Manipulation complexity and gender neutrality in stable marriage procedures

Maria Silvia Pini[1], Francesca Rossi[1], K. Brent Venable[1], and Toby Walsh[2]

[1] Dipartimento di Matematica Pura ed Applicata, Università di Padova, Italy
E-mail: {mpini,frossi,kvenable}@math.unipd.it
[2] NICTA and UNSW, Sydney, Australia. Email: Toby.Walsh@nicta.com.au

Abstract. The stable marriage problem is a well-known problem of matching men to women so that no man and woman who are not married to each other both prefer each other. Such a problem has a wide variety of practical applications, ranging from matching resident doctors, to hospitals to matching students to schools. A well-known algorithm to solve this problem is the Gale-Shapley algorithm, which runs in quadratic time in the number of men/women.
It has been proven that stable marriage procedures can always be manipulated. Whilst the Gale-Shapley algorithm is computationally easy to manipulate, we prove that there exist stable marriage procedures which are NP-hard to manipulate. We also consider the relationship between voting theory and stable marriage procedures, showing that voting rules which are NP-hard to manipulate can be used to define stable marriage procedures which are themselves NP-hard to manipulate. Finally, we consider the issue that stable marriage procedures like Gale-Shapley favour one gender over the other, and we show how to use voting rules to make any stable marriage procedure gender neutral.

1 Introduction

The stable marriage problem (SMP) [12] is a well-known problem of matching the elements of two sets. Given n men and n women, where each person expresses a strict ordering over the members of the opposite sex, the problem is to match the men to the women so that there are no two people of opposite sex who would both rather be matched with each other than with their current partners. If there are no such people, all the marriages are said to be *stable*. Gale and Shapley [8] proved that it is always possible to solve the SMP and make all marriages stable, and provided an algorithm which is quadratic in the number of men/women, that can be used to find one of two particular but extreme stable marriages, the so-called *male optimal* or *female optimal* solution. The Gale-Shapley algorithm has been used in many real-life applications, as shown in [22], such as in systems for matching hospitals to resident doctors [24] and the assignment of primary school students in Singapore to secondary schools [28]. Variants of the stable marriage problem turn up in many domains. For example, the US Navy has a web-based multi-agent system for assigning sailors to ships [19].

One important issue is whether agents have an incentive to tell the truth or can manipulate the result by mis-reporting their preferences. However, Roth [23] has proved that *all* stable marriage procedures can be manipulated. He demonstrated a stable marriage problem with 3 men and 3 women which can be manipulated whatever stable marriage procedure we use. This result is in some sense analogous to the classical Gibbard-Satterthwaite [11, 27] theorem for voting theory, which states that all voting procedures are manipulable under modest assumptions provided we have 3 or more voters. For voting theory, Bartholdi, Tovey and Trick [3] proposed that computational complexity might be an escape: whilst manipulation is always possible, there are voting rules where it is NP-hard to find a manipulation. We might hope that computational complexity might also be a barrier to manipulate stable marriage procedures. Unfortunately, the Gale-Shapley algorithm is computationally easy to manipulate [28].

Another drawback of many stable marriage procedures such as the one proposed by Gale-Shapley is their bias towards one of the two genders. The stable matching returned by the Gale-Shapley algorithm is either male optimal (that is, the best possible for every man) but female pessimal (that is, the worst possible for every woman), or female optimal but male pessimal. It is often desirable to use stable marriage procedures that are gender neutral [20]. Such procedures return a stable matching that is not affected by swapping the men with the women. The goal of this paper is to study both the complexity of manipulation and gender neutrality in stable marriage procedures, and to design gender neutral procedures that are difficult to manipulate.

It is known that the Gale-Shapley algorithm is computationally easy to manipulate [28]. Our first contribution is to prove that if the male and female preferences have a certain form, it is computationally easy to manipulate *any* stable marriage procedure. We provide a universal polynomial time manipulation scheme that, under certain conditions on the preferences, guarantees that the manipulator marries his optimal stable partner irrespective of the stable marriage procedure used. On the other hand, our second contribution is to prove that, when the preferences of the men and women are unrestricted, there exist stable marriage procedures which are NP-hard to manipulate.

Our third contribution is to show that any stable marriage procedure can be made gender neutral, if applied twice (to the profile and its swapped version) and then choosing among the two stable marriages via a score-based method which exploits a bijection between scores and marriages. However, this may give a gender neutral stable matching procedure which is easy to manipulate.

Our final contribution is a stable matching procedure which is both gender neutral and NP-hard to manipulate. This procedure uses a voting rule that, considering the male and female preferences, helps to choose between stable matchings. In fact, it picks the stable matching that is most preferred by the most popular men and women. We prove that, if the voting rule is Single Transferable Vote (STV) or the hybrid plurality rule [1], which are NP-hard to manipulate, then the resulting stable matching procedure is both gender neutral and NP-hard to manipulate. We conjecture that other voting rules which are NP-hard to manipulate will give rise to stable matching procedures which are also gender neutral and NP-hard to manipulate. Thus, our approach shows how combining voting rules and stable matching procedures can be beneficial in two ways: by using preferences to discriminate among stable matchings and by providing a possible computational shield against manipulation.

2 Background

We now give some basic notions on stable marriage problems.

Definition 1 (profile). *Given n men and n women, a profile is a sequence of $2n$ strict total orders, n over the men and n over the women.*

Given a profile, the stable marriage problem (SMP) [8] is the problem of finding a matching between men and women. The goal is to marry the men to the women such that there are no two people of opposite sex who would both rather be married to each other than their current partners. If there are no such people, the matching is said to be stable.

Definition 2 (feasible partner). *Given an SMP P, a feasible partner for a man m (resp., a woman w) is a woman w (resp., a man m) such that there is a stable marriage for P where m and w are married.*

Stable marriages for an SMP form a lattice w.r.t. the men's or women's preferences. The top of this lattice is the stable matching where men (resp., women) are mostly satisfied. Conversely, the bottom is the stable matching where men's (resp., women's) preferences are least satisfied.

Definition 3 (male (resp., female) optimal matching). *Given an SMP P, a matching is male (resp., female) optimal iff every man (resp., woman) is paired with his (resp., her) highest ranked feasible partner in P.*

Definition 4 (male (resp., female) pessimal matching). *Given an SMP P, a matching is male (resp., female) pessimal iff every man (resp., woman) is paired with his (resp., her) lowest ranked feasible partner.*

Notice that male (resp., female) optimal and female (resp., male) pessimal stable matchings coincide [17].

2.1 The Gale-Shapley algorithm

The *Gale-Shapley algorithm* [8] is a well-known algorithm to solve the SMP problem. At the start of the algorithm, each person is free and becomes engaged during the execution of the algorithm. Once a woman is engaged she never becomes free again (although to whom she is engaged may change), but men can alternate between being free and being engaged. The following step is iterated until all men are engaged: choose a free man m, and let m propose to the most preferred woman w on his preference list, such that w has not already rejected m. If w is free, then w and m become engaged. If w is engaged to man m', then she rejects the man (m or m') that she least prefers, and becomes, or remains, engaged to the other man. The rejected man becomes, or remains, free. When all men are engaged, the engaged pairs are said to be paired and form the male optimal stable matching.

This algorithm needs a number of steps that is quadratic in n (that is, the number of men), and it guarantees that, if the number of men and women coincide, and all participants express a linear order over all the members of the other group, everyone gets married, and the returned matching is stable. Since the input is a profile, the algorithm is linear in the size of the input. Note that the pairing generated by the Gale-Shapley algorithm is male optimal and female pessimal. It would be the reverse, of course, if the roles of male and female participants in the algorithm were interchanged.

Example 1. Assume $n = 3$. Let $W = \{w_1, w_2, w_3\}$ and $M = \{m_1, m_2, m_3\}$ be respectively the set of women and men. The following sequence of strict total orders defines a profile:

- $m_1 : w_1 > w_2 > w_3$ (i.e., the man m_1 prefers the woman w_1 to w_2 to w_3),
- $m_2 : w_2 > w_1 > w_3$,
- $m_3 : w_3 > w_2 > w_1$,
- $w_1 : m_1 > m_2 > m_3$,
- $w_2 : m_3 > m_1 > m_2$,
- $w_3 : m_2 > m_1 > m_3$

For this profile, the Gale-Shapley algorithm returns the male optimal solution $\{(m_1, w_1), (m_2, w_2), (m_3, w_3)\}$. On the other hand, the female optimal solution is $\{(w_1, m_1), (w_2, m_3), (w_3, m_2)\}$. □

2.2 Gender neutrality and non-manipulability

Against this background, a desirable property of a stable marriage procedure is gender neutrality.

Definition 5 (gender neutral procedure [20]). *A stable marriage procedure is gender neutral if and only if, when we swap the men with the women, we get the same result.*

A related property is called peer indifference.

Definition 6 (peer indifferent property [20]). *A stable marriage procedure has the peer indifference property if the result is not affected by the order in which the members of the same sex are considered.*

The Gale-Shapley procedure is peer indifferent but it is not gender neutral. In fact, if we swap men and women in Example 1, we obtain the female optimal solution rather than the male optimal one.

Another useful property of a stable marriage procedure is its resistance to manipulation. Indeed, it would be desirable that lying would not lead to better results for the liar. A stable marriage procedure is *manipulable* if there is a way for one person to mis-report their preferences and obtain a result which is better than the one they would have obtained with everybody using his or her true preference.

Roth [23] has proven that stable marriage procedures can always be *manipulated*, i.e, that no stable marriage procedure always yields a stable outcome and gives agents the incentive to reveal their true preferences. He demonstrated a 3 men, 3 women profile which can be manipulated whatever stable marriage procedure we use. A similar result in a different context is the one by Gibbard and Satterthwaite [11, 27], that proves that all voting procedures [1] are manipulable when there are three or more candidates and the voting procedure satisfies non-imposition (for every candidate, there exist votes that would make that candidate win) and non-dictatorship (the rule does not simply always choose the most-preferred candidate of a single fixed voter).

In this context, Bartholdi, Tovey and Trick [3] proposed that computational complexity might be an escape: whilst manipulation is always possible, there are voting rules like, for example, Single Transferable Vote (STV) where it is NP-hard to find a manipulation [2]. This resistance to manipulation arises from the difficulty of inverting the voting rule and does not depend on other assumptions like the difficulty of discovering the preferences of the other voters. In this paper, we study whether computational complexity may also be an escape from the manipulability of stable marriage procedures. Our results are only initial steps to a more complete understanding of the computational complexity of manipulating stable matching procedures. As mentioned before, NP-hardness results only address the worst case and may not apply to preferences met in practice.

3 Manipulating stable marriage procedures

In the following, we will call a manipulation attempt by a participant p the mis-reporting of p's preferences.

Definition 7 (non-worsening manipulation). *A manipulation attempt by a participant p in a stable marriage procedure is non-worsening if the resulting marriage for p is better than or equal to the marriage obtained by using the true preferences of p.*

In what follows, it will be useful to distinguish the manipulation attempts that strictly improve the manipulator's satisfaction, which will be called successful.

Definition 8 (successful manipulation). *A manipulation attempt by a participant p in a stable marriage procedure is successful if the resulting marriage for p is better than the marriage obtained by using the true preferences of p.*

Definition 9 (manipulable). *A stable marriage procedure is manipulable if there is a profile with a successful manipulation attempt from a participant.*

The Gale-Shapley procedure, which, depending on which set of participants it is applied to, returns either the male optimal or the female optimal solutions, is computationally easy to manipulate [28]. However, besides these two extreme solutions, there may be many other stable matchings. Several procedures have been defined to return some of these other stable matchings [13]. Our first contribution is to show that, under certain conditions on the shape of the male and female preferences, *any* stable marriage procedure is computationally easy to manipulate.

Definition 10 (universally manipulable profile). *Consider a profile p and a woman w in such a profile. Let m be the male optimal partner with w in p, and r be the female optimal partner for w in p. Profile p is universally manipulable by w if the following conditions hold:*

- *in the men-proposing Gale-Shapley algorithm, w receives more than one proposal;*
- *there exists a woman v such that r is the male optimal partner for v in p;*
- *v prefers m to r;*
- *r's preferences are $\ldots > v > w > \ldots$;*
- *m's preferences $\ldots w > v > \ldots$.*

We now show that, if the profile is universally manipulable by a woman, then it is easy for this woman to give false preferences to get her female optimal partner with any stable matching procedure.

Theorem 1. *Consider any stable marriage procedure and any woman w. There is a polynomial non-worsening manipulation scheme that, for any profile p which is universally manipulable by w, produces the female optimal partner for w.*

Proof. Consider the manipulation attempt that moves the male optimal partner m of w to the lower end of w's preference ordering, obtaining the new profile p'. Consider now the behaviour of the men-proposing Gale-Shapley algorithm on p and p'. Since the profile is universally manipulable, w receives a proposal from m, that is, her male optimal partner, and from some other man o. In p she will decide for m, while in p' she will decide for o. At this point, in p', m will have to propose to the next best woman for him, that is, v, and she will accept because of the assumptions on her preference ordering. This means that r (who was married to v in p) now in p' has to propose to his next best choice, that is, w, who will accept, since w prefers r to o. So, in p', the male optimal partner for w, as well as her female optimal partner, is r, since, with respect to p, w has only lowered m who wasn't her female optimal anyway, and the others' preferences remained unchanged. This means that there is only one stable partner for w in p'. Therefore, any stable marriage procedure applied to p' must return r as the partner for w. Hence, if the profile is universally manipulable, then w can give false preferences p', by moving m to the far right, in a way such that the female pessimal and the female optimal partner for w in p' becomes the same as the female optimal partner for w in p, and thus she will get her female optimal partner, whatever stable marriage procedure is used. The procedure is polynomial since the Gale-Shapley algorithm is polynomial. $\square$

If a woman wants to manipulate a stable marriage procedure to obtain her female optimal partner, she can check if the profile is universally manipulable by her. This involves checking some conditions on the shape of male and female optimal preferences and simulating the Gale-Shapley algorithm to see whether she receives a proposal only from her male optimal or also from some other man. In the former case, she will not do the manipulation. Otherwise, she can give false preferences, by moving her male optimal partner to the far right, thus obtaining her female optimal partner, whatever stable marriage procedure is used.

Example 2. In a setting with 3 men and 3 women, consider the profile $\{m_1 : w_1 > w_2 > w_3;\ m_2 : w_2 > w_1 > w_3;\ m_3 : w_1 > w_2 > w_3;\}$ $\{w_1 : m_2 > m_1 > m_3;\ w_2 : m_1 > m_2 > m_3;\ w_3 : m_1 > m_2 > m_3;\}$. In this profile, the male optimal solution is $\{(m_1, w_1), (m_2, w_2), (m_3, w_3)\}$. This profile is universally manipulable by w_1. It is easy to see that w_1 can obtain her female optimal partner, that is m_2, by moving m_1 after m_3. Notice that this holds no matter what stable marriage procedure is used. In fact, if the male and female optimal partner for w_1 coincide, then there is only one stable partner for w_1 and thus every stable marriage procedure produces this partner.

The same profile is not universally manipulable by w_2 or w_3, since they receive just one proposal in the men-proposing Gale-Shapley algorithm. It is possible to show that woman w_2 cannot manipulate: trying to move m_2 after m_3 she gets a worse result. Also, woman w_3 cannot manipulate since her male optimal partner is her least preferred man. $\square$

We now give an example of an arbitrary-sized universally manipulable profile.

Example 3. Assume to have a profile over k men, denoted by $m_1, \ldots, m_k$, and k women, denoted by $w_1, \ldots, w_k$, such that:

- m_1's preferences are $w_1 > w_2 > \ldots$;
- m_2's preferences are $w_2 > w_1 > \ldots$;
- m_3's preferences are $w_1 > w_3 \ldots$;
- m_i's preferences are $w_i > \ldots\ ; \forall i, 4 \leq i \leq k$,
- w_1's preferences are $m_2 > m_1 > \ldots$;
- w_2's preferences $m_1 > m_2 \ldots$;
- w_j's preferences $m_j > \ldots;\ \ \forall j, 3 \leq j \leq k$.

It is easy to see that m_1 (resp., m_2) is the male optimal of w_1 (resp., w_2), m_2 is the female optimal of w_1, and that the profile satisfies all the requirements to be universally manipulable by w_1. Since the profile is universally manipulable by w_1, by Theorem 1, w_1 can give false preferences and obtain her female optimal partner. In fact, she can move m_1 at the end of her preference ordering. Doing so, in any stable matching procedure, she is matched with m_2, that is, her female optimal partner. $\square$

Restricting to universally manipulable profiles makes manipulation computationally easy. On the other hand, if we allow all possible profiles, there are stable marriage procedures that are NP-hard to manipulate. The intuition is simple. We construct a stable marriage procedure that is computationally easy to compute but NP-hard to invert. To manipulate, a man or a woman will essentially need to be able to invert the procedure to choose between the exponential number of possible preference orderings. Hence, the constructed stable marriage procedure will be NP-hard to manipulate.

Theorem 2. *There exist stable marriage procedures for which deciding the existence of a successful manipulation is NP-complete.*

Proof. We construct a stable marriage procedure which chooses between the male and female optimal solution based on whether the profile encodes a NP-complete problem and its polynomial witness. The manipulator's preferences define the witness. The other people's preferences define the NP-complete problem. Hence, the manipulator needs to be able to solve a NP-complete problem to be able to manipulate successfully. Deciding if there is a successful manipulation for this stable marriage procedure is clearly in NP since we can compute male and female optimal solutions in polynomial time, and we can check a witness to a NP-complete problem also in polynomial time.

Our stable marriage procedure is defined to work on $n+3$ men (m_1, m_2 and p_1 to p_{n+1}) and $n+3$ women (w_1, w_2 and v_1 to v_{n+1}). It returns the female optimal solution if the preferences of woman w_1 encode a Hamiltonian path in a directed graph encoded by the other women's preferences, otherwise it returns the male optimal solution. The 3rd to $n+2$th preferences of woman w_1 encode a possible Hamiltonian path in a n node graph. In particular, if the $2+i$th man in the preference ordering of woman w_1 for $i > 0$ is man p_j, then the path goes from vertex i to vertex j. The preferences of the women v_i for $i \leq n$ encode the graph in which we find this Hamiltonian path. In particular, if man p_j for $j < n+1$ and $j \neq i$ appears before man p_{n+1} in the preference list of woman v_i, then there is a directed edge in the graph from i to j. It should be noticed that any graph can be produced using this construction.

Given a graph which is not complete in which we wish to find a Hamiltonian path, we now build a special profile. Woman w_1 will be able to manipulate this profile successfully iff the graph contains a Hamiltonian path. In the profile, woman w_1 most prefers to marry man m_1 and then man m_2. Consider any pair of vertices (i, j) not in the graph. Woman w_1 puts man p_j at position $2+i$ in her preference order. She puts all other men in any arbitrary order. This construction will guarantee that the preferences of w_1 do not represent a Hamiltonian path. Woman w_2 most prefers to marry man m_2. Woman v_i most prefers to marry man p_i, and has preferences for the other men according to the edges from vertex i. Man m_1 most prefers woman w_2. Man m_2 most prefers woman w_1. Finally, man p_i most prefers woman v_i. All other unspecified preferences can be chosen in any way. By construction, all first choices are different. Hence, the male optimal solution has the men married to their first choice, whilst the female optimal solution has the women married to their first choice.

The male optimal solution has woman w_1 married to man m_2. The female optimal solution has woman w_1 married to man m_1. By construction, the preferences of woman w_1 do not represent a Hamiltonian path. Hence our stable matching procedure returns the male optimal solution: woman w_1 married to man m_2. The only successful manipulation then for woman w_1 is if she can marry her most preferred choice, man m_1. As all first choices are different, woman w_1 cannot successfully manipulate the male or female optimal solution. Therefore, she must manipulate her preferences so that she spells out a Hamiltonian path in her preference ordering, and our stable marriage procedure therefore returns the female optimal solution. This means she can successfully manipulate iff there is a Hamiltonian path. Hence, deciding if there is a successful manipulation is NP-complete. □

Note that we can modify the proof by introducing $O(n^2)$ men so that the graph is encoded in the tail of the preferences of woman w_2. This means that it remains NP-hard to manipulate a stable marriage procedure even if we collude with all but one of the women. It also means that it is NP-hard to manipulate a stable marriage procedure when the problem is imbalanced and there are just 2 women but an arbitrary number of men. Notice that this procedure is not peer indifferent, since it gives special roles to different men and women. However, it is possible to make it peer indifferent, so that it computes the same result if we rename the men and women. For instance, we just take the men's preferences and compute from them a total ordering of the women (e.g. by running an election with these preferences). Similarly, we take the women's preferences and compute from them a total ordering of the men. We can then use these orderings to assign indices to men and women. Notice also this procedure is not gender neutral. If we swap men and women, we may get a different result. We can, however, use the simple procedure proposed in the next section to make it gender neutral.

4 Gender neutrality

As mentioned before, a weakness of many stable marriage procedures like the Gale-Shapley procedure and the procedure presented in the previous section, is that they are not gender neutral. They may greatly favour one sex over the other. We now present a simple and universal technique for taking any stable marriage procedure and making it gender neutral. We will assume that the men and the women are named from 1 to n.

We can convert any stable marriage procedure μ into one that is gender neutral μ' by applying Algorithm gn to μ. This algorithm runs μ twice, once over a given profile and once over the profile in which men and women are swapped. This will give us two (not necessarily different) stable matchings, say M_1 and M_2. We then choose among these two stable matchings by exploiting the following bijection between matchings and vectors of integers: given a matching M, the corresponding vector $V(M) = [v_1, \ldots, v_n]$ is defined by $v_i = j$ if $(i, j) \in M$.

Between M_1 and M_2, we return the matching with the lexicographically smallest vector. If the two vectors coincide, $M_1 = M_2$ and we return this matching.

Algorithm 1: gn

Input: μ: a stable matching procedure;
Output: μ': a stable matching procedure;
$p \leftarrow$ a profile;
$M_1 \leftarrow \mu(p)$;
$p' \leftarrow$ SwapMenWomen(p);
$M_2 \leftarrow \mu(p')$;
if $V(M_1) \leq_{lex} V(M_2)$ **then**
  $\mu' \leftarrow M_1$;
else
  $\mu' \leftarrow M_2$;
return μ'

Theorem 3. *Let μ be a stable marriage procedure. Then, for any profile p, the procedure $\mu' = gn(\mu)$ returns a gender neutral stable matching.*

Proof. Let us assume that running procedure μ' on profile p we get the stable matching M. Let us now swap men and women in p, obtaining the new profile p'. Notice that M is either $\mu(p)$ or $\mu(p')$. W.l.o.g., assume $M = \mu(p)$. This means that $V(\mu(p)) \leq_{\text{lex}} V(\mu(p'))$. If we now apply μ' to p', we get two matchings: $\mu(p')$ and $\mu(p)$, since swapping men and women on p' generates p. Thus the returned matching will be $\mu(p)$, since we assumed that $V(\mu(p)) \leq_{\text{lex}} V(\mu(p'))$. Therefore μ' is gender neutral. $\square$

5 Voting Rules and Stable Marriage Procedures

We now show how to use voting rules to build stable marriage procedures which are both gender neutral and difficult to manipulate. The intuition behind the procedure is to choose between stable matchings according to the preferences of the most preferred men or women. In particular, we will pick the stable matching that is most preferred by the most popular men and women. A voting rule that ranks candidates is a social welfare function. Given such a social welfare function, we order the men using the women's preferences and order the women using the men's preferences. Using this ordering of the men (where a more preferred man is before a less preferred one), we construct a male score vector for a matching. The ith element of the male score vector is the integer j iff the ith man in this order is married to his jth most preferred woman. A large male score vector is a measure of dissatisfaction with the matching from the perspective of the more preferred men. A female score vector is computed in an analogous manner.

The overall score for a matching is the lexicographically largest of its male and female score vectors. A large overall score corresponds to dissatisfaction with the matching from the perspective of the more preferred men or women. Using the Gale Shapley algorithm, we compute both the male and female optimal stable matchings. We then choose between these two matchings according to which has the lexicographically least overall score. Since we compare scores lexicographically, this is guaranteed to be the stable matching from the lattice of stable matchings with the lowest score. In the event of a tie, we can use any gender neutral tie-breaking rule, such as the one described in Section 4, which relies on vectors V to choose between the two matchings. Let us call this stable matching procedure the *lexicographical minimal regret* (LMR) stable marriage procedure. In particular, when the social welfare function v is used to rank the men and women we will call it a *v-based* lexicographical minimal regret stable marriage procedure. It is easy to see that this procedure is gender neutral.

Theorem 4. *The lexicographical minimal regret stable marriage procedure is gender neutral.*

Proof. Given a set of stable matchings, let M be the matching returned by the LMR procedure. If we swap men and women, for each marriage in the set, what was the male score vector now is the female score vector, and vice versa. There are two cases. In the first case, the male and female score vectors are different. However, the larger score vector is always the same as before, although denoted by a different name. In the second case, the male and female score vectors are the same. We now tie-break with a gender-neutral procedure. For example, we can use the one described in Section 4. Therefore the chosen matching will remain the same. □

There exist voting rules, based on which this stable procedure is computationally hard to manipulate. Here we use STV [1] and the hybrid plurality rule [5] to rank the men and women. Both of these voting rules are NP-hard to manipulate. We conjecture that similar results hold for stable matching procedures using other voting rules which are NP-hard to manipulate. In the STV rule, each voter provides a total order on candidates and, initially, an individual's vote is allocated to his most preferred candidate. If no candidate has a majority, the candidate with the fewest votes is eliminated thus the votes are transferred to the second choices of the voters who had selected him as first choice. This step is repeated until some candidate has a majority. We rank the candidates according to the order in which they are eliminated.

Theorem 5. *It is NP-complete to decide if an agent can manipulate the STV-based lexicographical minimal regret stable marriage procedure.*

Proof. We adapt the reduction used to prove that constructive manipulation of the STV rule by a single voter is NP-hard [2]. We construct a profile in which every man and every woman have different first choices. Thus, in the male optimal stable matching, every man will marry his first choice whilst in the female optimal stable matching, every woman will marry her first choice. The (male) manipulator will try to choose between these two matchings by ordering the women appropriately.

In our proof, we need to consider how the STV rule treats ties. As all men and women have different first choices, the first round of STV to rank the women needs to tie break between all the men. In such tie breaks, we suppose that the alphabetically last woman is eliminated. Note that this means that we are not peer indifferent in general.

To prove membership in NP, we observe that a manipulation has a polynomial witness which is simply the preferences reported by the manipulator which give the desired result. To prove NP-hardness, we give a reduction from 3-COVER (also called X3C). Given a set S with $|S| = n$, subsets S_i with $i \in [1, m]$, $|S_i = 3|$ and $S_i \subset S$, we ask if there exists an index set I with $|I| = n/3$ and $\bigcup_{i \in I} S_i = S$.

We will construct a profile of preferences for the men so that the only possibility for STV to rank first one of only two women, w or y. The (male) manipulator h will try to vote strategically so that woman y is ranked first. In this case, we return the male optimal stable matching in which the manipulator marries his first choice. On the other hand, if w is ranked first, we return the female optimal stable matching in which the manipulator is not married to his first choice.

The following women participate in the problem:

- the two possible winners of the first STV election, w and y;
- "second line" in any election, a_i and b_i for $i \in [1, m]$;
- "third line" in any election, c_i and d_i for $i \in [1, m]$;
- "e-bloc", e_i for $i \in [0, n]$;
- "garbage collectors", g_i for $i \in [1, m]$;
- "dummy women", $z_{i,j,k}$ where $i \in [1, 19]$ and j and k depend on i as outlined in the description given shortly for the men's preferences (e.g. for $i = 1$, $j = 1$ and $k \in [1, 12m - 1]$ but for $i \in [6, 8]$, $j \in [1, m]$ and $k \in [1, 6m + 4j - 6]$).

The men's preferences are such that the dummy women are the first women eliminated by the STV rule, and that a_i and b_i are $2m$ out of the next $3m$ woman eliminated. In addition, let $I = \{i : b_i \textit{ is eliminated before } a_i\}$. Then the men's preferences will be constructed so that STV ranks woman y first if and only if I is a 3-COVER. The manipulator can ensure b_i is eliminated by the STV rule before a_i for $i \in I$ by placing a_i in the $i + 1$th position and b_i otherwise.

The men's preferences are constructed as follows (where preferences are left unspecified, they can be completed in any order):

- a man n with preference $(y, \ldots)$ and $\forall k \in [1, 12m-1]$ a man with $(z_{1,1,k}, y, \ldots)$;
- a man p with preference $(w, y, \ldots)$ and $\forall k \in [1, 12m - 2]$ a man with $(z_{2,1,k}, w, y, \ldots)$;
- a man q with preference $(e_0, w, y, \ldots)$ and $\forall k \in [1, 10m + 2n/3 - 1]$ a man with $(z_{3,1,k}, e_0, w, y, \ldots)$;
- $\forall j \in [1, n]$, a man with preference $(e_j, w, y, \ldots)$ and $\forall k \in [1, 12m - 3]$ a man with preference $(z_{4,j,k}, e_j, w, y, \ \ldots)$;

- $\forall j \in [1, m]$, a man r_j with preference $(g_j, w, y, \ldots)$ and $\forall k \in [1, 12m-1]$ a man with preference $(z_{5,j,k}, g_j,\ w, y, \ldots)$;
- $\forall j \in [1, m]$, a man with preference $(c_j, d_j, w, y, \ldots)$ and $\forall k \in [1, 6m+4j-6]$ a man with preference $(z_{6,j,k}, c_j, d_j,\ w, y, \ldots)$, and for each of the three k s.t. $k \in S_j$, a man with preference $(z_{7,j,k}, c_j, e_k, w, y, \ldots)$, and one with preference $(z_{8,j,k}, c_j, e_k, w, y, \ldots)$;
- $\forall j \in [1, m]$, a man with preference $(d_j, c_j, w, y, \ldots)$ and $\forall k \in [1, 6m + 4j - 2]$ a man with preference $(z_{9,j,k}, d_j, c_j,\ w, y, \ldots)$, one with preference $(z_{10,j,k}, d_j, e_0,\ w, y, \ldots)$, and one with $(z_{11,j,k}, d_j, e_0, w, y, \ldots)$;
- $\forall j \in [1, m]$, a man with preference $(a_j, g_j, w, y, \ldots)$ and $\forall k \in [1, 6m+4j-4]$ a man with preference $(z_{12,j,k}, a_j, g_j,\ w, y, \ldots)$, one with preference $(z_{13,j,k}, a_j, c_j,\ w, y, \ldots)$, one with preference $(z_{14,j,k}, a_j, b_j, w, y, \ldots)$, and one with preference $(z_{15,j,k}, a_j, b_j, w, y, \ldots)$.
- $\forall j \in [1, m]$, a man with preference $(b_j, g_j, w, y, \ldots)$ and $\forall k \in [1, 6m+4j-4]$ a man with preference $(z_{16,j,k}, b_j, g_j,\ w, y, \ldots)$, one with preference $(z_{17,j,k}, b_j, d_j,\ w, y, \ldots)$, one with preference $(z_{18,j,k}, b_j, a_j, w, y, \ldots)$, and one with preference $(z_{19,j,k}, b_j, a_j, w, y, \ldots)$.

Note that each woman is ranked first by exactly one man. We suppose that the manipulator's most preferred women is $z_{1,1,1}$ and then $z_{2,1,1}$. The women's preference are such that the manipulator is assured that he at least marries his second choice woman as this will be his female optimal partner. To manipulate the election, the manipulator needs to put $z_{1,1,1}$ first in his preferences and to report the rest of his preferences so that the result returned is the male optimal solution. As all woman are ranked first by exactly one man, the male optimal matching marries h with his first choice, $z_{1,1,1}$.

When we use STV to rank the women, $z_{i,j,k}$ are alphabetically last so are eliminated first by the tie-breaking rule. This leaves the following profile:

- $12m$ men with preference $(y, \ldots)$;
- $12m - 1$ men with preference $(w, y, \ldots)$;
- $10m + 2n/3$ men with preference $(e_0, w, y, \ldots)$;
- $\forall j \in [1, n]$, $12m - 2$ men with preference $(e_j, w, y,\ \ldots)$;
- $\forall j \in [1, m]$, $12m$ men with preference $(g_j, w, y,\ \ldots)$;
- $\forall j \in [1, m]$, $6m + 4j - 5$ men with preference $(c_j, d_j, w, y,\ \ldots)$, and for each of the three k such that $k \in S_j$, two men with preference $(c_j, e_k, w, y, \ldots)$;
- $\forall j \in [1, m]$, $6m + 4j - 1$ men with preference $(d_j, c_j, w, y,\ \ldots)$, and two men with preference $(d_j, e_0, w, y, \ldots)$,
- $\forall j \in [1, m]$, $6m + 4j - 3$ men with preference $(a_j, g_j, w, y,\ \ldots)$, a man with preference $(a_j, c_j, w, y, \ldots)$, and two men with preference $(a_j, b_j, w, y, \ldots)$;
- $\forall j \in [1, m]$, $6m + 4j - 3$ men with preference $(b_j, g_j, w, y,\ \ldots)$ a man with preference $(b_j, d_j, w, y, \ldots)$, and two men with preference $(b_j, a_j, w, y, \ldots)$.

At this point, the votes are identical (up to renaming of the men) to the profile constructed in the proof of Theorem 1 in [2]. Using the same argument as there, it follows that the manipulator can ensure that STV ranks woman y first instead of w if and only if there is a 3-COVER. The manipulation will place $z_{1,1,1}$ first in h's preferences. Similar to the proof of Theorem 1 in [2], the manipulation puts woman a_j in $j+1$th place and b_j otherwise where $j \in J$ and J is any index set of a 3-COVER.

The women's preferences are as follows:

- the woman y with preference $(n, \ldots)$;
- the woman w with preference $(q, \ldots)$;
- the woman $z_{1,1,1}$ with preference $(p, \ldots)$;
- the woman $z_{2,1,1}$ with preference $(h, \ldots)$;
- the women g_i with preference $(r_i, \ldots)$;
- the other women with any preferences which have all different first choices, and which ensure STV ranks r_0 first and r_1 second overall.

Each man is ranked first by exactly one woman. Hence, the female optimal stable matching is the first choice of the women.

The male score vector of the male optimal stable matching is $(1, 1, \ldots, 1)$. The female score vector of the male optimal stable matching is $(1, 2, \ldots)$ if the manipulation is successful and $(2, 1, \ldots)$ if it is not. Hence, the overall score vector of the male optimal stable matching is $(1, 2, \ldots)$ if the manipulation is successful and $(2, 1, \ldots)$ if it is not. On the other hand, the overall score vector of the female optimal stable matching is $(1, 3, \ldots)$. Hence the lexicographical minimal regret stable marriage procedure will return the male optimal stable matching iff there is a successful manipulation of the STV rule. Note that the profile used in this proof is not universally manipulable. The first choices of the man are all different and each woman therefore only receives one proposal in the men-proposing Gale-Shapley algorithm. □

After having shown that the lexicographical minimal regret stable marriage procedure, when based on STV, is NP hard to manipulate and gender neutral, we will do the same for another voting rule, called the hybrid plurality rule. The hybrid plurality rule uses one round of the cup rule and then applies the plurality rule [5]. In other words, candidates are paired and from each pair a winner remains, according to what the majority says about the two candidates in the pair. Thus, half of the candidates remain after this first phase. Then, a winner among these candidates is chosen by majority. This simple hybridization has been shown to make manipulation NP-hard. For simplicity, we suppose that the schedule for the cup rule is fixed in advance. Fixing the schedule after the manipulator has voted will only increase the computational complexity.

Theorem 6. *It is NP-complete to decide if an agent can manipulate the lexicographical minimal regret stable marriage procedure based on the hybrid plurality rule.*

Proof. We adapt the reduction used in the proof of Theorem 2 of [5], which shows that constructive manipulation of the hybrid plurality rule by a single voter is NP-hard. We again construct a profile in which every man and every woman have different first choices. Thus, in the male optimal stable matching, every man will marry his first choice whilst in the female optimal stable matching, every woman will marry her first choice. The (male) manipulator will try to choose between these two matchings by ordering the women appropriately.

To prove membership in NP, we again observe that a manipulation has a polynomial witness. To prove NP-hardness, we give a reduction from 3-SAT. Let L be the literals in the 3-SAT problem (e.g. v and $\neg v$), K be the clauses, $l = |L|$ and $k = |K|$. We introduce the following women: p, $C_L = \{c_i \mid i \in L\}$, $D_K = \{d_i \mid 0 \leq i < k\}$, and four dummy sets of women, $X = \{x_i, x_{-i} \mid 1 \leq i \leq 4k+1\}$, $Y = \{y_{i,j}, y_{-i,-j} \mid 1 \leq i \leq k, 1 \leq j \leq 4k-1\}$, $Z = \{z_i, z_{-i} \mid 1 \leq i \leq 3k\}$ and $W = \{w_i \mid 1 \leq i \leq k+1\}$.

The manipulating man h will try to vote strategically so that woman p is ranked first when the hybrid plurality rule is applied to the men's preferences. In this case, we return the male optimal stable matching. On the other hand, if p is not ranked first, we will return the female optimal stable matching (which is less good for h). Similar to [5], we write (A, b) if, for all $a \in A$, the agent prefers a to b, but the order between the elements of A is not important. To simplify notation, we write A' for every remaining element of A. For instance, (a, A') stands for $(a, A - \{a\})$. This states that a is preferred to every other element in A but there is no constraint between the other elements in A.

The men's preferences are constructed as follows:

- a man a with preference $(p, C_L, D_K, X, Y, Z, W)$ and for each $i \in [1, 4k+1]$ a man e_i with preference $(x_i, p, C_L, D_K, X', Y, Z, W)$;
- for each $i \in [1, k]$, a man b_i with preference $(d_i, D'_k, C_L, p, X, Y, Z, W)$ and $4k-1$ men with preferences $(y_{i,j}, d_i, D'_k, C_L, p, X, Y', Z, W)$;
- for each $i \in [1, k]$ and $\varphi_i \in K$, a man with preference $(\{c_l \mid l \in \varphi_i\}, d_k, \{c_l \mid l \in L, l \notin \varphi_i\}, D'_k, p, X, Y, Z, W)$ and 3 men with preferences $(z_{i,j}, \{c_l \mid l \in \varphi_i\}, d_k, \{c_l \mid l \in L, l \notin \varphi_i\}, D'_k, p, X, Y, Z', W)$.

Note that there are an even number of men. We therefore are able to rank within the otherwise unconstrained preferences for C_L so that c_l and c_{-l} are tied for every pair of literals in L. The manipulator's job will be to decide if c_l is ranked ahead of c_{-l}. As c_l plays against c_{-l} in the initial preround, the manipulator will get to decide which of these two literals survives the preround.

Each woman is ranked first by exactly one man. Similarly, each man is ranked first by exactly one woman. Hence, the male optimal stable matching will be the first choice of the men, and the female optimal stable matching will be the first choice of the women.

The schedule of the cup puts c_l against c_{-l} for every pair of literals in L, d_i against w_i, p against w_{k+1}, x_i against x_{-i}, $y_{i,j}$ against $y_{-i,-j}$, and z_i against z_{-i}. We arrange the (otherwise unconstrained) votes within the sets X, Y and Z so that x_i necessarily loses to x_{-i} in the preround, $y_{i,j}$ loses to $y_{-i,-j}$, and z_i loses to z_{-i}. Note that, as W is last place in each vote, W is also necessarily eliminated in the preround.

After half the women are eliminated by the preround, we are left with the following profile:

- $4k + 2$ men with preference $(p, C_L, D_K, X^-, Y^-, Z^-)$
- for each $i \in [1, k]$, $4k$ men with preferences $(d_i, D'_k, C_L, p, X^-, Y^-, Z^-)$;
- for each $i \in [1, k]$ and $\varphi_i \in K$, 4 men with preferences $(\{c_l \mid l \in \varphi_i\}, d_k, \{c_l \mid l \in L, l \notin \varphi_i\}, D'_k, p, X^-, Y^-, Z^-)$.

where $X^- = \{x_{-i} \mid 1 \leq i \leq 4k+1\}$, $Y^- = \{y_{-i,-j} \mid 1 \leq i \leq k, 1 \leq j \leq 4k-1\}$, and $Z^- = \{z_{-i} \mid 1 \leq i \leq 3k\}$.

At this point, the preferences are identical (up to renaming of the men and the addition of X^-, Y^- and Z^-) to the profile constructed in the proof of Theorem 2 in [5]. Note that p gets $4k + 2$ votes. If, for some $\varphi_i \in K$, all $c_l \in \varphi_i$ are eliminated in the preround, then d_k gets $4k+4$ votes and p must lose. Otherwise d_k gets just $4k$ votes and p wins. The manipulator must order the literals c_l and c_{-l} to ensure this. The manipulator puts c_l ahead of c_{-l} iff l is true in the corresponding satisfying assignment. Thus, the manipulator can ensure that p wins iff the original 3SAT problem is satisfiable.

We now turn to the women's preferences. These are as follows:

- woman p with preference $(a, \ldots)$;
- women d_{i+1} with preferences $(b_i, b_{i+1}, \ldots)$ for $1 \in [1, k)$ and woman d_1 with preference $(b_k, b_1, \ldots)$;
- woman c_1 with preference $(e_1, \ldots)$;
- the other women with any preferences in which the first choices are all different, and which ensure that the hybrid plurality rule ranks a first and then e_1.

As in the case for the STV rule (see proof of Theorem 5), the overall score vector of the male optimal stable matching is $(1, 2, \ldots)$ if the manipulation is successful and $(2, \ldots)$ if it is not. Similarly, the overall score vector of the female optimal stable matching is $(1, i, \ldots)$ where $i > 2$. Hence the lexicographical minimal regret stable marriage procedure will return the male optimal stable matching iff there is a successful manipulation of the hybrid plurality rule. □

We can thus see how the proposed matching procedures (the one based on STV and the one based on the hybrid plurality rule) are reasonable and appealing. In fact, they allow us to discriminate among stable matchings according to the men and women's preferences, and they are difficult to manipulate while ensuring gender neutrality.

6 Related work

In [20] fairness of a matching procedure is defined in terms of four axioms, two of which are gender neutrality and peer indifference. Then, the existence of a matching procedures satisfying all or a subset of the axioms is considered in terms of restrictions on preference orderings. Here, instead, we propose a preprocessing step that allows to obtain a gender neutral matching procedure from any matching procedure without imposing any restrictions on the preferences in the input.

A detailed description of results about manipulation of stable marriage procedures can be found in [14]. In particular, several early results [6, 7, 9, 23] indicated the futility of men lying, focusing later work mostly on strategies in which the women lie. Gale and Sotomayor [10] presented the manipulation strategy in which women truncate their preference lists. Roth and Vande Vate [26] discussed strategic issues when the stable matching is chosen at random, proposed a truncation strategy and showed that every stable matching can be achieved as an equilibrium in truncation strategies. We instead do not allow the elimination of men from a woman's preference ordering, but permit reordering of the preference lists.

Teo et al. [28] suggested lying strategies for an individual woman, and proposed an algorithm to find the best partner with the male optimal procedure. We instead focus on the complexity of determining if the procedure can be manipulated to obtain a better result. Moreover, we also provide a universal manipulation scheme that, under certain conditions on the profile, assures that the female optimal partner is returned.

Also Kobayashi et al. [18] study a strategic issue in the stable marriage model with complete preference lists. Given complete preference lists of men over women and a matching, they present a necessary and sufficient condition for the existence of a set of preference lists of women over men, such that the men-proposing Gale-Shapley algorithm applied to these lists produces the given matching. We are instead interested in situations where there exists a specific woman, that is the only deceitful agent, and that she knows the preferences of all the other agents.

Coalition manipulation is considered in [14]. Huang shows how a coalition of men can get a better result in the men-proposing Gale-Shapley algorithm. By contrast, we do not consider a coalition but just a single manipulator, and do not consider just the Gale-Shapley algorithm.

7 Conclusions

We have studied the manipulability and gender neutrality of stable marriage procedures. We first looked at whether, as with voting rules, computationally complexity might be a barrier to manipulation. It was known already that one prominent stable marriage procedure, the Gale-Shapley algorithm, is computationally easy to manipulate. We proved that, under some simple, yet demanding, restrictions on agents' preferences, *all* stable marriage procedures are in fact easy to manipulate. Our proof provides a universal manipulation method that an agent can use to improve his result. On the other hand, when preferences are unrestricted, we proved that there exist stable marriage procedures which are NP-hard to manipulate.

We also showed how to use a voting rule to choose between stable matchings. In particular, we gave a stable marriage procedure which picks the stable matching that is most preferred by the most popular men and women. This procedure inherits the computational complexity of the underlying voting rule. Thus, when the STV voting rule, or the hybrid plurality rule (which are both NP-hard to manipulate), is used to compute the most popular men and women, the corresponding stable marriage procedure is NP-hard to manipulate. Another desirable property of stable marriage procedures is gender neutrality. Our procedure of turning a voting rule into a stable marriage procedure is gender neutral.

This study of stable marriage procedures is only an initial step to understanding if computational complexity might be a barrier to manipulation. Many questions remain to be answered. For example, the preferences in practice may be highly correlated. Men may have similar preferences for many of the women. Are such profiles computationally difficult to manipulate? As a second example, it has been recently recognised (see, for example, $[29, 31, 21, 4]$) that worst-case results may represent an insufficient barrier against manipulation since they may only apply to problems that are rare. Are there stable marriage procedures which are difficult to manipulate on average? There are also many interesting and related questions connected with privacy and mechanism design. For instance, how do we design a decentralised stable marriage procedure which is resistant to manipulation and in which the agents do not share their preference lists? As a second example, how can side payments be used in stable marriage procedures to prevent manipulation?

References

1. K.J. Arrow, A.K. Sen and K. Suzumura. Handbook of Social Choice and Welfare. North Holland, Elsevier, 2002
2. J. Bartholdi and J. Orlin. Single transferable vote resists strategic voting. In *Social Choice and Welfare, 8(4):341-354*, 1991.
3. J.J. Bartholdi, C. A. Tovey and M. A. Trick, The computational difficulty of manipulating an election. In *Social Choice and Welfare 6(3):227-241*, 1989.
4. V. Conitzer and T. Sandholm. Nonexistence of voting rules that are usually hard to manipulate. In*Proc. AAAI'06*, AAAI Press, 2006.
5. V. Conitzer and T. Sandholm. Universal Voting Protocol Tweaks to Make Manipulation Hard. In *Proc. IJCAI'03*, pp. 781-788, 2003.
6. G. Demange, D. Gale, and M. Sotomayor. A further note on the stable matching problem. In *Discrete Applied Mathematics, 16:217-222*, 1987.
7. L. Dubins and D. Freedman. Machiavelli and the Gale-Shapley algorithm. In *American Mathematical Monthly, 88:485-494*, 1981.
8. D. Gale, L. S. Shapley. College Admissions and the Stability of Marriage. In *Amer. Math. Monthly, 69:9-14*, 1962.
9. D. Gale and M. Sotomayor. Some remarks on the stable matching prob- lem. *Discrete Applied Mathematics, 11:223-232*, 1985.
10. D. Gale and M. Sotomayor. Machiavelli and the stable matching problem. *American Mathematical Monthly, 92:261-268*, 1985.
11. A. Gibbard. Manipulation of Voting Schemes: A General Result. In *Econometrica, 41(3):587-601*, 1973.
12. D. Gusfield and R. W. Irving. The Stable Marriage Problem: Structure and Algorithms. MIT Press, Boston, Mass.,1989.
13. D. Gusfield. Three fast algorithms for four problems in stable marriage. In *SIAM J. of Computing, 16(1)*, 1987.

14. C.-C. Huang. Cheating by men in the Gale-Shapley stable matching algorithm. In *Proc. ESA'06*, pp. 418-431, Springer-Verlag, 2006.
15. Robert W. Irving, P. Leather and D. Gusfield, An efficient algorithm for the "optimal" stable marriage, In *JACM 34(3):532-543*,1987.
16. B. Klaus and F. Klijin. Procedurally fair and stable matching. In *Economic Theory, 27:431-447*, 2006.
17. D. E. Knuth. Mariages stables et leurs relations avec d'autres problmes combinatoires. (French) Introduction l'analyse mathmatique des algorithmes. Collection de la Chaire Aisenstadt. Les Presses de l'Universit de Montral, Montreal, Que., 1976.
18. H. Kobayashi and T. Matsui. Successful manipulation in stable marriage model with complete preference lists. *In Proc. MATCH-UP workshop: a satellite workshop of ICALP'08*, pp. 17-22, 2008.
19. J. Liebowitz and J. Simien. Computational efficiencies for multi-agents: a look at a multi-agent system for sailor assignment. In *Electronic government: an International Journal 2 (4):384-402*, 2005.
20. F. Masarani and S.S. Gokturk. On the existence of fair matching algorithms. In *Theory and Decision, 26:305-322*, Kluwer, 1989.
21. A.D. Procaccia and J.S. Rosenschein. Junta Distributions and the Average-Case Complexity of Manipulating Elections. In *JAIR 28:157-181*, 2007.
22. A. Roth. Deferred Acceptance Algorithms: History, Theory, Practice, and Open Questions. In *International Journal of Game Theory, Special Issue in Honor of David Gale on his 85th birthday, 36:537-569* 2008.
23. A. Roth. The Economics of Matching: Stability and Incentives. In *Mathematics of Operations Research, 7:617-628*, 1982.
24. A. Roth. The evolution of the labor market for medical interns and residents: a case study in game theory. In *Journal of Political Economy 92:991-1016*, 1984.
25. A. Roth and M. Sotomayor. Two-sided matching: A study in game-theoretic modeling and analysis. *Cambridge University Press*, 1990.
26. A. Roth and J. H. Vande Vate. Incentives in two-sided matching with random stable mechanisms. In *Economic Theory, 1(1):31-44*, 1991.
27. M. A. Satterthwaite. Strategy-proofness and Arrow's conditions: Existence and correspondence theorems for voting procedures and social welfare function. In *Economic Theory, 10(3):187-217*, 1975.
28. C.-P. Teo, J. Sethuraman and W.-P. Tan. Gale-Shapley Stable Marriage Problem Revisited: Strategic Issues and Applications. In *Management Science, 47(9):1252-1267*, 2001.
29. L. Xia, V. Conitzer. Generalized scoring rules and the frequency of coalitional manipulability. In *Proc. ACM Conference on Electronic Commerce 2008*, pp. 109-118, ACM Press, 2008.
30. T. Walsh. Where are the really hard manipulation problems? The phase transition in manipulating the veto rule. In *Proc. IJCAI'09*, 2009
31. M. Zuckerman, A. D. Procaccia, J. S. Rosenschein. Algorithms for the coalitional manipulation problem. In *Artificial Intelligence 173(2):392-412*, 2009.