online computation and competitive analysis*. Cambridge
1043 University Press, 1998.
- 1044 5 Joanna Drummond and Craig Boutilier. Elicitation and approximately stable matching with
1045 partial preferences. In *IJCAI*, pages 97–105. IJCAI/AAAI, 2013.
- 1046 6 Joanna Drummond and Craig Boutilier. Preference elicitation and interview minimization in
1047 stable matchings. In *AAAI*, pages 645–653. AAAI Press, 2014.
- 1048 7 Christoph Dürr, Thomas Erlebach, Nicole Megow, and Julie Meißner. An adversarial model
1049 for scheduling with testing. *Algorithmica*, 82(12):3630–3675, 2020.
- 1050 8 Lars Ehlers and Jordi Massó. Matching markets under (in)complete information. *J. Econ.*
1051 *Theory*, 157:295–314, 2015.
- 1052 9 T. Erlebach and M. Hoffmann. Query-competitive algorithms for computing with uncertainty.
1053 *Bulletin of the EATCS*, 116:22–39, 2015. URL: [http://bulletin.eatcs.org/index.php/
1054 beatcs/article/view/335](http://bulletin.eatcs.org/index.php/beatcs/article/view/335).
- 1055 10 Thomas Erlebach, Murilo S. de Lima, Nicole Megow, and Jens Schlöter. Sorting and hypergraph
1056 orientation under uncertainty with predictions. In *IJCAI*, pages 5577–5585. ijcai.org, 2023.
- 1057 11 Guy Even, Joseph Naor, Baruch Schieber, and Madhu Sudan. Approximating minimum
1058 feedback sets and multicuts in directed graphs. *Algorithmica*, 20(2):151–174, 1998.
- 1059 12 D. Gale and L. S. Shapley. College admissions and the stability of marriage. *The American*
1060 *Mathematical Monthly*, 69(1):9–15, 1962. doi:10.1080/00029890.1962.11989827.
- 1061 13 Yannai A. Gonczarowski, Noam Nisan, Rafail Ostrovsky, and Will Rosenbaum. A stable
1062 marriage requires communication. *Games Econ. Behav.*, 118:626–647, 2019. doi:10.1016/j.
1063 geb.2018.10.013.
- 1064 14 Dan Gusfield and Robert W. Irving. *The Stable Marriage Problem – Structure and Algorithms*.
1065 Foundations of computing series. MIT Press, 1989.
- 1066 15 Guillaume Haeringer and Vincent Iehlé. Two-sided matching with one-sided preferences. In
1067 *EC*, page 353. ACM, 2014.
- 1068 16 Guillaume Haeringer and Vincent Iehlé. Two-sided matching with (almost) one-sided prefer-
1069 ences. *American Economic Journal: Microeconomics*, 11(3):155–190, 2019.
- 1070 17 Guillaume Haeringer and Vincent Iehlé. Enjeux stratégiques du concours de recrutement des
1071 enseignants chercheurs. *Revue Economique*, 61(4):697–721, 2010.
- 1072 18 M. M. Halldórsson and M. S. de Lima. Query-competitive sorting with uncertainty. In *MFCS*,
1073 volume 138 of *LIPIcs*, pages 7:1–7:15, 2019. doi:10.4230/LIPIcs.MFCS.2019.7.
- 1074 19 Michael Hoffmann, Thomas Erlebach, Danny Krizanc, Matús Mihalák, and Rajeev Raman.
1075 Computing minimum spanning trees with uncertainty. In *STACS*, volume 1 of *LIPIcs*, pages
1076 277–288. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Germany, 2008.
- 1077 20 Hadi Hosseini, Vijay Menon, Nisarg Shah, and Sujoy Sikdar. Necessarily optimal one-sided
1078 matchings. In *AAAI*, pages 5481–5488. AAAI Press, 2021.
- 1079 21 S. Kahan. A model for data in motion. In *STOC’91: 23rd Annual ACM Symposium on Theory*
1080 *of Computing*, pages 265–277, 1991. doi:10.1145/103418.103449.
- 1081 22 Richard M. Karp. Reducibility among combinatorial problems. In *50 Years of Integer*
1082 *Programming*, pages 219–241. Springer, 2010.
- 1083 23 Thomas Ma, Vijay Menon, and Kate Larson. Improving welfare in one-sided matchings using
1084 simple threshold queries. In *IJCAI*, pages 321–327. ijcai.org, 2021.
- 1085 24 David F. Manlove. *Algorithmics of Matching Under Preferences*, volume 2 of *Series on*
1086 *Theoretical Computer Science*. WorldScientific, 2013. doi:10.1142/8591.

- 1087 25 N. Megow, J. Meißner, and M. Skutella. Randomization helps computing a minimum spanning
1088 tree under uncertainty. *SIAM Journal on Computing*, 46(4):1217–1240, 2017. doi:10.1137/
1089 16M1088375.
- 1090 26 Cheng Ng and Daniel S. Hirschberg. Lower bounds for the stable marriage problem and its
1091 variants. *SIAM J. Comput.*, 19(1):71–77, 1990. doi:10.1137/0219004.
- 1092 27 Jannik Peters. Online elicitation of necessarily optimal matchings. In *AAAI*, pages 5164–5172.
1093 AAAI Press, 2022.
- 1094 28 Baharak Rastegari, Anne Condon, Nicole Immorlica, Robert W. Irving, and Kevin Leyton-
1095 Brown. Reasoning about optimal stable matchings under partial information. In *EC*, pages
1096 431–448. ACM, 2014.
- 1097 29 Baharak Rastegari, Anne Condon, Nicole Immorlica, and Kevin Leyton-Brown. Two-sided
1098 matching with partial information. In *EC*, pages 733–750. ACM, 2013.
- 1099 30 Alvin E. Roth and Marilda A. Oliveira Sotomayor. *Two-Sided Matching: A Study in Game-
1100 Theoretic Modeling and Analysis*. Econometric Society Monographs. Cambridge University
1101 Press, 1990. doi:10.1017/CCOL052139015X.
- 1102 31 Andrew Chi-Chih Yao. Probabilistic computations: Toward a unified measure of complexity.
1103 In *FOCS*, pages 222–227. IEEE Computer Society, 1977.