Minimax and Preferential Almost-Stable Matchings

Frederik Glitzner
School of Computing Science
University of Glasgow
Glasgow, United Kingdom
f.glitzner.1@research.gla.ac.uk

David Manlove
School of Computing Science
University of Glasgow
Glasgow, United Kingdom
david.manlove@glasgow.ac.uk

ABSTRACT

In the fundamental Stable Marriage and Stable Roommates problems, there are inherent trade-offs between the size and the stability of solutions. While the existence of desirable stable matchings can famously be determined in linear time when considering strict ordinal preferences, the computation of matchings that minimise the instability, either due to the presence of additional constraints on the size of the matching, or due to restrictive preference cycles, gives rise to a collection of infamously intractable *almost-stable matching* problems.

To better understand and deal with individual and collective incentives in such settings, we introduce two new perspectives on these problems. The first applies a minimax principle, seeking a matching that minimises the maximum number of blocking pairs that any single agent is involved in, thus limiting individual incentives to deviate. The second requires that a given set of agents is in few blocking pairs, or even entirely free of blocking pairs, motivated by contexts where some agents are unwilling or unable to initiate deviations even in the presence of such opportunities.

Surprisingly, both of these directions prove computationally intractable in strong ways: for example, it is NP-complete to decide whether a matching exists where no agent is in more than one blocking pair, even under bounded preference lists. On the positive side, we identify polynomial-time and fixed-parameter tractable cases, providing practical algorithmic tools for multi-agent systems where stability cannot be fully guaranteed, and offering new insights into the structure of almost-stable matchings.

KEYWORDS

Stable Marriage; Stable Roommates; Almost-Stable Matching

ACM Reference Format:
Frederik Glitzner and David Manlove. 2026. Minimax and Preferential Almost-Stable Matchings. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 9 pages. https://doi.org/10.65109/PCDE6577

1 INTRODUCTION

The two foundational stable matching problems Stable Marriage and Stable Roommates were first studied by Gale and Shapley [18]. In the former, there are two disjoint sets of agents A_1 and A_2, where each agent ranks a subset of the opposite set. The latter problem generalises this: there is a single set A of agents, each ranking a subset of the remaining agents. We will denote these problems by SMI and SRI, respectively, to indicate that incomplete preference lists, i.e., preferences over a strict subset of agents (implying that all unranked agents are not acceptable), are permissible. In both settings, the goal is to find a *matching* consisting of disjoint pairs of mutually acceptable agents, such that no two agents prefer each other to their partners (either or both of which could be none). If such a pair does exist, we refer to it as a *blocking pair*. The problem of finding a *stable matching* (i.e., a matching without blocking pairs) underpins many real-world allocation tasks, such as matching doctors to hospitals [35], students to dormitories [2], or users to communication channels [4]. Much of the literature focuses on efficient algorithms for centralised schemes [36]. While a stable matching for an SMI instance can always be found in linear time using the Gale-Shapley algorithm [18], not every SRI instance admits a stable matching [29]. If at least one exists, one can be found in linear time using Irving's algorithm [34].

Two key challenges arise in practice. Incomplete preference lists can prevent stable matchings from maximising the number of matched agents, resulting in unfilled positions and unmatched agents despite possible improvements to the size [8]. Secondly, in non-bipartite matching settings, stable matchings often do not exist at all [22]. Gale and Sotomayor [19] and Gusfield and Irving [29] observed that, under strict preferences, the set of matched agents is identical in all stable matchings. Consequently, there is no flexibility to increase the size of the matching while maintaining stability. A large number of unmatched agents can often be considered a significant weakness, and experimental works have highlighted the practical importance of relaxing stability for a substantial improvement in this regard [5]. Eriksson and Häggström [14] reviewed different measures of instability and concluded that counting blocking pairs is often the most indicative of instability, as each such pair has both the incentive and the opportunity to deviate [36]. In this work, we adopt a fine-grained, agent-centric view of blocking pairs and continue research on problems concerned with matchings and maximum-cardinality matchings that are "as stable as possible", also referred to as "almost-stable" matchings [1, 8].

1.1 Our Contributions

Given the practical importance of matchings that minimise instability (and possibly maximise size), strong intractability results for various problems involving the computation of such matchings are discouraging. Hence, we aim to understand the complexity frontier for related problems better in search of efficiently solvable variants.

In this paper, we consider two directions that are somewhat orthogonal to those studied previously. In the first one, we study

---

The full versions of this paper can be found at references [24, 25].

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/PCDE6577

almost-stability through the classical economic and optimisation principle of *minimax fairness.* Here, we investigate the computation of (any or maximum-cardinality) matchings that minimise the maximum number of blocking pairs that any agent is in. Hence, rather than minimising the aggregate instability, we minimise *individuals'* incentives to deviate. Surprisingly, even deciding whether a matching (respectively a maximum-cardinality matching) exists in which no agent is in more than one blocking pair is NP-complete for SRI (SMI). On the positive side, we present polynomial-time algorithms for instances with preference lists of length at most 2.

In the second direction, we study the computation of (any or maximum-cardinality) matchings that minimise the number of blocking pairs for a given subset of *preferential agents*, i.e., when giving *preferential treatment* to a subset of the agents. For example, agents that are bound by law, contracts, or regulations may be unable to initiate deviations, so there may be less of a practical need to minimise their instability. Hence, we aim to minimise the deviation incentive of the remaining (preferential) agents. Unfortunately, here we can also conclude strong intractability results: even deciding whether there exists a matching (respectively maximum-cardinality matching) in which no preferential agent is in any blocking pair is NP-complete for SRI (SMI). On the positive side, we present efficient parametrised algorithms.

1.2 Formal Definitions

We start by rigorously defining the models that we already introduced informally.

Definition 1.1 (SRI Instance). *Let $I = (A, \succ)$ be a* Stable Roommates with Incomplete lists *(*SRI*)* instance *consisting of a set $A = \{a_1, a_2, \dots, a_n\}$ of agents and a tuple of strict ordinal* preference rankings *$\succ$, where each agent a_i has a ranking $\succ_i$ over a subset of $A \setminus \{a_i\}$. Agent a_i* prefers *a_j to a_k if $a_j \succ_i a_k$. We write $a_j \succeq_i a_k$ if $a_j \succ_i a_k$ or $j = k$. Agents a_i and a_j are* acceptable *if a_i appears in $\succ_j$ and vice versa. Throughout, we assume that acceptability is symmetric.*

We can equivalently treat an SRI *instance $I = (A, \succ)$ as a graph $G = (A, E)$: let every agent $a_i \in A$ be a vertex in the graph and let all acceptable pairs of agents $\{a_i, a_j\}$ be in E. We refer to G as the* acceptability graph *and say that I has* complete preferences *if G is isomorphic to the complete graph K_n, and* incomplete preferences *if not. We also assume that every agent ranks themselves last.*

Definition 1.2 (SMI Instance). *Let $I = (A, \succ)$ be an* SRI *instance and let G be the acceptability graph of I. If G is bipartite, i.e., there exist disjoint sets A_1 and A_2 such that $A = A_1 \dot\cup A_2$ and no agents within A_1 or A_2 are mutually acceptable, then we refer to I as a* Stable Marriage with Incomplete lists *(*SMI*) instance. Furthermore, we say that* SMI *instance I has* complete preferences *if G is isomorphic to $K_{|A_1|,|A_2|}$, and* incomplete preferences *if not.*

We adopt the following notation and assume the following concepts for matchings.

Definition 1.3 (Matchings). *Let $I = (A, \succ)$ be an* SRI *(or* SMI*) instance. A* matching *M of I is a set of unordered pairs of agents such that no agent a_i is contained in more than one pair. If $\{a_i, a_j\} \in M$, then we refer to a_i as* matched *and to a_j as the* partner *of a_i, denoted by $M(a_i)$, and vice versa. If a_i is not paired with any other agent, we say that a_i is* unmatched *and set $M(a_i) = a_i$.*

We denote the set of all matchings of I by $\mathcal{M}$ and refer to some $M \in \mathcal{M}$ as a maximum-cardinality matching *if, for any $M' \in \mathcal{M}$, $|M| \geq |M'|$. We denote the set of all maximum-cardinality matchings by $\mathcal{M}^+$. We refer to $M \in \mathcal{M}^+$ as* perfect *if every $a_i \in A$ is matched to some agent in $A \setminus \{a_i\}$, and denote the set of such matchings by $\mathcal{M}^p$.*

We adopt the following classical stability notion.

Definition 1.4 (Stability). *Let $I = (A, \succ)$ be an* SRI *(or* SMI*) instance and let M be a matching of I. A* blocking pair *admitted by M is a pair of distinct agents $a_i, a_j \in A$ such that $a_i \succ_j M(a_j)$ and $a_j \succ_i M(a_i)$. We denote the* set of blocking pairs *by $bp(M)$, and the subset thereof that involves an agent a_i by $bp_{a_i}(M)$. Furthermore, the set of* blocking agents *$ba(M)$ consists of all agents in blocking pairs.*

If $bp(M) = \varnothing$, then M is referred to as stable*, and as* unstable *otherwise. We refer to I as* solvable *if it admits at least one stable matching and as* unsolvable *otherwise.*

In the absence of stable matchings, we previously noted that one might want to find an "almost-stable" matching that minimises the number of blocking pairs [1]. We define this problem as follows.

MinBP-AlmostStable-SRI
Input: SRI instance I.
Output: Matching M of I s.t. $|bp(M)| = \min_{M' \in \mathcal{M}} |bp(M')|$.

Even when considering only SMI instances, the set of stable matchings might not intersect with $\mathcal{M}^+$. Thus, we may want to find a maximum-cardinality matching that minimises the number of blocking pairs [8]. This problem is defined as follows.

MinBP-AlmostStable-Max-SMI
Input: SMI instance I.
Output: A maximum-cardinality matching M of I such that $|bp(M)| = \min_{M' \in \mathcal{M}^+} |bp(M')|$.

1.3 Related Work

MinBP-AlmostStable-SRI was shown to be NP-hard regardless of whether preference lists are of length at most 3 or complete [1, 7], NP-hard to approximate within $n^{1/2-\varepsilon}$ (for any $\varepsilon > 0$, where n is the number of agents) [1], and W[1]-hard with respect to $\min_{M \in \mathcal{M}} |bp(M)|$ even when preference lists are of length at most 5 [10]. Biró, Manlove and Mittal [8] conducted a thorough study of the computational complexity of the trade-off between matching size and stability in SMI, where stability is measured in terms of the number of blocking pairs or blocking agents. They showed that MinBP-AlmostStable-Max-SMI (and similarly for blocking agents) is NP-hard to approximate within $n^{1-\varepsilon}$ (for any $\varepsilon > 0$) [8]. Hamada et al. [30] strengthened this result to preference lists of length at most 3. On the positive side, the problem is solvable efficiently for preference lists of length at most 2 [8].

Gupta et al. [27] showed that MinBP-AlmostStable-Max-SMI is W[1]-hard for the combined parameter (β, t, d), where β is the number of blocking pairs, t is the increase in size to a stable matching, and d is the maximum length of a preference list. In a very recent preprint, Chen, Roy and Simola [11] proved that it is W[1]-hard with respect to the optimal value to even approximate MinBP-AlmostStable-SRI and MinBP-AlmostStable-Max-SMI within any computable function depending only on this optimal value.

We will denote the decision problems "Does there exist a matching M of I such that $|bp(M)| \leq k$?" and "Does there exist a perfect matching M of bipartite instance I such that $|bp(M)| \leq k$?" by k-BP-AlmostStable-SRI and k-BP-AlmostStable-Perfect-SMI, respectively. Both problems are known to be NP-complete [1, 7–9], but in XP with respect to k [1, 8]. Minimising the number of blocking agents is also computationally intractable [8–10].

Different definitions of almost-stability have also been considered [16, 38]. Work on almost-stable matchings for other problem models has also been conducted [12, 20, 21, 27, 37]. Other alternative solution concepts for SRI have also been explored [15, 23, 26, 28, 31, 39, 40]. An alternative perspective on minimising instability in SMI was offered by Biermann [6], who argued that, instead of counting blocking pairs, one should count a subset thereof that comprises a possible and economically reasonable transformation. Similar perspectives have also been studied for SRI [3].

1.4 Structure of the Paper

We consider the computation of minimax almost-stable matchings in Section 2, giving guarantees with respect to the maximum number of blocking pairs for individual agents. Preferential almost-stability is considered in Section 3, targeting subsets of agents that are known in advance to be likely to deviate. We summarise the complexity results in Table 1 and pose open questions in Section 4.

An extended version of Section 2 with all omitted and shortened proofs, as well as further structural and intractability results, approximation and inapproximability results, and integer linear programs (ILP) and experiments, is available in [24]. Similarly, an extended version of Section 3 with all omitted and shortened proofs, as well as further results for minimising blocking preferential agents and generalised algorithms, is available in [25] (where we adopt the terminology "deviator agents" instead of "preferential agents").

2 MINIMAX ALMOST-STABLE MATCHINGS

In many real-world matching markets, agents contained in multiple blocking pairs not only have an increased opportunity to deviate and potentially cause unravelling, but they may also feel unfairly treated, knowing that they are in a position in which they could be better off after a decentralised reorganisation. To address this, we initiate the study of minimax almost-stable matchings, which optimise the solution for the worst-off agent by minimising the maximum number of blocking pairs any agent is involved in. This approach provides strong individual-level fairness guarantees and is particularly relevant in high-stakes multi-agent settings. We first study the following natural problem, which aims for an SRI matching with a minimax number of blocking pairs.

Minimax-AlmostStable-SRI
Input: SRI instance $I = (A, \succ)$.
Output: A matching $M \in \mathcal{M}$ such that $\max_{a_i \in A} |bp_{a_i}(M)| = \min_{M' \in \mathcal{M}} \max_{a_i \in A} |bp_{a_i}(M')|$.

We will also study the following related problem, in which we are given an SMI instance and aim for a maximum-cardinality matching with a minimax number of blocking pairs.

Minimax-AlmostStable-Max-SMI
Input: SMI instance $I = (A, \succ)$.
Output: A matching $M \in \mathcal{M}^+$ such that $\max_{a_i \in A} |bp_{a_i}(M)| = \min_{M' \in \mathcal{M}^+} \max_{a_i \in A} |bp_{a_i}(M')|$.

Minimax almost-stability differs fundamentally from aggregate stability metrics (like minimising the total number of blocking pairs) by prioritising the worst-off agent. Furthermore, it contrasts the minimisation of blocking agents [8, 10], which likely requires fewer agents with a stronger incentive to deviate, therefore causing a greater risk of unravelling.

2.1 Minimax Matchings in General Graphs

For our complexity analysis, we define a decision problem associated with Minimax-AlmostStable-SRI as follows.

k-Max-AlmostStable-SRI
Input: SRI instance $I = (A, \succ)$ and non-negative integer k.
Question: Does there exist a matching $M \in \mathcal{M}$ such that $\max_{a_i \in A} |bp_{a_i}(M)| \leq k$?

It is easy to see that if k-Max-AlmostStable-SRI cannot be solved in polynomial time, then Minimax-AlmostStable-SRI cannot be solved in polynomial time. Notice that I is a yes-instance to 0-Max-AlmostStable-SRI if and only if I is solvable, so this case is tractable [34]. Notice, furthermore, that k-Max-AlmostStable-SRI is trivial if k exceeds the maximum preference list length.

We will now highlight, though, that k-Max-AlmostStable-SRI is NP-complete even in the very restricted case where $k = 1$, and regardless of whether preference lists are bounded or complete. Specifically, the result below states that it is impossible to distinguish in polynomial time even between instances that admit a matching in which every agent is in at most one blocking pair, and instances where every matching requires at least one agent to be in more than one blocking pair (unless P=NP), even when preference lists are of bounded length or complete.

Theorem 2.1. *k-Max-AlmostStable-SRI is NP-complete, even if $k = 1$ and regardless of whether all preference lists are of length at most 10 or complete.*

This strong intractability result highlights the computational limitations of enforcing individual-level stability guarantees in these multi-agent settings. Intractability of Minimax-AlmostStable-SRI follows immediately, but we will state it again formally below. Note that a problem is para-NP-hard with respect to a parameter κ if it is NP-hard already for a constant value of κ, and a problem is in XP with respect to κ if there exists an $O(n^{f(\kappa)})$ algorithm. If a problem is para-NP-hard with respect to a parameter κ, then it is not in XP with respect to κ unless P=NP (we refer to [17] for an introduction to parametrised complexity theory).

Corollary 2.2. *Minimax-AlmostStable-SRI is para-NP-hard with respect to $\kappa = \min_{M \in \mathcal{M}} \max_{a_i \in A} |bp_{a_i}(M)|$, regardless of whether preference lists are of bounded length at most 10 or whether preferences* *are complete. Furthermore, the problem admits no polynomial-time approximation algorithm with a performance guarantee better than 2, unless* P=NP.

We provide further computational intractability results in [24].

On the positive side, we provide an efficient algorithm for SRI instances with preference lists of length at most 2 that returns an optimal solution to Minimax-AlmostStable-SRI. Although the setting where preference lists are of length at most 2 is very constrained, we highlight that these algorithms apply whenever choices are very limited, or when computing a matching in a two-stage approach, where first agents are asked to provide their top two choices, and then the remaining agents are matched arbitrarily. The procedure is given in Algorithm 1 and works as follows: determine whether the instance I admits a stable matching M_S using Irving's algorithm [34]. If yes, immediately return M_S. If not, compute a maximum-cardinality matching M_C as follows. Notice that because preference lists are of length at most 2, the acceptability graph of I has maximum degree at most 2, i.e., it consists only of paths and cycles. Thus, to compute M_C, we can simply include every second edge of every path, starting from one of the endpoints, and every second edge of every cycle (such that no included edges have a vertex in common). It is easy to see that this results in a maximum-cardinality matching for graphs with maximum degree at most 2. Next, we perform a rematching procedure on M_C that ensures that no agent is contained in more than one blocking pair, and return this modified matching.

Algorithm 1 Exact algorithm for Minimax-AlmostStable-SRI with lists of length at most 2

```
Input: $I$ : an SRI instance
Output: $M$ : a matching
$M_S \leftarrow \mathsf{Irving}(I)$        ▷ Compute a stable matching if one exists
if $M_S$ exists then
    return $M_S$
end if
$M_C \leftarrow \mathsf{MaxCard}(I)$        ▷ Compute a maximum-cardinality matching
while $\exists\ a_i$ such that $|bp_{a_i}(M_C)| = 2$ do        ▷ Rematching procedure
    pick one agent $a_r$ that is in a blocking pair with $a_i$
    $a_k \leftarrow M_C(a_r)$
    $M_C(a_i) \leftarrow a_r$ ; $M_C(a_r) \leftarrow a_i$ ; $M_C(a_k) \leftarrow a_k$
end while
return $M_C$
```

Theorem 2.3. *Let I be an* SRI *instance with n agents. If preference lists are of length at most 2, then* $\min_{M \in \mathcal{M}} \max_{a_i \in A} |bp_{a_i}(M)| \leq 1$ *and Algorithm 1 solves* Minimax-AlmostStable-SRI *in $O(n)$ time.*

Proof. Clearly, if M_S exists, then $\max_{a_i \in A} |bp_{a_i}(M_S)| \leq 0$, so M_S is optimal. Furthermore, if M_S does not exist, then it must be the case that $\min_{M \in \mathcal{M}} \max_{a_i \in A} |bp_{a_i}(M)| > 0$.

Now suppose that $\max_{a_i \in A} |bp_{a_i}(M_C)| > 1$, then, because all preference lists are of length at most 2, there exists an agent a_i who is unmatched in M_C and who is contained in exactly two blocking pairs, say $bp_{a_i}(M_C) = \{\{a_i, a_r\}, \{a_i, a_s\}\}$ (and suppose that $a_r \succ_i a_s$). Clearly, both a_r and a_s must be matched in M_C, otherwise M_C would not be a maximum-cardinality matching as $M_C \cup \{\{a_i, a_r\}\}$ or $M_C \cup \{\{a_i, a_s\}\}$ would be a matching of larger size. Also, without loss of generality, let a_k denote the current partner of a_r in M_C, then $a_i \succ_r a_k$, otherwise a_i, a_r would not block. Now we update M_C such that a_i and a_r are matched to each other (instead of a_r to a_k). Due to the bound on the preference list lengths, a_k's preference list either contains just a_r, or two agents a_r and a_l (for some agent a_l). Note, though, that a_l must be matched in M_C, otherwise $M'_C = (M_C \backslash \{\{a_r, a_k\}\}) \cup \{\{a_i, a_r\}, \{a_k, a_l\}\}$ would exceed M_C in size and thus M_C would not be of maximum-cardinality.

After this rematching procedure terminates, the size of M_C remains the same as previously, but $|bp_{a_i}(M_C)| \leq 1$. Furthermore, because $a_i \succ_r a_k$, $|bp_{a_r}(M_C)| = 0$. Finally, $|bp_k(M_C)| \leq 1$ because while a_k can generally block with at most two agents due to its preference list of length at most two, it cannot block with a_r because $a_i \succ_r a_k$. Thus, this rematching procedure strictly decreases the number of agents who are contained in more than one blocking pair at each iteration of the main while loop. While at least one such agent remains, we repeat this rematching procedure with this agent taking the role of a_i in the argument above.

Asymptotically, Irving's algorithm requires at most linear time in the sum of the preference list lengths of I [34], which, due to bounded preference list lengths, is linear in the number of agents of I. The maximum-cardinality matching algorithm we described simply walks over paths and cycles, so it requires $O(m)$ time for a graph with m edges. We execute this on the acceptability graph of I, which has n vertices and, again due to bounded preference list lengths, $m = O(n)$ edges. Hence, computing M_C requires $O(n)$ time. The rematching procedure, if invoked, is executed at most $O(n)$ times because the number of agents in more than one blocking pair strictly decreases with each iteration (as argued above), and each execution requires constant time. We can implement the while loop efficiently by building a stack of unmatched agents (every agent involved in two blocking pairs must clearly be unmatched) and, at every iteration of the loop, popping an agent from this stack and verifying (in constant time) whether this agent is indeed blocking with both agents on its preference list. Because we do the computation of the initial M_C and the rematching sequentially, and assuming that matching size comparisons can be implemented in constant time, we arrive at an overall time complexity of $O(n)$. □

This result identifies a complexity frontier. Even though minimax almost-stability is generally intractable (for instances with preference lists of length at most 10), instances with very short preference lists (of length at most 2) allow an efficient computation.

2.2 Minimax Maximum Matchings in Bipartite Graphs

We will now turn to Minimax-AlmostStable-Max-SMI. For our complexity analysis, we will consider a restricted decision version of this problem, which we define as follows.

k-Max-AlmostStable-Perfect-SMI
Input: SMI instance $I = (A, \succ)$ and non-negative integer k.
Question: Does there exist a matching $M \in \mathcal{M}^p$ such that $\max_{a_i \in A} |bp_{a_i}(M)| \leq k$?

It is easy to see that if k-Max-AlmostStable-Perfect-SMI cannot be solved in polynomial time, then Minimax-AlmostStable-Max-SMI cannot be solved in polynomial time. Note that I is a yes-instance to 0-Max-AlmostStable-Perfect-SMI if and only if there exists a perfect stable matching of I. This can be decided in polynomial time by computing an arbitrary stable matching M of I using the classical Gale-Shapley algorithm [18], and checking whether M is a perfect matching. All stable matchings are known to match the same number of agents [19], so if M is not perfect, then I does not admit any perfect stable matching.

We now highlight, though, that k-Max-AlmostStable-Perfect-SMI is NP-complete even in the very restricted case where $k = 1$ and preference lists are very short. The result below states that it is impossible to distinguish in polynomial time even between instances that admit a perfect matching in which every agent is in at most one blocking pair and instances where every perfect matching requires at least one agent to be in more than one blocking pair (unless P=NP).

Theorem 2.4. *k-Max-AlmostStable-Perfect-SMI is NP-complete, even if $k = 1$ and all preference lists are of length at most 3.*

Corollary 2.5. *Minimax-AlmostStable-Max-SMI is para-NP-hard with respect to $\kappa = \min_{M \in \mathcal{M}^+} \max_{a_i \in A} |bp_{a_i}(M)|$ even if all preference lists are of length at most 3. Furthermore, there exists no approximation algorithm for Minimax-AlmostStable-Max-SMI with a performance guarantee strictly better than 2 and a polynomial runtime, unless P=NP.*

Notice that this result cannot be extended to SMI instances with complete preference lists, as all stable matchings of such instances are maximum-cardinality matchings. However, we provide further negative complexity results in [24].

Interestingly, the high-level technique of Algorithm 1 for low-degree acceptability graphs in SRI carries over to this problem. Specifically, the procedure we give in Algorithm 2 solves Minimax-AlmostStable-Max-SMI to optimality when preference lists are of length at most 2. It works as follows: compute a stable matching M_S using the Gale-Shapley algorithm [18] and a maximum-cardinality matching M_C using the same procedure as the one that we described for Algorithm 1 (as the acceptability graph still only consists of paths and cycles). If $|M_S| = |M_C|$ then return M_S, otherwise perform the same rematching procedure on M_C as in Algorithm 1 and return this modified matching.

Theorem 2.6. *Let I be an SMI instance with n agents. If all preference lists are of length at most 2, then $\min_{M \in \mathcal{M}^+} \max_{a_i \in A} |bp_{a_i}(M)| \leq 1$ and Algorithm 2 solves Minimax-AlmostStable-Max-SMI in $O(n)$ time.*

Algorithm 2 Exact algorithm for Minimax-AlmostStable-Max-SMI with lists of length at most 2

```
Input: $I$ : an SMI instance
Output: $M$ : a matching
$M_S \leftarrow$ GaleShapley($I$)          ▷ Compute a stable matching
$M_C \leftarrow$ MaxCard($I$)          ▷ Compute a maximum-cardinality matching
if $|M_S| = |M_C|$ then
    return $M_S$
end if
while $\exists\ a_i$ such that $|bp_{a_i}(M_C)| = 2$ do          ▷ Rematching procedure
    pick one agent $a_r$ that is in a blocking pair with $a_i$
    $a_k \leftarrow M_C(a_r)$
    $M_C(a_i) \leftarrow a_r$ ; $M_C(a_r) \leftarrow a_i$ ; $M_C(a_k) \leftarrow a_k$
end while
return $M_C$
```

Proof sketch. The approach remains the same as in the proof of Theorem 2.3. We highlight some differences below and refer to [24] for further details. First, recall that all stable matchings of I have the same size [19]. Clearly, $|M_C| \geq |M_S|$ and $\max_{a_i \in A} |bp_{a_i}(M_S)| = 0$, so if $|M_S| = |M_C|$, then M_S is also a maximum-cardinality matching and a stable matching of I, so it is also a solution to Minimax-AlmostStable-Max-SMI. However, if $|M_S| \neq |M_C|$, then necessarily $|M_S| < |M_C|$, so no solution to Minimax-AlmostStable-Max-SMI is stable, i.e., for all $M' \in \mathcal{M}^+$, $\max_{a_i \in A} |bp_{a_i}(M')| > 0$.

The Gale-Shapley algorithm requires at most linear time in the preference list lengths [18], which, here, is linear in the number of agents. The maximum-cardinality matching procedure we described also requires $O(n)$ time. The rematching procedure, if invoked, is executed at most $O(n)$ times because the number of agents in more than one blocking pair strictly decreases, and each execution requires constant time. □

This positive result implies a tight complexity frontier: Minimax-AlmostStable-Max-SMI can be solved very efficiently when preference lists are of length at most 2. However, the problem is NP-hard when preference lists are of length at most 3, as previously established in Corollary 2.5.

In [24], we also study the approximability and exact ILP formulations for these minimax problems.

3 PREFERENTIAL ALMOST-STABLE MATCHINGS

It is clear from previous work and our new results that minimising either the aggregate or the individual instability among the agents is generally intractable. However, not every agent might be able to deviate, even when a beneficial opportunity to do so presents itself. For example, some agents might be prevented from initiating deviations by external constraints. Hence, in this section, we study optimisation and decision problems that aim for matchings that give preferential treatment to the remaining agents, or generally, any designated subset of agents. Specifically, we seek matchings that minimise the blocking pairs involving these agents, or ideally, exclude them from any blocking pairs altogether.

We first propose the following natural optimisation problem, which seeks a matching that minimises the number of blocking pairs involving agents from a designated subset A'. Intuitively, the problem asks: how stable can a matching be made for a privileged subset of agents?

Preferential-AlmostStable-SRI
Input: SRI instance $I = (A, \succ)$ and a subset $A' \subseteq A$.
Output: A matching $M \in \mathcal{M}$ such that $|\bigcup_{a_i \in A'} bp_{a_i}(M)| = \min_{M' \in \mathcal{M}} |\bigcup_{a_i \in A'} bp_{a_i}(M')|$.

We will also study the maximum-cardinality bipartite matching analogue, which we define formally below.

Preferential-AlmostStable-Max-SMI
Input: SMI instance $I = (A, \succ)$ and a subset $A' \subseteq A$.
Output: A matching $M \in \mathcal{M}^+$ such that $|\bigcup_{a_i \in A'} bp_{a_i}(M)| = \min_{M' \in \mathcal{M}^+} |\bigcup_{a_i \in A'} bp_{a_i}(M')|$.

Of course, both of these problems are NP-hard in general, because they are equivalent to MinBP-AlmostStable-SRI and MinBP-AlmostStable-Max-SMI when $A' = A$. However, we will investigate restrictions and algorithms for the case when A' is small.

Rather than beginning with the general SRI case, we first examine maximum-cardinality matchings with these properties in bipartite acceptability graphs, before extending our techniques to SRI.

3.1 Preferential Maximum Matchings in Bipartite Graphs

We start with the following parametrised version of Preferential-AlmostStable-Max-SMI.

k-Preferential-AlmostStable-Max-SMI
Input: SMI instance $I = (A, \succ)$ and a subset $A' \subseteq A$.
Output: A matching $M \in \mathcal{M}^+$ such that $|\bigcup_{a_i \in A'} bp_{a_i}(M)| \leq k$, if one exists

We will denote the problem of deciding whether there exists a solution to k-Preferential-AlmostStable-Max-SMI when $\mathcal{M}^+ = \mathcal{M}^p$ by k-Preferential-AlmostStable-Perfect-SMI-Dec.

Notice that when $A' = A$ and $k = 0$, then (I, A') is a yes-instance to 0-Preferential-AlmostStable-Max-SMI if and only if there exists a stable maximum-cardinality matching of I. This can be decided efficiently by computing a stable matching M of I using the Gale-Shapley algorithm [18] and checking whether M is a maximum-cardinality matching. If it is not, then no solution exists.

Somewhat surprisingly, though, when $A' \subset A$, then even 0-Preferential-AlmostStable-Perfect-SMI-Dec is intractable, i.e., even when we merely ask whether a perfect matching exists such that no A' agent is in a blocking pair.

Theorem 3.1. *0-Preferential-AlmostStable-Perfect-SMI-Dec is NP-complete, even if all preference lists are of length at most 3.*

Corollary 3.2. *k-Preferential-AlmostStable-Max-SMI is NP-hard, even when $k = 0$ and all preference lists are of length at most 3. Preferential-AlmostStable-Max-SMI is para-NP-hard with respect to $\kappa = \min_{M \in \mathcal{M}^+} |\bigcup_{a_i \in A'} bp_{a_i}(M)|$.*

Note that, as in our discussion of Minimax-AlmostStable-Max-SMI, these hardness results no longer apply when preference lists are complete, since then every stable matching is a maximum-cardinality matching.

On the positive side, we show that the problem k-Preferential-AlmostStable-Max-SMI is fixed-parameter tractable (in the complexity class FPT [17]) with respect to the combined parameter $(|A'|, d_{\max})$, as well as in XP with respect to just $|A'|$, where $d_{\max}$ is the maximum preference list length of the instance. We show this by giving an algorithm in the following proof, which is efficient for small $|A'|$, k, and/or $d_{\max}$. Note that $d_{\max}$ is often naturally small: as preference lists usually need to be elicited manually, there is often a limit on how many choices a participant ranks. Furthermore, when maximising the matching size as a primary objective, it is often beneficial for agents to submit short preference lists. Of course, $|A'|$ need not be small in general. However, one possible strategy is to fix $k = 0$, in which case the algorithm terminates in $O(d_{\max}^{O(|A'|)} n^{3/2} \log(n))$ time, and rank agents based on their likelihood to deviate. Then, one could start by including only the strongest deviators in A', and then increasing the size of A' up until either no more solution exists, or the computation takes too long.

Intuitively, our algorithm enumerates possible configurations of partners and blocking pairs for the preferential agents, and tries to extend each (feasible) configuration to a desirable matching using a maximum-weight matching subroutine.

Theorem 3.3. *Let $(I = (A, \succ), A')$ be an instance of k-Preferential-AlmostStable-Max-SMI with n agents and maximum preference list length $d_{\max}$. Then k-Preferential-AlmostStable-Max-SMI is solvable in $O(|A'|^k d_{\max}^{|A'|+k+1} n^{3/2} \log(n))$ time.*

Proof sketch. We sketch the algorithm here; full details are given in [25]. We start with a maximum-cardinality matching M_S, which can be computed using the classical Hopcroft-Karp algorithm [32], to determine the target size of our matching. Recall that we aim to find a maximum-cardinality matching M_P of I such that the agents in A' are collectively involved in at most k blocking pairs in M_P. We will denote such blocking pairs by $bp(M_P)|_{A'}$. Clearly, every agent in $a_i \in A'$ is either matched to one of the at most $d_{\max}$ many agents on their preference list or remains unmatched. Thus, there are $(d_{\max} + 1)^{|A'|} = O(d_{\max}^{|A'|})$ possible combinations of choices of partner, including the possibility of being unmatched, for the agents in A'. We can discard any combinations of choices that are not matchings, and we denote the remaining set of candidate matchings by $\mathcal{M}_C$.

Next, we consider each candidate matching sequentially and aim to extend it to a solution. Let us fix a candidate matching $M_C \in \mathcal{M}_C$. Now, we consider each possible set of blocking pairs involving some agent in A' of size at most k. There are at most $O((|A'| d_{\max})^k)$ many such sets. Let us fix a candidate set of blocking pairs B. Clearly, no pair of agents can be simultaneously matched and blocking, so if $M_C \cap B \neq \varnothing$, then we reject this configuration (M_C, B).

Then, for every agent $a_i \in A'$ and every $a_r \in A$ such that $a_r \succ_i M_C(a_i)$ and $\{a_i, a_r\} \notin B$, a_r must end up with a partner better than a_i (according to $\succ_r$). Thus, we truncate $\succ_r$, discarding every agent worse than and including a_i (if multiple agents in A' satisfy this criterion with respect to a_r, we truncate at the best-ranked such agent on $\succ_r$). We must keep track of these agents a_r to later verify that they are matched, so we denote the set of such agents by Q. If, after carrying out these truncations, some agent $a_i \in A'$

has a partner in M_C that is no longer in their preference list (or vice versa), we can reject this configuration (M_C, B). Otherwise, we look for a maximum matching among the agents not yet matched in M_C such that every agent in Q is matched.

We do this as follows: we construct a maximum-weight matching instance (G, w) (where G is a graph and w is a weight function from the set of edges to the positive integers) consisting of all agents $A \setminus A(M_C)$ (where $A(M_C)$ denotes the set of agents matched in M_C) as vertices, and all acceptable matches among these agents as edges. Furthermore, we construct w as follows: for every acceptable pair of agents a_r, a_s in the instance, we let $w(\{a_r, a_s\}) = n + |\{a_r, a_s\} \cap Q|$, i.e., we assign a weight of at least n to every edge, plus one extra point for each endpoint in Q. We can construct (G, w) in linear time in the size of I, so in $O(d_{\max} n)$ time. Now we find a maximum-weight matching M_{mw} for (bipartite) instance (G, w) using the algorithm by Duan and Su [13], which runs in $O(d_{\max} n^{3/2} \log(n))$ time. M_{mw} is a maximum-cardinality matching of (G, w) in which every agent in Q is matched if and only if such a matching exists. Thus, if there exists an agent $a_r \in Q$ that is not matched in M_{mw}, we can reject (M_C, B). Otherwise, by construction, M_C and M_{mw} are disjoint, so if $|M_C \cup M_{mw}| = |M_S|$, we accept $M = M_C \cup M_{mw}$ as our solution. Otherwise, M_C does not extend to a solution with respect to B, and we reject (M_C, B).

Putting the complexity analysis together, we consider $O(d_{\max}^{|A'|})$ candidate matchings and, for every candidate matching, we consider $O((|A'|d_{\max})^k)$ sets of blocking pairs, leading to $O(|A'|^k d_{\max}^{|A'|+k})$ candidate configurations. For each configuration, we construct a maximum-weight matching instance in $O(d_{\max} n)$ time and compute a maximum-weight matching in $O(d_{\max} n^{3/2} \log(n))$ time. □

Corollary 3.4. *k-Preferential-AlmostStable-Max-SMI is in* FPT *with respect to* $(|A'|, d_{\max})$ *and in* XP *with respect to* $|A'|$. *Furthermore,* Preferential-AlmostStable-Max-SMI *is in* FPT *with respect to* $(|A'|, d_{\max})$ *and in* XP *with respect to* $(|A'|, OPT)$, *where OPT is the optimal value* $\min_{M \in \mathcal{M}^+} |\bigcup_{a_i \in A'} bp_{a_i}(M)|$.

Proof. From Theorem 3.3, we know that the problem is solvable in $O(|A'|^k d_{\max}^{|A'|+k+1} n^{3/2} \log(n))$ time. Given that k is part of the problem definition, not part of the input, we consider $k = O(1)$. Alternatively, we can justify this as follows: the optimal parameter value $|\sum_{a_i \in A'} bp_{a_i}(M)|$ is bounded from above by $|A'|d_{\max}$. Therefore, for any $k \geq |A'|d_{\max}$, the problem becomes trivial. Thus, we may assume without loss of generality that $k < |A'|d_{\max}$, in which case k does not need to be treated as a parameter when $|A'|, d_{\max}$ are fixed. Hence, the first FPT result follows directly from $O(|A'|^{O(1)} d_{\max}^{|A'|+O(1)} n^{O(1)})$. For the first XP result, notice that $d_{\max} \leq n - 1 < n$, so we can crudely upper-bound the complexity by $O(|A'|^{O(1)} n^{O(|A'|)})$ as required. Finally, to solve Preferential-Max-SMI, we can simply solve k-Preferential-Max-SMI with increasing values of k up until a solution is found, i.e., for instance (I, A'), we iterate $0 \leq k \leq \min_{M \in \mathcal{M}^+} |\bigcup_{a_i \in A'} bp_{a_i}(M)| \leq |A'|d_{\max}$ and find a solution in $O(|A'|d_{\max} |A'|^{|A'|d_{\max}} d_{\max}^{|A'|+|A'|d_{\max}+1} n^{O(1)}) = O(|A'|^{O(|A'|d_{\max})} d_{\max}^{O(|A'|d_{\max})} n^{O(1)})$ time. Furthermore, it holds that $O(k|A'|^k d_{\max}^{|A'|+k+1} n^{O(1)}) = O(k|A'|^{O(k)} n^{O(|A'|+k)})$. □

In [25], these results are generalised to Max-SRI problem variants and the minimisation of blocking A' agents. We also provide efficient algorithms for instances with preference lists of length at most 2.

3.2 Preferential Matchings in General Graphs

We now turn to the general case, where matchings may have arbitrary size and the acceptability graph need not be bipartite. We start by defining a restricted version of Preferential-AlmostStable-SRI as follows.

k-Preferential-AlmostStable-SRI
Input: SRI instance $I = (A, \succ)$ and a subset $A' \subseteq A$.
Output: A matching $M \in \mathcal{M}$ such that $|\bigcup_{a_i \in A'} bp_{a_i}(M)| \leq k$, if one exists.

Again, we will focus on the very restricted case of $k = 0$, where we require a matching in which no A' agent is in a blocking pair, and denote the problem of deciding whether there exists a solution to 0-Preferential-AlmostStable-SRI by 0-Preferential-AlmostStable-SRI-Dec.

We start with the following observation about a special class of instances to 0-Preferential-AlmostStable-SRI.

Theorem 3.5. *Let $(I = (A, \succ), A')$ be an instance with n agents of* 0-Preferential-AlmostStable-SRI. *If $G' = (A, E \setminus E[A \setminus A'])$ is bipartite, then we can solve* 0-Preferential-AlmostStable-SRI *in $O(n^2)$ time.*

Proof. Suppose that $G' = (A, E \setminus E[A \setminus A'])$ is bipartite. Then the instance $I' = (A, \succ')$ induced by the restricted acceptability graph G' is an SMI instance. Any SMI instance admits a stable matching M, and we can find one in $O(n^2)$ time [18]. Now suppose that M admits a blocking pair $\{a_r, a_s\} \in bp(M)$ with respect to instance I. Then, by construction of I', it must be the case that $\{a_r, a_s\} \in E[A \setminus A']$, in which case $\bigcup_{a_i \in A'} bp_{a_i}(M) = \varnothing$ as required. □

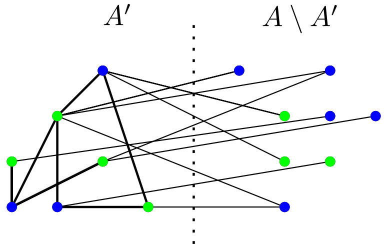

Figure 1: Illustration of restricted acceptability graph (colours indicate a possible bipartition)

An example of a restricted preference graph can be seen in Figure 1. Unfortunately, this structural insight does not generalise to an efficiently verifiable necessary and sufficient condition for 0-Preferential-AlmostStable-SRI-Dec, as there may be complex interactions among agents between and outside A'. The following result leaves little hope for such a condition.

Theorem 3.6. 0-Preferential-AlmostStable-SRI-Dec *is* NP-complete, *regardless of whether all preference lists are of length at most 5 or all preference lists are complete.*

Table 1: Overview of key complexity results. Here, d denotes the maximum preference list length (also denoted by $d_{\max}$ in other parts of the paper), and CL indicates complete preference lists. Our contributions are coloured in blue. We use $\top$ to indicate that the result holds trivially, and * to indicate that we introduced these problems in this paper.

	SRI	Max-SMI
MinBP ($\kappa = \min \lvert bp(M)\rvert$)	P ($d \leq 2$) [7], XP (κ) [1] NP-h ($d \leq 3$ and CL) [1, 7], W[1]-h (κ) [10]	P ($d \leq 2$ and CL) [8], XP (κ) [8] NP-h ($d \leq 3$) [8]
MinBA ($\kappa = \min \lvert ba(M)\rvert$)	P ($d \leq 2$) [25], XP (κ) [10] NP-h ($d \leq 5$ and CL [25]) [10], W[1]-h (κ) [10]	P ($d \leq 2$ and $d = C$) [8], XP (κ) [8] NP-h ($d \leq 3$) [8]
Minimax* ($\kappa = \min\max \lvert bp_i(M)\rvert$)	P ($d \leq 2$) [T2.3] NP-h ($\kappa = 1, d \leq 10$ and CL) [T2.1-2.1]	P ($d \leq 2$ and CL) [T2.6,×] NP-h ($\kappa = 1, d \leq 3$) [T2.4]
Preferential* ($\kappa_1 = \lvert A'\rvert$, $\kappa_2 = \min \lvert \bigcup_{a_i \in A'} bp_i(M)\rvert$)	P ($d \leq 2$) [25] FPT ((d, κ_1)) [C3.9] XP ((κ_1, κ_2)) [C3.9] NP-h ($\kappa_2 = 0, d \leq 5$ and CL) [T3.6]	P ($d \leq 2$ and CL) [[25],$\top$] FPT ((d, κ_1)) [C3.4] XP ((κ_1, κ_2)) [C3.4] NP-h ($\kappa_2 = 0, d \leq 3$) [T3.1]

Corollary 3.7. *k-Preferential-AlmostStable-SRI is* NP-hard, *even when $k = 0$ and regardless of whether preference lists are of length at most 5 or complete. Thus,* Preferential-AlmostStable-SRI *is* para-NP-hard *with respect to* $\kappa = \min_{M \in \mathcal{M}} \lvert \bigcup_{a_i \in A'} bp_{a_i}(M)\rvert$.

Having established intractability even under strong restrictions, we next identify a class of tractable special cases by adapting the FPT algorithm from the bipartite case. Notice that, although we consider more general preference systems, the fact that we do not require a maximum-cardinality matching allows us to arrive at a complexity independent of the total number of agents n. This is because any agent not in A' may be left unmatched, unless they block with an agent in A'. We performed the analysis assuming that $d_{\max} \ll n$, which leads to a larger exponent for $d_{\max}$, but it is easy to refine the analysis to a smaller exponent in exchange for a sub-quadratic dependency on n when this is not the case.

Theorem 3.8. *Let $(I = (A, \succ), A')$ be an instance of problem k-Preferential-AlmostStable-SRI with maximum preference list length $d_{\max}$. Then k-Preferential-AlmostStable-SRI is solvable in $O(\lvert A'\rvert^{k+3/2} d_{\max}^{\lvert A'\rvert+k+4})$ time.*

Proof sketch. The high-level approach is the same as in the proof of Theorem 3.3. We highlight critical differences below and refer to [25] for full details. When searching for a matching of agents in Q, we can drop the maximum-cardinality requirement. This simplifies the construction of (G, w): G now consists of all agents that are distance at most 2 away from some agent in A' (and not yet matched in M_C), and all acceptable matches among these agents as edges. We construct w such that, for every acceptable a_r, a_s, $w(\{a_r, a_s\}) = \lvert\{a_r, a_s\} \cap Q\rvert$, i.e., we assign a weight of 1 for every endpoint in Q. We can construct (G, w) in linear time in the graph. By the preference list lengths, there are at most $\lvert A'\rvert d_{\max}^2$ many agents with distance at most 2 away from some agent in A', and G has at most $\lvert A'\rvert d_{\max}^3$ edges. Thus, we can create (G, w) in $O(\lvert A'\rvert d_{\max}^3)$ time. Now we find a maximum-weight matching M_{mw} using the algorithm by Huang and Kavitha [33], which runs in $O(\lvert A'\rvert^{3/2} d_{\max}^4)$ time.

If there exists an agent $a_r \in Q$ that is not matched in M_{mw}, we can reject this configuration (M_C, B). Otherwise, by construction, M_C and M_{mw} are disjoint, so we can accept $M = M_C \cup M_{mw}$ as our solution. □

Again, due to a similar argument as in Corollary 3.4, we can immediately note the following complexity classification.

Corollary 3.9. *k-Preferential-AlmostStable-SRI is in* FPT *with respect to $(\lvert A'\rvert, d_{\max})$ and in* XP *with respect to $\lvert A'\rvert$. Furthermore, the problem* Preferential-AlmostStable-SRI *is in* FPT *with respect to $(\lvert A'\rvert, d_{\max})$ and in* XP *with respect to $(\lvert A'\rvert, OPT)$, where $OPT = \min_{M \in \mathcal{M}} \lvert \bigcup_{a_i \in A'} bp_{a_i}(M)\rvert$.*

In [25], we extend these results to the minimisation of blocking A' agents too. We also provide efficient algorithms for instances with preference lists of length at most 2.

4 CONCLUSION

We studied two natural classes of optimisation problems arising from stable matching theory, initiating the study of both minimax and preferential notions of almost-stability. Across bipartite and general settings, our results, summarised in Table 1, reveal a sharp contrast between strong intractability and tractable special cases, such as when preference lists are very short or there are few preferential agents. Beyond their algorithmic significance, these findings also carry implications for multi-agent systems. In many applications where centralised coordination mechanisms are sought, solutions need to be robust to agent incentives to deviate. We presented two new approaches and characterised the computational possibilities and limits of achieving such robustness.

Many interesting directions remain open. For instance, are there efficient approximation algorithms for Minimax-AlmostStable-Max-SMI? Furthermore, are Preferential-AlmostStable-Max-SMI and Preferential-AlmostStable-SRI in FPT with respect to just $\lvert A'\rvert$? We also conjecture that Minimax-AlmostStable-SRI cannot be solved in polynomial time even when $d_{\max} \leq 8$ (under standard complexity-theoretic assumptions), but that the optimal value of the problem under this restriction is at most 1. This would suggest a total-search type problem. What is the complexity of this problem under these restrictions?

ACKNOWLEDGMENTS

Frederik Glitzner is supported by a Minerva Scholarship from the School of Computing Science, University of Glasgow. We would like to thank the reviewers for their helpful comments.

REFERENCES

[1] David J. Abraham, Péter Biro, and David Manlove. 2006. "Almost stable" matchings in the roommates problem. *Proceedings of WAOA 2005* 3879 (2006), 1–14. https://doi.org/10.1007/11671411_1

[2] David J. Abraham, Ariel Levavi, David Manlove, and Gregg O'Malley. 2007. The stable roommates problem with globally-ranked pairs. *Proceedings of WINE 2007* 4858 (2007), 431–444. https://doi.org/10.1007/978-3-540-77105-0_48

[3] Ata Atay, Ana Mauleon, and Vincent Vannetelbosch. 2021. A bargaining set for roommate problems. *Journal of Mathematical Economics* 94 (2021), 102465. https://doi.org/10.1016/j.jmateco.2020.102465

[4] Mohammed W. Baidas, Zainab Bahbahani, and Emad Alsusa. 2019. User Association and Channel Assignment in Downlink Multi-Cell NOMA Networks: A Matching-Theoretic Approach. *EURASIP J. Wirel. Commun. Netw.* 2019, 1 (12 2019), 1–21. https://doi.org/10.1186/s13638-019-1528-8

[5] Dimitris Bertsimas and Carol Gao. 2025. Should We Relax Stability in Matching Markets? arXiv:2509.14475 [math.OC] https://arxiv.org/abs/2509.14475

[6] Florian M. Biermann. 2011. *A Measure to Compare Matchings in Marriage Markets.* Technical Report 592. Fondazione Eni Enrico Mattei. https://services.bepress.com/feem/paper592 Also published as Working Papers 005-11, International School of Economics at TSU.

[7] Péter Biró, David Manlove, and Eric J. McDermid. 2012. "Almost stable" matchings in the Roommates problem with bounded preference lists. *Theoretical Computer Science* 432 (5 2012), 10–20. https://doi.org/10.1016/J.TCS.2012.01.022

[8] Péter Biró, David Manlove, and Shubham Mittal. 2010. Size versus stability in the marriage problem. *Theoretical Computer Science* 411 (3 2010), 1828–1841. Issue 16-18. https://doi.org/10.1016/J.TCS.2010.02.003

[9] Jiehua Chen. 2019. Computational Complexity of Stable Marriage and Stable Roommates and Their Variants. arXiv:1904.08196 [cs.GT]

[10] Jiehua Chen, Danny Hermelin, Manuel Sorge, and Harel Yedidsion. 2018. How hard is it to satisfy (almost) all roommates? *Proceedings of ICALP 2018* 107 (7 2018), 35:1–35:15. https://doi.org/10.4230/LIPIcs.ICALP.2018.35

[11] Jiehua Chen, Sanjukta Roy, and Sofia Simola. 2025. FPT-Approximability of Stable Matching Problems. arXiv:2508.10129 [cs.GT] https://arxiv.org/abs/2508.10129

[12] Gergely Csáji, David Manlove, Iain McBride, and James Trimble. 2024. Couples Can Be Tractable: New Algorithms and Hardness Results for the Hospitals/Residents Problem with Couples. In *Proceedings of IJCAI 2024*, Kate Larson (Ed.). IJCAI, Jeju, South Korea, 2731–2739. https://doi.org/10.24963/ijcai.2024/302

[13] Ran Duan and Hsin-Hao Su. 2012. A scaling algorithm for maximum weight matching in bipartite graphs. In *Proceedings of SODA 2012*. Society for Industrial and Applied Mathematics, USA, 1413–1424.

[14] Kimmo Eriksson and Olle Häggström. 2008. Instability of Matchings in Decentralized Markets with Various Preference Structures. *International Journal of Game Theory* 36, 3 (March 2008), 409–420. https://doi.org/10.1007/s00182-007-0081-6

[15] Yuri Faenza, Telikepalli Kavitha, Vladlena Powers, and Xingyu Zhang. 2019. Popular matchings and limits to tractability. In *Proceedings of SODA 2019* (San Diego, California). Society for Industrial and Applied Mathematics, USA, 2790–2809.

[16] Patrik Floréen, Petteri Kaski, Valentin Polishchuk, and Jukka Suomela. 2010. Almost Stable Matchings by Truncating the Gale–Shapley Algorithm. *Algorithmica* 58, 1 (2010), 102–118. https://doi.org/10.1007/s00453-009-9353-9

[17] Jörg Flum and Martin Grohe. 2006. *Parameterized Complexity Theory.* Springer, Berlin, Germany. https://doi.org/10.1007/3-540-29953-X

[18] David Gale and Lloyd S. Shapley. 1962. College Admissions and the Stability of Marriage. *The American Mathematical Monthly* 69 (1 1962), 9. Issue 1. https://doi.org/10.2307/2312726

[19] David Gale and Marilda Sotomayor. 1985. Some remarks on the stable matching problem. *Discrete Applied Mathematics* 11, 3 (1985), 223–232. https://doi.org/10.1016/0166-218X(85)90074-5

[20] Frederik Glitzner. 2026. Near-Feasible Stable Matchings: Incentives and Optimality. arXiv:2602.10851 [cs.GT] https://arxiv.org/abs/2602.10851

[21] Frederik Glitzner. 2026. Near-Feasible Stable Matchings: Incentives and Optimality. In *Proceedings of AAMAS 2026*. IFAAMAS, Paphos, Cyprus, 3. https://doi.org/10.65109/RZEM3915

[22] Frederik Glitzner and David Manlove. 2025. Perspectives on Unsolvability in the Roommates Problem. arXiv:2505.06717 [cs.GT] https://arxiv.org/abs/2505.06717

[23] Frederik Glitzner and David Manlove. 2025. Unsolvability and Beyond in Many-to-Many Non-bipartite Stable Matching. In *Proceedings of SAGT 2025*. Springer Nature Switzerland, Cham, Switzerland, 267–285. https://doi.org/10.1007/978-3-032-03639-1_15

[24] Frederik Glitzner and David Manlove. 2026. A Minimax Perspective on Almost-Stable Matchings. arXiv:2601.14195 [cs.GT] https://arxiv.org/abs/2601.14195

[25] Frederik Glitzner and David Manlove. 2026. Stable Matching with Deviators and Conformists. arXiv:2601.18573 [cs.GT] https://arxiv.org/abs/2601.18573

[26] Frederik Glitzner and David Manlove. 2026. Structural and Algorithmic Results for Stable Cycles and Partitions in the Roommates Problem. *ACM Transactions on Economics and Computation* (1 2026). https://doi.org/10.1145/3789257

[27] Sushmita Gupta, Pallavi Jain, Sanjukta Roy, Saket Saurabh, and Meirav Zehavi. 2020. On the (Parameterized) Complexity of Almost Stable Marriage. In *Proceedings of FSTTCS 2020*, Vol. 182. Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 24:1–24:17. https://doi.org/10.4230/LIPIcs.FSTTCS.2020.24

[28] Sushmita Gupta, Pranabendu Misra, Saket Saurabh, and Meirav Zehavi. 2021. Popular Matching in Roommates Setting Is NP-hard. *ACM Transactions on Computation Theory (TOCT)* 13 (3 2021), 1–20. Issue 2. https://doi.org/10.1145/3442354

[29] Daniel Gusfield and Robert Irving. 1989. *The Stable Marriage problem: Structure and Algorithms.* MIT Press, Cambridge, MA, USA.

[30] Koki Hamada, Kazuo Iwama, and Shuichi Miyazaki. 2009. An improved approximation lower bound for finding almost stable maximum matchings. *Inform. Process. Lett.* 109 (8 2009), 1036–1040. Issue 18. https://doi.org/10.1016/J.IPL.2009.06.008

[31] P. Jean-Jacques Herings and Yu Zhou. 2025. Harmonious Equilibria in Roommate Problems. https://doi.org/10.2139/ssrn.5144979

[32] John E. Hopcroft and Richard M. Karp. 1973. An $n^{5/2}$ Algorithm for Maximum Matchings in Bipartite Graphs. *SIAM J. Comput.* 2, 4 (1973), 225–231. https://doi.org/10.1137/0202019

[33] Chien-Chung Huang and Telikepalli Kavitha. 2012. Efficient algorithms for maximum weight matchings in general graphs with small edge weights. In *Proceedings of SODA 2012*. Society for Industrial and Applied Mathematics, USA, 1400–1412.

[34] Robert W. Irving. 1985. An efficient algorithm for the "stable roommates" problem. *Journal of Algorithms* 6 (12 1985), 577–595. Issue 4. https://doi.org/10.1016/0196-6774(85)90033-1

[35] Robert W Irving. 1998. Matching medical students to pairs of hospitals: a new variation on a well-known theme. In *Proceedings of ESA 1998*. Springer, Berlin, Germany, 381–392.

[36] David Manlove. 2013. *Algorithmics of Matching Under Preferences.* Series on Theoretical Computer Science, Vol. 2. World Scientific, Singapore. https://doi.org/10.1142/8591

[37] David Manlove, Iain McBride, and James Trimble. 2017. "Almost-stable" matchings in the Hospitals / Residents problem with Couples. *Constraints* 22 (1 2017), 50–72. Issue 1. https://doi.org/10.1007/S10601-016-9249-7

[38] Rafail Ostrovsky and Will Rosenbaum. 2015. Fast Distributed Almost Stable Matchings. In *Proceedings of PODC 2015*. ACM, USA, 101–108. https://doi.org/10.1145/2767386.2767424

[39] Jimmy J.M. Tan. 1991. Stable matchings and stable partitions. *International Journal of Computer Mathematics* 39 (1 1991), 11–20. Issue 1-2. https://doi.org/10.1080/00207169108803975

[40] Elise Vandomme, Yves Crama, and Marie Baratto. 2025. Locally stable matchings for the roommates problem. https://orbi.uliege.be/bitstream/2268/330279/1/Local_stability_Complexity_ORBI.pdf