Algorithmica (2020) 82:1136–1188
https://doi.org/10.1007/s00453-019-00636-y

Stable Matchings with Covering Constraints: A Complete Computational Trichotomy

Matthias Mnich[1,2] · **Ildikó Schlotter**[3]

Received: 2 October 2017 / Accepted: 5 October 2019 / Published online: 6 February 2020
© The Author(s) 2020

Abstract
Stable matching problems with lower quotas are fundamental in academic hiring and ensuring operability of rural hospitals. Only few tractable (polynomial-time solvable) cases of stable matching with lower quotas have been identified; most such problems are NP-hard and also hard to approximate (Hamada et al. in Algorithmica 74(1):440–465, 2016). We therefore consider stable matching problems with lower quotas under a relaxed notion of tractability, namely fixed-parameter tractability. By cloning hospitals we focus on the case when all hospitals have upper quota equal to 1, which generalizes the setting of "arranged marriages" first considered by Knuth (Mariages stables et leurs relations avec d'autres problèmes combinatoires, Les Presses de l'Université de Montréal, Montreal, 1976). We investigate how a set of natural parameters, namely the maximum length of preference lists for men and women, the number of distinguished men and women, and the number of blocking pairs allowed determine the computational tractability of this problem. Our main result is a complete complexity trichotomy: for each choice of parameters we either provide a polynomial-time algorithm, or an NP-hardness proof and fixed-parameter algorithm, or NP-hardness proof and W[1]-hardness proof. As corollary, we negatively answer a question by Hamada et al. (Algorithmica 74(1):440–465, 2016) by showing fixed-parameter intractability parameterized by optimal solution size. We also classify all cases of one-sided constraints where only women may be distinguished.

Keywords Stable marriage · Lower quotas · Fixed-parameter algorithms

Mathematics Subject Classification 68R10 · 91B68

A preliminary version of the results presented in here appeared at the Proceedings of the 10th International Symposium on Algorithmic Game Theory [39]. Research of M. Mnich supported by ERC Starting Grant 306465 (BeyondWorstCase) and DFG Grant MN 59/4-1. Research of I. Schlotter supported by the Hungarian National Research Fund (OTKA Grants Nos. K-108383 and K-108947).

Extended author information available on the last page of the article

1 Introduction

The Stable Marriage (SM) problem is a fundamental problem first studied by Gale and Shapley [19]. An instance of SM consists of a set $\mathcal{M}$ of men, a set $\mathcal{W}$ of women, and a preference list for each person ordering members of the opposite sex. We aim to find a stable matching, i.e., a matching for which there exists no pair of a man and a woman who prefer each other to their partners given by the matching; such a pair is called a blocking pair. Gale and Shapley proved [19] that any instance of SM admits at least one stable matching, and gave a polynomial-time algorithm, known as the Gale–Shapley algorithm, to find one. Gale and Shapley also considered the many-to-one extension of SM, known as the Hospitals/Residents (HR) problem. In HR, the two sets $\mathcal{H}$ and $\mathcal{R}$ that correspond to men and women in the SM problem are residents and hospitals, respectively. Each hospital $H \in \mathcal{H}$ has an upper quota on the number of residents in $\mathcal{R}$ that it can accept. For HR it still holds true that a stable matching always exists, and can be found efficiently.

An extension of HR that is motivated by several real-world applications is the Hospitals/Residents with Lower Quota (HRLQ) problem, where hospitals declare both *lower* and *upper quotas* which bound the number of residents they can accept; as before, residents rank hospitals and vice versa. Now it is no longer true that a stable assignment always exists. The possible non-existence of stable assignments motivates the design of algorithms that find an assignment with a minimum number of blocking pairs; this is the task of the HRLQ problem. Indeed, the HRLQ problem and its variants have recently gained quite some interest from the algorithmic community [4, 8,16,22,24,26,30,40,44,45]. In his book, Manlove [36, Chapter 5.2] devotes an entire chapter to the algorithmics of different versions of the HRLQ problem.

The reason for this high interest in HRLQ is explained by its importance in several real-world matching markets [17,18,42] such as school admission systems, centralized assignment of residents to hospitals, or of cadets to military branches. Lower quotas are a common feature of such admission systems. Their purpose is often to remedy the effects of under-staffing that are explained by the well-known Rural Hospitals Theorem [20]: as an example, governments usually want to assign at least a small number of medical residents to each rural hospital to guarantee a minimum service level. Minimum quotas are also discussed in controlled school choice programs [13,33,43] where students are divided into a small number of types, and schools set lower bounds for each type. Such models can represent various forms of affirmative actions taken by schools to, e.g., admit a certain number of minority students [13]. Another example is the German university admission system for admitting students to highly oversubscribed subjects, where a certain percentage of study places is assigned according to high school grades or waiting time [43]. But lower quotas may also arise due to financial considerations: for instance, a business course with too few (tuition-paying) attendees may not be profitable. Certain aspects of airline preferences for seat upgrade allocations can be also modelled by lower quotas [33].

Much of the algorithmic research found that the HRLQ problem (in its different variants) is NP-hard, and thus considered intractable. A common approach then to identify tractable (polynomial-time solvable) cases of HRLQ; this avenue has lead to several beautiful algorithms [16,26,45]. However, restricting to polynomial-time solvability

necessarily means (if $\mathsf{P} \neq \mathsf{NP}$) that some original features of HRLQ must be restricted more or less, which may be undesirable in the application. Another approach aimed at addressing the intractability of HRLQ was the design of approximation algorithms. Unfortunately, it turns out that non-trivial approximation algorithms for HRLQ are highly unlikely to exist: Hamada et al. [24] showed that, unless $\mathsf{P} = \mathsf{NP}$, no algorithm with approximation guarantee $(|\mathcal{H}| + |\mathcal{R}|)^{1-\varepsilon}$ can exist for any $\varepsilon > 0$ (which they complement by an algorithm with approximation guarantee $(|\mathcal{H}| + |\mathcal{R}|)$). In light of their strong inapproximability bounds, Hamada et al. explicitly suggested to consider a more relaxed notion of tractability for HRLQ, namely, fixed-parameter tractability. They particularly asked whether HRLQ is fixed-parameter tractable parameterized by the minimum number b of blocking pairs over all matchings meeting all lower and upper quota requirements.

In this paper we follow this avenue, and study the fixed-parameter tractability of HRLQ. It allows us to provide a fine-grained analysis of HRLQ, and the design of efficient algorithms for NP-hard variants of HRLQ for small parameter values. Our main focus will be the case of HRLQ when each hospital has unit upper quota. The reason is that by the frequently applied method of "cloning" hospitals, stable instances of HRLQ reduce to the case where each hospital has unit upper quota—and thus we can reduce stable instances of HRLQ to the SM setting with lower quotas. In fact, this is equivalent to the special case of SM where only a subset of women (or, equivalently, men) are distinguished by having also a unit lower quota. From now on, we refer to HRLQ with unit upper quotas as SMC (where C means that we have to *cover* the women/men who have a unit lower quota), to the special case of SMC with one-sided covering constraints, linking SMC and HRLQ, as SMC-1. So formally, in SMC a set $\mathcal{W}^\star$ of women and a set $\mathcal{M}^\star$ of men are distinguished, and a feasible matching is one where each person in $\mathcal{W}^\star \cup \mathcal{M}^\star$ gets matched. By the Rural Hospitals Theorem [20] we know that the set of unmatched men and women is the same in all stable matchings, so clearly, feasible stable matchings may not exist. Thus, we define the task in SMC as finding a feasible matching with a minimum number of blocking pairs.

Apart from the recent interest in HRLQ, its reduction SMC also serves as a "modern version" of a classical problem first introduced by Donald Knuth. Knuth [31] considered SM with "arranged marriages", which are a set $\mathcal{X}$ of man-woman pairs that must be matched with each other. He showed that the Gale–Shapley algorithm can be extended to decide the existence of a stable matching still in time $O(n^2)$ for n-person instances, but in case of absence of a stable matching, minimizing the number of blocking pairs is NP-hard. Now in SMC, one does not prescribe any more which woman has to marry which man, but only requires certain women and men to marry without dictating their partner. There is a natural Turing reduction from SMC to the variant considered by Knuth. Coarse analysis into polynomial-time solvable and NP-hard cases has been studied by several researchers [1].

Another motivation for studying the SMC problem comes from the following scenario that we dub *Control for Stable Marriage*. Consider a two-sided market where each participant of the market expresses its preferences over members of the other party, and some central agent (e.g., a government) performs the task of finding a stable matching in the market. It might happen that this central agency wishes to apply a certain *control* on the stable matching produced: it may favour some participants by

trying to assign them a partner in the resulting matching. Such a behaviour might be either malicious (e.g., the central agency may accept bribes and thus favour certain participants) or beneficial (e.g., it may favour those who are at disadvantage, like handicapped or minority participants). However, there might not be a stable matching that covers all participants the agency wants to favour; thus arises the need to produce a matching that is *as stable as possible* among those that fulfil our constraints—the most natural aim in such a case is to minimize the number of blocking pairs in the produced matching, which yields exactly the SMC problem. Similar control problems have been extensively studied in the area of social choice for voting systems [7, Chapter 7] and recently also for fair division scenarios [2], but have not yet been considered in connection to stable matchings.

1.1 Our Results

We provide an extensive algorithmic analysis of the SMC problem and its special case SMC-1. In our analysis, we examine how different aspects of the input influence the tractability of these problems. To this end, we apply the framework of parameterized complexity, which deals with computationally hard problems and focuses on how certain *parameters* of a problem instance influence its tractability; for background, we refer to the book by Cygan et al. [11]. We aim to design so-called *fixed-parameter algorithms*, which perform well in practice if the value of the parameter on hand is small (for the precise definitions, see Sect. 2).

The parameters we consider are

- the number b of blocking pairs allowed,
- the number $|\mathcal{W}^\star|$ of women with covering constraint,
- the number $|\mathcal{M}^\star|$ of men with covering constraint,
- the maximum length $\Delta_{\mathcal{W}}$ of women's preference lists, and
- the maximum length $\Delta_{\mathcal{M}}$ of men's preference lists.

The choice of each of these parameters is motivated by the aforementioned applications. For instance, we seek matchings where ideally no blocking pairs at all or at least only few of them appear, to ensure stability of the matching and happiness of those getting matched. The number of women/men with covering constraints corresponds, for instance, to the number of rural hospitals for which a minimum quota specifically must be enforced, which we can expect to be small among the set of all hospitals accepting medical residents. Finally, preference lists of hospitals and residents can be expected to be small, as each hospital might not rank many more candidates than the number of positions it has to fill, whereas residents might rank only their top choices of hospitals.

We investigate in detail how these parameters influence the complexity of the SMC problem. A *parameterized restriction of SMC* with respect to the set $S = \{b, |\mathcal{W}^\star|, |\mathcal{M}^\star|, \Delta_{\mathcal{M}}, \Delta_{\mathcal{W}}\}$ means a (possibly parameterized) special case of SMC where each element of S is either restricted to be some constant integer, or regarded as a parameter, or left unbounded. Intuitively, these different choices for the elements of S correspond to their expected "range" in applications, from very small to mid-range to large (compared to the size of the entire system). By considering all combinations,

Table 1 Summary of our results for Stable Marriage with Covering Constraints

Constants	Parameters	Complexity
$\lvert\mathcal{M}^\star\rvert = 0, \lvert\mathcal{W}^\star\rvert = 0$		in P (Gale–Shapley algorithm)
$\lvert\mathcal{M}^\star\rvert = 0, \lvert\mathcal{W}^\star\rvert, \Delta_\mathcal{M}$		in P (Theorem 7)
$\lvert\mathcal{M}^\star\rvert, \lvert\mathcal{W}^\star\rvert, \Delta_\mathcal{M}, \Delta_\mathcal{W}$		in P (Theorem 7)
$\lvert\mathcal{M}^\star\rvert = 0, \Delta_\mathcal{M} \leq 2$		in P (Theorem 8)
$\Delta_\mathcal{W} \leq 2, \Delta_\mathcal{M} \leq 2$		in P (Observation 2)
b		in P (Observation 1)
$\lvert\mathcal{M}^\star\rvert = 0, \Delta_\mathcal{W} = 2, \Delta_\mathcal{M} \geq 3$		NP-hard (Theorem 9)
$\lvert\mathcal{W}^\star\rvert = 1, \Delta_\mathcal{W} = 2$		NP-hard (Theorem 10)
$\lvert\mathcal{M}^\star\rvert = 0, \Delta_\mathcal{W} \geq 3, \Delta_\mathcal{M} \geq 3, \Delta^\star = 1$	$b + \lvert\mathcal{W}^\star\rvert$	W[1]-hard (Theorem 2)
$\lvert\mathcal{M}^\star\rvert = 0, \lvert\mathcal{W}^\star\rvert \geq 1, \Delta_\mathcal{W} \geq 3, \Delta^\star = 1$	$b + \Delta_\mathcal{M}$	W[1]-hard (Theorem 5)
$\Delta_\mathcal{W} \leq 2$	$\lvert\mathcal{W}^\star\rvert + \lvert\mathcal{M}^\star\rvert$	FPT (Theorem 11)
$\Delta_\mathcal{W} \leq 2$	b	FPT (Corollary 3)
$\Delta_\mathcal{W} \leq 2$	$\lvert\mathcal{W}^\star\rvert + \Delta_\mathcal{M}$	FPT (Theorem 12)

Here, $\Delta^\star$ denotes the maximum length of the preference list of any distinguished person

we can flexibly model the whole range of applications mentioned above. We can even cover some cases of *master lists*, where all men's preference lists are restrictions of the exact same total order over the women, and likewise all women's preference lists are restrictions of the exact same total order over the men.

Theorem 1 *Any parameterized restriction of SMC with respect to* $\{b, |\mathcal{W}^\star|, |\mathcal{M}^\star|, \Delta_\mathcal{M}, \Delta_\mathcal{W}\}$ *is in P, or NP-hard and fixed-parameter tractable, or NP-hard and W[1]-hard with the given parameterization,*[1] *and is covered by one of the results shown in Table 1.*

In particular, SMC is W[1]-hard parameterized by $b + |\mathcal{W}^\star|$*, even if there are no distinguished men (i.e.,* $|\mathcal{M}^\star| = 0$*), there is a master list over men as well as one over women,* $\Delta_\mathcal{M} = 3$*,* $\Delta_\mathcal{W} = 3$ *and each distinguished woman finds only a single man acceptable.*

We give a decision diagram in Sect. 7 to show that the presented results indeed cover all restrictions of SMC with respect to $\{b, |\mathcal{W}^\star|, |\mathcal{M}^\star|, \Delta_\mathcal{M}, \Delta_\mathcal{W}\}$. Table 1 summarizes our results on the complexity of SMC. Note that some results are implied directly by the symmetrical roles of men and women in SMC, and thus are not stated explicitly. Here and later, we assume for simplicity that $\Delta_\mathcal{M} \geq 2$ and $\Delta_\mathcal{W} \geq 2$.

As a special case, we answer a question by Hamada et al. [24] who gave an exponential-time algorithm that in time $O(|I|^{b+1})$ decides for a given instance I of HRLQ whether it admits a feasible matching with at most b blocking pairs;[2] the authors asked whether HRLQ is fixed-parameter tractable parameterized by b. As shown by Theorem 1, SMC-1 and therefore also HRLQ is W[1]-hard when parameterized by b, already in a very restricted setting. Thus, the answer to the question by

[1] Restrictions without any parameters are simply classified as polynomial-time solvable or NP-hard.

[2] Hamada et al. claim only a run time $O((|\mathcal{W}||\mathcal{M}|)^{b+1})$, but their algorithm can easily be implemented to run in time $O(|I|^{b+1})$.

Hamada et al. [24] is negative: SMC-1, and hence HRLQ, admits no fixed-parameter algorithm with parameter b unless $\mathsf{FPT} = \mathsf{W}[1]$.

1.2 Related Work

There is a dynamically growing literature on matching markets with lower quotas [4,8,16–18,22,24,26,30,40,44,45]. These papers study several variants of HRLQ, adapting the general model to the various particularities of practical problems. However, there are only a few papers which consider the problem of minimizing the number of blocking pairs [17,24]. The most closely related work to ours is the paper by Hamada et al. [24]: they prove that the HRLQ problem is NP-hard and give strong inapproximability results; they also consider the SMC-1 problem directly and propose an $O(|I|^{b+1})$ time algorithm for it.

A different line of research connected to SMC is the problem of *arranged marriages*, an early extension of SM suggested by Knuth [31]. Here, a set $\mathcal{Q}^\star$ of man-woman pairs is distinguished, and we seek a stable matching that contains $\mathcal{Q}^\star$ as a subset. Thus, as opposed to SMC, we not only require that each distinguished person is assigned *some* partner, but instead prescribe its partner exactly. Initial work on arranged marriages [23,31] was extended by Dias et al. [12] to consider also *forbidden marriages*, and was further generalized by Fleiner et al. [15] and Cseh and Manlove [10]. Despite the similar flavour of the studied problems, none of these papers have a direct consequence on the complexity of SMC.

Our work also fits into the line of research that addresses computationally hard problems in the area of stable matchings by focusing on instances with bounded preference lists [6,27,29,32,41] or by applying the more flexible approach of parameterized complexity [1,3,5,37,38].

Organization After the preliminaries in Sect. 2, we start with the main intractability result in Sect. 3, which answers Hamada et al.'s question. This result shows $\mathsf{W}[1]$-hardness of SMC parameterized by $b+|\mathcal{W}^\star|$ even when $\mathcal{M}^\star = \emptyset$ and $\Delta_{\mathcal{M}} = \Delta_{\mathcal{W}} = 3$. Thus, we explore three directions to achieve tractability: (i) to lower b to be a constant, (ii) to lower $|\mathcal{W}^\star|$ to be a constant, or (iii) to lower either $\Delta_{\mathcal{W}}$ or $\Delta_{\mathcal{M}}$ to 2. We cover the cases (i) and (ii) in Sect. 5, and case (iii) in Sect. 6. In addition, Sect. 4 provides polynomial-time approximation results for HRLQ and SMC, used also in the polynomial-time algorithms of Sect. 5.

2 Preliminaries

An instance I of the Stable Marriage (SM) problem consists of a set $\mathcal{M}$ of men and a set $\mathcal{W}$ of women. Each person $x \in \mathcal{M} \cup \mathcal{W}$ has a preference list $L(x)$ that strictly orders the members of the other party acceptable for x. We thus write $L(x)$ as a vector $L(x) = (y_1, \ldots, y_t)$, denoting that y_i is (strictly) *preferred* by x over y_j for each i and j with $1 \leq i < j \leq t$. A *matching* M for I is a set of man-woman pairs appearing in each other's preference lists such that each person is contained in

at most one pair of M; some persons may be left unmatched by M. For each person x we denote by $M(x)$ the person assigned by M to x. For a matching M, a man m and a woman w included in each other's preference lists form a *blocking pair* if (i) m is either unmatched or prefers w to $M(m)$, and (ii) w is either unmatched or prefers m to $M(w)$. In the STABLE MARRIAGE WITH COVERING CONSTRAINTS (SMC) problem, we are given additional subsets $\mathcal{W}^\star \subseteq \mathcal{W}$ and $\mathcal{M}^\star \subseteq \mathcal{M}$ of *distinguished* people that must be matched; a matching M is *feasible* if it matches everybody in $\mathcal{W}^\star \cup \mathcal{M}^\star$. The objective of SMC is to find a feasible matching for I with minimum number of blocking pairs. If only people from one gender are distinguished, then without loss of generality, we assume these to be women; this special case will be denoted by SMC-1.

The many-to-one extension of SMC-1 is the HOSPITALS/RESIDENTS WITH LOWER QUOTAS (HRLQ) problem whose input consists of a set $\mathcal{R}$ of residents and a set $\mathcal{H}$ of hospitals that have ordered preferences over the acceptable members of the other party. Each hospital $h \in \mathcal{H}$ has a quota lower bound $\underline{\mathsf{q}}(h)$ and a quota upper bound $\overline{\mathsf{q}}(h)$, which bound the number of residents that can be assigned to h from below and above. One seeks an *assignment* M that maps a subset of the residents to hospitals that respects acceptability and is *feasible*, that is, $\underline{\mathsf{q}}(h) \leq |M(h)| \leq \overline{\mathsf{q}}(h)$ for each hospital h. Here, $M(h)$ is the set of residents assigned to some $h \in \mathcal{H}$ by M. We say that a hospital h is *under-subscribed* if $|M(h)| < \overline{\mathsf{q}}(h)$. For an assignment M of an instance of HRLQ, a pair $\{r, h\}$ of a resident r and a hospital h is *blocking* if (i) r is unassigned or prefers h to the hospital assigned to r by M, and (ii) h is under-subscribed or prefers r to one of the residents in $M(h)$. The task in HRLQ is to find a feasible assignment with minimum number of blocking pairs.

Some instances of SMC may admit a *master list* over women, which is a total ordering $L_{\mathcal{W}}$ of all women, such that for each man $m \in \mathcal{M}$, the preference list $L(m)$ is the restriction of $L_{\mathcal{W}}$ to those women that m finds acceptable. Similarly, we consider master lists over men.

With each instance I of SMC (or HRLQ) we can naturally associate a bipartite graph G_I whose vertex partitions correspond to $\mathcal{M}$ and $\mathcal{W}$ (or $\mathcal{R}$ and $\mathcal{H}$, respectively), and there is an edge between a man $m \in \mathcal{M}$ and a woman $w \in \mathcal{W}$ (or between a resident $r \in \mathcal{R}$ and a hospital $h \in \mathcal{H}$, respectively) if they appear in each other's preference lists. We may refer to entities of I as vertices, or a pair of entities as edges, without mentioning G_I explicitly. For a graph G, we denote its vertex set by $V(G)$ and its edge set by $E(G)$; furthermore, let $d_G(v)$ denote the degree of vertex $v \in V(G)$ in G. A *path* P in G is a series of vertices that contains each vertex at most once, with an edge of G connecting any two consecutive vertices of P. With a slight abuse of the notation, we will often identify paths with their edge sets; we will write $P' \subseteq P$ to express that P' is a subpath of P. A *matching* in G is a set of edges $M \subseteq E(G)$ such that no vertex in G is adjacent to more than one edges of M; note that a matching in an instance I of SMC indeed corresponds to a matching in the graph G_I. For a matching M in G, a path P is called *M-alternating*, if among any two consecutive edges along P exactly one belongs to M. We will use a few other notions from the theory of matchings about the symmetric difference of matchings, see e.g., the book by Lovász and Plummer [35] for an introduction to this topic.

Parameterized complexity The framework of parameterized complexity deals with computationally hard problems, examining their complexity in a more detailed way than classical complexity theory. In a parameterized problem problem Π, each input instance I is associated with an integer k called the *parameter*. An algorithm which decides instances I of Π in time $f(k) \cdot |I|^{O(1)}$ for some computable function f is called a *fixed-parameter algorithm*. Note that the dependence of the polynomial in the run time is constant, but the dependence on the parameter k can be arbitrary (and is typically exponential). However, if the parameter of a given instance is small, then such an algorithm can be useful in practice even if the overall size of the instance is large.

The class of problems admitting fixed-parameter algorithms is denoted by FPT. To argue that a problem is *not* in FPT, parameterized complexity provides a hardness theory. For two parameterized problems Π_1 and Π_2, a *parameterized reduction* from Π_1 to Π_2 is a function f, computable by a fixed-parameter algorithm, that maps each instance (I_1, k_1) of Π_1 to an instance $f(I_1, k_1) = (I_2, k_2)$ of Π_2 such that (i) (I_1, k_1) is a "yes"-instance of Π_1 if and only if (I_2, k_2) is a "yes"-instance of Π_2, and (ii) $k_2 \leq g(k_1)$ for some function g. The basic class of parameterized intractability is W[1]: proving a problem Π to be W[1]-hard is strong evidence that $\Pi \notin$ FPT. Given some problem Π that is known to be W[1]-hard, a parameterized reduction from Π to some parameterized problem Π' implies W[1]-hardness of Π' as well.

For more on parameterized complexity, we refer the reader to the book by Cygan et al. [11].

3 Strong Parameterized Intractability of SMC

This section provides parameterized intractability and inapproximability results for SMC showing the hardness of finding feasible matchings with minimum number of blocking pairs. Namely, we prove SMC-1 to be W[1]-hard parameterized by the number b of blocking pairs we aim for plus the number $|\mathcal{W}^\star|$ of distinguished women, even in a very restricted setting.

Theorem 2 *SMC-1 is W[1]-hard parameterized by $b + |\mathcal{W}^\star|$, even if there is a master list over men as well as one over women, all preference lists are of length at most 3, and $|L(w)| = 1$ for each woman $w \in \mathcal{W}^\star$.*

Before proving Theorem 2, let us quickly state a simple but useful claim.

Proposition 1 *Let M_1 and M_2 be two matchings in an instance of SMC. Let $v_0 v_1 \ldots v_p$ (with $p \geq 1$) be a maximal path in the symmetric differenc of M_1 and M_2, denoted by $M_1 \triangle M_2$. Then*

(a) *P contains an edge that blocks either M_1 or M_2, and*
(b) *if $i \in \{1, \ldots, p-1\}$ is such that v_i prefers v_{i-i} to v_{i+1}, then $P_i = v_0 v_1 \ldots v_i$ contains an edge that blocks either M_1 or M_2.*

Proof We call a person v_i a *leftist* if either $i = p$, or $i \in \{1, \ldots, p-1\}$ and v_i prefers v_{i-1} to v_{i+1}. Similarly, we call v_i a *rightist* if either $i = 0$, or $i \in \{1, \ldots, p-1\}$ and

v_i prefers v_{i+1} to v_{i-1}. Observe that any person on P is either a leftist or a rightist. Moreover, the path P, and under the conditions of (b) also the path P_i, must contain an edge $\{x, y\}$ such that x is a leftist and y is a rightist. Let M_i be the matching that does *not* contain $\{x, y\}$ (where $i \in \{1, 2\}$). Then both x and y prefer being matched to each other as opposed to their situation in M_i (where they may or may not be matched), proving both (a) and (b). □

Proof of Theorem 2 We give a reduction from the W[1]-hard MULTICOLORED CLIQUE parameterized by the size of the solution [14]. Let G be the input graph, with its vertex set partitioned into k sets $V_1, \ldots, V_k$; the task is to find a clique of size k in G containing exactly one vertex from each of the sets V_i. We let $E_{i,j}$ denote those edges that run between V_i and V_j for some $1 \leq i < j \leq k$. We fix an ordering on the vertices and edges of G that places vertices of V_i before vertices of V_j whenever $i < j$ (the ordering on the edges of G can be chosen arbitrarily). We will write $\mathrm{succ}(x)$ to denote the vertex following x in this ordering, and we let v_i^1 and v_i^∞ denote the first and last vertices in V_i, respectively. Similarly, we write $\mathrm{succ}(\{x, y\})$ for the edge following $\{x, y\}$, and we let $e_{i,j}^1$ and $e_{i,j}^\infty$ denote the first and last edges in $E_{i,j}$, respectively. We will also write $\mathrm{pred}(x)$ and $\mathrm{pred}(\{x, y\})$ for the predecessor of x or $\{x, y\}$, respectively. Also, we denote the h-th neighbor of some vertex x as $n(x, h)$. For simplicity, we assume that G has no isolated vertices.

We construct an instance I of SMC as follows; see Figs. 1 and 2 for an illustration.

We set the number of blocking pairs allowed for I to be $b = 2k + \binom{k}{2}$. Together with the instance I, we will define a stable (but not feasible) matching M_s for I as well. If a woman w of I is matched by M_s, we will denote the man $M_s(w)$ by $\hat{w}$. Some women will need "dummy" partners in their preference lists: we denote the dummy of w by $\widetilde{w}$. The dummy $\widetilde{w}$ will always appear as the last item on w's preference list, and its preference list will always be $L(\widetilde{w}) = (w)$.

For each i and j with $1 \leq i < j \leq k$, we construct an *edge selecting gadget* $G_{i,j}$ that involves women $s_{i,j}$ and $t_{i,j}$, together with women $a_{\{x,y\}}$, $b_{x\to y}$, and $b_{y\to x}$ for each edge $\{x, y\} \in E_{i,j}$. All women in $G_{i,j}$ are matched by M_s except for $s_{i,j}$, and $G_{i,j}$ contains the man $\hat{w}$ for each of these women w, together with additional dummies $\widetilde{b}_{y\to x}$ for each $\{x, y\} \in E_{i,j}$ with x preceding y.

For each $i \in \{1, \ldots, k\}$, we also construct a *node selecting gadget* G_i involving women s_i, t_i, and $u_i^1, \ldots, u_i^{b+1}$, together with women a_x, $b_x^1, \ldots, b_x^{d_G(x)}$, and $c_x^1, \ldots, c_x^{b+1}$ for each $x \in V_i$. The men in G_i include $\hat{w}$ for each woman w of G_i except for s_i, and additional dummies $\widetilde{b}_x^1, \ldots, \widetilde{b}_x^{d_G(x)}$ and $\widetilde{c}_x^{b+1}$ for each $x \in V_i$.

We define the following sets of women:

$$
\begin{aligned}
A &= \{a_x \mid x \in V(G)\} & C &= \{c_x^h \mid x \in V(G), 1 \leq h \leq b+1\}\\
A' &= \{a_{\{x,y\}} \mid \{x, y\} \in E(G)\} & S &= \{s_i \mid 1 \leq i \leq k\} \cup \{s_{i,j} \mid 1 \leq i < j \leq k\}\\
B &= \{b_x^h \mid x \in V(G), 1 \leq h \leq d_G(x)\} & T &= \{t_i \mid 1 \leq i \leq k\} \cup \{t_{i,j} \mid 1 \leq i < j \leq k\}\\
B' &= \{b_{x\to y}, b_{y\to x} \mid \{x, y\} \in E(G)\} & U &= \{u_i^h \mid 1 \leq i \leq k, 1 \leq h \leq b+1\}
\end{aligned}
$$

To define the set $\mathcal{W}^\star$ of women in I with covering constraint we let $\mathcal{W}^\star = S \cup T \cup U$; note $|\mathcal{W}^\star| = 2\binom{k}{2} + 2k + k(\binom{k}{2} + 2k + 1)$. To finish the definition of I, we define the precise structure of these gadgets as well as the connections between them by the

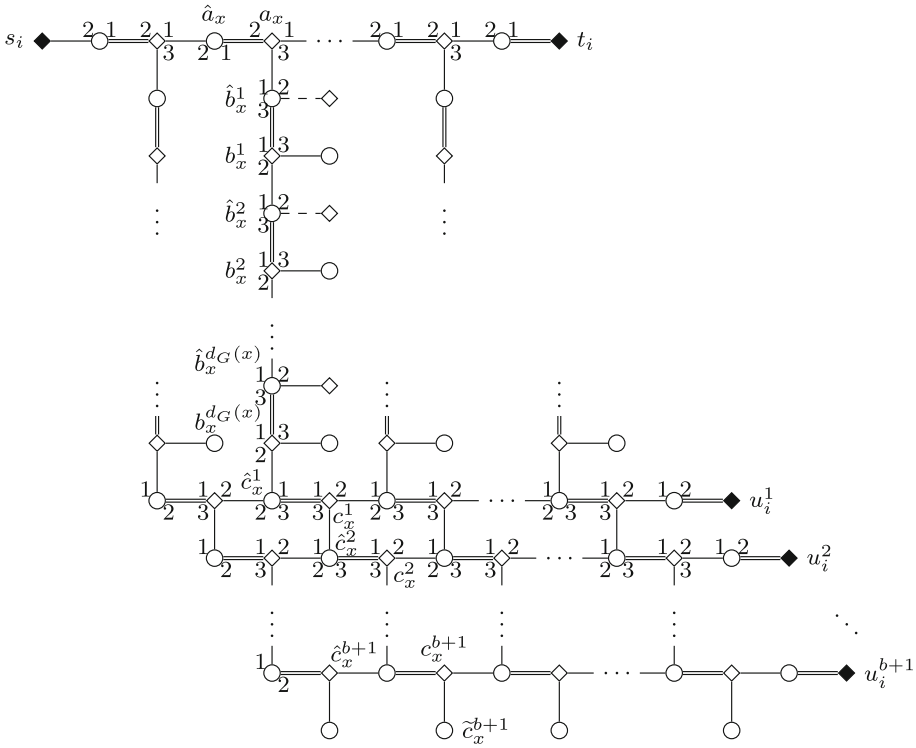

Fig. 1 Node selecting gadget G_i in the proof of Theorem 2. Throughout the paper, we use squares for women, circles for men; distinguished persons are denoted by filled squares/circles. The numbering of edges incident to some vertex (or, sometimes, arrows between edges) indicate preferences. Double edges denote edges of the stable matching M_s, and dashed edges are those leaving some gadget

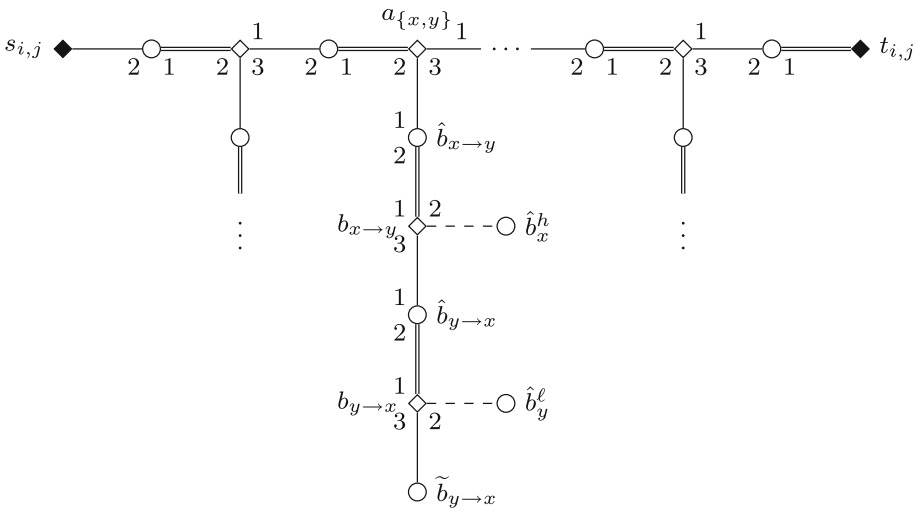

Fig. 2 Edge selecting gadget $G_{i,j}$ in the reduction of Theorem 2

Table 2 Preference lists of women and men in node selecting gadgets

$L(a_x) = (\hat{a}_{\mathrm{succ}(x)}, \hat{a}_x, \hat{b}^1_x)$,	where $x \in V_i \backslash \{v_i^\infty\}$,
$L(a_x) = (\hat{t}_i, \hat{a}_x, \hat{b}^1_x)$,	where $x = v_i^\infty$,
$L(b^h_x) = (\hat{b}^h_x, \hat{b}^{h+1}_x, \widetilde{b}^h_x)$,	where $1 \le h \le d_G(x)$,
$L(c^h_x) = (\hat{c}^h_x, \hat{c}^h_{\mathrm{succ}(x)}, \hat{c}^{h+1}_x)$,	where $1 \le h \le b, x \in V_i \backslash \{v_i^\infty\}$,
$L(c^h_x) = (\hat{c}^h_x, \hat{u}^h_i, \hat{c}^{h+1}_x,)$,	where $1 \le h \le b$ and $x = v_i^\infty$,
$L(c^{b+1}_x) = (\hat{c}^{b+1}_x, \hat{c}^h_{\mathrm{succ}(x)}, \widetilde{c}^{b+1}_x)$,	where $x \in V_i \backslash \{v_i^\infty\}$,
$L(c^{b+1}_x) = (\hat{c}^{b+1}_x, \hat{u}^{b+1}_i, \widetilde{c}^{b+1}_x,)$,	where $x = v_i^\infty$,
$L(s_i) = (\hat{a}_x)$,	where $x = v_i^1$,
$L(t_i) = (\hat{t}_i)$,	
$L(u^h_i) = (\hat{u}^h_i)$,	
$L(\hat{a}_x) = (a_x, a_{\mathrm{pred}(x)})$,	where $x \in V_i \backslash \{v_i^1\}$,
$L(\hat{a}_x) = (a_x, s_i)$,	where $x = v_i^1$,
$L(\hat{b}^h_x) = (b^{h-1}_x, b_{x\to y}, b^h_x)$,	where $y = n(x,h), x \in V_i, y \in V_j$ and $i < j$
$L(\hat{b}^h_x) = (b^{h-1}_x, b_{y\to x}, b^h_x)$,	where $y = n(x,h), x \in V_i, y \in V_j$ and $i > j$
$L(\hat{c}^h_x) = (c^{h-1}_x, c^h_{\mathrm{pred}(x)}, c^h_x)$,	where $x \in V_i \backslash \{v_i^1\}$,
$L(\hat{c}^h_x) = (c^{h-1}_x, c^h_x)$,	where $x = v_i^1$,
$L(\hat{t}_i) = (t_i, a_x)$,	where $x = v_i^\infty$,
$L(\hat{u}^h_i) = (c^h_x, u^h_i)$,	where $x = v_i^\infty$,
$L(\widetilde{w}) = (w)$,	for any dummy woman $\widetilde{w}$.

Table 3 Preference lists of women and men in edge selecting gadgets

$L(a_{\{x,y\}}) = (\hat{a}_{\mathrm{succ}(\{x,y\})}, \hat{a}_{\{x,y\}}, \hat{b}_{x\to y})$,	where $\{x,y\} \in E_{i,j} \backslash \{e^\infty_{i,j}, \}$ and x precedes y,
$L(a_{\{x,y\}}) = (\hat{t}_{i,j}, \hat{a}_{\{x,y\}}, \hat{b}_{x\to y})$,	where $\{x,y\} = e^\infty_{i,j}$ and x precedes y,
$L(b_{x\to y}) = (\hat{b}_{x\to y}, \hat{b}^h_x, \hat{b}_{y\to x})$,	where $y = n(x,h)$ and x precedes y in $V(G)$,
$L(b_{y\to x}) = (\hat{b}_{y\to x}, \hat{b}^h_y, \widetilde{b}_{y\to x})$,	where $y = n(x,h)$ and x precedes y in $V(G)$,
$L(s_{i,j}) = (\hat{a}_{\{x,y\}})$,	where $\{x,y\} = e^1_{i,j}$,
$L(t_{i,j}) = (\hat{t}_{i,j})$,	
$L(\hat{a}_{\{x,y\}}) = (a_{\{x,y\}}, a_{\mathrm{pred}(\{x,y\})})$,	where $\{x,y\} \in E_{i,j} \backslash \{e^1_{i,j}\}$,
$L(\hat{a}_{\{x,y\}}) = (a_{\{x,y\}}, s_{i,j})$,	where $\{x,y\} = e^1_{i,j}$,
$L(\hat{b}_{x\to y}) = (a_{\{x,y\}}, b_{x\to y})$,	where x precedes y in $V(G)$,
$L(\hat{b}_{y\to x}) = (b_{x\to y}, b_{y\to x})$,	where x precedes y in $V(G)$,
$L(\hat{t}_{i,j}) = (t_{i,j}, a_{\{x,y\}})$,	where $\{x,y\} = e^\infty_{i,j}$,
$L(\widetilde{w}) = (w)$,	for any dummy woman $\widetilde{w}$.

preference lists shown in Tables 2 and 3; when not stated otherwise, indices take all possible values. For simplicity, we write $b^0_x = a_x$, $b^{d_G(x)+1}_x = c^1_x$, and $c^0_x = b^{d_G(x)}_x$ for any vertex $x \in V(G)$.

Let us define a master list $L_{\mathcal{W}}$ over all women as follows. The first women in $L_{\mathcal{W}}$ are those in T, in any ordering. They are followed by women in A, ordered according to the reversed ordering over $V(G)$, that is, a_x precedes a_y exactly if y precedes x. Next follow women of A', ordered according to the reversed ordering over $E(G)$. Next come women in $B \cup B'$. To order them, we first order those in B by putting b_x^h before b_y^ℓ in $L_{\mathcal{W}}$ if and only if x precedes y or $x = y$ and $h < \ell$, then for each edge $\{x, y\} \in E(G)$ with x preceding y, $y = n(x, h)$ and $x = n(y, \ell)$ we insert $b_{x\to y}$ just before b_x^h, and we insert $b_{y\to x}$ just before b_y^ℓ, thus determining the ordering of $B \cup B'$. After women in $B \cup B'$ come women of C, with c_x^h preceding c_y^ℓ exactly if $h < \ell$ or $h = \ell$ and x precedes y. We finish the definition of the master list $L_{\mathcal{W}}$ by putting all women in $S \cup U$ at the end of $L_{\mathcal{W}}$ in an arbitrary order.

The master list over men is derived from $L_{\mathcal{W}}$ by letting $\hat{w}_1$ precede $\hat{w}_2$ whenever w_1 precedes w_2 in $L_{\mathcal{W}}$, and adding all dummies at the end in an arbitrary order. It is easy to check that the preference lists given in Tables 2 and 3 are indeed compatible with these master lists. This completes the construction of the instance.

We are going to prove that the constructed instance I admits a feasible assignment with at most b blocking pairs if and only if there is a clique of size k in the graph G.

"$\Rightarrow$": Suppose there is a feasible matching M of men to women with at most b blocking pairs. Let G_Δ be the symmetric difference $M \triangle M_s$. Notice that for each woman $s \in S$, the difference G_Δ must contain exactly one maximal path containing s as its endpoint, since the women in S must be matched in M, but are unmatched in M_s. Similarly, no path of G_Δ can contain a woman in $T \cup U$, because these women are matched by M_s to their only possible partners, and they must be matched by M as well, since $T \cup U$ is contained in $\mathcal{W}^\star$. We call a maximal path P in G_Δ with an endpoint s in S an *improving path*. We say that P *starts* at s and *ends* at its other endpoint, and we refer to the path starting at s_i (or $s_{i,j}$) as P_i (or $P_{i,j}$, respectively).

We define the *cost* of some path P of G_Δ as the number of blocking pairs $\{m, w\}$ for M involving a woman w that appears on P. By Proposition 1, each improving path contains at least one edge that is blocking for M, because no edge can block M_s. Therefore, each path in G_Δ has cost at least 1.

As there are exactly $k + \binom{k}{2}$ improving paths (as all women in S must be matched by M), we get a minimum cost of $k + \binom{k}{2}$. Note also that the total cost of all paths in G_Δ cannot exceed $b = 2k + \binom{k}{2}$. Claim 1 is therefore crucial.

Claim 1 *The following holds for any improving path P of G_Δ:*

(a) *P cannot end at a dummy $\tilde{c}_x^{b+1}$ for some $x \in V(G)$.*
(b) *P contains an edge $\{a, \hat{a}\}$ for some $a \in A \cup A'$ that blocks M.*
(c) *If P is not disjoint from G_i for some i, then P has cost at least 2.*

Proof of Claim 1 To prove (a), suppose for contradiction that P ends at $\tilde{c}_x^{b+1}$, where $x \in V_i$. Clearly, P must contain at least one woman from each of the $b + 1$ sets $\{c_v^h \mid v \in V_i\}$, $h = 1, \ldots, b+1$. Fix h, and let us consider the last $v \in V_i$ for which c_v^h is incident to an edge of G_Δ. Let $w = c^h_{\mathrm{succ}(v)}$ if $v \neq v_i^\infty$, or otherwise let $w = u_i^h$. Then the edge $\{c_v^h, \hat{w}\}$ yields a blocking pair in M, as $M(w) = M_s(w) = \hat{w}$, and thus $\hat{w}$ prefers c_v^h to w. This reasoning gives us $b + 1$ different blocking pairs for M, one for each index h, contradicting our assumption on M.

To prove (b), let us consider the case when $P = P_i$ for some i; the argument goes the same way for the case where $P = P_{i,j}$ for some i and j. If P ends at a_x for some $x \in V_i$, then a_x forms a blocking pair with $\hat{a}_x$ in M. If P does not end at a woman in A, then it must contain the edge $\{a_x, b_x^1\}$ for some x, in which case $\{a_x, \hat{a}_x\}$ is again blocking in M, showing (b).

To see (c), first observe that if P is not disjoint from G_i, then P ends in G_i, simply because of its property that it contains edges from M and M_s in an alternating fashion. Therefore, the last woman w on P must be in $B \cup C$. If $w = b_x^h$ for some $b \in B$, then the edge $\{b_x^h, \hat{b}_x^{h+1}\}$ is blocking M, as b_x^h cannot get its first choice $\hat{b}_x^h$ in M (and $\hat{b}_x^{h+1}$ cannot be on P, as that would imply that b_x^{h+1} is on P, contradicting the choice of w). If, by contrast, $w = c$ for some $c \in C$, then P must end at w by (a), and then c forms a blocking pair with the third man in its preference list (for whom c is the first choice). In either case, w is involved in a blocking pair, which together with the blocking pair guaranteed by (b) implies that P has cost at least 2.

This completes the proof of Claim 1. □

Claim 1 proves that for each $i \in \{1, \ldots, k\}$ the improving path P_i has cost at least 2. Since all the remaining $\binom{k}{2}$ improving paths have cost at least 1, and the total cost of these paths must be at most $b = 2k + \binom{k}{2}$, we get that any path P_i (or $P_{i,j}$) must have cost *exactly* 2 (or 1, respectively). Furthermore, it also follows that no other path of G_Δ can enter or start in G_i, for any i, as that would imply that the number of blocking pairs for M is more than b. In addition, it is not hard to see that G_Δ does not contain any cycle, because all cycles in the graph underlying I contain two consecutive edges not in M_s. Hence, it follows that the only connected component in G_Δ that is not disjoint from G_i is P_i.

To deal with the possible courses the path P_i may take in the graph for some $i \in \{1, \ldots, k\}$, let x_i denote the vertex in V_i for which $\{a_{x_i}, \hat{a}_{x_i}\}$ is the blocking edge guaranteed by statement (b) of Claim 1. Observe that P_i either ends at a_{x_i} or contains the edge $\{a_{x_i}, \hat{b}_{x_i}^1\}$. In either case, we say that P_i *selects* x_i from V_i; clearly, there can be only one vertex in V_i selected by P_i.

Consider now $P_{i,j}$ for some $1 \le i < j \le k$. Recall that $P_{i,j}$ has cost 1. Therefore, statement (b) of Claim 1 proves that the only blocking edge incident to some woman on $P_{i,j}$ must be $\{a_{\{x,y\}}, \hat{a}_{\{x,y\}}\}$ for some $\{x, y\} \in E_{i,j}$. We say that $P_{i,j}$ *selects* the edge $\{x, y\}$; without loss of generality, let us assume that x precedes y. By statement (c) of Claim 1, we also know that $P_{i,j}$ cannot leave $G_{i,j}$, which means that it can only have cost 1 if it ends at $\widetilde{b}_{y\to x}$. In particular, it contains the edges $\{b_{x\to y}, \hat{b}_{y\to x}\}$ and $\{b_{y\to x}, \widetilde{b}_{y\to x}\}$. Observe that the edge $\{b_{x\to y}, \hat{b}_x^h\}$ where h is such that $y = n(x, h)$ cannot be blocking in M (as this would indicate a cost of 2 for $P_{i,j}$), yielding that $\hat{b}_x^h$ must be matched to b_x^{h-1} in M. By the arguments of the previous paragraph, this means that P_i must contain the subpath $(a_x, \hat{b}_x^1, b_x^1, \ldots, \hat{b}_x^h, b_x^h)$. Hence, we obtain that x must be selected by P_i. Similarly, from the fact that the edge $\{b_{y\to x}, \hat{b}_y^\ell\}$ where $x = n(y, \ell)$ is not blocking in M we get that y must be selected by P_j.

Thus, we obtain that if an edge is selected by $P_{i,j}$ for some i and j, then its endpoints must be selected by P_i and P_j. As this must hold for each pair of indices with $1 \le i < j \le k$, we obtain that there must be $\binom{k}{2}$ edges in G whose endpoints are

among the k selected vertices. This can only happen if these edges are the edges of a clique of size k.

"$\Leftarrow$": Suppose now that G has a clique of size k formed by the vertices $x_1, \ldots, x_k$, with $x_i \in V_i$ for each $i \in \{1, \ldots, k\}$. Instead of directly defining the required matching M that is feasible and admits at most b blocking pairs, we give $M_s \triangle M$ as the union of paths P_i for $i \in \{1, \ldots, k\}$, and paths $P_{i,j}$ for $1 \leq i < j \leq k$, defined as follows.

We set P_i as the path

$$P_i = (s_i, \hat{a}_{v_i^1}, a_{v_i^1}, \ldots, \hat{a}_{x_i}, a_{x_i}, \hat{b}^1_{x_i}, b^1_{x_i}, \ldots, \hat{b}^{d_G(x_i)}_{x_i}, b^{d_G(x_i)}_{x_i}, \widetilde{b}^{d_G(x_i)}_{x_i}) \ .$$

Similarly, we define

$$P_{i,j} = (s_{i,j}, \hat{a}_{e^1_{i,j}}, a_{e^1_{i,j}}, \ldots, \hat{a}_{\{x_i,x_j\}}, a_{\{x_i,x_j\}}, \hat{b}_{x_i \to x_j}, \\ b_{x_i \to x_j}, \hat{b}_{x_j \to x_i}, b_{x_j \to x_i}, \widetilde{b}_{x_j \to x_i}) \ .$$

It is straightforward to verify that the blocking pairs for M are then the k edges $\{a_{x_i}, \hat{a}_{x_i}\}, i \in \{1, \ldots, k\}$, the k edges $\{b^{d_G(x_i)}_{x_i}, c^1_{x_i}\}$, and the $\binom{k}{2}$ edges $\{a_{\{x_i,x_j\}}, \hat{a}_{\{x_i,x_j\}}\}$, $1 \leq i < j \leq k$. The feasibility of M is trivial; this completes the proof of Theorem 2. □

A fundamental hypothesis about the complexity of NP-hard problems is the *Exponential Time Hypothesis* (ETH), which stipulates that algorithms solving all SATISFIABILITY instances in subexponential time cannot exist [28]. Assuming ETH, the fundamental CLIQUE problem parameterized by solution size k was shown not to admit any algorithm giving the correct answer in time $f(k) \cdot n^{o(k)}$ for all n-vertex instances and any computable function f [9, Thm. 5.4]. The known reduction from CLIQUE to MULTICOLORED CLIQUE does not change the parameter [14]. Finally, in the proof of Theorem 2, an instance of MULTICOLORED CLIQUE with solution size k is reduced to an instance of SMC-1 with parameter $b = O(k^2)$.

Corollary 1 *Assuming ETH, SMC-1 cannot be solved in time $f'(b) \cdot n^{o(\sqrt{b})}$ for any computable function f', even if there is a master list over men and over women, all preference lists have length at most* 3*, and each woman in $\mathcal{W}^\star$ finds only a single man acceptable.*

4 Polynomial-Time Approximation

Here we first provide a polynomial-time algorithm that yields an approximation for HRLQ with factor $(\Delta_{\mathcal{R}} - 1)\mathsf{q}_{\underline{\Sigma}}$, where $\Delta_{\mathcal{R}}$ is the maximum length $\Delta_{\mathcal{R}}$ of residents' preference lists and $\mathsf{q}_{\underline{\Sigma}}$ is the total sum of all lower quotas. Then we use this result to propose an exact polynomial-time algorithm for HRLQ for the case where both $\Delta_{\mathcal{R}}$ and $\mathsf{q}_{\underline{\Sigma}}$ are constant. Recall that in HRLQ, our objective is to find an assignment that satisfies all quota lower and upper bounds and minimizes the number of blocking pairs.

Theorem 3 *Let I be an instance of HRLQ. Let $\Delta_{\mathcal{R}}$ denote the maximum length of residents' preference lists, and let $\underline{q}_{\Sigma}$ denote the sum of lower quota bounds taken over all hospitals in I. There is an algorithm that in polynomial time either outputs a feasible assignment for I with at most $(\Delta_{\mathcal{R}} - 1)\underline{q}_{\Sigma}$ blocking pairs, involving only $\underline{q}_{\Sigma}$ residents, or concludes that no feasible assignment exists.*

Proof Let $\mathcal{H}^\star$ denote the set of hospitals with positive lower quotas. We start by finding an assignment M_q that assigns $\underline{\mathsf{q}}(h)$ residents to each hospital $h \in \mathcal{H}^\star$, and has the following property:

for each hospital $h \in \mathcal{H}^\star$, all residents that are not in $M_q(h)$ but preferred by h to the least preferred resident in $M_q(h)$ are contained in $\bigcup_{h' \in \mathcal{H}^\star \setminus \{h\}} M_q(h')$. (†)

Such an assignment can be obtained as follows. We start from an arbitrary assignment M that assigns $\underline{\mathsf{q}}(h)$ residents to each $h \in \mathcal{H}^\star$ (if no such assignment exists, then we can stop and reject); such an assignment, if existent, can be found in polynomial time by an algorithm of Hopcroft and Karp [25]. Then we greedily re-assign residents to hospitals of $\mathcal{H}^\star$, one-by-one: at each step, we take a hospital $h \in \mathcal{H}^\star$, and if there exists a resident r not assigned to any other hospital in $\mathcal{H}^\star$ that h prefers to the least preferred resident r' in $M(h)$, then we replace r' with r in $M(h)$. If this step cannot be applied anymore, then we arrive at an assignment M_q with the desired property (†).

Given M_q, we reduce the upper quotas of each hospital $h \in \mathcal{H}^\star$ by $\underline{\mathsf{q}}(h)$, set all lower quotas to 0, and delete all residents in $\mathcal{R}^\star := M_q(\mathcal{H}^\star)$. We then find a stable assignment M_s in the resulting instance I'; note that I' is an instance of HR, so we can find M_s in polynomial time [19]. Finally, we output $M^{\text{out}} = M_s \cup M_q$. Clearly, M^{out} is feasible. Also, any blocking pair that M^{out} admits must involve either a hospital from $\mathcal{H}^\star$ or a resident from $\mathcal{R}^\star = M_q(\mathcal{H}^\star)$ by the stability of M_s with respect to I'. Observe that if some $h \in \mathcal{H}^\star$ is involved in some blocking pair $\{r, h\}$ of M^{out}, then we must have $r \in \mathcal{R}^\star$. To see this, recall that each resident that is preferred by h to its least preferred resident in $M_q(h)$ must be in $\mathcal{R}^\star$ because of property (†), and furthermore, h is under-subscribed in M^{out} (within I) if and only if h is under-subscribed in M_s (within I'). Therefore, we can conclude that each blocking pair for M^{out} must involve some resident in $\mathcal{R}^\star$; observe that $|\mathcal{R}^\star| \leq \sum_{h \in \mathcal{H}} \underline{\mathsf{q}}(h) = \underline{\mathsf{q}}_{\Sigma}$. Since each resident in $\mathcal{R}^\star$ is incident to at most $\Delta_{\mathcal{R}} - 1$ edges not in M^{out}, we also have that M^{out} admits at most $(\Delta_{\mathcal{R}} - 1)|\mathcal{R}^\star| \leq (\Delta_{\mathcal{R}} - 1)\underline{\mathsf{q}}_{\Sigma}$ blocking pairs. □

If both $\Delta_{\mathcal{R}}$ and $\underline{\mathsf{q}}_{\Sigma}$ are constant, then Theorem 3 implies that HRLQ becomes polynomial-time solvable. Indeed, we can use the following simple strategy, depending on the number b of blocking pairs allowed: if $b \geq (\Delta_{\mathcal{R}} - 1)\underline{\mathsf{q}}_{\Sigma}$, then we apply Theorem 3 directly; if $b < (\Delta_{\mathcal{R}} - 1)\underline{\mathsf{q}}_{\Sigma}$, then we use the algorithm by Hamada et al. [24] running in time $O(|I|^{b+1})$ which is polynomial, since b is upper-bounded by a constant.

Corollary 2 *If both the maximum length $\Delta_{\mathcal{R}}$ of residents' preference lists and the total sum $\underline{q}_{\Sigma}$ of all lower quotas is constant, then* HRLQ *is polynomial-time solvable.*

Another application of Theorem 3 is an approximation algorithm that works regardless of whether $\Delta_{\mathcal{R}}$ or $\mathsf{q}_{\underline{\Sigma}}$ is a constant. In fact, the algorithm of Theorem 3 can be turned into a $(\Delta_{\mathcal{R}} - 1)\mathsf{q}_{\underline{\Sigma}}$-factor approximation algorithm as follows. First, we find a stable assignment M_s for I in polynomial time using the extension of the Gale–Shapley algorithm for the HOSPITALS/RESIDENTS problem. If M_s is not feasible, then by the Rural Hospitals Theorem [20], we know that any feasible assignment for I must admit at least one blocking pair; hence, the algorithm presented in Theorem 3 clearly yields an approximation with (multiplicative and also additive) factor $(\Delta_{\mathcal{R}} - 1)\mathsf{q}_{\underline{\Sigma}}$.

To close this section, we also state an analogue of Theorem 3 that deals with SMC: it can handle covering constraints on both sides, but assumes that all quota upper bounds are 1.

Theorem 4 *There is an algorithm that in polynomial time either outputs a feasible matching for an instance I of SMC with at most $(\Delta_{\mathcal{W}} - 1)|\mathcal{M}^\star| + (\Delta_{\mathcal{M}} - 1)|\mathcal{W}^\star|$ blocking pairs, or concludes that I admits no feasible matching.*

Proof The proof uses the same ideas as those used in our proof for Theorem 3, so the reader may skip the proof below, which we only include for completeness.

We start by finding an arbitrary matching M that covers each distinguished person (if no such matching exists, then we can stop and reject); such a matching, if existent, can be found in polynomial time by standard flow techniques. We assume, without loss of generality, that each edge in M is incident to some distinguished person. Let us define $\mathcal{X}^\star = \mathcal{W}^\star \cup \mathcal{M}^\star$, and let $\mathcal{U}^\star$ be the set of those persons $x \in \mathcal{X}^\star$ whose partner $M(x)$ is also in $\mathcal{X}^\star$.

We proceed by modifying M into a matching M_q that covers $\mathcal{X}^\star$ and has the following property:

If a person $x \in \mathcal{X}^\star \backslash \mathcal{U}^\star$ belongs to a blocking pair $\{x, y\}$ for M_q, then $M_q(y) \in \mathcal{X}^\star$. (✠)

Such an assignment can be obtained as follows. We greedily assign partners to the men and women in $\mathcal{X}^\star \backslash \mathcal{U}^\star$, one-by-one: at each step, we take a person $x \in \mathcal{X}^\star \backslash \mathcal{U}^\star$, and if x forms a blocking pair (with respect to the current matching) with some y that is not the partner of a distinguished person, then we replace the partner of x with y: we add the edge $\{x, y\}$ to the matching, and delete all the other edges incident to x or y. Observe that the obtained matching is still feasible. If this step cannot be applied anymore, then we arrive at a matching M_q with the desired property (✠); note also that each edge in M_q is incident to some distinguished person.

Given M_q, we delete all men and women covered by M_q. We then find a stable matching M_s in the resulting instance I'; note that I' is an instance of STABLE MARRIAGE, so we can find M_s in polynomial time [19]. Finally, we output $M^{\text{out}} = M_s \cup M_q$. Clearly, M^{out} is feasible. Also, any blocking pair that M^{out} admits must involve a person covered by M_q due to the stability of M_s with respect to I'.

We claim that any blocking pair $\{x, y\}$ involves a person whose partner by M_q is distinguished, so either $M_q(x) \in \mathcal{X}^\star$ or $M_q(y) \in \mathcal{X}^\star$. We can assume that x is covered by M_q (because this holds for at least one of x and y). To see the claim, first note that if x is not distinguished, then $M_q(x)$ must be distinguished, because each

edge of M_q contains a distinguished person. Second, if $x \in \mathcal{X}^\star$, then either $x \in \mathcal{U}^\star$ (in which case $M_q(x) \in \mathcal{X}^\star$) or $M_q(y) \in \mathcal{X}^\star$ because of property (✠). Therefore, we can conclude that each blocking pair for M^{out} must involve the partner of some distinguished resident. The partners of distinguished women can be incident to at most $|\mathcal{W}^\star|(\Delta_{\mathcal{M}} - 1)$ blocking pairs, and similarly, the partners of distinguished men can be incident to at most $|\mathcal{M}^\star|(\Delta_{\mathcal{W}} - 1)$ blocking pairs, proving the theorem. □

5 SMC with Bounded Number of Distinguished Persons or Blocking Pairs

In Theorem 2 we proved W[1]-hardness of SMC-1 for the case where $\Delta_{\mathcal{M}} = \Delta_{\mathcal{W}} = 3$, with parameter $b + |\mathcal{W}^\star|$. Here we investigate those instances of SMC and SMC-1 where the length of preference lists may be unbounded, but either b, or the number of distinguished persons is constant.

First, if the number b of blocking pairs allowed is constant, then SMC can be solved by simply running the extended Gale–Shapley algorithm after guessing and deleting all blocking pairs. This complements the result by Hamada et al. [24].

Observation 1 *SMC can be solved in time $O(|I|^{b+1})$, where b denotes the number of blocking pairs allowed in the input instance I.*

In Theorem 5 we prove hardness of SMC-1 even if only one woman must be covered. If we require preferences to follow master lists, then a slightly weaker version of Theorem 5, where $|\mathcal{W}^\star| = 2$, still holds.

Theorem 5 *SMC-1 is W[1]-hard parameterized by $b + \Delta_{\mathcal{M}}$, even if $\mathcal{W}^\star = \{s\}$, $\Delta_{\mathcal{W}} = 3$, and $|L(s)| = 1$.*

Proof We present a reduction based on the one from Multicolored Clique given in the proof of Theorem 2. Given some graph G and an integer k as inputs, we are going to re-use the instance I constructed in the proof of Theorem 2. Recall that I has a feasible matching with at most $b = \binom{k}{2} + 2k$ blocking pairs if and only if G has a clique of size k. Recall also that the set of women that must be covered in I is $S \cup T \cup U$; here we denote this set by $\mathcal{W}_I^\star$. We define a modified instance I' of SMC as follows. For each $w \in \mathcal{W}_I^\star$, we create a *forcing gadget* F_w which apart from w contains the newly introduced women a_w, b_w, c_w, d_w and men a'_w, b'_w, c'_w, d'_w. We also add the distinguished woman s, who must be covered in I', and the unique man t in $L(s)$. See Fig. 3 for an illustration.

Let $n(w)$ denote the unique man acceptable for some $w \in \mathcal{W}_I^\star$ in I. Additionally, we let $Y = \{a_w, c_w \mid w \in \mathcal{W}_I^\star\}$, and we write $[Y]$ for an arbitrarily fixed ordering of the elements of Y. The preferences of the newly introduced men and women, as well as the modified preferences of those agents that find them acceptable, is given below. Here, again, indices take all possible values, and w can be any woman in $\mathcal{W}_I^\star$. We let I' contain all other women and men defined in I, having the same preferences as in I.

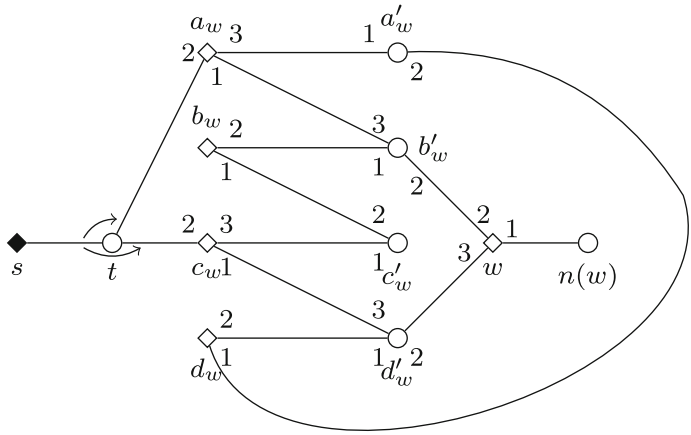

Fig. 3 Illustration depicting the forcing gadget F_w in the proof of Theorem 5

$$
\begin{aligned}
&L(s) = (t), && L(t) = ([Y], s),\\
&L(a_w) = (b'_w, t, a'_w), && L(a'_w) = (a_w, d_w),\\
&L(b_w) = (c'_w, b'_w), && L(b'_w) = (b_w, w, a_w),\\
&L(c_w) = (d'_w, t, c'_w), && L(c'_w) = (c_w, b_w),\\
&L(d_w) = (a'_w, d'_w), && L(d'_w) = (d_w, w, c_w),\\
&L(w) = (n(w), b'_w, d'_w).
\end{aligned}
$$

We will show that I' has a feasible matching with at most b blocking pairs if and only if I has such a matching; this clearly proves the theorem.

First observe that any feasible matching M' for I' contains the edge $\{s, t\}$. Thus, if some woman y in Y is not matched by M' to her first choice, then $\{y, t\}$ is blocking in M'. Consider now F_w for some $w \in \mathcal{W}_I^\star$. It is straightforward to check that if $M'(w) \neq n(w)$, then there are at least two blocking pairs incident to a woman in F_w. Indeed, assume first that $\{t, a_w\}$ is the only blocking pair in F_w; this quickly implies $M'(c_w) = d'_w$ and $M'(a_w) = a'_w$, which in turn leads to $\{d_w, d'_w\}$ blocking M', a contradiction. Second, assume that $\{t, a_w\}$ does not block M'; from this follows $M'(a_w) = b'_w$ and we have that $\{b'_w, w\}$ is a blocking pair for M'. Now either $\{t, c_w\}$ is blocking (in which case our claim holds), or we get $M'(c_w) = d'_w$, which implies that $\{d'_w, w\}$ blocks M', again a contradiction.

Now, let $\mathcal{W}_i$ be the women in G_i that must be covered in I, i.e., $\mathcal{W}_i = \{s_i, t_i, u_i^1, \ldots, u_i^{b+1}\}$. Consider the number β_i of blocking pairs for M' that involve a woman either in the gadget G_i or in a gadget F_w for some $w \in \mathcal{W}_i$. On the one hand, if some $w \in \mathcal{W}_i$ is not matched by M' to $n(w)$, then $\beta_i \geq 2$ because of the blocking pairs in F_w. On the other hand, if each $w \in \mathcal{W}_i$ is matched by M' to $n(w)$, then using the arguments of the proof for Theorem 2, we again know $\beta_i \geq 2$ because of the blocking pairs in G_i. Also, $\beta_i = 2$ can only be achieved if (i) $M'(t_i) = n(t_i)$, as otherwise $\{t_i, n(t_i)\}$ would be blocking for M', in addition to the two blocking pairs in F_{t_i}, and (ii) $M'(u_i^h) = n(u_i^h)$ for each $h \in \{1, \ldots, b+1\}$, as otherwise we would

have $M'(s_i) = n(s_i)$ (so as to avoid having four blocking pairs due to women in $F_{u_i^h}$ and F_{s_i}), implying at least one blocking pair in G_i in addition to those in $F_{u_i^h}$.

Analogously, let $\beta_{i,j}$ denote the number of blocking pairs for M' that involve a woman either in the gadget $G_{\{i,j\}}$ or in a gadget F_w for some $w \in \{s_{i,j}, t_{i,j}\}$. Then either $\beta_{i,j} \geq 2$, or we know that $M'(w) = n(w)$ for both women $w \in \{s_{i,j}, t_{i,j}\}$; in this case, from the proof of Theorem 2 we get $\beta_{i,j} \geq 1$. However, supposing that M' has at most $b = 2k + \binom{k}{2}$ blocking pairs, it follows that $\beta_i = 2$ and $\beta_{i,j} = 1$ must hold for each $i \in \{1, \ldots, k\}$ and each i, j with $1 \leq i < j \leq k$, respectively.

Along the same lines as in the proof of Theorem 2, it can also be verified that $\beta_{i,j} = 1$ for each pair of indices i, j can only be achieved if $M' \triangle M_s$ contains a path in each gadget G_i. From $M'(w) = n(w)$ for each $w \in \mathcal{W}_i \setminus \{s_i\}$ we get that such a path contains at least one blocking pair. This implies $M'(s_i) = n(s_i)$, as otherwise we would end up with $\beta_i \geq 3$ because of the blocking pairs incident to women of F_{s_i}.

Altogether, we have proved that $M'(w) = n(w)$ for each $w \in \mathcal{W}_I^\star$. Hence, the restriction of M' to I yields a feasible matching for I that admits at most b blocking pairs.

For the other direction, suppose that I has a feasible matching M. Then it is easy to see that adding the edges $\{a_w, b'_w\}$, $\{b_w, c'_w\}$, $\{c_w, d'_w\}$, and $\{d_w, a'_w\}$ for each $w \in \mathcal{W}_I^\star$ together with the edge $\{s, t\}$ to M yields a feasible matching for I' that contains exactly the same number of blocking pairs in I' as M does in I. □

Theorem 6 *SMC-1 is W[1]-hard parameterized by $b + \Delta_{\mathcal{M}}$, even if there is a master list over men as well as one over women, $|\mathcal{W}^\star| = 2$, $\Delta_{\mathcal{W}} \leq 3$, and $|L(w)| = 1$ for each $w \in \mathcal{W}^\star$.*

Proof The proof is very similar to the one for Theorem 5, so we will only sketch it. Again, we are going to re-use the instance I constructed in the proof of Theorem 2, and construct a modified instance I' of SMC, adding only two new women z_1 and z_2 and two men m_1 and m_2 to I. We append z_1 and z_2, in this order, to the master list over women, and similarly, we append m_1 and m_2 to the master list over men. We define the women to be covered in I' as z_1 and z_2.

Again, we denote the set of women to be covered in I by $\mathcal{W}_I^\star$, and we denote by $n(w)$ the unique man acceptable for some $w \in \mathcal{W}_I^\star$ in I. The preferences of the newly introduced men and women, as well as the modified preferences of those agents that find them acceptable, is given below (here, $[\mathcal{W}_I^\star]_\prec$ denotes the ordering of $\mathcal{W}_I^\star$ given by the master list). We let I' contain all other women and men defined in I, having the same preferences as in I.

$$\begin{aligned} &L(z_1) = (m_1), && L(m_1) = ([\mathcal{W}_I^\star]_\prec, z_1), \\ &L(z_2) = (m_2), && L(m_2) = ([\mathcal{W}_I^\star]_\prec, z_2), \\ &L(w) = (n(w), m_1, m_2) && \forall w \in \mathcal{W}_I^\star. \end{aligned}$$

Arguing analogously as before in the proof of Theorem 5, one can show that I' has a feasible matching with at most b blocking pairs if and only if I has such a matching; this suffices to prove the theorem. □

To contrast our intractability results, we show next that if each of the four parameters $|\mathcal{W}^\star|$, $|\mathcal{M}^\star|$, $\Delta_\mathcal{W}$, and $\Delta_\mathcal{M}$ is constant, then SMC becomes polynomial-time solvable. Our algorithm relies on the observation that in this case, the number of blocking pairs in an optimal solution is at most $(\Delta_\mathcal{M}-1)|\mathcal{W}^\star|+(\Delta_\mathcal{W}-1)|\mathcal{M}^\star|$ by Theorem 4. Note that for instances of SMC-1, Theorem 7 yields a polynomial-time algorithm already if both $|\mathcal{W}^\star|$ and $\Delta_\mathcal{M}$ are constant.

Theorem 7 *SMC can be solved in time* $O(|I|^{(\Delta_\mathcal{M}-1)|\mathcal{W}^\star|+(\Delta_\mathcal{W}-1)|\mathcal{M}^\star|+1})$.

Proof By Theorem 4, there is a matching with at most $b_{\max} = (\Delta_\mathcal{M} - 1)|\mathcal{W}^\star| + (\Delta_\mathcal{W} - 1)|\mathcal{M}^\star|$ blocking pairs. Hence, if the number b of blocking pairs allowed is at least $b_{\max}$, then we can simply run the algorithm of Theorem 4. Otherwise, we can use Observation 1, which gives us the required run time. □

Importantly, restricting only three of the values $|\mathcal{W}^\star|$, $|\mathcal{M}^\star|$, $\Delta_\mathcal{W}$, and $\Delta_\mathcal{M}$ to be constant does not yield tractability for SMC, showing that Theorem 7 is tight in this sense. Indeed, Theorem 5 implies immediately that restricting the maximum length of the preference lists on only one side still results in a hard problem: SMC remains W[1]-hard with parameter $b+\Delta_\mathcal{M}$, even if $\Delta_\mathcal{W} = 3$, $|\mathcal{W}^\star| = 1$, and $|\mathcal{M}^\star| = 0$. On the other hand, Theorem 2 shows that the problem remains hard even if $\Delta_\mathcal{W} = \Delta_\mathcal{M} = 3$ and $|\mathcal{M}^\star| = 0$.

6 SMC with Preference Lists of Length at Most Two

In this section we investigate the computational complexity of SMC where the maximum length of preference lists is bounded by 2 on one side. This restriction leads to important tractable special cases: we obtain both polynomial-time algorithms and fixed-parameter tractability results for various parameterizations.

Let I be an instance of SMC with underlying graph G. Let M_s be a stable matching in I, and let $\mathcal{M}_0^\star$ and $\mathcal{W}_0^\star$ denote the set of distinguished men and women, respectively, unmatched by M_s. Furthermore, let $\mathcal{M}_0$ and $\mathcal{W}_0$ denote the set of all men and women, respectively, unmatched by M_s. A path P in G is called an *augmenting path*, if $M_s \triangle P$ is a matching, and either both endpoints of P are in $\mathcal{M}_0^\star \cup \mathcal{W}_0^\star$, or one endpoint of P is in $\mathcal{M}_0^\star \cup \mathcal{W}_0^\star$, and its other endpoint is not distinguished. This definition ensures that for an augmenting path P, the set of distinguished men and women that are matched in $M_s \triangle P$ strictly contains the set of distinguished men and women matched in M_s.[3] We will call an augmenting path P *masculine* or *feminine* if it contains a man in $\mathcal{M}_0^\star$ or a woman in $\mathcal{W}_0^\star$, respectively; if P is both masculine and feminine, then we call it *neutral*. If P is not neutral, then we say that it *starts* at the (unique) person from $\mathcal{M}_0^\star \cup \mathcal{W}_0^\star$ it contains, and *ends* at its other endpoint.

[3] We remark that our concept of an augmenting path is analogous, but not identical, to the standard notion of an augmenting path in general matching theory. According to the standard definition, an augmenting path for a given matching M is an M-alternating path P' such that $M \triangle P'$ is a matching containing more edges than M. In our case, however, instead of increasing the number of edges in the matching, we aim for a path which can be used to increase the number of distinguished men and women that are matched.

6.1 Covering Constraints on One Side

Here we deal with the SMC-1 problem where only women need to be covered. We first give a polynomial-time algorithm for SMC-1 when each man finds at most two women acceptable, and then show NP-hardness of SMC-1 for instances where each woman finds at most two men acceptable. We start by considering the special case of SMC-1 where $\Delta_{\mathcal{M}} \leq 2$.

Theorem 8 *There is a polynomial-time algorithm for the special case of SMC-1 where each man finds at most two women acceptable.*

High-Level Description The main observation behind Theorem 8 is that if $\Delta_{\mathcal{M}} \leq 2$, then any two augmenting paths starting from different women in $\mathcal{W}_0^\star$ are almost disjoint, namely they can only intersect at their endpoints. Thus, we can modify the stable matching M_s by selecting augmenting paths starting from each woman in $\mathcal{W}_0^\star$ in an almost independent fashion: intuitively, we simply need to take care not to choose paths sharing an endpoint—a task which can be managed by finding a bipartite matching in an appropriately defined auxiliary graph. To ensure that the number of blocking pairs in the output is minimized, we will assign costs to the augmenting paths. Roughly speaking, the cost of an augmenting path P determines the number of blocking pairs introduced when modifying M_s along P (though certain special edges need not be counted); hence, our problem reduces to finding a bipartite matching with minimum weight in the auxiliary graph.

To present the algorithm of Theorem 8 in detail, we start with the following properties of augmenting paths which are easy to prove using that $\Delta_{\mathcal{M}} \leq 2$:

Proposition 2 *Suppose $\Delta_{\mathcal{M}} \leq 2$. Let P_1 and P_2 be augmenting paths starting at women w_1 and w_2, respectively.*

(a) *If $w_1 \neq w_2$, then P_1 and P_2 are either vertex-disjoint, or they both end at some $m \in \mathcal{M}_0$, with $V(P_1) \cap V(P_2) = \{m\}$.*
(b) *If there is an edge $\{m, w\}$ of G (with $m \in \mathcal{M}$ and $w \in \mathcal{W}$) connecting P_1 and P_2, then $m \in \mathcal{M}_0$ and P_1 or P_2 must end at m.*
(c) *If $w_1 = w_2$ and P is the maximal common subpath of P_1 and P_2 starting at w_1, then either $V(P_1) \cap V(P_2) = V(P)$, or P_1 and P_2 both end at some $m \in \mathcal{M}_0$ and $V(P_1) \cap V(P_2) = V(P) \cup \{m\}$.*

With a set P of edges (typically a set of augmenting paths) where $M_s \triangle P$ is a matching, we associate a *cost*, which is the number of blocking pairs that $M_s \triangle P$ admits. A pair $\{m, w\}$ for some $m \in \mathcal{M}$ and $w \in \mathcal{W}$ is *special*, if $m \in \mathcal{M}_0$ and w is the second (less preferred) woman in $L(m)$. As it turns out, such edges can be ignored during certain steps of the algorithm; thus, we define the *special cost* of P as the number of non-special blocking pairs in $M_s \triangle P$.

Lemma 1 *For vertex-disjoint augmenting paths P_1 and P_2 with cost c_1 and c_2, respectively, the cost of $P_1 \cup P_2$ is at most $c_1 + c_2$. Further, if the cost of $P_1 \cup P_2$ is less than $c_1 + c_2$, then the following holds for $\{i_1, i_2\} = \{1, 2\}$: there is a special edge $\{m, w\}$ with P_{i_1} ending at m and w appearing on P_{i_2}; moreover, $\{m, w\}$ is blocking in $M_s \triangle P_{i_2}$, but not in $M_s \triangle (P_1 \cup P_2)$.*

Proof First observe that if some edge $\{m, w\}$ has a common vertex with only one of the paths P_1 and P_2, say P_1, then $\{m, w\}$ is blocking in $M_s \triangle P_1$ if and only if it is blocking in $M_s \triangle (P_1 \cup P_2)$.

Consider now the case when $\{m, w\}$ connects P_1 and P_2. By Proposition 2, this implies that one of the paths, say P_1, ends at $m \in \mathcal{M}_0$ (and w lies on P_2). Clearly, $\{m, w\}$ is not blocking in $M_s \triangle P_1$, by the stability of M_s. If, on the one hand, w is the first choice of m, then $\{m, w\}$ is blocking in $M_s \triangle P_2$ exactly if it is blocking in $M_s \triangle (P_1 \cup P_2)$. If, on the other hand, $\{m, w\}$ is special, then it cannot be blocking in $M_s \triangle (P_1 \cup P_2)$, but it might be blocking in $M_s \triangle P_2$. Putting all these facts together, the lemma follows immediately. □

We are ready to provide the algorithm, in a sequence of four steps.

Step 1: Computing all augmenting paths By Proposition 2, if we delete $\mathcal{M}_0$ from the union of all augmenting paths starting at some $w \in \mathcal{W}_0^\star$, then we obtain a tree. Furthermore, these trees are mutually vertex-disjoint for different starting vertices of $\mathcal{W}_0^\star$. This allows us to compute all augmenting paths in linear time, e.g., by an appropriately modified version of a depth-first search algorithm (so that only augmenting paths are considered). During this process, we can also compute the special cost of each augmenting path in a straightforward way.

Step 2: Constructing an auxiliary graph Using the results of the computation of Step 1, we construct an edge-weighted single bipartite graph $G_{\text{path}} = (U, V; E)$ as follows. To define the vertices of G_{path} we set $U := \mathcal{W}_0^\star$ and $V := \mathcal{M}_0 \cup \{w' \mid w \in \mathcal{W}_0^\star\}$, so for each woman $w \in \mathcal{W}_0^\star$ we create a corresponding new vertex w'. The edge set E contains an edge $\{w, w'\}$ for each $w \in U$, as well as an edge $\{w, m\}$ whenever $w \in U$, $m \in \mathcal{M}_0$ and there exists an augmenting path with endpoints w and m. We define the weight of an edge $\{w, w'\}$ as the minimum special cost $c_w^{\min}$ of any augmenting path starting at w and *not* ending in $\mathcal{M}_0$, and we define the weight of an edge $\{w, m\}$ with $w \in U$ and $m \in \mathcal{M}_0$ as the minimum special cost of any augmenting path with endpoints w and m.

Step 3: Computing a minimum weight matching We compute a matching M_P in G_{path} covering U and having minimum weight; this can be done in polynomial time by, e.g., the Hungarian method [34]. Observe that the matching M_P corresponds to a set of augmenting paths $\mathcal{P} = \{P_w \mid w \in \mathcal{W}_0^\star\}$ that are mutually vertex-disjoint by Proposition 2. Recall that the special cost of P_w is the weight of the edge in M_P incident to w.

Step 4: Eliminating blocking special edges In this step, we modify $\mathcal{P}$ iteratively. We start by setting $\mathcal{P}_{\text{act}} = \mathcal{P}$. At each iteration we modify $\mathcal{P}_{\text{act}}$ as follows. We check whether there exists a special edge $\{m^*, w^*\}$ that is blocking in $M_s \triangle \mathcal{P}_{\text{act}}$. If yes, then notice that m^* is not matched in $M_s \triangle \mathcal{P}_{\text{act}}$, because $\{m^*, w^*\}$ is special and thus $m^* \in \mathcal{M}_0$. Let P be the path of $\mathcal{P}_{\text{act}}$ containing w^*. We modify $\mathcal{P}_{\text{act}}$ by truncating P to its subpath between its starting vertex and w^*, and appending to it the edge $\{m^*, w^*\}$. This way, $\{m^*, w^*\}$ becomes an edge of the matching $M_s \triangle \mathcal{P}_{\text{act}}$. The iteration stops when there is no special edge blocking $M_s \triangle \mathcal{P}_{\text{act}}$. Note that once a special edge ceases to be blocking in $M_s \triangle \mathcal{P}_{\text{act}}$, it cannot become blocking again during this process, so

the algorithm performs at most $|\mathcal{M}_0|$ iterations. For each $w \in \mathcal{W}_0^\star$, let P_w^* denote the augmenting path in $\mathcal{P}_{\text{act}}$ covering w at the end of Step 4; we define $\mathcal{P}^* = \{P_w^* \mid w \in \mathcal{W}_0^\star\}$ and output the matching $M_s \triangle \mathcal{P}^*$.

This completes the description of the algorithm; we now provide its analysis.

Lemma 2 *$M_{\text{sol}} := M_s \triangle \mathcal{P}^*$ is a feasible matching for I, and the number of blocking pairs for M_{sol} is at most the weight of M_P.*

Proof Consider the situation when the iteration in Step 4 deals with a special edge $\{m^*, w^*\}$ blocking in $\mathcal{P}_{\text{act}}$. Notice that since w^* is the second woman in $L(m^*)$ (by the definition of a special edge), and since $\{w^*, m^*\}$ is blocking in $M_s \triangle \mathcal{P}_{\text{act}}$, we know that m^* is unmatched in $M_s \triangle \mathcal{P}_{\text{act}}$, that is, m^* does not lie on any of the augmenting paths in $\mathcal{P}_{\text{act}}$. From this follows that the augmenting paths in $\mathcal{P}_{\text{act}}$, and hence in $\mathcal{P}^*$, remain mutually vertex-disjoint. Therefore, M_{sol} is indeed a matching. As it covers $\mathcal{W}_0^\star$, and no augmenting path ends at a woman in $\mathcal{W}^\star \setminus \mathcal{W}_0^\star$, matching M_{sol} is feasible.

Clearly, Step 4 ensures that there are no blocking special edges in M_{sol}. Note that when the algorithm modifies P_w for some $w \in \mathcal{W}_0^\star$, at most one new blocking pair may arise with respect to $M_s \triangle \mathcal{P}_{\text{act}}$, and from the stability of M and Proposition 2 it follows that such an edge must be a special edge (incident to the man at which P_w ends before its modification). This means that Step 4 gets rid of all blocking special edges without introducing any non-special blocking edges. Hence, we obtain that the cost of P_w^* is at most the special cost of P_w, for each $w \in \mathcal{W}_0^\star$. By Lemma 1, the number of blocking pairs that M_{sol} admits is at most the sum of the costs of all augmenting paths in $\mathcal{P}^*$; this finishes the proof. □

To show that our algorithm is correct and M_{sol} is optimal, by Lemma 2 it suffices to prove that the weight of M_P is at most the number of blocking pairs in M^{opt}, where M^{opt} denotes an optimal solution in I. To this end, we are going to define a matching covering $\mathcal{W}_0^\star$ in G_{path} whose weight is at most the number of blocking pairs in M^{opt}.

Clearly, $M_s \triangle M^{\text{opt}}$ contains an augmenting path Q_w covering w for each $w \in \mathcal{W}_0^\star$. If some Q_w ends at a man $m \in \mathcal{M}_0$, then clearly no other path in $M_s \triangle M^{\text{opt}}$ can end at m. So let us take the matching M_Q in G_{path} that includes all pairs $\{m, w\}$ where Q_w ends at $m \in \mathcal{M}_0$ for some $w \in \mathcal{W}_0^\star$. Also, we put $\{w, w'\}$ into M_Q if Q_w does not end at a man of $\mathcal{M}_0$. Note that M_Q is indeed a matching.

It remains to show that the weight of M_Q is at most the number of blocking pairs in M^{opt}. By definition, the weight of M_Q is at most the sum of the special costs of the paths Q_w for every $w \in \mathcal{W}_0^\star$. By Lemma 1, any non-special blocking pair in $M_s \triangle Q_w$ remains a blocking pair in $M_s \triangle (\bigcup_{w \in \mathcal{W}_0^\star} Q_w)$, and hence in M^{opt} as well. Hence, there is a matching in G_{path} with weight at most the number of blocking pairs in an optimal solution, implying the correctness of our algorithm. As the algorithm runs in polynomial time, Theorem 8 follows.

By contrast to Theorem 8, if men may have preference lists of length 3, then SMC-1 (and hence SMC) is NP-hard even if each woman finds at most two men acceptable.

Theorem 9 *SMC-1 is NP-hard even if $\Delta_{\mathcal{W}} = 2$ and $\Delta_{\mathcal{M}} = 3$.*

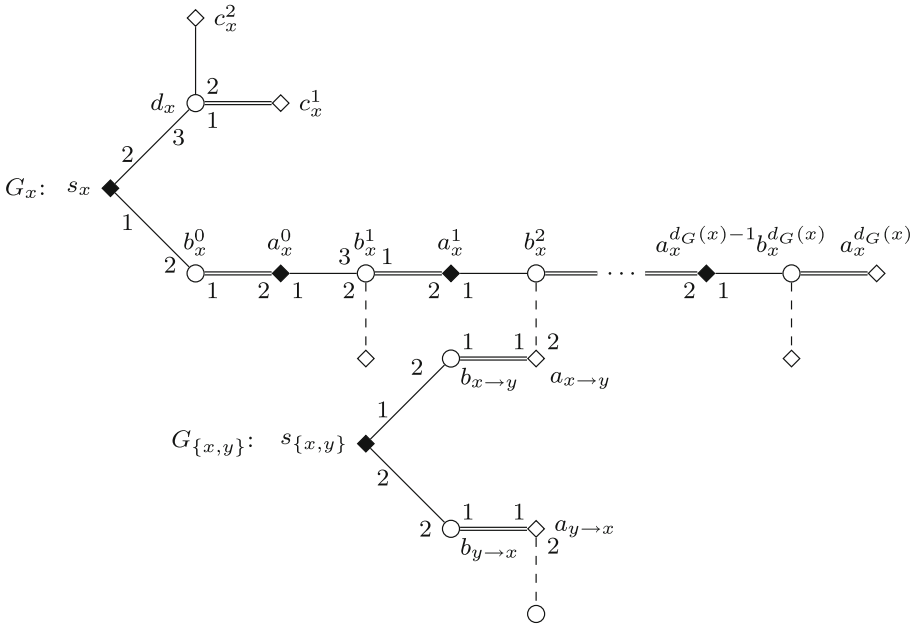

Fig. 4 Illustration of a node gadget G_x and an edge gadget $G_{\{x,y\}}$ constructed in the proof of Theorem 9. Double edges denote edges of a stable matching for I, and dashed edges are those leaving some gadget. The example depicted assumes $y = n(x, 2)$

Proof We give a reduction from the NP-hard VERTEX COVER problem, asking whether the input graph G has a vertex cover of size at most k. We order the vertices of G arbitrarily, and denote the h-th neighbor of some vertex x by $n(x, h)$ for any $h \in \{1, \ldots, d_G(x)\}$.

Let us construct an instance I of SMC as follows; see Fig. 4 for an illustration. For each vertex $x \in V(G)$ we construct a *node gadget* G_x which contains women $s_x, a_x^0, \ldots, a_x^{d_G(x)}$, c_x^1 and c_x^2, and men $b_x^0, \ldots, b_x^{d_G(x)}$, and d_x. For each edge $\{x, y\} \in E(G)$ we also construct an *edge gadget* $G_{\{x,y\}}$ involving women $s_{\{x,y\}}$, $a_{x\to y}$ and $a_{y\to x}$, and men $b_{x\to y}$ and $b_{y\to x}$. Furthermore, there are two edges in the underlying graph connecting $G_{\{x,y\}}$ to G_x and G_y, namely $\{a_{x\to y}, b_x^h\}$ and $\{a_{y\to x}, b_y^\ell\}$ where $y = n(x, h)$ and $x = n(y, \ell)$.

The preference lists of I are given in Table 4. Let the set of women with covering constraints be

$$\mathcal{W}^\star = \bigcup_{x\in V(G)} \{s_x, a_x^0, \ldots, a_x^{d_G(x)-1}\} \cup \bigcup_{\{x,y\}\in E(G)} \{s_{\{x,y\}}\},$$

and set the number of blocking pairs allowed to be $|V(G)| + k$. We are going to prove that I admits a feasible matching with at most $|V(G)| + k$ blocking pairs if and only if there is a vertex cover of size k in the graph G.

"$\Rightarrow$": Let M be a feasible matching with at most $|V(G)| + k$ blocking pairs. We say that the *cost* of some gadget G_x (or $G_{\{x,y\}}$) is the number of edges blocking M

Table 4 Preference lists of women and men in the proof of Theorem 9

$L(s_x) = (b_x^0, d_x)$,	
$L(b_x^0) = (a_x^0, s_x)$,	
$L(b_x^h) = (a_x^h, a_{x\to y}, a_x^{h-1})$	where $1 \le h \le d_G(x)$ and $y = n(x, h)$,
$L(a_x^h) = (b_x^{h+1}, b_x^h)$	where $0 \le h < d_G(x)$,
$L(a_x^{d_G(x)}) = (b_x^{d_G(x)})$,	
$L(c_x^h) = (d_x)$,	where $h \in \{1, 2\}$,
$L(d_x) = (c_x^1, c_x^2, s_x)$,	
$L(s_{\{x,y\}}) = (b_{x\to y}, b_{y\to x})$	where x precedes y,
$L(b_{x\to y}) = (a_{x\to y}, s_{\{x,y\}})$,	
$L(a_{x\to y}) = (b_{x\to y}, b_x^h)$	where $y = n(x, h)$.

When not stated otherwise, indices take all possible values

which are incident to some man of G_x (or $G_{\{x,y\}}$, respectively.) We will prove that the set S of vertices x for which G_x has cost at least 2 is a vertex cover of G.

First, let us consider some x for which $M(s_x) = d_x$. In this case, both c_x^1 and c_x^2 form a blocking pair for M with d_x, implying $x \in S$. Second, let us consider some x with $M(s_x) = b_x^0$. Since each a_x^h with $0 \le h < d_G(x)$ must be matched by M (because it is contained in $\mathcal{W}^\star$), we obtain $M(a_x^h) = b_x^{h+1}$ for each such h. Hence, $a_x^{d_G(x)}$ and $b_x^{d_G(x)}$ form a blocking pair for M. Moreover, if the woman $a_{x\to y}$ is unmatched in M for some y, then $\{a_{x\to y}, b_x^h\}$ is also a blocking pair in M (where $y = n(x, h)$), and implies a cost of at least 2 for G_x. Therefore, we can observe that if $x \notin S$, then $a_{x\to y}$ must be matched by M to $b_{x\to y}$ for each neighbor y of x in G.

However, for any $\{x, y\} \in E(G)$, M must match $s_{\{x,y\}}$ either to $b_{x\to y}$ or to $b_{y\to x}$ (because $s_{\{x,y\}}$ is contained in $\mathcal{W}^\star$), which means that $x \in S$ or $y \in S$. This proves that S is indeed a vertex cover for G. Moreover, the number of vertices in S can be at most k, since each G_x with $x \in S$ has cost at least 2, each G_x with $x \notin S$ has cost at least 1, and the total cost of all gadgets cannot exceed our budget $|V(G)| + k$.

"$\Leftarrow$": Given a vertex cover S of size at most k for G, we define a matching M with the desired properties. Namely, for each $x \in S$ we set $M(s_x) = d_x$ and $M(a_x^h) = b_x^h$ for each $h \in \{0, \ldots, d_G(x)\}$. In this case, c_x^1, c_x^2 are unmatched by M, both forming a blocking pair with d_x. By contrast, all of the men $b_x^0, \ldots, b_x^{d_G(x)}$ get their first choices.

Next, for each $x \in V(G)\backslash S$ we set $M(s_x) = b_x^0$, $M(c_x^1) = d_x$, and $M(a_x^h) = b_x^{h+1}$ for each $h \in \{0, \ldots, d_G(x) - 1\}$. Note that $a_x^{d_G(x)}$ is unmatched by M, and thus forms a blocking pair with $b_x^{d_G(x)}$. Observe also that d_x is not contained in any blocking pair.

Finally, for some $\{x, y\} \in E(G)$, let us assume $y \in S$ (since S is a vertex cover, it contains x or y). We set $M(s_{\{x,y\}}) = b_{y\to x}$ and $M(a_{x\to y}) = b_{x\to y}$. Note that $a_{x\to y}$ gets her first choice, so it cannot be involved in a blocking pair. Although $a_{y\to x}$ is unmatched by M, we know that it cannot form a blocking pair with b_y^ℓ where $x = n(y, \ell)$, because $y \in S$ and hence b_y^ℓ is assigned her first choice by M. Thus, no man or woman of some edge gadget participates in a blocking pair, and therefore we obtain that the total number of blocking pairs for M is exactly $|V(G)| + k$.

Since M is feasible, the theorem follows. □

6.2 Covering Constraints on Both Sides

Let us now investigate the complexity of SMC with covering constraints both for men and women. If we restrict the maximum length of preference lists on both sides to be at most 2, SMC becomes linear-time solvable. To see this, observe that by $\max(\Delta_{\mathcal{W}}, \Delta_{\mathcal{M}}) \leq 2$, the underlying graph G must be a collection of paths and cycles. Thus, we can process the connected components of G one by one. Applying dynamic programming on each component K, we can determine the minimum number b_K of blocking pairs for any matching that is feasible for K, together with a matching M_K for K that admits b_K blocking pairs. If K is a path, then this can be done in a straightforward manner, traversing the edges of K one by one in linear time. If K is a cycle, then we can pick any edge e of K, guess whether it is contained in M_K, or blocks M_K, or neither of the two; we can then process the remainder of K (which is a path) taking into account our guess for e. To compute the minimum number of blocking pairs that a feasible matching admits in our instance, we can simply sum up the values b_K over each connected component K of G. Hence, we arrive at the following.

Observation 2 *Instances of SMC with* $\max(\Delta_{\mathcal{W}}, \Delta_{\mathcal{M}}) \leq 2$ *are linear-time solvable.*

Recall that the case where $\Delta_{\mathcal{W}} = 2$ and $\Delta_{\mathcal{M}} = 3$ is NP-hard by Theorem 9, even if there are no distinguished men to be covered. However, switching the roles of men and women in Theorem 8, we obtain that if there are no women to be covered, then $\Delta_{\mathcal{W}} \leq 2$ guarantees polynomial-time solvability for SMC. This raises the natural question whether SMC with $\Delta_{\mathcal{W}} \leq 2$ can be solved efficiently if the number of distinguished women is bounded. Next we show that this is unlikely, as the problem turns out to be NP-hard for $|\mathcal{W}^\star| = 1$.

Theorem 10 *SMC is NP-hard, even if* $\Delta_{\mathcal{W}} = 2$, $|\mathcal{W}^\star| = 1$ *and there is only one man* m *with* $|L(m)| > 3$.

Proof We present a reduction from the following special case of Exact- 3- Cover. We are given a set $U = \{u_1, \ldots, u_n\}$, a family $\mathcal{S}$ of subsets $S_1, \ldots, S_m$ of U, each having size 3, such that each element of U occurs in at most three sets of $\mathcal{S}$. The task is to decide whether there exists a collection of $n/3$ sets in $\mathcal{S}$ whose union covers U; such a collection of subsets is called an *exact cover* for U. This problem is NP-complete [21, GT2]. We construct an equivalent instance I of SMC as follows.

The set $\mathcal{W}$ of women in I contains the women s_j, p_j^1, p_j^2, p_j^3, and q_j for each $j \in \{1, \ldots, m\}$, women x and y, as well as two women $a_{i,j}, b_{i,j}$ for each element u_i contained in S_j for each $j \in \{1, \ldots, m\}$. The men defined in I are $\hat{p}_j^1$, $\hat{p}_j^2$, $\hat{p}_j^3$, $\hat{q}_j$, and t_j for each $j \in \{1, \ldots, m\}$, a man c_i for each $u_i \in U$, a man $\hat{b}_{i,j}$ for each element u_i contained in S_j for each $j \in \{1, \ldots, m\}$, plus one additional man $\hat{y}$. (The pairs $\{w, \hat{w}\}$ form a stable matching in I.) The only distinguished woman in I is x, and the set of distinguished men is $\mathcal{M}^\star = \{c_i \mid i = 1, \ldots, n\} \cup \{t_j \mid j = 1, \ldots, m\}$. The preferences of each person are as shown in Table 5. Note that since each subset S_j contains three elements, and each element u_i is contained in at most three subsets from $\mathcal{S}$, we get that all men except for $\hat{y}$ have a preference list of length at most 3, as promised. To finish the construction, we set the number of allowed blocking pairs to be $b = 2m + 2n/3 + 1$.

Table 5 Preference lists of women and men in the proof of Theorem 10

$L(x) = (\hat{y})$,	
$L(y) = (\hat{y})$,	
$L(s_j) = (\hat{p}_j^3, \hat{y})$,	
$L(p_j^1) = (\hat{p}_j^1, t_j)$,	
$L(p_j^h) = (\hat{p}_j^{h-1}, \hat{p}_j^h)$	for $h \in \{2, 3\}$,
$L(q_j) = (\hat{q}_j, t_j)$,	
$L(a_{i,j}) = (\hat{b}_{i,j}, \hat{p}_j^h)$	for the unique $h \in \{1, 2, 3\}$ that satisfies $i = \mathrm{ind}(j, h)$,
$L(b_{i,j}) = (\hat{b}_{i,j}, c_i)$,	
$L(\hat{y}) = (y, s_1, s_2, \ldots, s_m, x)$,	
$L(t_j) = (p_j^1, q_j)$,	
$L(\hat{p}_j^h) = (p_j^h, a_{i,j}, p_j^{h+1})$	for $h \in \{1, 2\}$ and $i = \mathrm{ind}(j, h)$,
$L(\hat{p}_j^3) = (p_j^3, a_{i,j}, s_j)$	where $i = \mathrm{ind}(j, 3)$,
$L(\hat{q}_j) = (q_j)$,	
$L(\hat{b}_{i,j}) = (b_{i,j}, a_{i,j})$,	
$L(c_i) = ([B_i])$	where $B_i = \{b_{i,j} \mid u_i \in S_j\}$ and $[B_i]$ is some fixed ordering of B_i

We denote by $\mathrm{ind}(j, h)$ the index i for which u_i is the h-th element in S_j. When not stated otherwise, indices take all possible values

We claim that I admits a feasible matching with at most b blocking pairs if and only if $(U, \mathcal{S})$ is a "yes"-instance of EXACT- 3- COVER.

"$\Rightarrow$": Suppose that M is a feasible matching for I with at most b blocking pairs. Clearly, as x is distinguished, M must contain the edge $\{x, \hat{y}\}$. Thus, $\{y, \hat{y}\}$ is blocking in M. Second, since t_j is distinguished for each $j \in \{1, \ldots, m\}$, we get that M matches t_j either to q_j or to p_j^1, which in turn implies that either $\{q_j, \hat{q}_j\}$ or $\{p_j^1, \hat{p}_j^1\}$ blocks M, leading to m additional blocking pairs for M. Third, consider now any man c_i, $i \in \{1, \ldots, n\}$: as c_i is distinguished, we know $M(c_i) = b_{i,j}$ for some j such that S_j contains u_i. In this case, $\{b_{i,j}, \hat{b}_{i,j}\}$ is also a blocking pair for M, yielding n blocking pairs of such form. Thus, if b_U denotes the number of blocking pairs among the edges $\{b_{i,j}, \hat{b}_{i,j}\}$ for indices i and j with $u_i \in S_j$, then we get $b_U \geq n$. This adds up to $m + b_U + 1 \geq m + n + 1$ blocking pairs so far.

Let us define the set E_j of those edges that are incident to s_j, $\hat{p}_j^3$, $\hat{p}_j^2$, or $\hat{p}_j^1$, but not to p_j^1 for some $j \in \{1, \ldots, m\}$; note that these sets are pairwise disjoint, and none of them contains any of the (possibly) blocking edges mentioned in the previous paragraph.

Let k be the number of indices j for which E_j contains no blocking pairs for M; we call such indices (and the subsets S_j corresponding to them) *selected*. The $m - k$ non-selected indices clearly correspond to at least $m - k$ blocking pairs for M (each contained in E_j for some j).

Suppose now that j is selected. Then, since $\{s_j, \hat{y}\}$ is not blocking, we get $M(s_j) = \hat{p}_j^3$, since $\hat{y}$ prefers s_j to its partner x. This implies that $M(p_j^3) = \hat{p}_j^2$, as oth-

erwise $\{p_j^3, \hat{p}_j^3\}$ would be blocking in M. Similarly, from this we obtain $M(p_j^2) = \hat{p}_j^1$. Moreover, for each $h \in \{1, 2, 3\}$, to ensure that $\{\hat{p}_j^h, a_{i,j}\}$ does not block M where u_i is the h-th element of S_j, we must have $M(a_{i,j}) = \hat{b}_{i,j}$. Hence, $\{b_{i,j}, \hat{b}_{i,j}\}$ must be blocking in M. Since this holds for each $h \in \{1, 2, 3\}$ and each selected j, we get $b_U \geq 3k$.

Summing up the blocking pairs identified so far, we know that M admits at least $1 + m + (m - k) + \max(n, 3k)$ blocking pairs. Using that this must be upper-bounded by $b = 1 + m + n + (m - n/3)$, it is easy to show that only $k = n/3$ is possible. This yields that there exist exactly $n/3$ selected indices, and for all such indices j all the edges $\{b_{i,j}, \hat{b}_{i,j}\}$ for which $u_i \in S_j$ are blocking with respect to M. Moreover, we also must have $b_U = n$, as otherwise the number of blocking pairs would exceed b.

However, observe that for each $i \in \{1, \ldots, n\}$, there must exist some j with $u_i \in S_j$ for which the pair $\{b_{i,j}, \hat{b}_{i,j}\}$ is blocking in M (because c_i is distinguished), implying that for each $u_i \in U$ there must exist some selected S_j that contains u_i. Since there are exactly $n/3$ selected sets in $\mathcal{S}$, we get that they form an exact covering of U.

"$\Leftarrow$": Suppose that $(U, \mathcal{S})$ is a "yes"-instance of Exact- 3- Cover. Let J be the set of indices describing a solution, meaning that the subsets $S_j \in \mathcal{S}$ with $j \in J$ form an exact covering of U; clearly, $|J| = n/3$. We define $\sigma(i)$ as the unique index j in J for which $u_i \in S_j$. We define a feasible matching M for I with exactly b blocking pairs as follows (indices take all possible values, if not stated otherwise).

$$
\begin{array}{ll}
M(\hat{y}) = x, & \\
M(t_j) = p_j^1 & \text{if } j \in J, \\
M(t_j) = q_j & \text{if } j \notin J, \\
M(\hat{p}_j^h) = p_j^{h+1} & \text{if } j \in J, h \in \{1, 2\}, \\
M(\hat{p}_j^3) = s_j & \text{if } j \in J, \\
M(c_i) = b_{i,\sigma(i)}, & \\
M(\hat{b}_{i,\sigma(i)}) = a_{i,\sigma(i)}, & \\
M(\hat{w}) = w & \text{if } w \in \mathcal{W} \text{ and neither } M(w) \text{ nor } M(\hat{w}) \text{ is defined yet.}
\end{array}
$$

It is easy to check that M indeed is feasible, and the blocking pairs it admits are exactly the pairs $\{y, \hat{y}\}$, $\{b_{i,\sigma(i)}, \hat{b}_{i,\sigma(i)}\}$ for each $i \in \{1, \ldots, n\}$, $\{p_j^1, \hat{p}_j^1\}$ for each $j \in J$, $\{q_j, \hat{q}_j\}$ for each $j \notin J$, and $\{s_j, \hat{y}\}$ for each $j \notin J$. This proves the lemma. □

Contrasting Theorem 10, we establish fixed-parameter tractability of the case $\Delta_{\mathcal{W}} \leq 2$ with three different parameterizations. Considering our five parameters, the relevant cases (whose tractability or intractability does not follow from our results obtained so far) are as follows, assuming $\Delta_{\mathcal{W}} \leq 2$ throughout. Since letting the number b of blocking pairs and the number $|\mathcal{W}^\star|$ of distinguished women to be unbounded (while assuming $\Delta_{\mathcal{W}} \leq 2$) results in NP-hardness by Theorem 9, in order to obtain fixed-parameter tractability, we need to take either b or $|\mathcal{W}^\star|$ as a parameter. However, taking *only* $|\mathcal{W}^\star|$ as a parameter is not likely to result in tractability, as the case $|\mathcal{W}^\star| = 1$ is still NP-hard by Theorem 10. Thus we need to take either $|\mathcal{M}^\star|$ or $\Delta_{\mathcal{M}}$

as an additional parameter. Altogether, this results in the following parameterizations of the SMC problem, each case subject to the assumption $\Delta_{\mathcal{W}} \leq 2$:

- taking b as the parameter,
- taking $|\mathcal{W}^\star| + |\mathcal{M}^\star|$ as the parameter, and
- taking $|\mathcal{W}^\star| + \Delta_{\mathcal{M}}$ as the parameter.

We show fixed-parameter tractability for each three of these parameterizations. The first two parameterizations can be dealt with an algorithm whose properties are stated in Theorem 11 (see also Corollary 3), while the third parameterization will be considered by Theorem 12.

Theorem 11 *There is a fixed-parameter algorithm for the special case of* SMC *where each woman finds at most two men acceptable (i.e., $\Delta_{\mathcal{W}} \leq 2$), parameterized by the number $|\mathcal{W}_0^\star| + |\mathcal{M}_0^\star|$ of distinguished men and women left unmatched by some stable matching.*

Let M^{opt} denote an optimal solution for our instance I such that $M_s \triangle M^{\text{opt}}$ contains the minimum number of edges; recall that M_s is a fixed stable matching for I. Without loss of generality, we can further assume that there does not exist another optimal solution M' such that (i) M' has the same number of common edges with M_s as M^{opt}, and (ii) for each man m, either $M'(m) = M^{\text{opt}}(m)$ or m prefers $M'(m)$ to $M^{\text{opt}}(m)$.[4] Indeed, as long as such a "superior" matching M' exists, we can simply pick that instead of M^{opt}, until this is no longer possible. This way, we eventually end up with an optimal matching that satisfies our requirement.

We denote by b the number of blocking pairs in M^{opt}.

High-Level Description Let us remark first that simply guessing the optimal partners for each woman in $\mathcal{W}_0^\star$ and then using the polynomial-time algorithm presented in Sect. 6.1 (after exchanging the roles of men and women) does not work, since that algorithm heavily relies on the assumption that we start with a stable matching. In fact, the main difficulty to overcome is that feminine and masculine augmenting paths may "interact" in the sense that certain blocking pairs introduced by a feminine augmenting path can be "eliminated" (i.e., made non-blocking again) by an appropriately chosen masculine path. Therefore, we apply the following strategy. In Phase I, we find all feminine paths (as well as all cycles) in $M_s \triangle M^{\text{opt}}$, and in Phase II we proceed with choosing the masculine paths carefully. Note that in Phase I it does not suffice to find a cheapest set of feminine augmenting paths, since we may not be able to eliminate as many blocking pairs afterwards as it is possible after an optimal choice of feminine paths. Instead, we need to find the exact feminine augmenting paths (and cycles) present in $M_s \triangle M^{\text{opt}}$; this can be accomplished by guessing certain properties of M^{opt}.

In Phase II, the main obstacle is that we do not know which blocking edges should be eliminated in an optimal solution, nor can we guess these edges efficiently. We deal with this problem by guessing the sets of those men in $\mathcal{M}_0^\star$ whose augmenting paths in $M_s \triangle M^{\text{opt}}$ contribute to the elimination of a blocking pair; this information allows us to find these masculine paths. Finally, we apply the algorithm of Theorem 8.

[4] Note that there may be optimal solutions that are incomparable to M^{opt} in the sense that they are preferred by some of the men while not preferred by some other men, but the existence of such a matching is of no importance to us: we can pick any optimal matching that fulfils our requirement above.

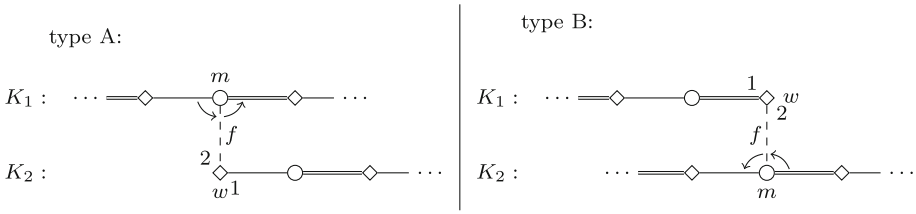

Fig. 5 Illustration of a dependent edge f, running between two connected components K_1 and K_2 of $M_s \triangle M^{\text{opt}}$ where K_1 relies on K_2. Double lines here denote edges of M_s, single lines denote edges of M^{opt}, and f is drawn with a dashed line

For the detailed description of our algorithm, we need a couple of simple observations and some additional notation. We begin with the following implications of the fact that each woman finds at most two men acceptable.

Proposition 3 *Suppose $\Delta_{\mathcal{W}} \leq 2$. Let P_1 and P_2 be two augmenting paths.*

(a) *If P_1 and P_2 start at some $w \in \mathcal{W}_0^\star$ through the same edge, then one of them is a subpath of the other.*
(b) *If P_1 and P_2 start at different women w_1 and w_2, respectively, and P_1 and P_2 are not disjoint, then the set of their common vertices induces a suffix of either P_1 or P_2 (or both); their first common vertex is a man.*
(c) *If P_1 and P_2 are disjoint and e is an edge incident to both, then one of the paths starts or ends at a women w, and e connects w with a man on the other path.*

Observations (a) and (b) of Proposition 3 immediately suggest that feminine paths are easy to find, since once we decide which edge to start with, all possible augmenting paths lead in the same direction; the only difficulty arises in deciding when to stop. Observation (c) describes the limited ways in which two augmenting path can interact; we next look closer at such interactions.

We say that an edge $f = \{m, w\}$ of G (with $m \in \mathcal{M}$ and $w \in \mathcal{W}$) is *dependent* if it connects two different connected components K_1 and K_2 of $M_s \triangle M^{\text{opt}}$ and, in addition, it holds that $M_s \triangle K_1$ admits more blocking pairs than $M_s \triangle (K_1 \cup K_2)$. We will say that f, and with a slight abuse of the notation, also K_1 *relies* on K_2. We say that

– f has *type A*, if w is the endpoint of K_2 (which is a path), f connects w with a man m on K_1 that prefers $M_s(m)$ to w, and w to $M^{\text{opt}}(m)$, and w is unmatched by M_s and prefers $M^{\text{opt}}(w)$ to m;
– f has *type B*, if w is the endpoint of K_1 (which is a path), unmatched by M^{opt}, and f connects w with a man m on K_2 that prefers $M^{\text{opt}}(m)$ to w, and w to $M_s(m)$.

See Fig. 5 for an illustration of the above definitions.

Lemma 3 *Let $\mathcal{K}$ be the set of connected components of $M_s \triangle M^{\text{opt}}$, and let $H \subset \mathcal{K}$ and $K \in \mathcal{K} \backslash H$. Then any edge f that is blocking in $M_s \triangle H$ but is not blocking in $M_s \triangle (H \cup K)$ connects K with a connected component of H, is a dependent edge relying on K, and has either type A or type B.*

Proof Let f be an edge that blocks $M_s \triangle H$ but does not block $M_s \triangle (H \cup K)$. Since M_s is stable, f must have an endpoint in a connected component of H, because it blocks $M_s \triangle H$. However, as it ceases to be blocking in $M_s \triangle (H \cup K)$, it also must have an endpoint in K. Let w and m be the woman and the man connected by f, respectively. We distinguish now between two cases.

First, let us assume that w is contained in K. Since m is not contained in K, and w can only be connected to two men, it follows that the degree of w in K is at most (and hence exactly) 1. Since each connected component of $M_s \triangle M^{\text{opt}}$ is either a cycle or a path, this implies that K is a path with w being one of its endpoints.

Since f blocks $M_s \triangle H$ and the partner of m under $M_s \triangle H$ is $M^{\text{opt}}(m)$, we know that m prefers w to $M^{\text{opt}}(m)$. But as M_s is stable, m must prefer $M_s(m)$ to w.

Notice now that w is unmatched either in M_s or in M^{opt}. However, as f blocks $M_s \triangle H$ (where w is matched as in M_s) but does not block $M_s \triangle (H \cup K)$ (where w is matched as in M^{opt}), it must be the case that w is matched by M^{opt} but is unmatched in M_s. Further, as f is not blocking in $M_s \triangle (H \cup K)$ even though m prefers w to $M^{\text{opt}}(m)$, we get that w prefers $M^{\text{opt}}(w)$ to m. This proves that f has type A.

Second, let us assume that w is contained in a connected component K' of H. As in the previous case, we quickly get that K' must be a path, with w being an endpoint of K'. Since f blocks $M_s \triangle H$, we know that w is matched by M_s but is unmatched by M^{opt}. The stability of M_s implies also that w prefers $M_s(w)$ to m.

Regarding m, the fact that f blocks $M_s \triangle H$ implies that m prefers w to $M_s(m)$. However, since f ceases to be blocking in $M_s \triangle (H \cup K)$, we get that m prefers $M^{\text{opt}}(m)$ to w. This proves that f has type B. □

We are now ready to present our algorithm, which is a branching algorithm: throughout its course, we make several "guesses" for which all possibilities have to be explored. When certain guesses turn out to be trivially wrong, such guesses are discarded, and we might not explicitly mention this in the algorithm. (In Step 1, we describe such issues in detail for illustration, but later we omit them.) Phases I and II consist of Steps 1-5 and Steps 6-8, respectively.

Step 1: Guessing the first edges of feminine augmenting paths First, for each $w \in \mathcal{W}_0^\star$ with $|L(w)| = 2$, we guess the edge of M^{opt} incident to w. This results in at most $2^{|\mathcal{W}_0^\star|}$ possibilities, all of which must be explored. Naturally, we discard those guesses where the edges $\{w, M^{\text{opt}}(w)\}$, $w \in \mathcal{W}_0^\star$, do not form a matching. From now on we assume that we know $M^{\text{opt}}(w)$ for each $w \in \mathcal{W}_0^\star$.

Additionally, we delete those edges $\{m, w\}$ for which $w \in \mathcal{W}_0^\star$ and w prefers $M^{\text{opt}}(w)$ to m. Such edges are neither needed in M^{opt}, nor can they block any matching that contains all the edges $\{w, M^{\text{opt}}(w)\}$, $w \in \mathcal{W}_0^\star$, guessed in this step.

Before proceeding to Step 2, we state an important lemma about augmenting paths.

Lemma 4 *Each connected component of $M_s \triangle M^{\text{opt}}$ that is not a cycle is an augmenting path. Further, assume that Step 1 has already been performed, and K_1 and K_2 are connected components of $M_s \triangle M^{\text{opt}}$ such that K_1 relies on K_2 via a dependent edge f. Then*

(a) *if f has type A, then K_2 is a masculine path and not a feminine path;*
(b) *if f has type B, then K_1 is a feminine path and not a masculine path, and K_2 is either a cycle or a feminine path.*

Proof We begin by proving the first sentence of the lemma. Let us first suppose that Q is a connected component of $M_s \triangle M^{\text{opt}}$ that is not a cycle (thus is a path) but is not an augmenting path. The feasibility of M^{opt} implies that if Q has a distinguished person p as its endpoint, then p must be unmatched by M_s. This means that Q can only be non-augmenting if neither of its endpoints is distinguished. This implies that $M_Q := M^{\text{opt}} \triangle Q$ is a feasible matching. Recall that b is the number of blocking pairs M^{opt} admits. If M_Q admits at most b blocking pairs as well, then this contradicts the choice of M^{opt}, because there are strictly less edges in $M_s \triangle M_Q$ than in $M_s \triangle M^{\text{opt}}$.

Hence, $M^{\text{opt}} \triangle Q$ admits at least $b+1$ blocking pairs. Since Q is a maximal path in $M_s \triangle M^{\text{opt}}$ and M_s is stable, Proposition 1 implies that there must be an edge along Q that blocks M^{opt}. Obviously, there is no edge on Q that blocks M_s, since M_s is stable. Hence, modifying M^{opt} by switching the edges of M_s and M^{opt} along Q *decreases* the number of blocking edges *among the edges of* Q. However, the same operation *increases* the *total* number of blocking pairs (from b to at least $b+1$), thus we get that there must exist at least two edges that become blocking only as a result of this switch along Q. In other words, there exist at least two edges that are blocking in $M^{\text{opt}} \triangle Q$ but are not blocking in M^{opt}. By Lemma 3 we obtain that these edges must be dependent edges relying on Q.[5] By the definition of a type A edge, at most one type A edge can rely on Q (as only one endpoint of Q can be a woman unmatched by M_s, and this woman may be incident to at most one edge outside Q), so we get that at least one of these dependent edges relying on Q must have type B.

Let us call the man endpoint of a type B dependent edge a *joiner*; by the previous paragraph, Q contains at least one joiner. Our aim is to fix an "outer-most" joiner m on Q. However, for technical reasons we also have to take into account a special case where some man m_ℓ lying on Q, called the *looper* for Q, fulfils the following properties: (i) one endpoint of Q is a woman w_ℓ that is adjacent to m_ℓ in G but prefers $M_s(w_\ell)$ to m_ℓ, and (ii) m_ℓ prefers w_ℓ to $M_s(m_\ell)$ but prefers $M^{\text{opt}}(m_\ell)$ to w_ℓ.[6] Now, we choose a man m that is either a joiner or the looper for Q so that the following holds: if m splits Q into two subpaths Q_1 and Q_2 with Q_1 containing $M^{\text{opt}}(m)$, then $V(Q_1) \setminus \{m\}$ contains neither joiners nor the looper for Q.

Case for a Joiner First, let us assume that m is a joiner. In this case, there might be several women who form a dependent edge of type B with m, so let w denote the one that is most preferred by m. Let f be the edge $\{m, w\}$, and let P be the path of $M_s \triangle M^{\text{opt}}$ that has w as its endpoint. We illustrate these concepts in Fig. 6.

We claim that $M_f = M^{\text{opt}} \triangle (Q_1 \cup \{f\}) = (M^{\text{opt}} \triangle Q_1) \cup \{f\}$ is an optimal solution. Observe that $M_s \triangle M_f$ can be obtained from $M_s \triangle M^{\text{opt}}$ by deleting Q and substituting P by the path $P + f + Q_2$ where the plus sign means concatenation. Let x denote

[5] Formally, we need to apply Lemma 3 with setting H to contain all connected component of $M_s \triangle M^{\text{opt}}$ except for Q, and setting K as Q.

[6] Note that there cannot exist two loopers for Q, because only one woman endpoint of Q can be matched by M_s and such a woman can only be connected to at most one man on Q other than its partner in M_s.

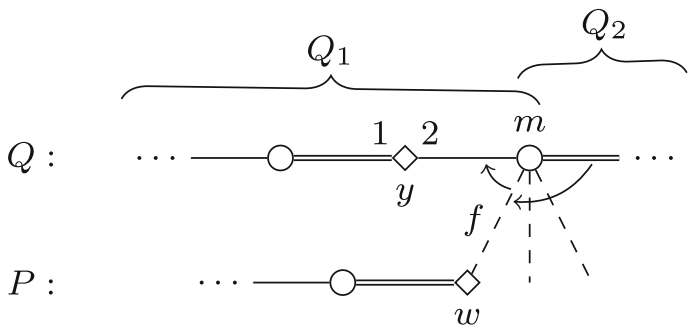

Fig. 6 The joiner m splits Q into subpaths Q_1 and Q_2. Here and in later figures, double lines denote edges of M_s. Single lines denote edges of M^{opt}, and dashed lines are for dependent edges

the endpoint of Q_1 that is not m. First, M_f is clearly feasible, since x is not distinguished (by our assumption that Q is non-augmenting). Next, suppose that some edge e is blocking in M_f but is not blocking in M^{opt}. It is easy to see that by our choice of w, e cannot be incident to w or m. As only vertices in $Q_1 \cup \{f\}$ are matched differently in M_f as in M^{opt}, we obtain that one endpoint of e must lie on Q_1. Using the stability of M_s, we also get that the other endpoint of e must lie either on a connected component of $M_s \triangle M^{\text{opt}}$ other than Q or on Q_2. In the former case, Lemma 3 implies[7] that e is a dependent edge. By our choice of m, e cannot be of type B, hence by Lemma 3, e must have type A, and thus it is incident to x which must be a woman not covered by M_s. In the latter case, supposing that the man endpoint m' of e lies on Q_1 we quickly get that it must be the looper for Q. To see this, first note that the only woman that m' can be adjacent to on Q_2 must be the common endpoint of Q_2 and Q (by $\Delta_{\mathcal{W}} \leq 2$); let w' denote this woman. Second, since e blocks M_f but it does not block M^{opt}, we know that m' prefers w' to $M_s(m')$, but prefers $M^{\text{opt}}(m')$ to w'. Third, since m' prefers w' to $M_s(m')$, the stability of M_s implies that w' is matched by M_s and w' prefers $M_s(w')$ to m'. Hence, m' is indeed the looper for Q, contradicting our choice of m. Thus, we know that it must be the woman endpoint of e that lies on Q_1. Since the two endpoints of Q_1 are $m \in \mathcal{M}$ and x, by $\Delta_{\mathcal{W}} \leq 2$ we obtain that the woman endpoint of e lying on Q_1 can only be x, which therefore must be a woman not covered by M_s. In either case, we can conclude that there can only exist at most one such edge e (because x has degree at most 2 in G). So the number of edges that are blocking in M_f but not blocking in M^{opt} is at most one.

Furthermore, since f is a dependent edge of type B, by definition we know that m prefers $y = M^{\text{opt}}(m)$ to $M_s(m)$, and hence, y must prefer $M_s(y)$ to m (as otherwise $\{m, y\}$ would be blocking in M_s, which is not possible). Note that, in particular, $M_s(y)$ exists and Q_1 contains at least two edges. By Proposition 1, Q_1 must contain at least one edge that blocks M^{opt}, and this edge is not blocking in M_f, simply because any edge of Q_1 that blocks M^{opt} but be an edge contained in M_s and hence in M_f. Thus, the number of edges blocking M_f cannot be more than b. Hence, M_f is an optimal

[7] Again, we need to apply Lemma 3 with setting H to contain all connected component of $M_s \triangle M^{\text{opt}}$ except for Q, and setting K as Q.

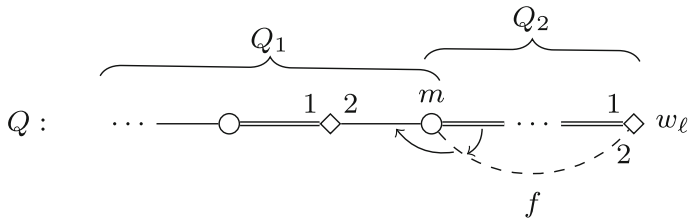

Fig. 7 The looper m of Q splits Q into subpaths Q_1 and Q_2

solution such that there are less edges in $M_s \triangle M_f$ than in $M_s \triangle M^{\text{opt}}$, a contradiction. This proves the first statement of the lemma for the case of m being a joiner.

Case for a Looper Let us now assume that $m = m_\ell$ is the looper for Q. In this case let f be the edge connecting m_ℓ to the woman endpoint w_ℓ of Q that is matched by M_s; it is easy to see that w_ℓ is the endpoint of Q_2 as well. See Fig. 7 for an illustration. Let us again define $M_f = M^{\text{opt}} \triangle (Q_1 \cup \{f\})$. As in the previous case, it is immediate that M_f is feasible. Arguing similarly as before and using that Q_1 does not contain any joiners, we get that any edge that blocks M_f but does not block M^{opt} can only be a type A dependent edge incident to the (woman) endpoint x of Q_1 that is not m_ℓ. Hence, there can be at most one edge blocking M_f but not M^{opt}. From this point on, we can use the same reasoning as in the previous case to arrive at the conclusion that M_f is an optimal solution that has more edges common with M_s than M^{opt}, a contradiction. This proves the first statement of our lemma.

Let us prove (b) now. Suppose that K_1 and K_2 are two connected components of $M_s \triangle M^{\text{opt}}$ such that K_1 relies on K_2 via an edge $f = \{w, m\}$ of type B. By the definition of a type B edge, we know that K_1 is a path with an endpoint that is a woman covered by M_s, so its other endpoint is either a woman not covered by M_s, or a man covered by M_s. Thus, K_1 cannot be a masculine path, so by the first statement of the lemma, it is feminine. It remains to show that if K_2 is a path, then it is feminine. Assume for contradiction that K_2 is a non-feminine path Q in $M_s \triangle M^{\text{opt}}$ and some other path P in $M_s \triangle M^{\text{opt}}$ relies on Q via a type B edge f. In this case we can argue exactly as above to show that there must exist an optimal matching M_f (defined the same way as we did while proving the first statement of the lemma) for which $M_s \triangle M_f$ contains less edges than $M_s \triangle M^{\text{opt}}$, a contradiction.

To show (a), suppose that K_1 and K_2 are two connected components of $M_s \triangle M^{\text{opt}}$ such that K_1 relies on K_2 via an edge $f = \{w, m\}$ of type A. By the definition of a type A edge, the woman endpoint w of f is unmatched in M_s, and K_2 is a path that has w as an endpoint. Also, m is the second choice of w and $M^{\text{opt}}(w)$ is the first choice of w. Hence $w \in \mathcal{W}_0^\star$ is not possible, as in that case the edge f would have been deleted in Step 1 of the algorithm. Furthermore, since K_2 is an M_s-alternating path with an endpoint in $\mathcal{W}_0$, the other endpoint of K_2 cannot be a woman in $\mathcal{W}_0$. Hence, K_2 is not a feminine augmenting path. Since we already know that K_2 is an augmenting path, (a) follows. □

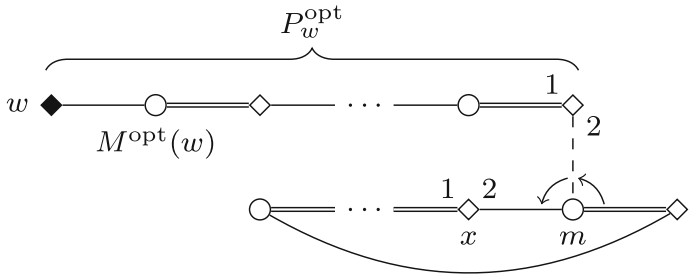

Fig. 8 Step 2 of the algorithm for Theorem 11, for finding cycles in $M \triangle M^{\mathrm{opt}}$

By Lemma 4, each connected component in $M_s \triangle M^{\mathrm{opt}}$ that is a path must be an augmenting path; we let P_x^{opt} denote the augmenting path in $M_s \triangle M^{\mathrm{opt}}$ that contains some $x \in \mathcal{W}_0^\star \cup \mathcal{M}_0^\star$ as an endpoint.

Consider now any woman $w \in \mathcal{W}_0^\star$. Let us define a *candidate path* for w as an augmenting path that starts through the edge $\{w, M^{\mathrm{opt}}(w)\}$. Observe that P_w^{opt} itself is a candidate path for w. Furthermore, although we cannot determine P_x^{opt} directly, Proposition 3 implies that there exists a unique maximal candidate path A_w for w that contains P_w^{opt} as a subpath. We can compute A_w easily: starting from the edge $\{w, M^{\mathrm{opt}}(w)\}$, we can build A_w by adding edges of M_s and edges not in M_s in an alternating manner, always appending a new edge to the last vertex of the current subpath of A_w. Observe that after an addition of an edge of M_s, the resulting subpath of A_w ends at a woman, and thus $\Delta_{\mathcal{W}} \leq 2$ implies that there exists at most one edge that we can add to our path. This shows the uniqueness of A_w.

Step 2: Finding cycles in $M_s \triangle M^{\mathrm{opt}}$ We make one more guess for each $w \in \mathcal{W}_0^\star$ by guessing whether P_w^{opt} relies on some cycle of $M_s \triangle M^{\mathrm{opt}}$ or not; this again yields $2^{|\mathcal{W}_0^\star|}$ possibilities. If P_w^{opt} relies on some cycle C, then by Proposition 3, both P_w^{opt} and C can be found in time $O(|P_w^{\mathrm{opt}}| + |C|)$ as follows. We compute the longest possible candidate path A_w for w. Since P_w^{opt} relies on cycle C, the last person on A_w must be a woman x incident to an edge $\{x, m\}$ for which m is already on A_w. Then the subpath of A_w between m and x together with the edge $\{x, m\}$ form the cycle C. For an illustration, see Fig. 8.

Lemma 5 *Assuming that the guesses made by the algorithm are correct, all cycles in $M_s \triangle M^{\mathrm{opt}}$ and all paths relying on some cycle in $M_s \triangle M^{\mathrm{opt}}$ are found in Step 2.*

Proof Let C be a cycle in $M_s \triangle M^{\mathrm{opt}}$. By the stability of M_s, no edge of C blocks the matching $M^{\mathrm{opt}} \triangle C$. Since $M^{\mathrm{opt}} \triangle C$ has more edges in common with M_s than M^{opt} does, we know by the choice of M^{opt} that there must be an edge incident to C that is blocking in $M^{\mathrm{opt}} \triangle C$ but is not blocking in M^{opt}. By Lemma 3, such an edge must be a dependent edge e relying on C. This proves that there are no cycles in $M_s \triangle M^{\mathrm{opt}}$ without augmenting paths relying on them.

By Lemma 4 we know that dependent edges of type A rely on (masculine) paths,[8] so e must have type B. From this, Lemma 4 yields that all paths relying on a cycle must be feminine paths. Hence, Step 2 indeed finds all paths that rely on some cycle, together with the cycles that these paths rely upon, which (as pointed out above) means that Step 2 finds all cycles of $M_s \triangle M^{\text{opt}}$, proving our claim. □

Step 3: Finding neutral paths In this step, for each $w \in \mathcal{W}_0^\star$ we guess whether w lies on a neutral path in $M_s \triangle M^{\text{opt}}$, yielding at most $2^{|\mathcal{W}_0^\star|}$ possibilities. Clearly, if w lies on a neutral path, then by Proposition 3, P_w^{opt} is the maximal candidate path starting with the edge $\{w, M^{\text{opt}}(w)\}$, that is, $P_w^{\text{opt}} = A_w$.

Step 4: Finding feminine paths relying on other feminine paths In this step, we find all feminine paths that rely on some other feminine path; note that feminine paths that do not depend on other feminine paths will be taken care of by Step 5.

In Step 4, we first guess for each $w \in \mathcal{W}_0^\star$ whether P_w^{opt} relies on another feminine path, and if so, we also guess on which one. This means at most $|\mathcal{W}_0^\star|$ possibilities for each $w \in \mathcal{W}_0^\star$, a total of $|\mathcal{W}_0^\star|^{|\mathcal{W}_0^\star|}$ possibilities.

Supposing that, according to our guesses, P_w^{opt} relies on P_y^{opt} for some w and y in $\mathcal{W}_0^\star$, we can find P_w^{opt} easily as follows. Consider the maximal candidate paths A_w and A_y for w and y, respectively. We know that P_w^{opt} and P_y^{opt} are disjoint, but supposing that our guesses were made correctly, Lemma 4 implies that they are connected by a dependent edge f of type B relying on P_y^{opt}. Hence, A_w contains P_w^{opt}, followed by the dependent edge f connecting P_w^{opt} and P_y^{opt}, followed possibly by a suffix of P_y^{opt}. Thus, by Proposition 3, A_w and A_y must share the man endpoint m of f as their first common vertex, and they coincide after m. Therefore, P_w^{opt} can be obtained by deleting all edges incident to or occurring after m on A_w. This way, f and P_w^{opt} can be found in $O(|P_w^{\text{opt}}| + |P_y^{\text{opt}}|)$ time.

Notice that this method only enables us to locate P_w^{opt} but not P_y^{opt}; we only know that P_y^{opt} must contain the man m incident to f, but we do not know where P_y^{opt} ends. To remind us that P_y^{opt} contains m, we store m as an *obligatory* man for y.

Step 5: Finding all remaining feminine paths In this step, we find all feminine paths that rely neither on a cycle, nor on another feminine path. To do so, we will use a surprisingly simple method. Intuitively, it makes sense to construct augmenting paths to be as short as possible, since the longer such a path gets, the more opportunities arise for a blocking pair to appear along the way. This intuition is almost entirely correct, if no other path (or cycle) relies on the augmenting path P_w^{opt} we are looking for: it turns out that P_w^{opt} is either indeed the shortest possible, or if this is not the case, then we can "get rid of" only one blocking pair by constructing a longer augmenting path—and this, too, we can keep as short as it is possible, by stopping whenever we encounter a decrease in the number of blocking pairs along the augmenting path. In

[8] We remark that it is still possible that C itself relies on a masculine path via a dependent edge of type A, but such edges do not play a role in Step 2.

the case when some other path relies on P_w^{opt}, we need to be somewhat more careful, and take into account the obligatory men defined in Step 4.

After this high-level explanation, let us define Step 5 formally. Let $\mathcal{W}_r$ be the set of those women $w \in \mathcal{W}_0^\star$ for which P_w^{opt} has not been found yet (that is, where P_w^{opt} is not neutral, and it does not rely on any other path or cycle in $M_s \triangle M^{\text{opt}}$). Let us fix any woman $w \in \mathcal{W}_r$. With any candidate path Q for w, we associate a matching $M_Q = (M^{\text{opt}} \triangle P_w^{\text{opt}}) \triangle Q$; we can think of M_Q as the matching that results from replacing P_w^{opt} with Q in the symmetrical difference of M_s and the solution. Let us state a few useful properties of candidate paths and their associated matchings.

Proposition 4 *The following holds for any candidate path Q for w:*

(a) *$M_s \triangle Q$ is a matching.*
(b) *M_Q is a feasible matching.*
(c) *Q must be one of the following two forms:*

  (i) *Q ends at a man in $\mathcal{M}_0$, in which case $Q = A_w = P_w^{\text{opt}}$, or*
  (ii) *Q ends at a non-distinguished woman.*

Proof Recall that the definition of an augmenting path (see the beginning of Sect. 6) directly requires that $M_s \triangle Q$ be a matching, proving (a). This immediately implies that if Q ends at a man, then this man cannot be matched by M_s, from which $Q = A_w$ and thus $Q = P_w^{\text{opt}}$ follows. Using again the definition of an augmenting path, we get that the last person on Q is not distinguished (because P_w^{opt} is not neutral by Step 3), showing (c).

To see that M_Q is a matching, let $\mathcal{K}_{-w}$ denote the union of all connected components of $M_s \triangle M^{\text{opt}}$ except for P_w^{opt}. Then $M^{\text{opt}} \triangle P_w^{\text{opt}} = M_s \triangle \mathcal{K}_{-w}$ is a matching. But as $\mathcal{K}_{-w}$ is vertex-disjoint from Q, from (a) we get that $M_Q = (M_s \triangle \mathcal{K}_{-w}) \triangle Q$ must be a matching too. The properties of Q described in (c) show also that M_Q is feasible. □

Now, we define two candidate paths for w as follows:

- First, we let Q_w^1 be the shortest candidate path for w that contains all obligatory men for w. In particular, if there are no obligatory men for w, then Q_w^1 is the shortest candidate path for w; see Fig. 9 for an illustration.
- Second, let Q_w^2 be the shortest candidate path containing Q_w^1 such that $M_s \triangle Q_w^2$ admits less blocking pairs than $M_s \triangle Q_w^1$. Such a path, however, may not exist, in which case we leave Q_w^2 undefined.

Lemma 6 *If P_w^{opt} does not rely on any other path or cycle in $M_s \triangle M^{\text{opt}}$, then either $P_w^{\text{opt}} = Q_w^1$ or $P_w^{\text{opt}} = Q_w^2$.*

Proof First observe that for any $w \in \mathcal{W}_r$, if the guesses made by the algorithm are correct, then P_w^{opt} contains Q_w^1 as a subpath. To see this, recall that if the guesses made in Step 4 of the algorithm are correct, then P_w^{opt} indeed must contain all obligatory men for w. Since P_w^{opt} is a candidate path for w, it indeed must contain Q_w^1.

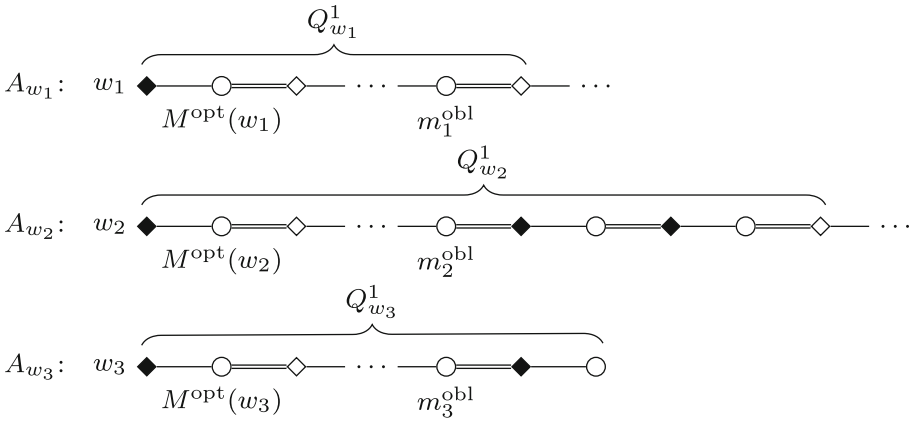

Fig. 9 An example of the paths Q^1_w for three women $w \in \{w_1, w_2, w_3\} \subseteq \mathcal{W}_r$, as defined in Step 5. The man m_i^{obl} for each $i \in \{1, 2, 3\}$ is the last obligatory man for w_i on the maximal augmenting path A_{w_i}. In cases where there is no obligatory man for w_i, one can identify m_i^{obl} with $M^{\mathrm{opt}}(w_i)$ to get an appropriate figure

Next, suppose that $P_w^{\mathrm{opt}} \neq Q^1_w$, and consider the matching $M_1 := M_{Q^1_w}$. By Proposition 4, M_1 is a feasible matching, and $P_w^{\mathrm{opt}} \neq Q^1_w$ implies that Q^1_w ends at a non-distinguished woman x.

Claim 2 *Let Q be a candidate path for w such that $Q^1_w \subset Q \subseteq P_w^{\mathrm{opt}}$. If e is an edge blocking in M_1 but not blocking in M_Q, then e is adjacent to x.*

Proof of Claim 2 Suppose for contradiction that e is not incident to x. First note that e must be adjacent to a person in $V(Q)\backslash V(Q^1_w)$, as all other persons (except for x) have the same partner in M_1 as in M_Q. Second, as M_s is stable, at least one endpoint of e must be matched differently in M_1 than in M_s. Thus, e must connect a person in $V(Q)\backslash V(Q^1_w)$ either (i) with a connected component of $\mathcal{K}_{-w}$, denoting the union of all connected components of $M_s \triangle M^{\mathrm{opt}}$ except for P_w^{opt}, or (ii) with Q^1_w.

In case (i), e must be blocking in $M_s \triangle \mathcal{K}_{-w}$ (as e blocks M_1). As e is not blocking in $M_Q = M_s \triangle (\mathcal{K}_{-w} \cup Q)$, it is not blocking in $M^{\mathrm{opt}} = M_s \triangle (\mathcal{K}_{-w} \cup P_w^{\mathrm{opt}})$ either (because any person in $V(Q)\backslash V(Q^1_w)$ has the same partner in M_Q as in M^{opt}, except possibly for the last person on Q who might be left unmatched in M_Q). Therefore, Lemma 3 implies that e is a dependent edge relying on P_w^{opt}. By Lemma 4, only feminine paths may rely on P_w^{opt}. However, in Step 4 we have stored the man endpoint of all edges that connect P_w^{opt} with a feminine path relying on P_w^{opt} as an obligatory man, so assuming correct guesses in Step 4, e must have an obligatory man as its endpoint. But Q^1_w contains all the obligatory men for w, which contradicts our assumption of case (i).

In case (ii), let w_e and m_e be the woman and man endpoint of e, respectively; we assumed $w_e \neq x$. By $\Delta_{\mathcal{W}} \leq 2$, w_e can only be adjacent to both $V(Q)\backslash V(Q^1_w)$ and $V(Q^1_w)$ if either $w_e = w$, or Q and also P_w^{opt} ends at w_e. In the former case, $m_e \in V(P_w^{\mathrm{opt}})\backslash V(Q^1_w)$, yielding $M_1(m_e) = M_s(m_e)$. As e blocks M_1, m_e prefers w

to $M_s(m_e)$; however, w is not matched in M_s, so this means that e blocks M_s as well, a contradiction. In the latter case, m_e lies on Q_w^1, yielding $M_1(m_e) = M^{\text{opt}}(m_e)$. As e blocks M_1, m_e prefers w_e to $M^{\text{opt}}(m_e)$; however, w_e is not matched by M^{opt}, so this means that e blocks M^{opt} as well, a contradiction. Hence the claim follows. □

Let c_1 denote the number of blocking pairs with respect to M_1 that are incident to a man or woman on Q_w^1.

Claim 3 *If Q is a candidate path for w such that $Q_w^1 \subset Q \subseteq P_w^{\text{opt}}$, and c is the number of blocking pairs with respect to M_Q that are incident to a man or woman on Q, then $c \geq c_1 - 1$.*

Proof of Claim 3 By Claim 2, $c \leq c_1 - 2$ can only happen if both edges incident to x are blocking in M_1, and none of them is blocking in M_Q. However, this cannot happen, as we are going to show now. Let $m_1 = M_s(x)$ and $m_2 = M^{\text{opt}}(x)$ be the two men in x's preference list; note that Q_w^1 contains m_1 but not m_2. If both $\{x, m_1\}$ and $\{x, m_2\}$ are blocking in M_1, then m_1 prefers x to $M^{\text{opt}}(m_1)$, and m_2 either prefers x to $M_s(m_2)$ or m_2 is unmatched by M_s. Therefore the stability of M_s implies that x prefers m_1 to m_2, as otherwise $\{x, m_2\}$ would block M_s. But this shows that $\{x, m_1\}$ is a blocking pair in M^{opt} and thus in M_Q as well, proving our claim. □

Now, let c^{opt} denote the number of blocking pairs with respect to M^{opt} that are incident to a man or woman on P_w^{opt}. By our assumption $Q_w^1 \neq P_w^{\text{opt}}$, we get $c^{\text{opt}} < c_1$, as otherwise M_1 would also be an optimal solution which would contradict our choice of M^{opt} (as $M_s \triangle M_1$ contains less edges than $M_s \triangle M^{\text{opt}}$). Hence, Claim 3 implies $c^{\text{opt}} = c_1 - 1$.

By the definition of Q_w^2, we know that $Q_w^1 \subset Q_w^2 \subseteq P_w^{\text{opt}}$, and Proposition 4 yields that $M_2 := M_{Q_w^2}$ is a feasible solution. The definition of Q_w^2 and Claim 3 together imply that the number of blocking pairs with respect to M_2 that are incident to a man or woman on Q_w^2 is exactly $c_1 - 1$. Hence, M_2 is an optimal solution. Therefore, $M_s \triangle M_2$ cannot contain less edges than $M_s \triangle M^{\text{opt}}$, because assuming otherwise would contradict our choice of M^{opt}. From this, $Q_w^2 = P_w^{\text{opt}}$ follows. □

According to Lemma 6, we can find all remaining feminine paths in $M_s \triangle M^{\text{opt}}$ by guessing for each $w \in \mathcal{W}_r$ whether P_w^{opt} equals Q_w^1 or Q_w^2. This means at most $2^{|\mathcal{W}_0^\star|}$ guesses; the computations needed for each guess take linear time.

Step 6: Computing elimination paths. Let $\mathcal{F}$ be the union of cycles, feminine and neutral paths found in Steps 1 to 5. When searching for masculine paths, we will have to deal with edges that might be type A dependent edges in the optimum solution. We call an edge *volatile* if it connects a woman in $\mathcal{W}_0$ with her second choice. The importance of this definition is shown by the following two lemmas.

Lemma 7 *Let f be an edge that is blocking in $M_s \triangle \mathcal{F}$, but not blocking in M^{opt}. Then f is volatile, its woman endpoint is contained on a non-feminine path of $M_s \triangle M^{\text{opt}}$, and its man-endpoint is contained in $\mathcal{F}$.*

Proof Observe that since f does not block M^{opt}, there must exists a connected component Q of $M_s \triangle M^{\text{opt}}$ not in $\mathcal{F}$ such that f does not block $M_s \triangle (\mathcal{F} \cup Q)$. Clearly, Q is a masculine but not feminine augmenting path by the definition of $\mathcal{F}$. Lemma 3 then

implies that f is a dependent edge relying on Q, and Lemma 4 yields also that f has type A and Q is masculine and not feminine. The remaining statements follow from the definition of a type A edge. □

Lemma 8 *If f is a volatile edge incident to some non-feminine path in $M_s \triangle M^{\text{opt}}$, then f is not blocking in M^{opt}.*

Proof Recall that M^{opt} is an optimal solution for which $M_s \triangle M^{\text{opt}}$ has as few edges as possible, and there does not exist another optimal solution M' such that (i) M' has the same number of common edges with M_s as M^{opt}, and (ii) for each man m, either $M'(m) = M^{\text{opt}}(m)$ or m prefers $M'(m)$ to $M^{\text{opt}}(m)$.

Suppose for the sake of contradiction, that f is a volatile edge incident to a masculine path P_x^{opt} in $M_s \triangle M^{\text{opt}}$ for some $x \in \mathcal{M}_0^\star$, and f blocks M^{opt}. Let m and w be the man and woman connected by f.

Since m is the second choice of w and f is blocking in M^{opt}, we know that M^{opt} does not cover w. However, any woman on a masculine augmenting path is matched by M^{opt} with a man adjacent to her on this path, so w does not lie on P_x^{opt}. Therefore, m lies on P_x^{opt}.

Define P as the subpath of P_x^{opt} from x to m plus the edge f. Consider now $M_f = M^{\text{opt}} \triangle (P_x^{\text{opt}} \triangle P)$; M_f is a matching because $w \in \mathcal{W}_0$. Clearly, M_f is feasible, since P_x^{opt} is not a neutral path.

Claim 4 *If f' is an edge that blocks M_f but does not block M^{opt}, then P_x^{opt} ends at a woman y, and f' is a volatile edge incident to y.*

Proof of Claim 4 Let w' and m' be the woman and man connected by f'. The only persons that are matched differently in M_f than in M^{opt} are those in $(V(P_x^{\text{opt}}) \setminus V(P)) \cup \{m, w\}$; recall that $f = \{m, w\}$ with $w \in \mathcal{W}_0$. Since both m and w are better off in M_f than in M^{opt}, and f' blocks M_f but not M^{opt}, we get that one endpoint of f' must be in $V(P_x^{\text{opt}}) \setminus V(P)$.

Suppose first that $m' \in V(P_x^{\text{opt}}) \setminus V(P)$. Then $M_f(m') = M_s(m')$ and since f' blocks M_f, we know that m' prefers w' to $M_s(m')$. But M_s is stable, so w' prefers $M_s(w')$ to m', and in particular is matched by M_s. As f' blocks M_f, we get that $M_f(w') \neq M_s(w')$. Hence, as $w' \neq w$ (because $w \in \mathcal{W}_0$ but $w' \notin \mathcal{W}_0$), it must be the case that $M_f(w') = M^{\text{opt}}(w')$ and w' is on a connected component of $\mathcal{K}_{-x}$, denoting the union of all connected components of $M_s \triangle M^{\text{opt}}$ other than P_x^{opt}. This implies that f' is blocking in $M_s \triangle \mathcal{K}_{-x}$ as well. Since f' is not blocking in $M^{\text{opt}} = M_s \triangle (\mathcal{K}_{-x} \cup P_x^{\text{opt}})$, Lemma 3 implies that f' must be a dependent edge relying on P_x^{opt}. Note that since $m' \in V(P_x^{\text{opt}})$, f' cannot be of type A (by the definition of a type A edge), so it must be of type B. However, Lemma 4 tells us that a type B edge relies either on a cycle or on a feminine path, but as P_x^{opt} is neither of the two, we arrive at a contradiction.

Thus, we get $w' \in V(P_x^{\text{opt}}) \setminus V(P)$. Observe that $m' = m$ is not possible, because it would imply $f' = \{m, M^{\text{opt}}(m)\}$ contradicting the assumption that f' blocks M_f (note that m prefers w to $M^{\text{opt}}(m)$ because f blocks M^{opt}). From this we obtain that f'

cannot be contained in P_x^{opt}: we have $f' \notin M_f$ by definition, and $f' \in M^{\text{opt}}$ is also not possible, since in that case $m' \notin V(P_x^{\text{opt}}) \setminus V(P)$ would imply $m' = m$. Thus, by $\Delta_{\mathcal{W}} \leq 2$ we know that $w' = y$, and since f' blocks M_f but not M^{opt}, we also get that m' must be the second choice of w'. Hence f' is volatile, proving our claim. □

By Claim 4, there can be at most one edge that is blocking in M_f but not in M^{opt}. Since f is an edge that blocks M^{opt} but not M_f, we get that M_f is an optimal matching. Moreover, either M_f has more common edges with M_s than M^{opt}, or m is the last man on P_x^{opt}, in which case every man has the same partner in M_f as in M^{opt} except for m who is better off in M_f than in M^{opt}. This contradicts our choice of M^{opt}. □

For any volatile edge f we can decide in linear time if there exists a masculine augmenting path disjoint from $\mathcal{F}$ that contains the woman endpoint w of f, but not f itself. Indeed, we can build such a path starting from w by taking edges not in M_s and edges in M_s in an alternating manner; by $\Delta_{\mathcal{W}} \leq 2$ and since we need to start with an edge different from f, we always have at most one possibility to pick our next edge. This process may or may not result in a masculine augmenting path, but if it does, then the path is unique, proving the following claim.

Proposition 5 *Let f be a volatile edge. If there exists a masculine augmenting path disjoint from $\mathcal{F}$ that contains the woman endpoint w of f but not f itself, then this path is unique; we denote it by Q_f^{elim}.*

Let f be a volatile edge that is blocking in $M_s \triangle \mathcal{F}$. We say that a set $\mathcal{P}_f$ of masculine augmenting paths *eliminates* f if (i) Q_f^{elim} exists and $Q_f^{\text{elim}} \in \mathcal{P}_f$, and (ii) for any path $Q \in \mathcal{P}_f$, if there is a volatile blocking edge f' in $M_s \triangle Q$, then $Q_{f'}^{\text{elim}}$ exists and is contained in $\mathcal{P}_f$. We refer to the (inclusion-wise) minimal set of masculine paths eliminating f as the *elimination paths* for f, and denote it by $\mathcal{P}_f^{\text{elim}}$. Further, we refer to the starting vertices of these paths as the *elimination set* for f. The next lemma illuminates the role of elimination paths.

Lemma 9 *If f is an edge blocking in $M_s \triangle \mathcal{F}$ but not blocking in M^{opt}, then $M_s \triangle M^{\text{opt}}$ contains all paths in $\mathcal{P}_f^{\text{elim}}$.*

Proof By Lemma 7, the woman endpoint w of f lies on a masculine path. By Proposition 5, this path can only be Q_f^{elim}, and thus Q_f^{elim} must be contained in $M_s \triangle M^{\text{opt}}$. For an inductive reasoning, assume that Q is a masculine path in $\mathcal{P}_f^{\text{elim}}$ that is contained in $M_s \triangle M^{\text{opt}}$ and f' is a volatile edge that is blocking in $M_s \triangle Q$. We claim that $Q_{f'}^{\text{elim}}$ is contained in $M_s \triangle M^{\text{opt}}$, which by induction proves our lemma.

To see our claim, note that by Lemma 8, f' cannot be blocking in M^{opt}. Since the woman endpoint of f' is unmatched by M_s, it can only be blocking in $M_s \triangle Q$ if its man endpoint lies on Q. Thus, since f' is not blocking in M^{opt} (but is blocking in $M_s \triangle Q$), it must be the case that its woman endpoint lies on a masculine path in $M_s \triangle M^{\text{opt}}$. Since such a path can only be $Q_{f'}^{\text{elim}}$ by Proposition 5, we get that $Q_{f'}^{\text{elim}} \in \mathcal{P}_f^{\text{elim}}$ and the lemma follows. □

Step 7: Guessing relevant elimination sets in M^{opt} We call an edge *relevant* in M^{opt}, if it is a volatile edge blocking $M_s \triangle \mathcal{F}$, but it does not block M^{opt}. By Lemma 9, if f

is a relevant edge in M^{opt}, then $M_s \triangle M^{\text{opt}}$ must contain all paths in $\mathcal{P}_f^{\text{elim}}$. Since there may be several volatile edges blocking in $M_s \triangle \mathcal{F}$, we cannot determine the relevant ones among them by simply guessing them. Instead, we only guess the elimination sets for all relevant edges. Clearly, these sets must be pairwise disjoint subsets of $\mathcal{M}_0^\star$, so guessing them results in at most $(|\mathcal{M}_0^\star| + 1)^{|\mathcal{M}_0^\star|}$ possibilities. Let us denote by $R_1, \ldots, R_\ell$ the guessed elimination sets.

Step 8: Computing cheapest elimination paths For each set $R_i \subseteq \mathcal{M}_0^\star$ with $i = 1, \ldots, \ell$ that, according to our guesses made in Step 7, forms the elimination set for a volatile edge relevant in M^{opt}, we determine *some* volatile edge f incident to $\mathcal{F}$ that is blocking in $M_s \triangle \mathcal{F}$ and whose elimination set is exactly R_i. Namely, we pick an edge f among all such edges in a way that the number of blocking pairs in $M_s \triangle (\mathcal{F} \cup \mathcal{P}_f^{\text{elim}})$ is as small as possible. Let f_i be the volatile edge chosen this way, and let $\mathcal{P}^{\text{elim}} = \bigcup_{1 \le i \le \ell} \mathcal{P}_{f_i}^{\text{elim}}$.

Step 9: Computing remaining masculine paths We define $\mathcal{M}_r = \mathcal{M}_0^\star \setminus (R_1 \cup \cdots \cup R_\ell)$ as the set of distinguished men that are neither covered by M_s nor contained in any of the sets $R_1, \ldots, R_\ell$. For each such m we are going to compute an augmenting path P_m disjoint from $\mathcal{F}$ such that the number of edges that block $M_s \triangle (\mathcal{F} \cup \bigcup_{m \in \mathcal{M}_r} P_r)$ but not $M_s \triangle \mathcal{F}$ is minimized. We set $\mathcal{P}_r = \{P_m \mid m \in \mathcal{M}_r\}$.

To compute $\mathcal{P}_r$, we use the algorithm of Theorem 8 with some modifications.

Step 9.1 We compute all augmenting paths that start at a man $m \in \mathcal{M}_r$ and are disjoint from $\mathcal{F}$. For each such augmenting path P we define a set $C(P)$ containing those edges that block $M_s \triangle (\mathcal{F} \cup P)$ but not $M_s \triangle \mathcal{F}$, and are either non-volatile or have a woman endpoint in $\mathcal{F}$; we define the *contributing cost* of P as $|C(P)|$.

Step 9.2 We construct the auxiliary graph $G_{\text{path}} = (U, V; E)$ as follows: we set $U = \mathcal{M}_r$ and $V = \mathcal{W}_0 \cup \{m' \mid m \in \mathcal{M}_r\}$. The edge set E contains an edge $\{m, m'\}$ for each $m \in U$, as well as an edge $\{m, w\}$ whenever $m \in U$, $w \in \mathcal{W}_0$ and there exists an augmenting path disjoint from $\mathcal{F}$ with endpoints m and w. We define the weight of an edge $\{m, m'\}$ as the minimum contributing cost $c_m^{\min}$ of any augmenting path starting at m and *not* ending in $\mathcal{W}_0$, and we define the weight of an edge $\{m, w\}$ with $m \in U$ and $w \in \mathcal{W}_0$ as the minimum contributing cost of any augmenting path with endpoints m and w, disjoint from $\mathcal{F}$.

Step 9.3 We compute a minimum weight matching M_P in G_{path} covering U the same way as in the algorithm of Theorem 8; let $\mathcal{P}$ denote the set of augmenting paths corresponding to the edges of the matching M_P. Note that the paths in $\mathcal{P}$ are pairwise disjoint, by the construction of G_{path} and because M_P is a matching in G_{path}.

Step 9.4 We eliminate all volatile edges that block $M_s \triangle (\mathcal{F} \cup \mathcal{P})$ but not $M_s \triangle \mathcal{F}$ and are not incident to $\mathcal{F}$. We modify $\mathcal{P}$ iteratively. We start by setting $\mathcal{P}_{\text{act}} = \mathcal{P}$. At each iteration we modify $\mathcal{P}_{\text{act}}$ as follows. We check whether there exists a volatile edge $\{m^*, w^*\}$ with $w^* \in \mathcal{W}_0$ that is not incident to $\mathcal{F}$, and blocks $M_s \triangle (\mathcal{F} \cup \mathcal{P}_{\text{act}})$ but not $M_s \triangle \mathcal{F}$. If yes, then notice that w^* is not matched in $M_s \triangle (\mathcal{F} \cup \mathcal{P}_{\text{act}})$, because m^* is the second choice for w^*. Let P be the path of $\mathcal{P}_{\text{act}}$ containing w^*. We modify $\mathcal{P}_{\text{act}}$ by truncating P to its subpath between its starting vertex and m^*, and appending to it the edge $\{m^*, w^*\}$. This way, $\{m^*, w^*\}$ becomes an edge of the matching $M_s \triangle (\mathcal{F} \cup \mathcal{P}_{\text{act}})$. The iteration stops when there is no volatile edge disjoint from $\mathcal{F}$ blocking $M_s \triangle (\mathcal{F} \cup \mathcal{P}_{\text{act}})$ but not $M_s \triangle \mathcal{F}$. Note that once a volatile edge ceases to be blocking in $M_s \triangle (\mathcal{F} \cup$

$\mathcal{P}_{\text{act}}$), it cannot become blocking again during this process, so the algorithm performs at most $|\mathcal{W}_0|$ iterations. For each $m \in \mathcal{M}_r$, let P_m denote the augmenting path in $\mathcal{P}_{\text{act}}$ covering m at the end of this step; we set $\mathcal{P}_r := \{P_m \mid m \in \mathcal{M}_r\}$.
Step 9.5 Finally, we output the matching $M^{\text{out}} = M_s \triangle (\mathcal{F} \cup \mathcal{P}^{\text{elim}} \cup \mathcal{P}_r)$.

It is straightforward to verify that the number of guesses made are bounded by a function of $|\mathcal{W}_0^\star| + |\mathcal{M}_0^\star|$, and all computations in a branch can be performed in time polynomial in the size $|I|$ of the instance, yielding a fixed-parameter algorithm with parameter $|\mathcal{W}_0^\star| + |\mathcal{M}_0^\star|$. It remains to prove the correctness of the proposed algorithm.

To this end, we first prove a simple observation.

Proposition 6 *Suppose $\Delta_{\mathcal{W}} \leq 2$. Let M be a matching and $\mathcal{K}_M$ the set of connected components in $M_s \triangle M$. Let also $\mathcal{H}_1$ and $\mathcal{H}_2$ be two disjoint subsets of $\mathcal{K}_M$. Then there is no edge that blocks both $M_s \triangle \mathcal{H}_1$ and $M_s \triangle \mathcal{H}_2$*

Proof Suppose that e blocks both $M_1 := M_s \triangle \mathcal{H}_1$ and $M_2 := M_s \triangle \mathcal{H}_2$; by the stability of M_s, e connects a person a_1 in $\mathcal{H}_1$ with a person a_2 in $\mathcal{H}_2$. Since e blocks M_1, we know that a_2 prefers a_1 to its situation in M_s. Since e blocks M_2, we also get that a_1 prefers a_2 to its situation in M_s. This contradicts the stability of M_s. □

Next, let us prove that Step 9 works as promised.

Lemma 10 *For each $m \in \mathcal{M}_r$ let P'_m be an augmenting path disjoint from $\mathcal{F}$ and starting at m, and let $\mathcal{P}'_r = \{P'_m \mid m \in \mathcal{M}_r\}$. Then the number of edges that block $M_s \triangle (\mathcal{F} \cup \mathcal{P}_r)$ but not $M_s \triangle \mathcal{F}$ is at most the number of edges that block $M_s \triangle (\mathcal{F} \cup \mathcal{P}'_r)$ but not $M_s \triangle \mathcal{F}$.*

Proof Let c' be the number of edges that block $M_s \triangle (\mathcal{F} \cup \mathcal{P}'_r)$ but not $M_s \triangle \mathcal{F}$. Then Claim 5 implies that there can be at most c' edges in $\bigcup_{m \in \mathcal{M}_r} C(P'_m)$.

Claim 5 *If e is an edge such that $e \in C(P'_m)$ for some $m \in \mathcal{M}_r$, then e blocks $M_s \triangle (\mathcal{F} \cup \mathcal{P}'_r)$ but not $M_s \triangle \mathcal{F}$.*

Proof Let m^* and w^* be the man and woman connected by e. By $e \in C(P'_m)$ it does not block $M_s \triangle \mathcal{F}$, so we only need to show that e blocks $M_s \triangle (\mathcal{F} \cup \mathcal{P}'_r)$. Suppose for contradiction that e does not block $M_s \triangle (\mathcal{F} \cup \mathcal{P}'_r)$. Then it must clearly be adjacent to some path P'_x, $x \in \mathcal{M}_r$. By $\Delta_{\mathcal{W}} \leq 2$ we know that w^* is the endpoint of one of these paths, $w^* \in \mathcal{W}_0$, and since e blocks $M_s \triangle (\mathcal{F} \cup P'_m)$ but not $M_s \triangle (\mathcal{F} \cup \mathcal{P}'_r)$ we also get that m^* must be the less preferred choice of w^*; this yields that e is volatile. However, as e is not incident to $\mathcal{F}$ (because it connects P'_m and P'_x), this contradicts $e \in C(P'_m)$. □

Let m_1 and m_2 be two distinct men in $\mathcal{M}_r$; we will show $C(P'_{m_1}) \cap C(P'_{m_2}) = \emptyset$. Assuming otherwise, let e be an edge in $C(P'_{m_1}) \cap C(P'_{m_2})$. Then e does not block $M_s \triangle \mathcal{F}$ but blocks both $M_s \triangle (\mathcal{F} \cup P'_{m_1})$ and $M_s \triangle (\mathcal{F} \cup P'_{m_2})$. Hence e connects P'_{m_1} with P'_{m_1}, and thus blocks $M_s \triangle P'_{m_1}$ and $M_s \triangle P'_{m_2}$ as well, contradicting Proposition 6. Thus,

$$\left| \bigcup_{m \in \mathcal{M}_r} C(P'_m) \right| = \sum_{m \in \mathcal{M}_r} |C(P'_m)| \leq c'.$$

Then, by the definition of edge-weights in G_{path}, we get that there exists a matching M' in G_{path} with weight at most $\sum_{m \in \mathcal{M}_r} |C(P'_m)| \leq c'$ that covers $U = \mathcal{M}_r$. Thus, the matching computed by the algorithm in Step 9.3 has weight at most c' as well, yielding that the set $\mathcal{P}_r$ of augmenting path corresponding to this matching is such that $\sum_{m \in \mathcal{M}_r} |C(P_m)| \leq c'$, that is, their total contribution cost is at most c'. From this, Claim 6 below implies that there can be at most c' edges that block $M_s \triangle (\mathcal{F} \cup \mathcal{P}_r)$ but not $M_s \triangle \mathcal{F}$.

Claim 6 *If an edge e blocks $M_s \triangle (\mathcal{F} \cup \mathcal{P}_r)$ but not $M_s \triangle \mathcal{F}$, then $e \in C(P_m)$ for some $m \in \mathcal{M}_r$.*

Proof Let m^* and w^* be the man and woman connected by e. Clearly, e is adjacent to some path in $\mathcal{P}_r$.

First, suppose that e is adjacent to only one path P_m where $m \in \mathcal{M}_r$. Then e must block $M_s \triangle (\mathcal{F} \cup P_m)$ as well. In that case, to prove $e \in C(P_m)$ we only need to show that either w^* is in $\mathcal{F}$ or is e is non-volatile. Suppose for contradiction that e is volatile and w^* is not in $\mathcal{F}$. However, as a result of Step 9.4, there can be no volatile edge that blocks $M_s \triangle (\mathcal{F} \cup \mathcal{P}_r)$, unless it is incident to $\mathcal{F}$. Hence, we must have that m^* is in $\mathcal{F}$, and since e blocks $M_s \triangle (\mathcal{F} \cup P_m)$, we get that m^* prefers w^* to its partner in $M_s \triangle (\mathcal{F} \cup P_m)$, which is the same partner he has in $M_s \triangle \mathcal{F}$. However, this yields that e blocks $M_s \triangle \mathcal{F}$ as well: since $w^* \in \mathcal{W}_0$ and w^* is not in $\mathcal{F}$, she is unmatched in $M_s \triangle \mathcal{F}$. Hence, we arrive at a contradiction, proving $e \in C(P_m)$.

Second, suppose that e connects two paths P_{m_1} and P_{m_2} where m_1 and m_2 are distinct men in $\mathcal{M}_r$. By $\Delta_{\mathcal{W}} \leq 2$, we know that w^* is an endpoint of one of these paths; without loss of generality, we may assume that w^* is an endpoint of P_{m_1}. Since P_{m_1} is a masculine augmenting path, we get $w^* \in \mathcal{W}_0$. Now, we know that e blocks $M_s \triangle (\mathcal{F} \cup P_{m_1} \cup P_{m_2})$, so it must also block $M_s \triangle (\mathcal{F} \cup P_{m_2})$, because w^* is unmatched in $M_s \triangle (\mathcal{F} \cup P_{m_2})$ and m^* has the same partner in $M_s \triangle (\mathcal{F} \cup P_{m_2})$ as in $M_s \triangle (\mathcal{F} \cup P_{m_1} \cup P_{m_2})$. Again, e cannot be volatile as a result of Step 9.4, yielding that e is contained in $C(P_{m_2})$. □

This proves the optimality of $\mathcal{P}_r$ as stated by the lemma. □

Next we state the following useful lemma.

Lemma 11 *Let f be a volatile edge blocking in $M_s \triangle \mathcal{F}$, Q an elimination path in $\mathcal{P}_f^{\text{elim}}$, and Q' a non-feminine augmenting path not contained in $\mathcal{P}_f^{\text{elim}}$, disjoint from $\mathcal{F}$. Let e be an edge. Then e blocks $M_s \triangle (Q \cup Q')$ if and only if it blocks $M_s \triangle Q$ or $M_s \triangle Q'$.*

Proof First observe that the statement of the lemma is obviously true for any edge e that is incident to at most one of Q and Q'. So suppose that e connects Q with Q'. Since both Q and Q' are masculine augmenting paths and $\Delta_{\mathcal{W}} \leq 2$, by Proposition 3 we know that e connects the woman endpoint w of one of these paths with a man m on the other path; moreover, w must be unmatched by M_s as it is the endpoint of a masculine augmenting path. Let Q_w be the path (either Q or Q') that contains w, and let Q_m be the one containing m. By $w \in \mathcal{W}_0$ and the stability of M_s, e cannot be blocking in $M_s \triangle Q_w$.

Now, if e blocks $M_s \triangle (Q_m \cup Q_w)$, then it blocks $M_s \triangle Q_m$ as well, since m prefers w to its partner in $M_s \triangle Q_m$ and w is unmatched in $M_s \triangle Q_m$.

It remains to show that if e blocks $M_s \triangle Q_m$, then it also blocks $M_s \triangle (Q_w \cup Q_m)$. Supposing otherwise, it must be the case that (i) m prefers w to its partner in $M_s \triangle Q_m$, but (ii) w prefers its partner in $M_s \triangle Q_w$ to m. Consequently, e is a volatile edge.[9] We distinguish two cases.

Case (A) $Q_w = Q'$. Then Q' must be the unique masculine augmenting path containing the woman endpoint of e but not e itself, that is, $Q' = Q_e^{\text{elim}}$. Hence, since e is a volatile edge blocking $M_s \triangle Q$ and $Q \in \mathcal{P}_f^{\text{elim}}$, by the definition of elimination paths we get $Q' \in \mathcal{P}_f^{\text{elim}}$ as well, a contradiction.

Case (B) $Q_m = Q'$. Then Q' contains the man endpoint of e. However, again by the definition of elimination paths, as Q is a path of $\mathcal{P}_f^{\text{elim}}$, the man endpoint of e must be contained either in $\mathcal{F}$ (if $e = f$) or another path of $\mathcal{P}_f^{\text{elim}}$ that triggered the addition of Q into $\mathcal{P}_f^{\text{elim}}$; both possibilities contradict our conditions on Q'. □

Now we are ready to show that our algorithm is correct.

Proof of Theorem 11 To prove the correctness of the proposed algorithm, we first show that if all our guesses are true, then the paths and cycles in $\mathcal{F}$ are exactly the feminine paths and the cycles of $M_s \triangle M^{\text{opt}}$. The correctness of Step 2 is stated by Lemma 5. From the description of our algorithm, it should be clear that the correctness of Steps 3 and 4 follows directly from Proposition 3 and Lemma 4. Lemma 6 guarantees the correctness of Step 5, which proves that in Steps 1–5 the algorithm indeed finds all cycles and feminine paths of $M_s \triangle M^{\text{opt}}$.

Next, let us argue that M^{out} is indeed a matching. For this, apart from the correctness of Steps 1–5, we need that the masculine paths in $M_s \triangle M^{\text{out}}$ are disjoint from $\mathcal{F}$. Further, we also need that paths in $\mathcal{P}^{\text{elim}}$ are disjoint from all remaining masculine paths. To see this, observe that any path P in $\mathcal{P}^{\text{elim}}$ ends at a woman $w \in \mathcal{W}_0$ which is connected by a volatile edge (not on P) to either $\mathcal{F}$ or to another path in $\mathcal{P}^{\text{elim}}$. Hence, w cannot lie on any masculine path other than P by Proposition 5. Thus, M^{out} is a matching. Its feasibility is implied by the correctness of Steps 1–5, and the definition of augmenting paths.

It remains to argue that M^{out} admits at most as many blocking pairs as M^{opt}. First, Lemma 7 implies that all edges blocking in $M_s \triangle \mathcal{F}$ are either relevant volatile edges in M^{opt}, or they are also blocking in M^{opt}. Furthermore, if f_i^{opt} is a relevant volatile edge with elimination set R_i for some $i \in \{1, \ldots, \ell\}$, then by Lemma 9 we know that all elimination paths in $\mathcal{P}_{f_i^{\text{opt}}}^{\text{elim}}$ must be contained in $M_s \triangle M^{\text{opt}}$. As $\mathcal{F}$ is the set of feminine augmenting paths and cycles of $M_s \triangle M^{\text{opt}}$, in Step 8 the algorithm is bound to find *some* volatile edge f_i (though not necessarily f_i^{opt}) that is blocking in $M_s \triangle \mathcal{F}$ and whose elimination set is R_i. Furthermore, by our choice of f_i, there are at most as many blocking pairs in $M_s \triangle (\mathcal{F} \cup \mathcal{P}_{f_i}^{\text{elim}})$ as there are in $M_s \triangle (\mathcal{F} \cup \mathcal{P}_{f_i^{\text{opt}}}^{\text{elim}})$.

Let the *contribution* of a volatile edge e, denoted by $C(e)$, be the set of edges that block $M_s \triangle (\mathcal{F} \cup \mathcal{P}_e^{\text{elim}})$ but not $M_s \triangle \mathcal{F}$; we extend this notion to any set E of volatile edges with pairwise disjoint elimination sets by defining the contribution $C(E)$ of E

[9] This case is analogous with e being a type A edge; however, the paths Q and Q' here need not be paths of $M_s \triangle M^{\text{opt}}$.

as the set of edges that block $M_s \triangle (\mathcal{F} \cup \bigcup_{e \in E} \mathcal{P}_e^{\text{elim}})$ but not $M_s \triangle \mathcal{F}$. By our choice of f_i, we know $|C(f_i)| \leq |C(f_i^{\text{opt}})|$.

We are going to show

$$|C(\{f_i \mid 1 \leq i \leq \ell\})| \leq |C(\{f_i^{\text{opt}} \mid 1 \leq i \leq \ell\})|,$$

which implies that $M_s \triangle (\mathcal{F} \cup \bigcup_{1 \leq i \leq \ell} \mathcal{P}_{f_i}^{\text{elim}})$ has at most as many blocking pairs as $M_s \triangle (\mathcal{F} \cup \bigcup_{1 \leq i \leq \ell} \mathcal{P}_{f_i^{\text{opt}}}^{\text{elim}})$ does. To prove this, it suffices to prove the following claim.

Claim 7 *Let $F \cup \{f\}$ be a set of volatile edges with pairwise disjoint elimination sets, $F \neq \emptyset$. Then $|C(f)| + |C(F)| = |C(F \cup \{f\})|$.*

Proof We need to prove that any edge e blocks $M_s \triangle (\mathcal{F} \cup \mathcal{P}_f^{\text{elim}} \cup \bigcup_{f' \in F} \mathcal{P}_{f'}^{\text{elim}})$ but not $M_s \triangle \mathcal{F}$ if and only if it blocks exactly one of $M_s \triangle (\mathcal{F} \cup \mathcal{P}_f^{\text{elim}})$ and $M_s \triangle (\mathcal{F} \cup \bigcup_{f' \in F} \mathcal{P}_{f'}^{\text{elim}})$ but not $M_s \triangle \mathcal{F}$. So let us assume that e does not block $M_s \triangle \mathcal{F}$.

Notice that if e is incident to only one of $\mathcal{P}_f^{\text{elim}}$ and $\bigcup_{f' \in F} \mathcal{P}_{f'}^{\text{elim}}$, then our claim is immediate. So suppose that e connects a path $Q \in \mathcal{P}_f^{\text{elim}}$ with a path $Q' \in \bigcup_{f' \in F} \mathcal{P}_{f'}^{\text{elim}}$. Observe that it suffices to show that e blocks $M_s \triangle (Q \cup Q')$ if and only if it blocks exactly one of $M_s \triangle Q$ and $M_s \triangle Q'$, because e cannot be incident to any connected component of $\mathcal{F} \cup \mathcal{P}_f^{\text{elim}} \cup \bigcup_{f' \in F} \mathcal{P}_{f'}^{\text{elim}}$ other than Q and Q'. Now the claim follows directly from Lemma 11 and Proposition 6. □

It remains to consider the blocking pairs contributed by the paths $\mathcal{P}_r^{\text{opt}} := \{P_m^{\text{opt}} \mid m \in \mathcal{M}_r\}$.

Claim 8 *If e is an edge that blocks $M_s \triangle (\mathcal{F} \cup \mathcal{P}_r^{\text{opt}})$ but not $M_s \triangle \mathcal{F}$, then e blocks M^{opt} but not $M^{\text{opt}} \triangle \mathcal{P}_r^{\text{opt}}$.*

Proof Recall that by definition $M^{\text{opt}} = M_s \triangle (\mathcal{F} \cup \mathcal{P}_r^{\text{opt}} \cup \bigcup_{1 \leq i \leq \ell} \mathcal{P}_{f_i^{\text{opt}}}^{\text{elim}})$ and $M^{\text{opt}} \triangle \mathcal{P}_r^{\text{opt}} = M_s \triangle (\mathcal{F} \cup \bigcup_{1 \leq i \leq \ell} \mathcal{P}_{f_i^{\text{opt}}}^{\text{elim}})$, so if e does not have an endpoint in $\bigcup_{1 \leq i \leq \ell} \mathcal{P}_{f_i^{\text{opt}}}^{\text{elim}}$, then the claim follows immediately.

So suppose that e is incident to a path $Q \in \mathcal{P}_{f_i^{\text{opt}}}^{\text{elim}}$ for some $i \in \{1, \ldots, \ell\}$. As e blocks $M_s \triangle (\mathcal{F} \cup \mathcal{P}_r^{\text{opt}})$ but not $M_s \triangle \mathcal{F}$, the other endpoint of e must be contained in a path P_m^{opt} for some $m \in \mathcal{M}_r$. Since Q and P_m^{opt} are the only connected components of $M_s \triangle M^{\text{opt}}$ that e is adjacent to, the condition of our claim yields that e blocks $M_s \triangle P_m^{\text{opt}}$, and we need to show that e blocks $M_s \triangle (Q \cup P_m^{\text{opt}})$ but not $M_s \triangle Q$. To see this, first notice that by Proposition 6, edge e cannot block $M_s \triangle Q$, because Q and P_m^{opt} are disjoint. Second, as e blocks $M_s \triangle P_m^{\text{opt}}$, Lemma 11 implies that it must block $M_s \triangle (Q \cup P_m^{\text{opt}})$ as well, proving our claim. □

Let c_r be the number of those edges that block M^{opt} but not $M^{\text{opt}} \triangle \mathcal{P}_r^{\text{opt}}$. Then Claim 8 implies that there are at most c_r edges that block $M_s \triangle (\mathcal{F} \cup \mathcal{P}_r^{\text{opt}})$ but not $M_s \triangle \mathcal{F}$. Since Steps 9.1–9.4 calculate a set $\mathcal{P}_r = \{P_m \mid m \in \mathcal{M}_r\}$ of augmenting

paths that minimizes the number of edges that block $M_s \triangle (\mathcal{F} \cup \mathcal{P}_r)$ but not $M_s \triangle \mathcal{F}$, we know that there are at most c_r such edges. This in turn shows that there can be at most c_r edges that block M^{out} but not $M^{\text{out}} \triangle \mathcal{P}_r$, as implied by Claim 9. This proves the optimality of M^{out}.

Claim 9 *If e is an edge that blocks M^{out} but not $M^{\text{out}} \triangle \mathcal{P}_r$, then e blocks $M_s \triangle (\mathcal{F} \cup \mathcal{P}_r)$ but not $M_s \triangle \mathcal{F}$.*

Proof Since $M^{\text{out}} = M_s \triangle (\mathcal{F} \cup \mathcal{P}_r \cup \mathcal{P}^{\text{elim}})$ and $M^{\text{out}} \triangle \mathcal{P}_r = M_s \triangle (\mathcal{F} \cup \mathcal{P}^{\text{elim}})$, we have that if e does not have an endpoint in $\mathcal{P}^{\text{elim}}$, then the claim follows immediately.

So suppose that e is incident to a path $Q \in \mathcal{P}^{\text{elim}}_{f_i}$ for some $i \in \{1, \ldots, \ell\}$. As e blocks M^{out} but not $M^{\text{out}} \triangle \mathcal{P}_r$, the other endpoint of e must be contained in a path P_m for some $m \in \mathcal{M}_r$. Since Q and P_m are the only connected components of $M_s \triangle M^{\text{out}}$ that e is adjacent to, the condition of our claim yields that e blocks $M_s \triangle (Q \cup P_m)$ but not $M_s \triangle Q$, and we only need to show that e blocks $M_s \triangle P_m$, which is directly implied by Lemma 11, proving our claim. □

This finishes our proof of correctness for Theorem 11. □

As each augmenting path contains at least one edge that blocks M^{opt}, the number of blocking pairs admitted by M^{opt} is at least $(|\mathcal{W}^\star_0| + |\mathcal{M}^\star_0|)/2$. Thus, we get Corollary 3.

Corollary 3 *The special case of* SMC *where each woman finds at most two men acceptable (i.e., $\Delta_{\mathcal{W}} \leq 2$) is fixed-parameter tractable for parameter b.*

It remains to deal with the third among the parameterizations we have to consider. So let us turn our attention to the complexity of the SMC with $\Delta_{\mathcal{W}} \leq 2$ where we take $|\mathcal{W}^\star| + \Delta_{\mathcal{M}}$ as the parameter.

Theorem 12 *There is a fixed-parameter algorithm for the special case of* SMC *where each woman finds at most two men acceptable (i.e., $\Delta_{\mathcal{W}} \leq 2$), and the parameter is $|\mathcal{W}^\star| + \Delta_{\mathcal{M}}$, the number of distinguished women plus the maximum length of men's preference lists.*

It turns out that Theorem 12 can be proved by a modified version of the algorithm described above to prove Theorem 11. In fact, we start with applying Steps 1 to 6. However, we can no longer apply Step 7, because now $|\mathcal{M}^\star|$ can be unbounded, and thus guessing the elimination sets would not yield a fixed-parameter tractable algorithm. To circumvent this problem, we rely on an observation stated by Lemma 12 which shows that "almost all" volatile edges that are blocking in $M_s \triangle \mathcal{F}$ are relevant edges, and thus have to be eliminated. This allows us to efficiently guess the set of relevant volatile edges, leading us to a fixed-parameter tractable algorithm. Hence, instead of guessing only the elimination sets (that is, the man endpoints of the elimination paths) we can now guess the exact set of edges that the optimal solution eliminates, so in fact the modified algorithm can be viewed as a more simple approach.

Let E^{vol} denote the set of all volatile blocking edges in $M_s \triangle \mathcal{F}$, and let $\{m, w\} \in E^{\text{vol}}$ for some man m. We say that $\{m, w\}$ is *loose*, if

– both M_s and M^{opt} assign a partner to m,

– m prefers $M_s(m)$ to w, and w to $M^{\text{opt}}(m)$, and
– $M^{\text{opt}}(m)$ is not distinguished.

Lemma 12 *All loose edges in E^{vol} are relevant, and at most $(\Delta_{\mathcal{M}} - 1)|\mathcal{W}^\star|$ edges in E^{vol} are* not *loose.*

Proof First suppose that $\{m, w\}$ for some man m is loose. Since $\{m, w\}$ is volatile, we have that m is the less preferred man acceptable to w, and since $\{m, w\}$ is blocking in $M_s \triangle \mathcal{F}$, we know that w is unmatched in $M_s \triangle \mathcal{F}$.

For the sake of contradiction, let us suppose that $\{m, w\}$ is not relevant, meaning that it is still blocking in M^{opt}. In this case, w must be unmatched in M^{opt} too. We define a matching M' obtained from M^{opt} by modifying only the partner of m to be $M'(m) = w$ and, i.e., $M' = (M^{\text{opt}} \backslash \{\{m, M^{\text{opt}}(m)\}\}) \cup \{\{w, m\}\}$. Observe that M' is indeed a matching, because w was unmatched in M^{opt}. It is also feasible, because only $M^{\text{opt}}(m)$ becomes unmatched by this modification, and $M^{\text{opt}}(m)$ is not distinguished (by the definition of a loose edge). Furthermore, while $\{m, w\}$ is not blocking in M', the only edge that might be blocking in M' but not in M^{opt} is the edge adjacent to $M^{\text{opt}}(m)$ but not to m, if existent (because both m and w are better off in M' than in M^{opt}). Hence, M' is also an optimal matching. It should also be clear that it has the same number of common edges with M_s as M^{opt}. Note also that each man has the same partner in M' as in M^{opt} except for m who is better off in M' than in M^{opt}. This contradicts our choice of M^{opt}, proving that $\{m, w\}$ must indeed be relevant.

To prove the second part of the lemma, let us consider any edge $\{m, w\} \in E^{\text{vol}}$ with $m \in \mathcal{M}$. Again, w must be unmatched in the matching $M_s \triangle \mathcal{F}$, since otherwise $\{m, w\}$ would not block it (because m is the less preferred man in w's preference list). By contrast, $\{m, w\}$ cannot block M_s, which can be explained in three different ways, giving rise to the following three cases:

– Case A: $\{m, w\} \in M_s$,
– Case B: $\{m, w\} \notin M_s$, but $M_s(w)$ exists and is preferred by w to m, or
– Case C: $\{m, w\} \notin M_s$, but $M_s(m)$ exists and is preferred by m to w.

In Case A, by $\{m, w\} \in M_s$, the edge $\{m, w\}$ must be contained in an augmenting path of $M_s \triangle \mathcal{F}$, with w being an endpoint (since w is not matched in $M_s \triangle \mathcal{F}$). In Case B, M_s matches w to its first choice, say x, so the edge $\{w, x\}$ must be contained in an augmenting path of $M_s \triangle \mathcal{F}$, with w being an endpoint. Thus, in both cases we know that w is an endpoint of a feminine augmenting path of $\mathcal{F}$. However, Cases A and B exclude each other, because in Case A the augmenting path ending at w contains the edge connecting w to her second choice, while in Case B it connects w to her first choice. Since there are at most $|\mathcal{W}^\star|$ feminine augmenting paths, it follows that the number of such edges $\{m, w\} \in E^{\text{vol}}$ where the conditions of Cases A or B hold is at most $|\mathcal{W}^\star|$.

In Case C, let us first observe that m gets matched in $M_s \triangle \mathcal{F}$. Indeed, if this were not the case, then m would be the man endpoint of a feminine augmenting path, with the last edge contained in M_s; however, this is not possible (simply because feminine paths start with women and with edges not in M_s). Because $\mathcal{F}$ is a collection of connected components of $M_s \triangle M^{\text{opt}}$, with m contained in $\mathcal{F}$, it is clear that the partner of m in $M_s \triangle \mathcal{F}$ is the same as his partner in M^{opt}. Thus, $M^{\text{opt}}(m)$ exists, and since $\{m, w\}$

blocks $M_s \triangle \mathcal{F}$, we know that m prefers w to $M^{\text{opt}}(m)$. Therefore, either $\{m, w\}$ is a loose edge, or $M^{\text{opt}}(m)$ is distinguished; in the latter case, we say that $\{m, w\}$ *belongs* to this distinguished woman.

Now, observe that all edges belonging to some woman $z \in \mathcal{W}^\star$ are adjacent to $M^{\text{opt}}(z)$, and $M^{\text{opt}}(z)$ is adjacent to at least two edges *not* belonging to z: namely, the edges $\{M^{\text{opt}}(z), z\}$ and $\{M^{\text{opt}}(z), M_s(M^{\text{opt}}(z))\}$. Hence, at most $\Delta_{\mathcal{M}} - 2$ edges may belong to z, implying that there are at most $(\Delta_{\mathcal{M}} - 2)|\mathcal{W}^\star|$ edges in total belonging to distinguished women. Taking into account all three cases, we get that there are at most $(\Delta_{\mathcal{M}} - 1)|\mathcal{W}^\star|$ edges in E^{vol} that are not loose. □

Let us now describe our algorithm proving Theorem 12 in detail. This algorithm starts with Steps 1 to 6 (as given for Theorem 11), and then applies Steps 7⋆ and 8⋆ below. The last step of our algorithm is Step 9 from the algorithm for Theorem 11.

Step 7⋆: *Guessing the relevant edges* To determine which edges among those volatile edges that block $M_s \triangle \mathcal{F}$ are relevant, we first compute the set E^{loose} of all loose edges in E^{vol}. Notice that when checking some edge $\{m, w\} \in E^{\text{vol}}$, we already know $M^{\text{opt}}(m)$: since m lies on a path or cycle of $\mathcal{F}$, it gets matched in $M_s \triangle \mathcal{F}$ to the same woman as in M^{opt}. Therefore, checking whether $\{m, w\}$ is loose is straightforward.

After determining all loose edges (which are all relevant by Lemma 12), we guess the remaining set of relevant edges in $E^{\text{vol}} \setminus E^{\text{loose}}$. By Lemma 12, this yields at most $2^{(\Delta_{\mathcal{M}}-1)|\mathcal{W}^\star|}$ possibilities, and in the branch where our guess is correct, we obtain the set E^{rel} of all relevant volatile edges blocking in $M_s \triangle \mathcal{F}$.

Step 8⋆: *Computing cheapest elimination path* This step is a simplification of Step 8 of the algorithm for Theorem 11. By Lemma 9, for each $f \in E^{\text{rel}}$ the elimination paths in $\mathcal{P}_f^{\text{elim}}$ (determined in Step 6) are all contained in $M_s \triangle M^{\text{opt}}$; let $\mathcal{P}^{\text{elim}} = \cup_{f \in E^{\text{rel}}} \mathcal{P}_f^{\text{elim}}$.

Step 9⋆: *Computing remaining masculine paths* Let $\mathcal{M}_r$ denote the set of distinguished men in $\mathcal{M}_0^\star$ not covered by any of the elimination paths in $\mathcal{P}^{\text{elim}}$. For each $m \in \mathcal{M}_r$ we compute an augmenting path P_m disjoint from $\mathcal{F}$ such that the number of edges that block $M_s \triangle (\mathcal{F} \cup \mathcal{P}_r)$ but not $M_s \triangle \mathcal{F}$ is minimized, where $\mathcal{P}_r$ is the union of all paths P_m, $m \in M_r$. We accomplish this with the exact same method as in Step 9 of the algorithm for Theorem 11. Finally, we output the matching $M^{\text{out}} = M_s \triangle (\mathcal{F} \cup \mathcal{P}^{\text{elim}} \cup \mathcal{P}_r)$.

Note that the number of guesses made in Steps 1 to 6 are bounded by a function of $|\mathcal{W}_0^\star|$, and the guesses made in Step 7⋆ result in at most $2^{(\Delta_{\mathcal{M}}-1)|\mathcal{W}^\star|}$ possibilities. Since all computations in a branch can be performed in time polynomial in the size $|I|$ of the instance, we obtain a fixed-parameter algorithm with parameter $|\mathcal{W}^\star| + \Delta_{\mathcal{M}}$. It remains to prove its correctness.

Proof of Theorem 12 The proof is a straightforward adaptation of our proof for Theorem 11. The only difference is that instead of finding *some* elimination paths that eliminate the necessary number of relevant volatile edges blocking in $M_s \triangle \mathcal{F}$, our algorithm now directly finds all of the relevant edges among those volatile edges that block $M_s \triangle \mathcal{F}$, together with the corresponding elimination paths; this is a consequence of Lemma 12 and the definition of Step 7⋆. As Step 8⋆ can also be viewed as a simplification of Step 8 as defined for the algorithm for Theorem 11, it can be verified that the exact same arguments that prove the correctness of Step 8 in our algorithm for

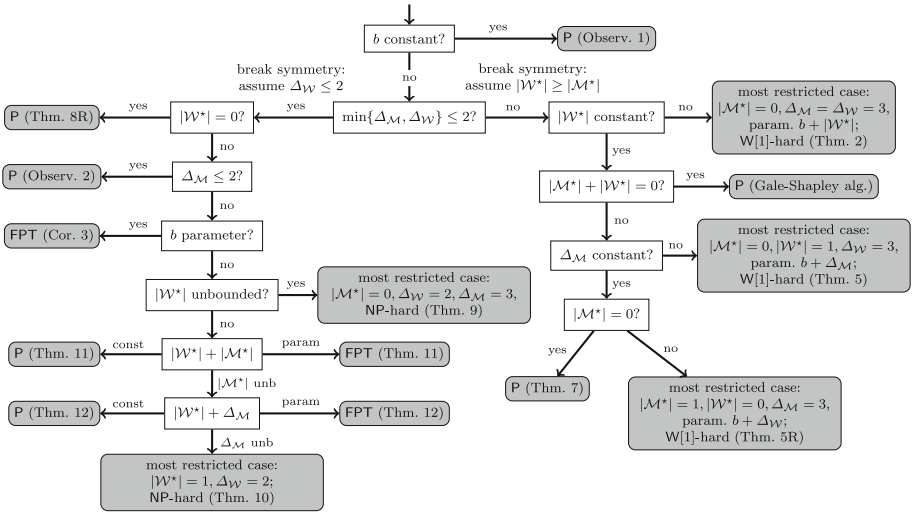

Fig. 10 Decision diagram for determining the complexity of SMC. We remark that if a certain restriction of SMC where one of the values $v \in S$ is a constant k proves to be NP-hard or W[1]-hard with some parameterization, then it is easy to see that the same hardness result also holds for the case where $v \geq k$ (and all other assumptions are the same). We refer to the "reflection" of a result by adding the postfix 'R' to its name (so Theorem xR denotes the reflection of Theorem x); here by reflection we mean the statement obtained by switching the roles of men and women

Theorem 11 also imply the correctness of Step 8⋆ for our algorithm for Theorem 12 in a straightforward manner. □

7 Discussion

We provided a systematic study of the computational complexity of Stable Marriage with Covering Constraints. Our main result is a complete computational complexity trichotomy into polynomial-time solvable cases, NP-hard and fixed-parameter tractable cases, and NP-hard and W[1]-hard cases, for all possible combinations of five natural parameters:

- $|\mathcal{M}^\star|$: the number of distinguished men,
- $|\mathcal{W}^\star|$: the number of distinguished women,
- $\Delta_{\mathcal{M}}$: the maximum length of preference lists for men,
- $\Delta_{\mathcal{W}}$: the maximum length of preference lists for women, and
- b: the number of blocking pairs allowed.

As a special case, we solved a problem by Hamada et al. [24].

Figure 10 provides a decision diagram showing that our results indeed fully determine the computational complexity of SMC with respect to the set $S = \{b, |\mathcal{W}^\star|, |\mathcal{M}^\star|, \Delta_{\mathcal{M}}, \Delta_{\mathcal{W}}\}$ of possible parameters. Going through this decision diagram should convince the reader that any parameterized restriction of SMC with respect to the set S is classified as either polynomial-time solvable (P) or NP-hard,

and in the latter case, either fixed-parameter tractable (FPT), or W[1]-hard with the given parameterization (if any). In particular, when we provide parameterized results, this means that the parameterized restriction of SMC in question is NP-hard without parameterization.

Given the strong polynomial-time inapproximability bounds, as well as the parameterized intractability results of this paper, we pose as an open question whether *fixed-parameter approximation algorithms* can beat either of these obstacles for solving SMC.

Another challenge for future research is to investigate possible adaptations of the proposed algorithms to the Hospitals/Residents model (note that, naturally, all our hardness results for SMC-1 apply to the HRLQ problem), or to a setting where ties are allowed in the preference lists.

Acknowledgements Open Access funding provided by Projekt DEAL. We thank the anonymous reviewers for their remarks which greatly helped to improve the paper; we are grateful for their patience and their rigorous reading of earlier drafts.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

References

1. Arulselvan, A., Cseh, Á., Groß, M., Manlove, D.F., Matuschke, J.: Matchings with lower quotas: algorithms and complexity. Algorithmica **8**(1), 185–208 (2018)
2. Aziz, H., Schlotter, I., Walsh, T.: Control of fair division. In: IJCAI 2016, pp. 67–73. IJCAI/AAAI Press (2016)
3. Aziz, H., Seedig, H.G., von Wedel, J.K.: On the susceptibility of the deferred acceptance algorithm. In: AAMAS 2015, pp. 939–947 (2015)
4. Biró, P., Fleiner, T., Irving, R.W., Manlove, D.F.: The college admissions problem with lower and common quotas. Theor. Comput. Sci. **411**(34–36), 3136–3153 (2010)
5. Biró, P., Irving, R.W., Schlotter, I.: Stable matching with couples: an empirical study. ACM J. Exp. Algorithmics **16**, 1.2:1.1–1.2:1.27 (2011)
6. Biró, P., Manlove, D.F., McDermid, E.J.: "Almost stable" matchings in the roommates problem with bounded preference lists. Theor. Comput. Sci. **432**, 10–20 (2012)
7. Brandt, F., Conitzer, V., Endriss, U., Lang, J., Procaccia, A.D.: Handbook of Computational Social Choice. Cambridge University Press, New York (2016)
8. Cechlárová, K., Fleiner, T.: Pareto optimal matchings with lower quotas. Math. Soc. Sci. **88**(C), 3–10 (2017)
9. Chen, J., Huang, X., Kanj, I.A., Xia, G.: Linear FPT reductions and computational lower bounds. In: STOC 2004, pp. 212–221 (2004)
10. Cseh, Á., Manlove, D.F.: Stable marriage and roommates problems with restricted edges: complexity and approximability. Discrete Optim. **20**, 62–89 (2016)
11. Cygan, M., Fomin, F.V., Kowalik, Ł., Lokshtanov, D., Marx, D., Pilipczuk, M., Pilipczuk, M., Saurabh, S.: Parameterized Algorithms. Springer, Cham (2015)
12. Dias, V.M., da Fonseca, G.D., de Figueiredo, C.M., Szwarcfiter, J.L.: The stable marriage problem with restricted pairs. Theor. Comput. Sci. **306**(1–3), 391–405 (2003)

13. Ehlers, L., Hafalir, I.E., Yenmez, M.B., Yildirim, M.A.: School choice with controlled choice constraints: hard bounds versus soft bounds. J. Econ. Theory **153**, 648–683 (2014)
14. Fellows, M.R., Hermelin, D., Rosamond, F., Vialette, S.: On the parameterized complexity of multiple-interval graph problems. Theor. Comput. Sci. **410**(1), 53–61 (2009)
15. Fleiner, T., Irving, R.W., Manlove, D.F.: Efficient algorithms for generalised stable marriage and roommates problems. Theor. Comput. Sci. **381**, 162–176 (2007)
16. Fleiner, T., Kamiyama, N.: A matroid approach to stable matchings with lower quotas. Math. Oper. Res. **41**(2), 734–744 (2016)
17. Fragiadakis, D., Iwasaki, A., Troyan, P., Ueda, S., Yokoo, M.: Strategyproof matching with minimum quotas. ACM Trans. Econ. Comput. **4**(1), 6:1–6:40 (2016)
18. Fragiadakis, D., Troyan, P.: Improving matching under hard distributional constraints. Theor. Econ. **12**, 863–908 (2017)
19. Gale, D., Shapley, L.S.: College admissions and the stability of marriage. Am. Math. Mon. **69**(1), 9–15 (1962)
20. Gale, D., Sotomayor, M.: Some remarks on the stable matching problem. Discrete Appl. Math. **11**(3), 223–232 (1985)
21. Garey, M., Johnson, D.S.: Computers and Intractability: A Guide to the Theory of NP-Completeness. W. H. Freeman and Co., New York (1979)
22. Goto, M., Iwasaki, A., Kawasaki, Y., Kurata, R., Yasuda, Y., Yokoo, M.: Strategyproof matching with regional minimum and maximum quotas. Artif. Intell. **235**, 40–57 (2016)
23. Gusfield, D., Irving, R.W.: The Stable Marriage Problem: Structure and Algorithms. Foundations of Computing Series. MIT Press, Cambridge (1989)
24. Hamada, K., Iwama, K., Miyazaki, S.: The hospitals/residents problem with lower quotas. Algorithmica **74**(1), 440–465 (2016)
25. Hopcroft, J.E., Karp, R.M.: An $n^{5/2}$ algorithm for maximum matchings in bipartite graphs. SIAM J. Comput. **2**(4), 225–231 (1973)
26. Huang, C.C.: Classified stable matching. In: SODA 2010, pp. 1235–1253 (2010)
27. Immorlica, N., Mahdian, M.: Marriage, honesty, and stability. In: SODA 2005, pp. 53–62 (2005)
28. Impagliazzo, R., Paturi, R., Zane, F.: Which problems have strongly exponential complexity? J. Comput. Syst. Sci. **63**(4), 512–530 (2001)
29. Irving, R.W., Manlove, D., O'Malley, G.: Stable marriage with ties and bounded length preference lists. J. Discrete Algorithms **7**, 213–219 (2009)
30. Kamiyama, N.: A note on the serial dictatorship with project closures. Oper. Res. Lett. **41**(5), 559–561 (2013)
31. Knuth, D.E.: Mariages stables et leurs relations avec d'autres problèmes combinatoires. Les Presses de l'Université de Montréal, Montreal (1976)
32. Kojima, F., Pathak, P.A., Roth, A.E.: Matching with couples: stability and incentives in large markets. Q. J. Econ. **128**(4), 1585–1632 (2013)
33. Kominers, S.D., Sönmez, T.: Matching with slot-specific priorities: theory. Theor. Econ. **11**(2), 683–710 (2016)
34. Kuhn, H.W.: The Hungarian method for the assignment problem. Nav. Res. Logist. Q. **2**, 83–97 (1955)
35. Lovász, L., Plummer, M.: Matching Theory. AMS Chelsea Publishing Series. North-Holland, Amsterdam (2009)
36. Manlove, D.F.: Algorithmics of Matching Under Preferences. Series on Theoretical Computer Science, vol. 2. World Scientific Publishing, Singapore (2013)
37. Marx, D., Schlotter, I.: Parameterized complexity and local search approaches for the stable marriage problem with ties. Algorithmica **58**(1), 170–187 (2010)
38. Marx, D., Schlotter, I.: Stable assignment with couples: parameterized complexity and local search. Discrete Optim. **8**, 25–40 (2011)
39. Mnich, M., Schlotter, I.: Stable marriage with covering constraints: a complete computational trichotomy. In: SAGT 2017. Lecture Notes Computer Science, vol. 10504, pp. 320–332 (2017)
40. Monte, D., Tumennasan, N.: Matching with quorums. Econ. Lett. **120**(1), 14–17 (2013)
41. Roth, A.E., Peranson, E.: The redesign of the matching market for American physicians: some engineering aspects of economic design. Am. Econ. Rev. **89**, 748–780 (1999)
42. Veskioja, T.: Stable marriage problem and college admission. Ph.D. thesis, Department of Informatics, Tallinn University of Technology (2005)

43. Westkamp, A.: An analysis of the German university admissions system. Econ. Theory **53**(3), 561–589 (2013)
44. Yokoi, Y.: Envy-free matchings with lower quotas. In: ISAAC 2017, Leibniz International Proceedings in Informatics, vol. 92, pp. 67:1–67:12 (2017)
45. Yokoi, Y.: A generalized polymatroid approach to stable matchings with lower quotas. Math. Oper. Res. **42**, 238–255 (2017)

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

Affiliations

Matthias Mnich[1,2] · Ildikó Schlotter[3]

✉ Matthias Mnich
matthias.mnich@tuhh.de

Ildikó Schlotter
ildi@cs.bme.hu

1 Universität Bonn, Bonn, Germany

2 Institute for Algorithms and Complexity, TU Hamburg, Hamburg, Germany

3 Budapest University of Technology and Economics, Budapest, Hungary