UvA-DARE (Digital Academic Repository)

Stable Matching with Uncertain Linear Preferences

Aziz, H.; Biró, P.; Gaspers, S.; de Haan, R.; Mattei, N.; Rastegari, B.

DOI

[10.1007/s00453-019-00650-0](https://doi.org/10.1007/s00453-019-00650-0)

Publication date

2020

Document Version

Final published version

Published in

Algorithmica

License

CC BY

[Link to publication](#)

Citation for published version (APA):

Aziz, H., Biró, P., Gaspers, S., de Haan, R., Mattei, N., & Rastegari, B. (2020). Stable Matching with Uncertain Linear Preferences. *Algorithmica*, 82(5), 1410–1433. <https://doi.org/10.1007/s00453-019-00650-0>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Stable Matching with Uncertain Linear Preferences

Haris Aziz^{1,2} · Péter Biró^{3,4} · Serge Gaspers¹ · Ronald de Haan⁵ ·
Nicholas Mattei⁶ · Baharak Rastegari⁷

Received: 1 March 2019 / Accepted: 29 October 2019 / Published online: 14 November 2019
© The Author(s) 2019

Abstract

We consider the two-sided stable matching setting in which there may be uncertainty about the agents' preferences due to limited information or communication. We consider three models of uncertainty: (1) lottery model—for each agent, there is a probability distribution over linear preferences, (2) compact indifference model—for each agent, a weak preference order is specified and each linear order compatible with the weak order is equally likely and (3) joint probability model—there is a lottery over preference profiles. For each of the models, we study the computational complexity of computing the stability probability of a given matching as well as finding a matching with the highest probability of being stable. We also examine more restricted problems such as deciding whether a certainly stable matching exists. We find a rich complexity landscape for these problems, indicating that the form uncertainty takes is significant.

Keywords Stable matchings · Stable marriage problem · Uncertain preferences · NP-hard problems · Polynomial-time algorithms

1 Introduction

We consider a *Stable Marriage problem (SM)* in which there is a set of men and a set of women. Each man has a linear order over the women, and each woman has a linear order over the men. For the purpose of this paper we assume that the preference lists are complete, i.e., each agent finds each member of the opposite side acceptable.¹ In the stable marriage problem, the goal is to compute a *stable matching*; a matching where no two agents prefer to be matched to each other rather than to be matched to their current partners. Unlike most of the literature on stable matching problems [13,19,27],

¹ We note that the complexity of all problems that we study is the same for complete and incomplete lists, where non-listed agents are deemed unacceptable—see Proposition 1 in Sect. 3.1.

A preliminary version of this paper has appeared in the proceedings of the 9th International Symposium on Algorithmic Game Theory (SAGT 2016) [2]. The majority of the claims appear without a proof or with only a short proof sketch in Aziz et al. [2].

we assume that men and women may have uncertainty in their preferences which can be captured by various probabilistic uncertainty models. We focus on *linear models* in which each possible deterministic preference profile is a set of linear orders.

Uncertainty in preferences could arise for a number of reasons, both practical or epistemological. For example, an agent could express a weak order because the agent did not invest enough time or effort to differentiate between potential matches and therefore one could assume that each linear extension of the weak order is equally likely; this maps to our *compact indifference model*. In many real applications, the ties are broken randomly with lotteries, e.g., in the school choice programs in New York and Boston as well as in centralized college admissions in Ireland. However, a central planner may also choose a matching that is optimal in some sense, without breaking the ties in the preference list. For instance, in Scotland they used to compute the maximum size (weakly) stable matching to allocate residents to hospitals [19]. Alternatively, there may be a cost associated with eliciting preferences from the agents, so a central planner may want to only obtain and provide a recommendation based on a subset of the complete orders [9,24]. Another cause of uncertainty could be that agents are certain about preferences over other agents according to specific criteria but there may be a probability distribution corresponding to the weight placed on different criteria. This motivates our lottery and joint probability models. In the *lottery model*, the agents have *independent* probabilities over possible linear orders, i.e., each linear order may correspond to a different criterion. In the *joint probability model*, the probability distribution is over possible preference profiles.

Uncertainty in preferences has already been studied in voting [15] and for cooperative games [18]. Ehlers and Massó [10] considers many-to-one matching markets under a Bayesian setting. Similarly, in auction theory, it is standard to examine Bayesian settings in which there is a probability distribution over the types of agents. In the two-sided matching setting, when preferences are uncertain, a natural solution is finding a matching which has the highest probability of being stable. We consider this fundamental computational problem and its related variants for the three uncertainty models discussed above.

To illustrate the problem, we describe a simple example with four agents. We write $b \succ_a c$ to say that agent a prefers b to c and assume the lottery model.

Example 1 We have two men m_1 and m_2 and two women w_1 and w_2 . Each agent assigns a probability to each strict preference ordering as follows. (i) $p(w_1 \succ_{m_1} w_2) = 0.4$ and $p(w_2 \succ_{m_1} w_1) = 0.6$ (ii) $p(w_1 \succ_{m_2} w_2) = 0.0$ and $p(w_2 \succ_{m_2} w_1) = 1.0$ (iii) $p(m_1 \succ_{w_1} m_2) = 1.0$ and $p(m_2 \succ_{w_1} m_1) = 0.0$ (iv) $p(m_1 \succ_{w_2} m_2) = 0.8$ and $p(m_2 \succ_{w_2} m_1) = 0.2$.

This setting admits two matchings that are stable with positive probability: $\mu_1 = \{(m_1, w_1), (m_2, w_2)\}$ and $\mu_2 = \{(m_1, w_2), (m_2, w_1)\}$. Notice that if each agent submits the preference list that s/he finds most likely to be true, then the setting admits a unique stable matching that is μ_2 . The probability of μ_2 being stable, however, is 0.48 whereas the probability of μ_1 being stable is 0.52 (Tables 1, 2).

Table 1 Pairwise probabilities for the agents in Example 1

Men	Women
$m_1 \quad p(w_1 \succ_{m_1} w_2) = 0.4$ $p(w_2 \succ_{m_1} w_3) = 0.6$	$w_1 \quad p(m_1 \succ_{w_1} m_2) = 1.0$ $p(m_2 \succ_{w_1} m_3) = 0.0$
$m_2 \quad p(w_1 \succ_{m_2} w_2) = 0.0$ $p(w_2 \succ_{m_2} w_3) = 1.0$	$w_2 \quad p(m_1 \succ_{w_2} m_2) = 0.8$ $p(m_2 \succ_{w_2} m_3) = 0.2$

Table 2 Stability probability for each matching in Example 1

	Matching	Stability probability
μ_1	$\{(m_1, w_1), (m_2, w_2)\}$	0.52
μ_2	$\{(m_1, w_2), (m_2, w_1)\}$	0.48

1.1 Uncertainty Models

In this article, we consider three different uncertainty models which assume that agents have *linear* preferences. In related work we have explored similar computational questions when agents define their uncertainty over *pairwise* preferences [3].

- *Lottery Model* For each agent, we are given a probability distribution over strict preference lists.
- *Compact Indifference Model* Each agent reports a single weak preference list that allows for ties. Each complete linear order extension of this weak order is assumed to be equally likely.
- *Joint Probability Model* A probability distribution over preference profiles is specified.

Note that for the Lottery Model and the Joint Probability Model the representation of the input preferences can be exponentially large (in the number of agents). However, in settings where similar models of uncertainty are used, including resident matching [9] and voting [15], a limited amount of uncertainty (i.e. small supports) is commonly expected and observed in real-world data [21,22]. Consequently, we consider special cases when the uncertainty is bounded in certain natural ways, including the existence of only a small number of uncertain preferences and/or uncertainty on only one side of the market.

Observe that the compact indifference model can be represented as a lottery model. This is a special case of the lottery model in which each agent expresses a weak order over the candidates, similar to the SMT setting [13,19]. However, the lottery model representation can be exponentially larger than the compact indifference model; for an agent that is indifferent among n agents on the other side of the market, there are $n!$ possible linearly ordered preferences. The uncertainty models considered in the paper have further been examined in the context of Pareto optimal assignment of items to agents [4–6]. In a subsequent paper, Chen et al. [7] consider additional problems related to the joint probability model.

Table 3 Summary of results

Problems	Lottery model	Compact indifference	Joint probability
STABILITYPROBABILITY	#P-complete in P for all three models if 1 side is certain	?	in P
ISSTABILITYPROBABILITYNON- ZERO	NP-complete	in P	in P
ISSTABILITYPROBABILITYONE	in P	in P	in P
EXISTSPOSSIBLYSTABLEMATCHING	in P	in P	in P
EXISTSCERTAINLYSTABLEMATCHING	in P	in P	NP-complete
MATCHINGWITHHIGHESTSTABILITYPROB	? in P for all models if 1 side is certain and there is $O(1)$ number of uncertain agents	NP-hard	NP-hard

1.2 Computational Problems

Given a stable marriage setting where agents have uncertain preferences, various natural computational problems arise. Let *stability probability* denote the probability that a matching is stable. We then consider the following two natural problems for each of our uncertainty models.

- **STABILITYPROBABILITY** Given a matching and uncertain preferences of the agents, what is the stability probability of the matching?
- **MATCHINGWITHHIGHESTSTABILITYPROBABILITY** Given uncertain preferences of the agents, compute a matching with the highest stability probability.

We also consider two specific problems that are simpler than STABILITYPROBABILITY:

- **ISSTABILITYPROBABILITYNON- ZERO** For a given matching, is its stability probability non-zero?
- **ISSTABILITYPROBABILITYONE** For a given matching, is its stability probability one?

We additionally consider problems connected to, and more restricted than, MATCHINGWITHHIGHESTSTABILITYPROBABILITY

- **EXISTSCERTAINLYSTABLEMATCHING** Does there exist a matching that has stability probability one?
- **EXISTSPOSSIBLYSTABLEMATCHING** Does there exist a matching that has non-zero stability probability?

Note that EXISTSPOSSIBLYSTABLEMATCHING is straightforward to answer for any of the three uncertainty models we consider here, since there exists a stable matching for each deterministic preference profile that is a possible realization of the uncertain preferences.

1.3 Results

Table 3 summarizes our main findings. Note that the complexity of each problem is considered with respect to the input size, and that under the lottery and joint probability

models the size of the input could be exponential in n , namely $O(n! \cdot 2n)$ for the lottery model and $O((n!)^{2n})$ for the joint probability model, where n is the number of agents on either side of the market.

We point out that STABILITYPROBABILITY is #P-complete for the lottery model even when each agent has at most two possible preferences, but in P if one side has certain preferences. Additionally, we show that ISSTABILITYPROBABILITYNON-ZERO is in P for the lottery model if each agent has at most two possible preferences. Note that STABILITYPROBABILITY is open for the compact indifference model when both sides may be uncertain, and we also do not know the complexity of MATCHINGWITHHIGHESTSTABILITYPROBABILITY in the lottery model, except when only a constant number of agents are uncertain on the same side of the market.

2 Preliminaries

In the Stable Marriage problem, there are two sets of agents. Let M denote a set of n men and W a set of n women. We use the term *agents* when making statements that apply to both men and women, and the term *candidates* to refer to the agents on the opposite side of the market to that of an agent under consideration. Each agent has a linearly ordered preference over the candidates. An agent may be uncertain about his/her linear preference ordering. Let L denote the *uncertain preference profile* for all agents.

We say that a given uncertainty model is *independent* if any uncertain preference profile L under the model can be written as a product of uncertain preferences L_a for all agents a , where all L_a 's are independent. Note that the lottery and the compact indifference models are both independent, but the joint probability model is not.

A *matching* μ is a pairing of men and women such that each man is paired with at most one woman and vice versa; defining a list of (man, woman) pairs (m, w) . We use $\mu(m)$ to denote the woman w that is matched to m and $\mu(w)$ to denote the match for w . Assume that each agent prefers being matched to remaining unmatched. Given linearly ordered preferences, a matching is *stable* if there is no pair (m, w) not in μ where m prefers w to his partner in μ , i.e., $w \succ_m \mu(m)$, and vice versa. If such a pair exists, it constitutes a *blocking pair*; as the pair would prefer to defect and match with each other rather than stay with their partner in μ . A matching μ is a *complete* matching if all agents are matched in μ . A matching is *certainly stable* if it is stable with probability 1. For an instance $I = (M, W, L)$ and matching μ , let $p(\mu, I)$ denote the probability of μ being stable, and let $p_S(I) = \max\{p(\mu, I) : \mu \text{ is a matching in } I\}$, that is the maximum probability of a matching being stable for I .

The following extensions of SM will come in handy in proving our results. The *Stable Marriage Problem with Partially Ordered Lists (SMP)* is an extension of SM in which agents' preferences are partial orders over the candidates. An instance $I = (M, W, p)$ of SMP is given by a set of n men M , a set of n woman W , and the partial preference ordering profile of agents p where p_a denotes the partial order that represents the preferences of agent a . If for a given agent a and two candidates b and c we have that neither b is related to c nor c is related to b then a cannot compare b and c . The *Stable Marriage problem with Ties (SMT)* is a special case of SMP in which

incomparability is transitive and is interpreted as indifference. Therefore, in SMT each agent partitions the candidates into different ties (equivalence classes), is indifferent between the candidates in the same tie, and has strict preference ordering over the ties. In some practical settings some agents may find some candidates unacceptable and prefer to remain unmatched than to be matched to the unacceptable ones. *SMP with Incomplete Lists (SMPI)* and *SMT with Incomplete lists (SMTI)* captures these scenarios where each agent's partially ordered list contains only his/her acceptable candidates. Three stability criteria, including the one we have already defined, have been introduced in the literature to capture different degrees of stability for these richer domains. The weakest criterion, that is the one we have already defined, is (*weak*) *stability*. A matching is (*weakly*) *stable* in an instance of SMPI if there is no pair $(m, w) \notin \mu$ where m (strictly) prefers w to his partner in μ and vice versa. The strongest criterion, referred to by *super-stability*, is closely related to our notion of certain stability. A matching is *super-stable* in an instance of SMPI if there is no pair (m, w) not in μ where m either prefers w to his partner in μ or finds them incomparable, and vice versa. If such a pair exists, it constitutes a *very weakly blocking pair*. It is easy to observe (see, e.g., [17]) that μ is super-stable if and only if it is stable w.r.t. all linear extensions of the partially ordered lists.

We define the *certainly preferred* relation $\succ_a^{\text{cert}}$ for agent a . We write $b \succ_a^{\text{cert}} c$ if and only if agent a prefers b over c with probability 1. Based on the certainly preferred relation, we can define a dominance relation D : $D_m(w) = \{w\} \cup \{w' : w' \succ_m^{\text{cert}} w\}$; $D_w(m) = \{m\} \cup \{m' : m' \succ_w^{\text{cert}} m\}$. Based on the notion of the dominance relation, we present a useful characterization of certainly stable matchings for independent uncertainty models.

Lemma 1 *A matching μ is certainly stable for an independent uncertainty model if and only if for each pair (m, w) , $\mu(m) \in D_m(w)$ or $\mu(w) \in D_w(m)$.*

Proof Assume that for each pair (m, w) we have that $\mu(m) \in D_m(w)$ or $\mu(w) \in D_w(m)$ for a given matching μ . This implies that for each unmatched pair (m, w) it is the case that $\mu(m) \succ_m^{\text{cert}} w$ or $\mu(w) \succ_w^{\text{cert}} m$ and hence (m, w) has zero probability of forming a blocking pair. It thus follows that μ is certainly stable.

Assume that a matching μ is certainly stable. Then no pair (m, w) has non-zero probability of forming a blocking pair. This is only possible if the pair (m, w) is part of the matching or one of m and w have zero probability of preferring the partner in (m, w) over their current partner in μ . In either case, $\mu(m) \in D_m(w)$ or $\mu(w) \in D_w(m)$. $\square$

We point out that the certainly preferred relation can be computed in polynomial time for all three models studied in this paper.

3 General Results

In this section, we first show that the complexity of all problems that we study is the same for complete and incomplete lists. We then present some general results that apply to multiple uncertainty models. We show that EXISTSCERTAINLYSTABLEMATCHING can be solved in polynomial time for a class of independent uncertainty models that

includes lottery and compact indifference. We then prove that, when the number of uncertain agents is constant and one side of the market is certain, we can solve MATCHINGWITHHIGHESTSTABILITYPROBABILITY efficiently for each of the linear models.

3.1 The Case for Incomplete Lists

The claims in this section explain that our efficient algorithms described for the case of complete lists can be extended to incomplete lists. Additionally, our hardness proofs for incomplete lists can be modified to extend to complete lists. In fact, all our hardness reductions, except Theorem 9, are for complete lists and so they trivially extend to the case of incomplete lists.

In the case of complete lists, we assumed that we have an equal number of men and women and everybody finds all candidates acceptable. When we consider the problem with incomplete lists we mean that the sizes of the two sets are not necessarily the same and not all the candidates are acceptable to all agents. However, we assume that in all realization of the preference profiles the same candidates are acceptable, so we only randomize on the preferences over the acceptable partners.

Proposition 1 *The computational complexity of STABILITYPROBABILITY is the same for complete and incomplete lists.*

Proof We show that if I is an instance of a linear probabilistic model with incomplete lists and μ is a given matching for I then we can, in linear time, construct an extended instance I' with complete lists and a complete matching μ' for I' such that the stability probability of μ under I , $p(\mu, I)$, is equal to the stability probability of μ' under I' , $p(\mu', I')$.

Assume, without loss of generality, that $|M| \geq |W|$. From I we create an extended instance I' with sets M and W' in the following way. First we ensure that $|M| = |W'|$ by adding enough agents to the woman's side of the market if necessary. Then we complete the preference lists of each agent by adding the previously unacceptable or nonexistent candidates to the end of her/his list according to a predetermined order, e.g., by the indices of the agents. Suppose now that μ is a matching in I and X is the set of matched men in M , whilst $\mu(X) = Y$. Let E denote the set of acceptable pairs in I . We assume that there is no pair $(m, w) \in (M \setminus X) \times (W \setminus Y)$ belonging to E , since in this case this pair would certainly block μ in I , thus $p(\mu, I) = 0$ trivially. Let us now extend μ to another matching μ' in I' by appending to μ the unique stable matching for the subinstance restricted to the unmatched agents. Namely, let μ_u be the stable matching that matches $M \setminus X$ to $W' \setminus Y$ in such a way that the k th pair contains the k th man and the k th woman from $M \setminus X$ and $W' \setminus Y$, respectively according to their indices, and let $\mu' = \mu \cup \mu_u$.

We claim that $p(\mu, I) = p(\mu', I')$. This is because there is no blocking pair in $(M \setminus X) \times (W' \setminus Y)$, and any other pair is blocking for some preference profile in I if and only if it is blocking for the corresponding preference profile in I' . To verify the latter statement, first consider a pair $(m, w) \in X \times Y$. If $(m, w) \in E$ then the partners of m and w are the same in μ and μ' and the ranks of m are the same in the corresponding preference profiles of w , and vice versa. If $(m, w) \notin E$ then this pair cannot block

in any preference profile. Similarly, let us consider a pair (m, w) when exactly one agent is matched in μ , say, m , i.e., $(m, w) \in M \times (W \setminus Y)$. If $(m, w) \in E$ then m has the same partner in μ and μ' , whilst w is unmatched in μ and has a previously unacceptable partner in μ' , namely $\mu'(w)$. Whenever this pair blocks for I then it also blocks for I' , since the rank of w is the same in the corresponding preference profiles of m , and w prefers m to $\mu'(w)$ in all profiles. Finally, if $(m, w) \notin E$ then this pair cannot block in any preference profile, as before. $\square$

The above proof implies the following Corollary.

Corollary 1 *The computational complexity of ISSTABILITYPROBABILITYNON-ZERO is the same for complete and incomplete lists, and the same holds for ISSTABILITYPROBABILITYONE.*

Proposition 2 *The computational complexity of MATCHINGWITHHIGHESTSTABILITYPROBABILITY is the same for complete and incomplete lists.*

Proof Suppose that we are given an instance of a linear probabilistic model with incomplete lists I , and we consider the extended instance I' with complete preference lists, as described in the proof of Proposition 1. We show that if μ' is one of the most stable complete matchings in I' then its restriction to E , μ , is one of the most stable matching for I .

First we note that $p(\mu, I) \geq p(\mu', I')$, since any pair in E that is blocking for μ under some preference profile in I is also blocking for μ' under the corresponding extended preference profile in I' . Suppose now for a contradiction that there is another matching ν for I such that $p(\nu, I) > p(\mu, I)$. But then, for its natural extension ν' we have $p(\nu, I) = p(\nu', I')$ by Proposition 1, contradicting with the maximum stability of μ' . $\square$

The above proof implies the following Corollary.

Corollary 2 *The computational complexity of EXISTSPOSSIBLYSTABLEMATCHING is the same for complete and incomplete lists, and the same holds for EXISTSCERTAINLYSTABLEMATCHING.*

3.2 An Efficient Algorithm for the Lottery and Compact Indifference Models

As pointed out earlier, certainly stable matchings are closely related to the notion of super-stable matchings [13, 16]. Deciding whether an instance of SMPI admits a super-stable matching or not can be done in polynomial time using algorithm SUPER-SMP in [25]. We next show that for a class of independent uncertainty models that includes the lottery and compact indifference, we can solve EXISTS CERTAINLY STABLE MATCHING in polynomial time by a straightforward reduction to the problem of deciding whether an instance of SMP admits a super-stable matching or not.

Theorem 1 *For any independent uncertainty model in which the certainly preferred relation is transitive and can be computed in polynomial time, EXISTS CERTAINLY STABLE MATCHING can be solved in polynomial time.*

Proof We prove this by reducing EXISTSCERTAINLYSTABLEMATCHING to the problem of deciding whether an instance of SMP admits a super-stable matching or not. Let $I = (M, W, L)$ be an instance of EXISTSCERTAINLYSTABLEMATCHING under an independent uncertainty model in which the certainly preferred relation is transitive and can be computed in polynomial time. We construct in polynomial time an instance $I' = (M, W, p)$ of SMP, where p is the agents' partial preference ordering profile, as follows. The set of men and women are unchanged. To create the partial preference ordering p_a for each agent a we do the following. Without loss of generality, assume that a is a man m . For every pair of women w_1 and w_2 (i) if $w_1 \succ_m^{\text{cert}} w_2$ then $(w_1, w_2) \in p_m$, denoting that m (strictly) prefers w_1 to w_2 in I' , (ii) if $w_2 \succ_m^{\text{cert}} w_1$ then $(w_2, w_1) \in p_m$, denoting that m (strictly) prefers w_2 to w_1 in I' . We claim that I' admits a super-stable matching if and only if I admits a certainly stable matching. It follows from the definition of super-stability that a matching μ in I' is super-stable if and only if there is no unmatched pair (m, w) where $(\mu(m), w) \notin p_m$ and $(\mu(w), m) \notin p_w$.

($\Leftarrow$) We first prove that if I' admits a super-stable matching μ then μ is certainly stable in I . Assume, for a contradiction, that μ is not certainly stable. It then follows Lemma 1 that $\mu(m) \notin D_m(w)$ and $\mu(w) \notin D_w(m)$, implying that $\mu(m) \not\succ_m^{\text{cert}} w$ and $\mu(w) \not\succ_w^{\text{cert}} m$, and thus $(\mu(m), w) \notin p_m$ and $(\mu(w), m) \notin p_w$. Therefore μ is not super-stable in I' , a contradiction.

($\Rightarrow$) Now we prove that if I admits a certainly stable matching μ then μ is super-stable in I' . Assume, for a contradiction, that μ is not super-stable in I' . Therefore there exists an unmatched pair (m, w) where $(\mu(m), w) \notin p_m$ and $(\mu(w), m) \notin p_w$, implying that $\mu(m) \not\succ_m^{\text{cert}} w$ and $\mu(w) \not\succ_w^{\text{cert}} m$. The latter statement, coupled with the fact that m and w are not matched together, implies that $\mu(m) \notin D_m(w)$ and $\mu(w) \notin D_w(m)$. Thus, by Lemma 1, μ is not certainly stable in I , a contradiction. $\square$

3.3 An Efficient Algorithm for the Case with a Constant Number of Uncertain Agents

Theorem 2 *When the number of uncertain agents is constant and one side of the market is certain, then MATCHINGWITHHIGHESTSTABILITYPROBABILITY is polynomial-time solvable for each of the linear models.*

Proof Let $I = (M, W, L)$ be an instance of MATCHINGWITHHIGHESTSTABILITYPROBABILITY and assume, without loss of generality, that uncertain agents are all men. Let $X \subseteq M$ be the set of uncertain agents with $|X| = k$ for a constant k . We consider all the possible matchings between X and W ; note that their total number is $K = n(n-1) \dots (n-k+1)$. Let $\mu_i, i \in \{1 \dots K\}$, be such a matching. The main idea of the proof is to show that there exists an extension of μ_i to $M \cup W$, which we denote by μ_i^* , that has stability probability at least as high as any other extension of μ_i . Furthermore, we can compute μ_i^* in polynomial time. Therefore, in order to compute a matching that has the highest stability probability, it is enough to generate μ_i^* 's, compute their stability probabilities, compare them and select the one with the highest stability probability. The total number of μ_i 's and hence μ_i^* 's is K and hence polynomial in n , we can compute each μ_i^* in polynomial time (as we will see later

in this proof) and computing the stability probability of a given μ_i^* can be done in polynomial time since all uncertain agents are on one side of the market (see Theorem 3 in Section 4, Theorem 8 in Section 5 and Theorem 10 in Section 6).

Take a matching μ_i between sets X and W . Let $Y = \mu_i(X)$ (i.e., the partners of X in W) and let $M' = M \setminus X$ and $W' = W \setminus Y$. Recall that all agents in $M' \cup W$ are certain. First, we compute the man-optimal matching μ_i^M for the sub-instance I' on $M' \cup W'$, that can be done efficiently by the Gale-Shapley algorithm [11]. Consider the matching $\mu'_i = \mu_i \cup \mu_i^M$ in I . If μ'_i admits a blocking pair (m', w) involving some (certain) agents $m' \in M'$ and $w \in Y$ (that we will refer to as a BP_Type1 blocking pair) then we can conclude that any extension of μ_i to a matching in I will have zero probability of being stable. This is because any extension of μ_i that has a positive probability of being stable in I must also be stable for the sub-instance I' . If (m', w) is a blocking pair for μ'_i then it also blocks any other extension of μ_i in I , since in all extensions w has the same partner and m' cannot have a better partner than in μ_i^M . Therefore we can exclude the extensions of μ_i from further consideration in this case.

Suppose now that $\mu'_i = \mu_i \cup \mu_i^M$ admits no BP_Type1 blocking pair. We truncate the preference lists of men in M' in the following way. For each man $m' \in M'$ we remove from his preference list all the women $w' \in W'$ that m' prefers less than some woman $w \in Y$ who prefers m' to her partner in μ_i . That is, we remove w' from the list of m' if there exists $w \in Y$ such that $w \succ_{m'} w'$ and $m' \succ_w \mu_i(w)$. To do this, it is enough to identify the first (i.e. the highest ranked) woman w in the preference list of m' such that $w \in Y$ and $m' \succ_w \mu_i(w)$, and then remove from the preference list of m' all women $w' \in W'$ that appear after w . Let us denote the sub-instance for $M' \cup W'$ with the truncated lists as I_i^r . Now we compute the woman-optimal matching μ_i^W in I_i^r . Let $\mu_i^* = \mu_i \cup \mu_i^W$. We claim that μ_i^* is a complete matching in I and is stable for the certain agents; that is, no blocking pair (m, w) exists where both m and w are certain agents. To see this, first note that since μ'_i admits no BP_Type1 blocking pair hence $\mu_i^M(m')$ remains in the truncated preference list of all $m' \in M'$. Therefore, the man-optimal matching for I_i^r is the same as the man-optimal matching for I' , μ_i^M , implying that a complete stable matching exists in I_i^r and that each $m' \in M'$ is either matched to the same woman in both μ_i^M and μ_i^W or prefers $\mu_i^M(m')$ to $\mu_i^W(m')$. As all stable matchings in I_i^r are of the same size (by the Rural-Hospital Theorem, see e.g. [26]), hence all agent in I_i^r are matched in μ_i^W and therefore μ_i^* is a complete matching for I . Now assume, for a contradiction, that μ_i^* is not stable for the certain agents and hence admits a blocking pair (m^*, w^*) where $m^* \in M'$ and $w^* \in W$. Note that $w^* \notin W'$ as otherwise (m^*, w^*) blocks μ_i^W in I_i^r , a contradiction. Therefore, $w^* \in Y$. But then we have that $w^* \in Y$ prefers $m^* \in M'$ to $\mu_i(w^*)$ and also that m^* prefers w^* to $\mu_i^W(m^*)$, implying that by the construction of the truncated preference lists of men in M' , woman $\mu_i^W(m^*)$ must have been removed from m^* 's preference list, a contradiction. Therefore μ_i^* is a complete matching in I and is stable for certain agents.

Finally, we show that for any matching μ_i^+ that is an extension of μ_i to I , the stability probability of μ_i^+ is less than, or equal to, the stability probability of μ_i^* . If μ_i^+ is not stable for the certain agents then μ_i^+ has zero probability of being stable, thus the statement holds. Otherwise, μ_i^+ is stable for the certain agents and hence for

any pair (m, w) that has a positive probability of being a blocking pair it must be that $m \in X$ and $w \in W'$. Moreover, the projection of μ_i^+ onto I_i^r must be stable in I_i^r , and as each woman $w \in W'$ is matched in μ_i^* to her optimal stable partner under I_i^r thus w either prefers her partner in μ_i^* to her partner in μ_i^+ , or is matched to the same partner under both matchings. Therefore, for each deterministic preference profile that has a positive probability of realization under I , if a pair (m, w) blocks μ_i^* then it must also block μ_i^+ , implying that if μ_i^+ is stable under the given preference profile then so is μ_i^* , and thus our statement follows. $\square$

4 Lottery Model

In this section, we turn to the lottery model. Recall that in the lottery model, each agent reports a probability distribution over a set of preference orders. We show that even with small supports, computing STABILITYPROBABILITY is still a computationally hard problem. However, other questions become more tractable with small supports or one side having certain preferences.

Theorem 3 *For the lottery model, if one side has certain preferences, STABILITYPROBABILITY is polynomial-time solvable.*

Proof Without loss of generality, assume that men have certain preferences so each man m has a deterministic preference relation $\succ_m$. We present a polynomial-time computable formula for the probability of μ being stable. For a woman w , we denote her set of possible preferences lists by $P_w = \{\succ_w^{p_w^1}, \dots, \succ_w^{p_w^{k_w}}\}$ with each preference list $\succ_w^{p_w^i}$ having corresponding probability p_w^i . Let q_w be the probability that woman w will not form a blocking pair. The term q_w is by definition as follows.

$$q_w = \sum_{\substack{\succ_w^{p_w^j} \in P_w: \nexists m \in M \text{ for which } w \succ_m \mu(m) \text{ and } m \succ_w^{p_w^j} \mu(w)}} p_w^j$$

So for each $w \in W$, q_w can be computed in time $O(k_w n)$ by scanning over the k_w preference lists of woman w and checking in which of them it is not the case that some m prefers w over $\mu(m)$ and w prefers m over $\mu(w)$. The probability that μ is stable is equal to the probability that no woman is in a blocking pair:

$$p(\mu \text{ is stable}) = \prod_{w \in W} q_w.$$

$\square$

Theorem 4 *For the lottery model, ISSTABILITYPROBABILITYONE can be solved in polynomial time.*

Proof The problem is equivalent to checking whether the given matching μ has non-zero probability of *not* being stable, i.e., we need only find a single possible blocking

pair. This can be checked as follows: for each pair (m, w) of agents that are not matched to each other, we check whether either member of this pair can form a blocking pair with a pair in μ with non-zero probability. To do this, we need to check whether m prefers w in some possible preference over $\mu(m)$ and whether w prefers m in some possible preference over $\mu(w)$. As this only involves checking $n^2 - n$ possible pairs, each of which can be checked in linear time, this can be done in polynomial time. $\square$

Theorem 5 *For the lottery model, ISSTABILITYPROBABILITYNON-ZERO is polynomial-time solvable when each agent has at most two possible preference orderings.*

Proof The problem is to decide whether there is some preference ordering for each agent (among the ones in their lottery) such that the given matching is stable. If each agent has at most two possible preference orderings in their lottery, we can reduce the problem to an instance φ of 2SAT [12], as follows. 2SAT is the problem of deciding whether a given propositional formula in 2CNF is satisfiable. A propositional formula φ is in 2CNF if it is the conjunction of clauses $(l_1 \vee l_2)$ of size 2, where l_1 and l_2 are propositional literals (propositional variables x or their negation $\neg x$).

Let $\{a_1, \dots, a_n\}$ and $\{b_1, \dots, b_n\}$ be the two sets of agents. Moreover, for each agent c and each $i \in \{1, 2\}$, let $\text{pref}(c, i)$ denote the i -th preference in the lottery for agent c .

We introduce a propositional variable for each preference $\text{pref}(c, i)$, which we also call $\text{pref}(c, i)$. Intuitively, these variables indicate which preference for the agents we choose to make the matching stable.

For each agent c , we add the following clauses to φ , to ensure that for each agent c there is exactly one preference that is selected: $(\text{pref}(c, 1) \vee \text{pref}(c, 2)) \wedge (\neg \text{pref}(c, 1) \vee \neg \text{pref}(c, 2))$.

Then, we add clauses to ensure that the selected matching is stable. For each agent c and each $i \in \{1, 2\}$, let $B_{c,i}$ be the set of preferences $\text{pref}(c', i')$ —for $c' \neq c$ and $i' \in \{1, 2\}$ —such that $\text{pref}(c, i)$ and $\text{pref}(c', i')$ together lead to the given matching being unstable (with (c, c') being a blocking pair). Then, for each c, i , we add the following clauses: $(\neg \text{pref}(c, i) \vee \neg \text{pref}(c', i'))$ for each $\text{pref}(c', i') \in B_{c,i}$.

The given matching is then stable if and only if φ is satisfiable. Since φ is a 2CNF formula, this can be decided in linear time [1]. $\square$

Theorem 6 *For the lottery model, STABILITYPROBABILITY is #P-complete, even when each agent has at most two possible preferences.*

Proof We show how to count the number of satisfying assignments for a 2CNF formula using the problem STABILITYPROBABILITY for the lottery model where each agent has two possible preferences. Since this problem is #P-hard, we get #P-hardness also for STABILITYPROBABILITY [23].

Let φ be a 2CNF formula over the variables $x_1, \dots, x_n$. We firstly transform φ to a 2CNF formula φ' over the variables $x_1, \dots, x_{2n}, y_1, \dots, y_{2n}$ that has exactly the same number of satisfying assignments, and that satisfies the property that each clause contains one variable among $x_1, \dots, x_{2n}$ and one variable among $y_1, \dots, y_{2n}$. We do so as follows. Firstly, for each $1 \leq i \leq n$, we add the clauses $(\neg x_i \vee y_i), (\neg y_i \vee x_{n+i}), (\neg x_{n+i} \vee y_{n+i})$ and $(\neg y_{n+i} \vee x_i)$, ensuring that in each satisfying assignment the

variables x_i, x_{n+i}, y_i and y_{n+i} get assigned the same truth value. Then, for each clause of φ , we replace one occurrence of a variable among $x_1, \dots, x_n$ by a corresponding variable among $y_1, \dots, y_{2n}$ —that is, we replace x_i either by y_i or by y_{n+i} —and we add the resulting clause to φ' . For example, if φ contains the clause $(x_1 \vee \neg x_3)$, we could add the clause $(x_1 \vee \neg y_3)$ to φ' . It is readily verified that φ' has the same number of satisfying assignments as φ .

Without loss of generality, we may assume that for each pair x_i, x_j of variables among $x_1, \dots, x_n$, the original formula φ contains at most two distinct clauses that contain both x_i and x_j . If φ were to contain three or more distinct clauses that contain both x_i and x_j , we know that at least one of the literals $x_i, \neg x_i, x_j, \neg x_j$ would be entailed by these clauses, and we could instantiate these entailed literals and simplify φ accordingly. For example, if φ contains the clauses $(x_1 \vee x_2)$, $(x_1 \vee \neg x_2)$ and $(\neg x_1 \vee \neg x_2)$, we can simplify φ by instantiating the entailed literals x_1 and $\neg x_2$. As a result, we know that we can construct φ' in such a way that for any two variables of φ' , there is at most one clause of φ' that contains both of these variables. For instance, if φ contains the clauses $(x_1 \vee x_2)$ and $(\neg x_1 \vee \neg x_2)$, we can construct φ' in such a way that it contains the clauses $(x_1 \vee y_2)$ and $(\neg y_1 \vee \neg x_2)$.

We now construct an instance of STABILITYPROBABILITY. The sets of agents that we consider are $\{x_1, \dots, x_{2n}, a_1, \dots, a_{2n}\}$ and $\{y_1, \dots, y_{2n}, b_1, \dots, b_{2n}\}$. The matching that we consider matches x_i to b_i and matches y_i to a_i , for each $1 \leq i \leq 2n$ as shown below.

Each agent b_i has only a single possible preference, namely one where they prefer x_i over all other agents. Similarly, each agent a_i has a single possible preference where they prefer y_i over all other agents. In other words, the agents a_i and b_i are perfectly happy with the given matching.

The agents x_i and y_i each have two possible preferences, that are each chosen with probability $\frac{1}{2}$. These two possible preferences are associated with setting these variables to true or false, respectively. We describe how these preferences are constructed for the agents x_i . The construction for the preferences of the agents y_i is then entirely analogous.

Take an arbitrary agent x_i . We show how to construct the two possible preferences for agent x_i , which we denote by p_{x_i} and $p_{\neg x_i}$. Both of these possible preferences are based on the following partial ranking: $b_1 > b_2 > \dots > b_{2n}$, and we add some of the agents among $y_1, \dots, y_{2n}$ to the top of this partial ranking, and the remaining agents to the bottom of this partial ranking.

To the ranking p_{x_i} we add exactly those agents y_j to the top where φ' contains a clause $(\neg x_i \vee y_j)$ or a clause $(\neg x_i \vee \neg y_j)$. All remaining agents we add to the bottom. Similarly, to the ranking $p_{\neg x_i}$ we add exactly those agents y_j to the top where φ'

contains a clause $(x_i \vee y_j)$ or a clause $(x_i \vee \neg y_j)$. The rankings p_{y_j} and $p_{\neg y_j}$, for the agents y_j , are constructed entirely similarly.

Consider a truth assignment $\alpha : \{x_1, \dots, x_{2n}, y_1, \dots, y_{2n}\} \rightarrow \{0, 1\}$, and consider the corresponding choice of preferences for the agents $x_1, \dots, x_{2n}, y_1, \dots, y_{2n}$, where for each agent x_i the preference p_{x_i} is chosen if and only if $\alpha(x_i) = 1$, and for each agent y_j the preference p_{y_j} is chosen if and only if $\alpha(y_j) = 1$. We show that α satisfies φ' if and only if the corresponding choice of preferences leads to the matching being stable. The only blocking pairs that can arise are between an agent x_i and an agent y_j . Take an arbitrary such pair (x_i, y_j) . We show that this is a blocking pair if and only if α falsifies at least one clause containing both variables x_i and y_j . We know that φ' contains exactly one clause containing x_i and y_j . We deal with the case where φ' contains the clause $(x_i \vee y_j)$ —the other possibilities are entirely analogous. By construction of p_{x_i} , $p_{\neg x_i}$, p_{y_j} , and $p_{\neg y_j}$, we get that (x_i, y_j) is a blocking pair if and only if both $p_{\neg x_i}$ and $p_{\neg y_j}$ are chosen. This is exactly the case where the clause $(x_i \vee y_j)$ is falsified. Thus, we can conclude that α falsifies φ' if and only if there is a blocking pair, and thus that α satisfies φ' if and only if the corresponding choice of preferences leads to the matching being stable.

Since each combination of preferences is equally likely to occur, and there are 2^{4n} many combinations of preferences, the probability that the given matching is stable is exactly $q = \frac{s}{2^{4n}}$, where s is the number of satisfying truth assignments for φ . Therefore, given q , s can be obtained by computing $s = q2^{4n}$. $\square$

If the agents are allowed to have more than two possible preferences as we have assumed in the last two Theorems, then even the following problem is NP-complete. The statement can be proved via a reduction from Exact Cover by 3-Sets (X3C).

Theorem 7 *For the lottery model, ISSTABILITYPROBABILITYNON-ZERO is NP-complete.*

Proof The problem is in NP, since we only need to provide one profile that occurs with non-zero probability for which the given matching is stable. We show NP-hardness by giving a reduction from Exact Cover by 3-Sets (X3C). Let (X, C) be an instance of X3C, where $|X| = 3n$ for some n , and $C = \{c_1, \dots, c_m\}$ is a collection of sets $c_i \subseteq X$, each of size 3. Moreover, let $c_i = \{x_{\ell_{i,1}}, x_{\ell_{i,2}}, x_{\ell_{i,3}}\}$, for each $1 \leq i \leq m$. The problem is to decide whether there is a subset $C' \subseteq C$ of size exactly n such that $\bigcup C' = X$.

We construct an instance of our problem as follows. We let $\{a_1, \dots, a_n, a'_1, \dots, a'_{3n}\}$ and $\{b_1, \dots, b_n, b'_1, \dots, b'_{3n}\}$ be the two sets of agents, we match a_i to b_i —for each $1 \leq i \leq n$ —and we match a'_j to b'_j —for each $1 \leq j \leq 3n$. This is depicted graphically below.

Each agent a_i prefers their matching to b_i over any other possible match, i.e., agent a_i has one preference, where b_i is ranked first, and the rest of the agents appear in some fixed order after b_i .

Similarly, each agent b'_j prefers their matching to a'_j over any other possible match. That is, agent b'_j has one preference ordering in which a'_j is ranked first and the rest of the agents appear in some fixed order after a'_j .

Then, for each agent b_i , we add the following $|C|$ possible preferences to the lottery:

$$\begin{aligned} P_{i,1} : & \quad a'_{\ell_{1,1}} > a'_{\ell_{1,2}} > a'_{\ell_{1,3}} > a_i > \dots \\ P_{i,2} : & \quad a'_{\ell_{2,1}} > a'_{\ell_{2,2}} > a'_{\ell_{2,3}} > a_i > \dots \\ & \quad \vdots \\ P_{i,m} : & \quad a'_{\ell_{m,1}} > a'_{\ell_{m,2}} > a'_{\ell_{m,3}} > a_i > \dots \end{aligned}$$

where in each preference the remaining agents appear in any (fixed) order after a_i . In other words, b_i prefers three agents a'_j to their current match, and these three form some set $c \in C$.

Finally, for each agent a'_j , we add the following n possible preferences to the lottery:

$$\begin{aligned} P'_{j,1} : & \quad b_2 > \dots > b_n > b'_j > b_1 > b'_1 > \dots > b'_{j-1} > b'_{j+1} > \dots > b'_{3n} \\ P'_{j,2} : & \quad b_1 > b_3 > \dots > b_n > b'_j > b_2 > b'_1 > \dots > b'_{j-1} > b'_{j+1} > \dots > b'_{3n} \\ P'_{j,3} : & \quad b_1 > b_2 > b_4 > \dots > b_n > b'_j > b_3 > b'_1 > \dots > b'_{j-1} > b'_{j+1} > \dots > b'_{3n} \\ & \quad \vdots \\ P'_{j,n} : & \quad b_1 > \dots > b_{n-1} > b'_j > b_n > b'_1 > \dots > b'_{j-1} > b'_{j+1} > \dots > b'_{3n} \end{aligned}$$

That is, each agent a'_j prefers each of the agents $b_1, \dots, b_n$, except one, to their current match (and they never prefer any of the agents $b'_{j'}$ for $j' \neq j$ over their current match).

We can show that there is a choice of preferences for the agents that makes this matching stable if and only if $(X, C) \in \text{X3C}$.

($\Rightarrow$) Firstly, suppose that there is a choice of preferences for the agents that makes this matching stable. That is, for each agent b_i there is some preference ordering P_{i,ℓ_i} , and for each agent a'_j there is some preference ordering P'_{j,k_j} , such that these orderings (together with the fixed preference orderings for the agents a_i and b'_j) make this matching stable. Now, consider the set $C' = \{c_\ell : i \in [n], \ell = \ell_i\}$. We show that $\bigcup C' = X$. To derive a contradiction, suppose that this is not the case, that is, suppose that $\bigcup C' \neq X$. Then, since $|C'| = n$, $|X| = 3n$ and for each $c \in C'$ it holds that $|c| = 3$, we know that there must be some $c_\ell, c_{\ell'} \in C'$ such that $c_\ell \cap c_{\ell'} \neq \emptyset$. Say that $x_j \in c_\ell \cap c_{\ell'}$. Therefore, there must be some $i, i' \in [n]$ such that both b_i and $b_{i'}$ prefer a'_j over their current match. On the other hand, a'_j will prefer either b_i or $b_{i'}$ over their current match. Therefore, either b_i and a'_j or $b_{i'}$ and a'_j will form a blocking pair. Thus, the matching is not stable. From this we can conclude that $\bigcup C' = X$.

($\Leftarrow$) Conversely, suppose that there exists some $C' \subseteq C$ of size exactly n such that $\bigcup C' = X$. Let $C' = \{c_{\ell_1}, \dots, c_{\ell_n}\}$. Now, for each agent b_i we pick some

preference ordering, and for each agent a'_j we pick some preference ordering, such that these orderings (together with the fixed preference orderings for the agents a_i and b'_j) make the matching stable. For each agent b_i , we pick the preference ordering P_{i,ℓ_i} , and for each agent a'_j we pick the preference ordering P'_{j,k_j} , where $k_j \in [n]$ is the unique value such that $x_j \in c_{\ell_{k_j}}$. It is straightforward to verify that these preferences make the matching stable. $\square$

Given that STABILITYPROBABILITY is computationally hard even in restricted settings, one may immediately wonder about approximation algorithms. However, we can state the following two corollaries, the first from Theorem 7 and the second from [29, Proposition 8] and Theorem 7, which show that even approximation remains intractable.

Corollary 3 *For the lottery model, unless $P=NP$, there exists no polynomial-time constant-factor approximation algorithm STABILITYPROBABILITY of a given matching.*

Corollary 4 *For the lottery model, unless $NP=RP$, there is no FPRAS for STABILITYPROBABILITY.*

5 Compact Indifference Model

In the compact indifference model we are given an instance of SMT and each linear order over candidates (each possible preference ordering) is achieved by breaking ties independently at random with uniform probabilities. It is easy to show that ISSTABILITYPROBABILITYNONZERO, ISSTABILITYPROBABILITYONE, and EXISTSCESTAINLYSTABLEMATCHING are all in P.

Proposition 3 *For the compact indifference model, ISSTABILITYPROBABILITYNONZERO is in P.*

Proof This is equivalent to checking whether a given matching μ is weakly stable in the given SMTI instance, which is polynomial-time solvable. To check this we only have to look for a blocking pair, which can be done in polynomial time: take every possible pair (m, w) who are not matched together and check whether they both (strictly) prefer each other to their current partner. $\square$

Proposition 4 *For the compact indifference model, ISSTABILITYPROBABILITYONE is in P.*

Proof This is equivalent to checking whether a given matching μ is super stable in the given SMTI instance, which is polynomial-time solvable. To check this we only have to look for a very weakly blocking pair, which can be done in polynomial time: take every possible pair (m, w) who are not matched together and check whether they each either prefers the other to his or her current partner or finds them incomparable. $\square$

Proposition 5 *For the compact indifference model, EXISTSCERTAINLYSTABLEMATCHING is in P.*

Proof Deciding whether there is a matching that is stable with probability one is equivalent to deciding whether there is a matching that is stable w.r.t. all refinements, i.e., a super-stable matching. Given an instance of SMTI one can decide in polynomial time whether it admits a super-stable matching [20]. $\square$

The complexity of computing the stability probability of a given matching remains open under the compact indifference model, but this problem can be shown to be in P if one side has certain preferences.

Theorem 8 *In the compact indifference model, if one side has certain preferences, STABILITYPROBABILITY is polynomial-time solvable.*

Proof Assume, without loss of generality, that men have certain preferences. The following procedure gives us the stability probability of any given matching μ . (1) For each uncertain woman w identify those men with whom she can potentially form a blocking pair. That is, those m such that $w \succ_m \mu(m)$ and w is indifferent between m and her partner in μ . Assume there are k such men. The probability of w not forming a blocking pair with any men is then the probability that her current match ranks first among these $k + 1$ men, which is $\frac{1}{k+1}$. (2) Multiply the probabilities from step 1. $\square$

We next show that MATCHINGWITHHIGHESTSTABILITYPROBABILITY is NP-hard. For an instance I of SMT and matching μ , let $p(\mu, I)$ denote the probability of μ being stable, and let $p_S(I) = \max\{p(\mu, I) \mid \mu \text{ is a matching in } I\}$, that is the maximum probability of a matching being stable. A matching μ is said to be weakly stable if there exists a tie-breaking rule where μ is stable. Therefore a matching μ has positive probability of being stable if and only if it is weakly stable. Furthermore, if the number of possible tie-breaking is N then any weakly stable matching has a probability of being stable at least $\frac{1}{N}$.

An extreme case occurs if we have one woman only with n men, where the woman is indifferent between all men. In this case any matching (pair) has a $\frac{1}{n}$ probability of being stable. An even more unfortunate scenario is when we have n men and n women, each woman is indifferent between all men, and each man ranks the women in a strict order in the same way, e.g., in the order of their indices. In this case, the probability that the first woman picks her best partner, and thus does not block any matching is $\frac{1}{n}$. Suppose that the first woman picked her best partner, the probability that the second woman also picks her best partner from the remaining $n - 1$ men is $\frac{1}{n-1}$, and so on. Therefore, the probability that an arbitrary complete matching is stable is $\frac{1}{n \cdot (n-1) \cdot \dots \cdot 2} = \frac{1}{n!}$.

Theorem 9 *For the compact indifference model, MATCHINGWITHHIGHESTSTABILITYPROBABILITY is NP-hard, even if only one side of the market has uncertain agents.*

Proof For an instance I of SMTI, let $\text{opt}(I)$ denote the maximum size of a weakly stable matching in I . [14] showed (in the proof of Corollary 3.4) that given an instance I of SMTI of size n , where only one side of the market has agents with indifferences

and each of these agents has a single tie of size two, and any arbitrary small positive ϵ , it is NP-hard to distinguish between the following two cases: (1) $\text{opt}(I) \geq \frac{21-\epsilon}{27}n$ and (2) $\text{opt}(I) < \frac{19+\epsilon}{27}n$.

When choosing ϵ so that $0 < \epsilon < \frac{1}{2}$ we can simplify the above cases to (1) $\text{opt}(I) > \frac{41}{54}n$, since $\text{opt}(I) \geq \frac{21-\epsilon}{27}n > \frac{41}{54}n$ and (2) $\text{opt}(I) < \frac{39}{54}n$, since $\text{opt}(I) < \frac{19+\epsilon}{27}n < \frac{39}{54}n$.

Therefore, the number of agents left unmatched on either side of the market is less than $\frac{13}{54}n$ in the first case and more than $\frac{15}{54}n$ in the second case. Let us now extend instance I to a larger instance of SMTI I' as follows. Besides the n men $M = \{m_1, \dots, m_n\}$ and n women $W = \{w_1, \dots, w_n\}$, we introduce $\frac{13}{54}n$ men $X = \{x_1, \dots, x_k\}$ and another $\frac{n}{27}$ men $Y = \{y_1, \dots, y_l\}$ and $\frac{n}{27}$ women $Z = \{z_1, \dots, z_l\}$. Furthermore, for each $y_j \in Y$, we introduce n men $Y^j = \{y_1^j, \dots, y_n^j\}$. We create the preferences of I' as follows. The preferences of men M remain the same. For each woman $w \in W$ we append the men X and then Y at the end of her list in the order of their indices. Each man $x_i \in X$ has only all the women W in his list in the order of their indices. Furthermore, each $y_j \in Y$ has all the women W first in his preference list in the order of their indices and then z_j . Let each $z_j \in Z$ has y_j as first choice and then all the men Y^j in one tie of size n . Each man in Y^j has only z_j in his list. We will show that in case one $p_S(I') \geq \frac{1}{2n}$, whilst in case two $p_S \leq (\frac{1}{n})^{\frac{n}{27}}$. Therefore, for $n > 2^{27}$, it is NP-hard to decide which of the two separate intervals contains the value $p_S(I')$.

To show the above statement, suppose first that we have the first case, so $\text{opt}(I) > \frac{41}{54}n$ and therefore less than $\frac{13}{54}n$ women are left unmatched in a maximum size weakly stable matching μ for I , denoted by $W_u \subset W$. We extend μ to μ' for I' as follows. We assign all the women in W_u to men in X in the unique stable way, namely we pair them in a mutually increasing order of their indices. Since $|X| > |W_u|$, we now matched all women in W , and left some men in X unmatched in μ' . We complete the matching by assigning y_j to z_j for each $j = 1, \dots, n$ and leaving all of the men in Y^j for all j unmatched. We shall see that no matter how we break the ties in I' , blocking pair can appear between the original I agents only, and therefore the probability of μ' being stable in I' is the same as the probability of μ being stable in I . Since we have at most n ties in I , each of length two, the number of different tie-breakings is at most 2^n , out of which at least one is stable. Therefore $p(\mu, I') = p(\mu, I) \geq \frac{1}{2^n}$.

In the second case, $\text{opt}(I) < \frac{39}{54}n$ and therefore more than $\frac{15}{54}n$ women are left unmatched in any weakly stable matching μ for I . Let μ' be one of the most stable matchings in I' . First we have to note that the restriction of μ' to I must be weakly stable in I , since otherwise $p(\mu', I') = 0$. Let W_u denote the set of women that are not matched to any man from M in μ' . According to our assumption $|W_u| > \frac{15}{54}n$, whilst $|X| + |Y| = \frac{15}{54}n$, therefore in order to avoid a certain blocking pair between W_u and $X \cup Y$ we shall match all the men in $X \cup Y$ to women in W_u in the only stable way (in the order of indices, where men in X are coming before men in Y), leaving some women in W_u unmatched in μ' . However, in this case no agent $z_j \in Z$ can be matched to y_j , and therefore, even if there was no potential blocking pair between agents of I , the probability that z_j is matched the best partner from Y^j is $\frac{1}{n}$ independently for each

$z_j \in Z$. Therefore the probability of μ' being stable is at most $(\frac{1}{n})^{\frac{n}{27}}$, which completes the proof of the first statement.

Regarding the NP-hardness of finding one of the most stable matchings, we shall prove that we can decide between the two cases according to the number of unmatched women in W in the restriction of μ' to I , where μ' is one of the most stable matchings in I' . To see this, let W_u denote again the set of women that are not matched to any man in M under μ' . In the first case, when $\text{opt}(I) > \frac{41}{54}n$, it must be the case that $|W_u| < \frac{15}{54}n$, since otherwise $p(\mu', I)$ would be less than $(\frac{1}{n})^{\frac{n}{27}}$ and could not achieve $\frac{1}{2^n}$, that is the minimum value for $p_S(I')$, as shown in the above argument. Whilst, in the second case $|W_u| > \frac{15}{54}n$ must hold, since $\text{opt}(I) < \frac{39}{54}n$ was assumed. $\square$

6 Joint Probability Model

In this section, we examine problems concerning the joint probability model. Recall that the input now consists of a set of preference profiles, and a probability distribution over these preference profiles. Our first result is that the stability probability of a matching can be computed in polynomial time in the size of the input.

Theorem 10 *For the joint probability model, STABILITYPROBABILITY can be solved in polynomial time.*

Proof The probability that a given matching is stable is equivalent to the sum of the probability weights of the preference profiles for which the matching is stable. This can be checked as follows. We check for which preference profiles the given matching is stable. For one profile, this can be checked in $O(n^2)$ time by checking for each (man, woman) pair, whether it is a blocking pair. Then we add up the probabilities of those profiles for which the matching is stable. The sum of the probabilities is the probability that the matching is stable. $\square$

Corollary 5 *For the joint probability model, ISSTABILITYPROBABILITYNON-ZERO and ISSTABILITYPROBABILITYONE can be solved in polynomial time.*

For the joint probability model, the problem EXISTSCERTAINLYSTABLEMATCHING is equivalent to checking whether the intersection of the sets of stable matchings of the different preference profiles is empty.

Theorem 11 *For the joint probability model, EXISTSCERTAINLYSTABLEMATCHING is NP-complete.*

Proof The problem is in NP, since computing STABILITYPROBABILITY can be done in polynomial time by Theorem 10, and so a matching that is stable with probability 1 is a certificate. The NP-hardness proof is by reduction from the NP-complete 3-COLORING problem. The input for 3-COLORING is a graph $G = (V, E)$, and the question is whether there exists a coloring of the vertex set with 3 colors, $\chi : V \rightarrow \{1, 2, 3\}$, such that no two adjacent vertices receive the same color. Such a coloring is called a *proper* coloring. Let $G = (V, E)$ be a graph specifying an instance of 3-COLORING,

where $V = \{v_1, \dots, v_n\}$. We construct an instance I of EXISTSCERTAINLYSTABLE-MATCHING in the joint probability model.

For each vertex $v_i \in V$, we introduce three men $m_{i,1}, m_{i,2}, m_{i,3}$ and three women $w_{i,1}, w_{i,2}, w_{i,3}$. We will add a sequence of preference profiles that have non-zero probability. The exact probability of the preference profiles is unimportant, so we may assume that there is a uniform probability associated with the preference profiles. The first preference profile P_0 will ensure that every certainly stable matching matches each $m_{i,j}$ to some $w_{i,j'}$ and each $w_{i,j}$ to some $m_{i,j''}$, where $j, j', j'' \in \{1, 2, 3\}$. Moreover, it ensures that for each $i \in \{1, \dots, n\}$, exactly one of three matchings between the men $m_{i,j}$ and the women $w_{i,j}$ must be used:

- (1) $m_{i,1}$ is matched to $w_{i,1}$, $m_{i,2}$ is matched to $w_{i,2}$, and $m_{i,3}$ is matched to $w_{i,3}$;
- (2) $m_{i,1}$ is matched to $w_{i,2}$, $m_{i,2}$ is matched to $w_{i,3}$, and $m_{i,3}$ is matched to $w_{i,1}$; or
- (3) $m_{i,1}$ is matched to $w_{i,3}$, $m_{i,2}$ is matched to $w_{i,1}$, and $m_{i,3}$ is matched to $w_{i,2}$.

Intuitively, choosing one of the matchings (1)–(3) for the agents $m_{i,j}, w_{i,j}$ corresponds to coloring vertex v_i with one of the three colors in $\{1, 2, 3\}$.

Then, for each edge $e = \{v_{i_1}, v_{i_2}\} \in E$, and for each color $c \in \{1, 2, 3\}$, we will introduce a preference profile $P_{e,c}$ that ensures that in each certainly stable matching, it cannot be the case that both the agents $m_{i_1,j}, w_{i_1,j}$ and the agents $m_{i_2,j}, w_{i_2,j}$ are matched to each other according to matching (c). As a result, any certainly stable matching directly corresponds to a proper 3-coloring of G .

A detailed description of the preference profiles P_0 and $P_{e,c}$ and a proof of correctness for this reduction follows.

In P_0 , for each $i \in [n]$, the preferences for $m_{i,j}, w_{i,j}$ are as follows:

$$\begin{array}{ll} m_{i,1} : & w_{i,1}, w_{i,2}, w_{i,3}, \text{---} \text{---} \text{---} & w_{i,1} : & m_{i,2}, m_{i,3}, m_{i,1}, \text{---} \text{---} \text{---} \\ m_{i,2} : & w_{i,2}, w_{i,3}, w_{i,1}, \text{---} \text{---} \text{---} & w_{i,2} : & m_{i,3}, m_{i,1}, m_{i,2}, \text{---} \text{---} \text{---} \\ m_{i,3} : & w_{i,3}, w_{i,1}, w_{i,2}, \text{---} \text{---} \text{---} & w_{i,3} : & m_{i,1}, m_{i,2}, m_{i,3}, \text{---} \text{---} \text{---} \end{array}$$

Observe that the matchings (1)–(3) are the only stable matchings restricted to the agents $\{m_{i,1}, m_{i,2}, m_{i,3}, w_{i,1}, w_{i,2}, w_{i,3}\}$.

Next, we continue with the preference profiles $P_{e,c}$. Take an arbitrary $e = \{v_{i_1}, v_{i_2}\} \in E$ and an arbitrary $c \in \{1, 2, 3\}$. In $P_{e,c}$, the preferences for $m_{i,j}, w_{i,j}$ for each $i \in [n] \setminus \{i_1, i_2\}$ are exactly the same as in P_0 . Only the preferences for $m_{i_1,j}, w_{i_1,j}$ and $m_{i_2,j}, w_{i_2,j}$ differ from P_0 ; namely, we construct these preferences as follows.

For $m_{i_1,j}, w_{i_1,j}$, we start with preferences that (i) for all $m_{i_1,j}$ have $\{w_{i_1,1}, w_{i_1,2}, w_{i_1,3}\}$ as top three choices, (ii) for all $w_{i_1,j}$ have $\{m_{i_1,1}, m_{i_1,2}, m_{i_1,3}\}$ as top three choices, (iii) admit only matchings (1), (2), and (3) as stable matchings between the agents $m_{i_1,j}, w_{i_1,j}$, and (iv) for the men $m_{i_1,j}$ the matching (c) is the worst option among the matchings (1), (2), and (3). Similarly, for $m_{i_2,j}, w_{i_2,j}$, we start with preferences that satisfy conditions (i), (ii) and (iii), and additionally satisfy the condition (iv') that for the women $w_{i_2,j}$ the matching (c) is the worst option among the matchings (1), (2), and (3). Then, we modify the preferences for $m_{i_1,1}$ and $w_{i_2,1}$ slightly. For $m_{i_1,1}$, we insert $w_{i_2,1}$ between his second and third preferred woman. Similarly,

for $w_{i_2,1}$, we insert $m_{i_1,1}$ between her second and third preferred man. As a result, $m_{i_1,1}$ and $w_{i_2,1}$ form a blocking pair in this preference profile if both the agents $m_{i_1,j}$, $w_{i_1,j}$ and the agents $m_{i_2,j}$, $w_{i_2,j}$ are matched to each other using matching (c) — and not if any set of agents is matched to each other using some other matching (c') .

For example, consider $e = \{v_{i_1}, v_{i_2}\}$ and $c = 2$. The preferences for the agents $m_{i_1,j}$, $w_{i_1,j}$ and $m_{i_2,j}$, $w_{i_2,j}$ in the preference profile $P_{e,c}$ are as follows:

$m_{i_1,1} :$	$w_{i_1,1}, w_{i_1,3}, \mathbf{w_{i_2,1}}, w_{i_1,2}, - - -$	$m_{i_2,1} :$	$w_{i_2,2}, w_{i_2,3}, w_{i_2,1}, - - -$
$m_{i_1,2} :$	$w_{i_1,2}, w_{i_1,1}, w_{i_1,3}, - - -$	$m_{i_2,2} :$	$w_{i_2,3}, w_{i_2,1}, w_{i_2,2}, - - -$
$m_{i_1,3} :$	$w_{i_1,3}, w_{i_1,2}, w_{i_1,1}, - - -$	$m_{i_2,3} :$	$w_{i_2,1}, w_{i_2,2}, w_{i_2,3}, - - -$
$w_{i_1,1} :$	$m_{i_1,3}, m_{i_1,2}, m_{i_1,1}, - - -$	$w_{i_2,1} :$	$m_{i_2,1}, m_{i_2,2}, \mathbf{m_{i_1,1}}, m_{i_2,3}, - - -$
$w_{i_1,2} :$	$m_{i_1,1}, m_{i_1,3}, m_{i_1,2}, - - -$	$w_{i_2,2} :$	$m_{i_2,2}, m_{i_2,3}, m_{i_2,1}, - - -$
$w_{i_1,3} :$	$m_{i_1,2}, m_{i_1,1}, m_{i_1,3}, - - -$	$w_{i_2,3} :$	$m_{i_2,3}, m_{i_2,1}, m_{i_2,2}, - - -$

We argue that G has a proper 3-coloring if and only if there is a certainly stable matching for the probability distribution over preference profiles that we constructed.

($\Rightarrow$) Firstly, suppose that G has a proper 3-coloring, say $\chi : V \rightarrow \{1, 2, 3\}$. We can then construct a certainly stable matching as follows. For each $i \in \{1, \dots, n\}$, we match the agents $m_{i,j}$, $w_{i,j}$ to each other using matching $(\chi(v_i))$. Clearly, this matching is stable for P_0 . Moreover, because χ is a proper 3-coloring of G , it is straightforward to verify that this matching is also stable for each $P_{e,c}$.

($\Leftarrow$) Conversely, suppose that there is a certainly stable matching. We know that in this matching, each man $m_{i,j}$ must be matched to some woman $w_{i',j'}$, and vice versa, each woman $w_{i,j}$ must be matched to some man $m_{i'',j''}$. If this were not the case, the matching would not be stable for P_0 , and thus not certainly stable. Moreover, by a similar argument, we know that for each $i \in \{1, \dots, n\}$, the matching between the men $m_{i,j}$ and the women $w_{i,j}$ must be one of the matchings (1), (2), or (3). We can then construct a 3-coloring $\chi : V \rightarrow \{1, 2, 3\}$ as follows. For each $i \in \{1, \dots, n\}$, we set $\chi(v_i)$ in such a way that $(\chi(v_i))$ is the matching used in the certainly stable matching to match the men $m_{i,j}$ and the women $w_{i,j}$ to each other.

We argue that χ is a proper 3-coloring of G . Suppose that this is not the case, that is, that there is some $e = \{v_{i_1}, v_{i_2}\}$ such that $\chi(v_{i_1}) = \chi(v_{i_2}) = c$. Now consider the preference profile $P_{e,c}$. By construction of χ , we know that in the certainly stable matching, agent $m_{i_1,j}$ is matched to agent $w_{i_1,j}$ and agent $m_{i_2,j}$ is matched to agent $w_{i_2,j}$, both times according to matching (c) . However, by construction of $P_{e,c}$ we will have that $m_{i_1,1}$ and $w_{i_2,1}$ form a blocking pair in $P_{e,c}$. This is a contradiction with our assumption that the matching we considered is certainly stable. From this, we can conclude that χ is a proper 3-coloring of G . $\square$

By reducing from 3-COLORING on 4-regular graphs and using Vizing's theorem to combine preference profiles, we can strengthen Theorem 11 and show that the result holds even when the input contains only a constant number of preference profiles.

Corollary 6 *For the joint probability model, EXISTSCERTAINLYSTABLEMATCHING is NP-complete, even when there are only 16 preference profiles in the lottery.*

Proof We show this by modifying the proof of Theorem 11. We know that 3-COLORING is NP-complete even when restricted to 4-regular graphs [8]. We use the reduction in the proof of Theorem 11, and we assume that in the given graph G , each vertex has degree 4. Then, by Vizing's Theorem [28], we know that we can give a proper edge coloring of G that uses at most 5 colors. Moreover, we can find such an edge coloring in polynomial time. Then, since in the proof of Theorem 11, in each preference profile $P_{e,c}$ with $e = \{v_{i_1}, v_{i_2}\}$, only the preferences for the agents $m_{i_1,j}, w_{i_1,j}, m_{i_2,j}, w_{i_2,j}$ differ from P_0 , we can, for each color $c \in \{1, 2, 3\}$, combine the preference profiles $P_{e,c}$ for all edges e that are colored with the same color. This results in only 16 preference profiles: P_0 , and a preference profile for each of the 5 edge colors and each of the 3 vertex colors. $\square$

7 Future work

First we note that we left open two outstanding questions, as described in Table 3. In this paper we focused on the problem of computing a matching with the highest stability probability. However, a similarly reasonable goal could be to minimize the expected number of blocking pairs. It would also be interesting to investigate some further realistic probability models, such as the situation when the candidates are ranked according to some noisy scores (like the SAT scores in the US college admissions). This would be a special case of the joint probability model that may turn out to be easier to solve.

One may also consider the questions we asked in this paper for two natural generalizations on the one-to-one matching problem: many-to-one and the many-to-many two-sided problems, where the agents on one or both sides of a two-sided market have capacities. These are reasonable extensions with regard to many applications, e.g., college admissions, and course allocation, respectively. One issue when considering capacities larger than one is that there are many ways of extending the preferences over candidates to preferences over sets of candidates. Another possible generalization of the one-to-one marriage model is the non-bipartite stable roommates problem. We conjecture that most of our polynomial tractability results can be extended to these more general models, but we leave this work for future research.

Acknowledgements The authors gratefully acknowledge the support from European Cooperation in Science and Technology (COST) action IC1205. Aziz is supported by a UNSW Scientia Fellowship. Biró acknowledges the support of the Hungarian Academy of Sciences under its Momentum Programme (LP2016-3/2018) and Cooperation of Excellences Grant (KEP-6/2018), and the Hungarian Scientific Research Fund, OTKA, Grant No. K128611. Gaspers is the recipient of an Australian Research Council (ARC) Future Fellowship (FT140100048) and acknowledges support under the ARC's Discovery Projects funding scheme (DP150101134). De Haan is supported by the Austrian Science Fund (FWF), Project J4047. Rastegari was supported by EPSRC Grant EP/K010042/1 at the time of the submission of the conference version of this manuscript.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Aspvall, B., Plass, M.F., Tarjan, R.E.: A linear-time algorithm for testing the truth of certain quantified boolean formulas. *Inf. Process. Lett.* **8**, 121–123 (1979)
2. Aziz, H., Biró, P., Gaspers, S., de Haan, R., Mattei, N., Rastegari, B.: Stable matching with uncertain linear preferences. In: *Proceedings of the 9th International Symposium on Algorithmic Game Theory (SAGT)*, pp. 195–206 (2016)
3. Aziz, H., Biró, P., Fleiner, T., Gaspers, S., de Haan, R., Mattei, N., Rastegari, B.: Stable matching with uncertain pairwise preferences. In: *Proceedings of the 16th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 344–352 (2017a)
4. Aziz, H., de Haan, R., Rastegari, B.: Pareto optimal allocation under uncertain preferences. In: *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 1472–1474 (2017b)
5. Aziz, H., Biró, P., de Haan, R., Rastegari, B.: Pareto optimal allocation under uncertain preferences: uncertainty models, algorithms, and complexity. *Artif. Intell.* **276**, 57–78 (2019a)
6. Aziz, H., de Haan, R., Rastegari, B.: Pareto optimal allocation under compact uncertain preferences. In: *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)* (2019b)
7. Chen, J., Niedermeier, R., Skowron, P.: Stable marriage with multi-modal preferences. In: *Proceedings of the 19th ACM Conference on Electronic Commerce (ACM-EC)*, pp. 269–286 (2018)
8. Dailey, D.P.: Uniqueness of colorability and colorability of planar 4-regular graphs are NP-complete. *Discrete Math.* **30**(3), 289–293 (1980)
9. Drummond, J., Boutilier, C.: Preference elicitation and interview minimization in stable matchings. In: *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 645–653 (2014)
10. Ehlers, L., Massó, J.: Matching markets under (in)complete information. *J. Econ. Theory* **157**, 295–314 (2015)
11. Gale, D., Shapley, L.S.: College admissions and the stability of marriage. *Am. Math. Mon.* **69**(1), 9–15 (1962)
12. Garey, M.R., Johnson, D.S.: *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, New York (1979)
13. Gusfield, D., Irving, R.W.: *The Stable Marriage Problem: Structure and Algorithms*. MIT Press, Cambridge (1989)
14. Halldórsson, M.M., Iwama, K., Miyazaki, S., Yanagisawa, H.: Improved approximation results for the stable marriage problem. *ACM Trans. Algorithms* **3**(3) (2007). <https://dl.acm.org/citation.cfm?id=1273346>
15. Hazon, N., Aumann, Y., Kraus, S., Wooldridge, M.: On the evaluation of election outcomes under uncertainty. *Artif. Intell.* **189**, 1–18 (2012)
16. Irving, R.W.: Stable marriage and indifference. *Discrete Appl. Math.* **48**, 261–272 (1994)
17. Irving, R.W., Manlove, D.F., Scott, S.: The hospitals/residents problem with ties. In: *Proceedings of the 7th Scandinavian Workshop on Algorithm Theory (SWAT)*, Volume 1851 of *Lecture Notes in Computer Science*, pp. 259–271. Springer (2000)
18. Li, Y., Conitzer, V.: Cooperative game solution concepts that maximize stability under noise. In: *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 979–985 (2015)
19. Manlove, D.: *Algorithmics of Matching Under Preferences*. World Scientific Publishing Company, Singapore (2013)
20. Manlove, D.F.: Stable marriage with ties and unacceptable partners. Technical Report TR-1999-29, University of Glasgow, Department of Computing Science (1999)
21. Mattei, N., Walsh, T.: PrefLib: A library for preferences, <http://www.preflib.org>. In: *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT)* (2013)
22. Mattei, N., Walsh, T.: A PREFLIB.ORG Retrospective: lessons learned and new directions. In: Endriss, U. (ed.), *Trends in Computational Social Choice*, chapter 15, pp. 289–309. AI Access Foundation (2017)
23. Papadimitriou, C.H.: *Computational Complexity*. Wiley, New York (2003)
24. Rastegari, B., Condon, A., Immorlica, N., Leyton-Brown, K.: Two-sided matching with partial information. In: *Proceedings of the 14th ACM Conference on Electronic Commerce (ACM-EC)*, pp. 733–750. ACM (2013)

25. Rastegari, B., Condon, A., Immorlica, N., Irving, R., Leyton-Brown, K.: Reasoning about optimal stable matching under partial information. In: Proceedings of the 15th ACM Conference on Electronic Commerce (ACM-EC), pp. 431–448. ACM (2014)
26. Roth, A.E.: The evolution of the labor market for medical interns and residents: a case study in game theory. *J. Polit. Econ.* **92**(6), 991–1016 (1984)
27. Roth, A.E., Sotomayor, M.A.O.: Two-Sided Matching: A Study in Game Theoretic Modelling and Analysis. Cambridge University Press, Cambridge (1990)
28. Vizing, V.G.: On an estimate of the chromatic class of a p-graph. *Diskretnyi Analiz* **3**(7), 25–30 (1964)
29. Welsh, D.J.A., Merino, C.: The Potts model and the Tutte polynomial. *J. Math. Phys.* **41**(3), 1127–1152 (2000)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Affiliations

Haris Aziz^{1,2} · **Péter Biró**^{3,4} · **Serge Gaspers**¹ · **Ronald de Haan**⁵ · **Nicholas Mattei**⁶ · **Baharak Rastegari**⁷

✉ Baharak Rastegari
b.rastegari@soton.ac.uk

Haris Aziz
haris.aziz@unsw.edu.au

Péter Biró
peter.biro@krtk.mta.hu

Serge Gaspers
sergeg@cse.unsw.edu.au

Ronald de Haan
me@ronalddhaan.eu

Nicholas Mattei
nsmattei@tulane.edu

- ¹ UNSW Sydney, Sydney, Australia
- ² Data61, Sydney, Australia
- ³ Hungarian Academy of Sciences, Budapest, Hungary
- ⁴ Corvinus University of Budapest, Budapest, Hungary
- ⁵ ILLC, University of Amsterdam, Amsterdam, The Netherlands
- ⁶ Tulane University, New Orleans, USA
- ⁷ University of Southampton, Southampton, UK