

האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

**MANIPULATION OF STABLE MATCHINGS
USING MINIMAL BLACKLISTS**

By

YANNAI A. GONCZAROWSKI

Discussion Paper # 643R (July 2013, Revised Nov 2013)

מרכז לחקר הרציונליות

**CENTER FOR THE STUDY
OF RATIONALITY**

Feldman Building, Givat-Ram, 91904 Jerusalem, Israel
PHONE: [972]-2-6584135 FAX: [972]-2-6513681
E-MAIL: ratio@math.huji.ac.il
URL: <http://www.ratio.huji.ac.il/>

Manipulation of Stable Matchings using Minimal Blacklists

Yannai A. Gonczarowski*

November 11, 2013

Abstract

Gale and Sotomayor (1985) have shown that in the Gale-Shapley matching algorithm (1962), the proposed-to side W (referred to as *women* there) can strategically force the W -optimal stable matching as the M -optimal one by truncating their preference lists, each woman possibly blacklisting *all but one man*. As Gusfield and Irving have already noted in 1989, no results are known regarding achieving this feat by means other than such preference-list truncation, i.e. by also permuting preference lists.

We answer Gusfield and Irving's open question by providing tight upper bounds on the amount of blacklists and their combined size, that are required by the women to force a given matching as the M -optimal stable matching, or, more generally, as the unique stable matching. Our results show that the coalition of all women can strategically force any matching as the unique stable matching, using preference lists in which at most half of the women have nonempty blacklists, and in which the average blacklist size is *less than 1*. This allows the women to manipulate the market in a manner that is far more inconspicuous, in a sense, than previously realized. When there are less women than men, we show that in the absence of blacklists for men, the women can force any matching as the unique stable matching *without blacklisting anyone*, while when there are more women than men, each to-be-unmatched woman may have to blacklist as many as all men. Together, these results shed light on the question of how much, if at all, do given preferences for one side *a priori* impose limitations on the set of stable matchings under various conditions. All of the results in this paper are constructive, providing efficient algorithms for calculating the desired strategies.

Keywords: Matching; Stability; Deferred Acceptance; Manipulation; Blacklist

*Einstein Institute of Mathematics and Center for the Study of Rationality, Hebrew University of Jerusalem, Israel; and Microsoft Research. *Email:* yannai@gonch.name.

1 Introduction

Gale and Shapley (1962), in their seminal paper, have introduced the question of finding a matching between a set of women W and a set of men M , that is, a one-to-one mapping between a subset of W (the women *matched* by this matching) and a subset of M (the men matched by this matching), that satisfies certain desirable properties.¹ Such W and M , along with preferences for each participant, comprise a two-sided *matching market*. Over the past few decades, applications of matching markets have become widespread, from assigning medical interns to hospitals (Roth, 1984) and students to high schools (Abdulkadiroğlu et al., 2005a,b), to centralizing kidney donation assignments (Roth et al., 2005).

In Gale and Shapley’s model, each woman has a strict order of preference over a subset of M . This subset, ordered according to her order of preference, is called the *preference list* of this woman, and is interpreted as those men with whom she would prefer to be matched over being matched with no-one. Each man similarly has a preference list of women.

Definition 1 (Blacklist). The *blacklist* of a woman w is the set of all men that are not in her preference list, i.e. the men that w finds unacceptable, even when the alternative is being unmatched (i.e. being matched with no-one). The blacklist of a man is defined analogously.

Definition 2 (Individual Rationality). A matching is *M-rational* if no man is matched with a woman from his blacklist; it is *W-rational* if no woman is matched with a man from her blacklist.

A matching is *unstable* under the given orders of preference if either (i) it is not *W-rational*, or not *M-rational*, or (ii) there exist a woman and a man, each of whom prefer being matched with the other over the partner (or lack thereof) assigned to them by the matching. A matching that is not unstable is *stable*; Gale and Shapley (1962) have proved that a stable matching always exists, and have presented an efficient algorithm for finding such a matching.

While in some cases only a single stable matching exists, in general there are many stable matchings, differing significantly in the outcome for the various participants (Pittel, 1989). Gale and Shapley have shown that their algorithm finds the *M-optimal* stable matching, i.e. a stable matching over which no man prefers any other stable matching. McVitie and Wilson (1971) have additionally shown that this matching is at the same time the *W-worst* stable matching, i.e. a stable matching that is worst for each woman. The benefits of the Gale-Shapley matching algorithm for the men have been demonstrated even further by Dubins and Freedman (1981), who have shown that no man can unilaterally strategically manipulate this algorithm (i.e. manipulate the *M-optimal* stable matching) to his advantage and, moreover, that no coalition of men can all benefit from jointly manipulating it.² These observations have led to the study by Gale and Sotomayor (1985) of strategic manipulations of the algorithm, and of stable matchings in general, by women.

Gale and Sotomayor (1985) have shown that if more than one stable matching exists, then at least one woman can unilaterally manipulate the *M-optimal* stable matching to her own advantage. Moreover, they have shown that the coalition of all women can force the *W-optimal* stable matching as the unique stable matching, and thus as the *M-optimal* stable matching (and the outcome of the Gale-Shapley matching algorithm), by truncating their preference lists, each woman blacklisting all men over whom she prefers her *W-optimal* stable partner. Note that this strategy may lead to each woman blacklisting all but one man (i.e. declaring a blacklist of size $|M| - 1$), which makes the women’s manipulation painfully obvious to any observer of

¹Gale and Shapley have studied one-to-many matchings that generalize the scenario described here, as well as one-to-one matchings under the condition $|W| = |M|$. The scenario that we describe, of one-to-one matchings between sets of possibly-unequal size, was first explicitly studied by McVitie and Wilson (1970).

²The interested reader is referred to Demange et al. (1987) for an even stronger non-manipulability theorem, and to Huang (2006) for the study of lying by men in a probabilistic setting.

their submitted preferences. Despite the extensive literature regarding stable matchings and manipulation, as Gusfield and Irving (1989, pp. 59,65) have already noted in 1989, no results are known regarding achieving this feat by means other than such preference-list truncation, i.e. by also permuting preference lists.

As it is widely believed that successful matching mechanisms should produce stable outcomes (Roth, 2002), many real-life matching mechanisms (even those not based on Gale and Shapley’s algorithm) are designed with the goal of producing stable matchings. Therefore, the ability to force some matching as the unique stable matching is of broad significance.

In Section 3, we answer Gusfield and Irving’s open question by tightly characterizing the worst-case amount of blacklists and their sizes, that are required by the women to force the W -optimal stable matching (or more generally, any M -rational matching) as the unique stable matching, and in particular as the outcome of the Gale-Shapley matching algorithm. We start by analysing *full* matchings, i.e. matchings in which no participant is unmatched. A corollary of our results for this case is as follows.

Theorem 1 (Weaker version of Theorem 4; the latter is shown to be tight in Theorem 5).

1. When $|W| = |M|$, the coalition of all women can force any M -rational full matching as the unique stable matching, using a profile of preference lists in which at most half of the women have blacklists, and in which the average blacklist size is less than 1.
2. This profile of preference lists can be computed in $O(n^3)$ time.

This low upper bound of less than 1 on the average blacklist size that is guaranteed by Theorem 1 should be contrasted with the previously-known upper bound of $|M| - 1$, which is attainable using truncation, as mentioned above. Indeed, our results provide a far more “inconspicuous” manipulation than the one suggested by Gale and Sotomayor using preference truncating, even though both manipulations result in the same matching. Thus, if women have to pay even a small cost for every man that they blacklist, then our solution beats preference truncation by an order of magnitude. The manipulation we suggest becomes even more inconspicuous in implementations where, by design, participants are required to submit preference lists of at most a certain length, as in the case in some real-life matching markets (see e.g. Abdulkadiroğlu et al., 2005a).

It turns out that the case of a *balanced market*, i.e. $|W| = |M|$, is a singular case at which a phase change occurs. When there are less women than men, we show (see Theorems 6 and 7) that in the absence of blacklists for men, the women can always force any M -rational matching, in which all women are matched, as the unique stable matching *without blacklisting anyone*. In contrast, when there are more women than men (or more generally, when not all women should be matched), each to-be-unmatched woman may have to blacklist as many as all men. Together, all of these results shed light on the question of how much, if at all, do given preferences for one side *a priori* impose limitations on the set of stable matchings under various conditions.

It is interesting to note that Ashlagi et al. (2013) have shown that a somewhat similar phase change occurs w.r.t. the expected ranking of each participant’s partner in the participant’s preference list. Both that paper and this one show, in different senses, that the preferences of the smaller side of the market, even if it is only slightly smaller, play a far more significant role than may be expected in determining the stable matchings, and those of the larger side — a considerably insignificant one.

Another implication of our results is in goods allocation problems, i.e. matching problems in which only one side (the *buyers*) has preferences, while the other side (the *goods*) does not. Abdulkadiroğlu and Sönmez (2003) and Abdulkadiroğlu et al. (2009) consider using a version of the (student-optimal) Gale-Shapley algorithm in order to allocate school seats to children; as school priorities are very coarse (and sometimes nonexistent, see e.g. Abdulkadiroğlu et al., 2005a), some tie-breaking rule is needed so that the algorithm can be run. Both of these papers

advocate the use of a single tie-breaking rule for all schools (e.g. a single lottery conducted once before the algorithm runs, to assign each student a number to be used for tie breaking at every school) over a different tie-breaking rule for each school (e.g. a different lottery for each school), arguing that the former results in higher social welfare. Our results strengthen this argument and make it even more concrete, showing that in goods allocation problems, an implementation using multiple arbitrary (e.g. randomized) tie-breaking rules may yield any allocation that is buyer-rational, regardless of the internal orders of the buyer’s preference lists, if even an amortized less-than-one-buyer blacklist per good is allowed (or even in the absence any blacklists for the goods whatsoever, if there are more buyers than goods); in contrast, it is easy to see that an implementation with a single tie-breaking rule is equivalent to serial dictatorship (Abdulkadiroğlu and Sönmez, 1998), and thus e.g. constructs Pareto-efficient allocations.

While Theorem 1 guarantees an average blacklist size of less than 1, it provides no guarantee regarding the size of each individual blacklist. Indeed, it may be the case (see Theorem 5) that in order to force a particular matching, one of the women is required to blacklist all but one man, while all the other women’s blacklists are empty. In Section 4, we introduce a natural variant of the Gale-Shapley algorithm, in which blacklisting is a dynamic process called *divorcing*; this variant, which we call the *Gale-Shapley algorithm with divorces*, is exactly as manipulable as the unaltered Gale-Shapley algorithm (see Proposition 8 for the precise sense). We characterize the worst-case pattern of divorces in this variant, and show that the coalition of all women can force any matching as the outcome by at most $n - 1$ women getting divorced, each exactly once (see Theorem 9). This result is shown to be tight as well (see Theorem 10).

All of the results in this paper are constructive, providing efficient algorithms for calculating the desired strategies; full proofs for all theorems are given in the appendix. While the proofs are involved, the resulting algorithms are relatively simple to describe.

2 Model and Preliminaries

In order to standardize the notation used throughout this paper, and for self-containment, let us quickly recapitulate the definitions from the previous section, as well as the Gale-Shapley algorithm (1962). Let W and M be disjoint finite sets, referred to as the sets of women and men, respectively.

Definition 3 (Preference Lists; Blacklists; Combined Blacklist Size).

1. A *preference list* for a woman $w \in W$ is a totally-ordered subset of M .
2. A *profile of preference lists* for W is a specification of a preference list for each $w \in W$.
3. Given a profile \mathcal{P}_W of preference lists for W , the *blacklist* of a woman $w \in W$, denoted by $B_w(\mathcal{P}_W)$, is the set of men in M who do not appear in w ’s preference list (i.e. all men who are declared unacceptable by w). We say that w *blacklists* the men in $B_w(\mathcal{P}_W)$.
4. The *combined size* of all blacklists in \mathcal{P}_W is defined to be $\sum_{w \in W} |B_w(\mathcal{P}_W)|$.

We define preference lists, profiles of preference list, and blacklists for M analogously.

Definition 4 (Matching). A *matching* is a one-to-one mapping between a subset of W and a subset of M . By slight abuse of notation, we denote the woman matched with a man m by a matching μ by $\mu(m)$ instead of $\mu^{-1}(m)$. Given a matching μ , we define the subset of W (resp. M) on which μ is defined as W_μ (resp. M_μ). If $W_\mu = W$ and $M_\mu = M$, then we say that μ is *full*; otherwise, we say that μ is *partial*.

Throughout the remainder of this section, let \mathcal{P}_W and \mathcal{P}_M be profiles of preference lists for W and M , respectively.

Definition 5 (Individual Rationality; Stability). Let μ be a matching.

1. μ is said to be *G-rational*, for a side $G \in \{W, M\}$, if $\mu(p) \notin B_p(\mathcal{P}_G)$ for every matched participant $p \in G_\mu$.
2. μ is said to be *unstable* w.r.t \mathcal{P}_W and \mathcal{P}_M , if either it is not *W-rational* or not *M-rational*, or there exist a woman $w \in W$ and a man $m \in M$ that do not blacklist each other, each of whom either unmatched in μ , or preferring the other over the partner matched to them by μ .
3. If μ is not unstable, then it is said to be *stable*.

Definition 6 (The Gale-Shapley Algorithm (1962)). The following algorithm is henceforth referred to as the *Gale-Shapley algorithm*.³ The algorithm is divided into steps, to which we refer as *nights*. On each night, each man serenades under the window of the woman he prefers most among all women who have not (yet) rejected him and who are not in his blacklist (if such a woman exists), and then each woman, under whose window more than one man serenades, rejects every man who serenades under her window, except for the man she prefers most among these men; any woman under whose window serenade only men from her blacklist, rejects all of these men. The algorithm stops on a night on which no man is rejected by any woman, and then each woman under whose window a man has serenaded on this night is matched to this man. Other women are unmatched (as are men who have not serenaded under any window on this night).

Theorem 2 (Gale and Shapley (1962)). *The Gale-Shapley algorithm stops and yields a stable matching between W and M (in particular, such a matching exists), and this stable matching is M -optimal, i.e. no stable matching is better for any man.*

Theorem 2 states that the stable matching given by the Gale-Shapley algorithm is optimal (out of all stable matchings) for each man. A conceptually-reverse claim holds regarding the women:

Theorem 3 (McVitie and Wilson (1971)). *No stable matching is worse for any woman than the matching given by the Gale-Shapley algorithm.*

In particular, if both the Gale-Shapley algorithm as described above, and the Gale-Shapley algorithm with reverse roles (i.e. with women serenading to men, yielding the *W*-optimal stable matching), yield the same matching, then this matching is the unique stable matching.

3 Blacklists

3.1 Full Matchings

Example 1 (Forcing by men of any matching as the *M*-optimal stable matching). Let W and M be equal-sized sets of women and men, respectively. Let μ be a matching and let \mathcal{P}_M be any profile of preference lists for M in which each man m 's top choice is $\mu(m)$. For any profile \mathcal{P}_W of preference lists for W according to which μ is *W*-rational, the *M*-optimal stable matching (yield by the Gale-Shapley algorithm) is μ .

Gale and Sotomayor (1985) have shown that women can strategically force the *W*-optimal stable matching as the *M*-optimal one by truncating their preference lists, each woman blacklisting all men over whom she prefers her *W*-optimal stable partner (without altering the order of the men remaining in her preference list). If the top choices of all women are distinct, then the *W*-optimal stable matching is for each woman to be matched with her top choice. Gale and Sotomayor's strategy in this case is for each woman to blacklisting all men but her top choice. Such a profile of preference lists indeed allows the women to dictate the *M*-optimal stable matching a strong sense, as demonstrated by the following example.

³The semantics of the algorithm that we describe w.r.t. blacklists are slightly different than those of Gale and Shapley's original algorithm. Nonetheless, the outcome is the same by a theorem of Dubins and Freedman (1981).

Example 2 (Naïve forcing by women of any matching as the M -optimal stable matching). Let W and M be equal-sized sets of women and men, respectively. Let μ be a matching and let \mathcal{P}_W be the profile of preference lists for W in which each woman blacklists all of M except for $\mu(m)$. For any profile \mathcal{P}_M of preference lists for M according to which μ is M -rational, the only stable matching (and thus the M -optimal stable matching) is μ .

As illustrated by Examples 1 and 2, each side can force any “other-side-rational” matching as the M -optimal stable matching, however in order to do so, the women may need to submit far more specific, and far shorter, preference lists. Indeed, a conspiracy as in Example 2, with $n \triangleq |W| = |M|$ nonempty blacklists with a combined size of $n \cdot (n - 1)$, would be painfully obvious to anyone examining the women’s submitted preference lists. Despite the extensive literature regarding stable matchings and manipulation, as Gusfield and Irving (1989, pp. 59,65) have already noted in 1989, no results are known regarding achieving this feat by means other than such preference-list truncation, i.e. by also permuting preference lists.⁴

The main result of this paper provides an answer to Gusfield and Irving’s open question, and shows that if the preferences of all men are known to the women, then the latter can force any given M -rational matching (and in particular the W -optimal stable matching) as the unique stable matching (and in particular as the M -optimal stable matching), using a far smaller combined size for all blacklists, and with significantly less nonempty blacklists than by using preference-list truncation, rendering the manipulation far less obvious in some sense. For implementations where, by design, one of the sides is required to submit preference lists of at most a certain length (e.g. Abdulkadiroğlu et al., 2005a), this result allows such a side to force any matching that is “other-side-rational”, in an even further-inconspicuous way.

Theorem 4 (Manipulation with Minimal Blacklists). *Let W and M be equal-sized sets of women and men, respectively. Define $n \triangleq |W| = |M|$. Let \mathcal{P}_M be a profile of preference lists for M . For every M -rational full matching μ , there exists a profile \mathcal{P}_W of preference lists for W , s.t. all of the following hold.*

1. *The only stable matching, given \mathcal{P}_W and \mathcal{P}_M , is μ .*
2. *The blacklists in \mathcal{P}_W are pairwise disjoint, i.e. no man appears in more than one blacklist.*
3. *n_b , the number of women who have nonempty blacklists in \mathcal{P}_W , is at most $\lfloor \frac{n}{2} \rfloor$.*
4. *The combined size of all blacklists in \mathcal{P}_W is at most $n - n_b$, i.e. at most the number of women who have empty blacklists.*

Furthermore, \mathcal{P}_W can be computed in worst-case $O(n^3)$ time, best-case $O(n^2)$ time and average-case (assuming μ is uniformly distributed given \mathcal{P}_M) $O(n^2 \log n)$ time.

We note that Theorem 4 guarantees a combined blacklist size no greater than the size of each of the individual n blacklists from Example 2 — an order-of-magnitude improvement.

Corollary 1 (Upper Bound on Combined Blacklist Size). *Under the conditions of Theorem 4, the combined size of all blacklists in \mathcal{P}_W is at most $n - 1$, i.e. the average blacklist size is less than 1.*

Theorem 4 and Corollary 1 also have implications in goods allocation problems, i.e. matching problems in which only one side (the *buyers*) has preferences, while the other side (the *goods*) does not. Abdulkadiroğlu and Sönmez (2003) and Abdulkadiroğlu et al. (2009) consider using a version of the (student-optimal) Gale-Shapley algorithm in order to allocate school seats to children; as school priorities are very coarse (and sometimes nonexistent, see e.g. Abdulkadiroğlu et al., 2005a), some tie-breaking rule is needed so that the algorithm can be run.

⁴Manipulations via preference-list truncations generally lend to easier analysis than general manipulations. As Gale and Sotomayor (1985) show, if a woman’s utility is determined solely by the identity of her partner, then preference-list truncation is a (weakly) dominant strategy (as we show in this paper, this is no longer the case if it is preferable to blacklist as few men as possible). The combination of these has led to most of the relevant literature focusing on manipulation via preference-list truncation. Nonetheless, some results do extend to arbitrary manipulation, most of them requiring nonstandard proof techniques. (See e.g. Gonczarowski and Friedgut, 2013).

Both of these papers advocate the use of a single tie-breaking rule for all schools (e.g. a single lottery conducted once before the algorithm runs, to assign each student a number to be used for tie breaking at every school) over a different tie-breaking rule for each school (e.g. a different lottery for each school), arguing that the former results in higher social welfare. Theorem 4 strengthens this argument and makes it even more concrete, showing that in goods allocation problems (in the absence of any priorities for the goods), an implementation using multiple arbitrary (e.g. randomized) tie-breaking rules may yield (for the case of randomization, with possibly-small, albeit positive, probability) any allocation that is buyer-rational, regardless of the internal orders of the buyer’s preference lists, if even an amortized less-than-one-buyer blacklist per good is allowed (and as we show in the next section, even if no blacklists are allowed whatsoever, if $|W| < |M|$); in contrast, it is easy to see that an implementation with a single tie-breaking rule is equivalent to serial dictatorship, or to random serial dictatorship (Abdulkadiroğlu and Sönmez, 1998) if the tie-breaking rule is randomized, and thus e.g. constructs Pareto-efficient allocations.

We note that since the combined size of the preference lists of W as in Theorem 4 is $\Theta(n^2)$, then as long as the algorithm for finding them must encode its output explicitly, its time complexity must be $\Omega(n^2)$. In Theorem 11 in the appendix, we show that attaining $\Theta(n^2)$ time complexity is possible in the special case of Theorem 4 in which the top choices of all men are distinct. Somewhat surprisingly, the analysis and proof for this scenario, in which the men attempt to force some matching as the M -optimal stable matching *à la* Example 1, are simpler than those of the general case, as is calculating the women’s “response” \mathcal{P}_W .⁵ Thus, in a sense the men actually inadvertently help the women whenever they try to manipulate the algorithm and force some matching as the M -optimal stable matching.

We now show that Theorem 4 is tight in a strong sense, w.r.t. the amount of blacklists and their sizes. We conclude that Theorem 4 describes an optimal strategy for the women in any model in which, all assigned partners being equal, the utility of the coalition of women decreases as the combined size of their blacklists increases (e.g. if women have to pay even a small cost for every man that they blacklist). Indeed, in such models Theorem 4 beats preference truncation by an order of magnitude.

Theorem 5 (Tightness of Theorem 4 and Corollary 1). *Let $n \in \mathbb{N}$ and let W and M be sets of women and men, respectively, s.t. $|W| = |M| = n$. For every $0 \leq n_b \leq \lfloor \frac{n}{2} \rfloor$ and for every $l_1, \dots, l_{n_b} > 0$ s.t. $l_1 + \dots + l_{n_b} \leq n - n_b$, there exist a profile \mathcal{P}_M of preference lists for M in which all blacklists are empty, and a full matching μ , s.t. both of the following hold.*

1. *There exists a profile \mathcal{P}_W of preference lists for W with precisely n_b nonempty blacklists, all pairwise disjoint and of sizes l_1, \dots, l_{n_b} , s.t. the only stable matching, given \mathcal{P}_W and \mathcal{P}_M , is μ .*
2. *For every profile \mathcal{P}'_W of preference lists for W s.t. the only stable matching, given \mathcal{P}'_W and \mathcal{P}_M , is μ , there exist n_b distinct women $w_1, \dots, w_{n_b} \in W$ s.t. $|B_{w_i}(\mathcal{P}'_W)| \geq l_i$ for every $0 \leq i \leq n_b$.*

Remark 1. Regarding Theorem 5(2),

1. As shown in our proof below, the requirement that the only stable matching, given \mathcal{P}'_W and \mathcal{P}_M , is μ , may be replaced by the weaker assumption that the M -optimal stable matching, given \mathcal{P}'_W and \mathcal{P}_M , be μ .
2. No limitations are imposed regarding the number of blacklists in \mathcal{P}'_W in which each man appears; in particular, the blacklists in \mathcal{P}'_W need not necessarily be pairwise disjoint.

⁵Moreover, in this special case, while \mathcal{P}_W naturally depends on \mathcal{P}_M , if the implementation of the Gale-Shapley algorithm that is used is the one proposed by Dubins and Freedman (1981), then there exists a scheduling of this implementation for which the decision of whom each woman prefers or blacklists on each step according to \mathcal{P}_W can be taken online. More precisely, both the choice of who acts next, and the action of that participant if that participant is a woman, depend solely the history of the run (and not on the not-yet-disclosed suffixes of the men’s preference lists). Thus, if participants are not required to submit their preference lists in advance, but rather only to dynamically act upon them, then a strategy for the women that forces μ as the only stable matching against *every* profile \mathcal{P}_M of preference lists for M can be constructed for Dubins and Freedman’s implementation of the algorithm, if the women can control its scheduling; this does not seem to be possible in the general case of Theorem 4.

3.1.1 Proof Outline of Theorem 4

A full proof of Theorem 4 is given in the appendix. The general flow of the proof is as follows. We construct a profile of preference lists for W s.t. the top choice of every $w \in W$ is $\mu(w)$; thus, the W -optimal stable matching is μ , and it is enough to make sure that the M -optimal stable matching is μ as well. We construct this profile iteratively.

Assume, for the time being, that the top choices of all men are distinct, i.e. each woman is serenaded-to on the first night of the Gale-Shapley algorithm by precisely one man (in this case, in the absence of any blacklists for women, the algorithm would stop on the first night). We pick some woman \tilde{w} who is not matched with $\mu(\tilde{w})$, and set her blacklist list so that she rejects the unique man m serenading under her window on the first night. We adjust the women's preferences (see Lemma 7), so that m is continually rejected until he serenades under the window of $\mu(m)$, who then accepts him, and in turn rejects the man m' serenading under her window, who is continually rejected until he serenades under the window of $\mu(m')$, and so fourth until this *rejection cycle* concludes with $\mu(\tilde{w})$ serenading under the window of \tilde{w} . We note that while \tilde{w} may have rejected (and thus has blacklisted) more than one man during this rejection cycle, all such men are matched with their μ -partners at the end of this rejection cycle, and in addition so is $\mu(\tilde{w})$, whom no woman blacklisted; in fact, the only woman who blacklists anyone so far is \tilde{w} , and so we have that more men are now matched with their μ -partners than have been blacklisted, so we are "on the right track", in a sense.

At this point, we would have naïvely liked to pick another woman \tilde{w}' who is unmatched with $\mu(\tilde{w}')$, and initiate a similar rejection cycle triggered by her, and so fourth until μ is obtained (arguing that this "cycle-by-cycle" simulation yields the M -optimal stable matching as well, due to the outcome of the Gale-Shapley algorithm being invariant under certain timing changes, an invariance established by Dubins and Freedman, 1981). Alas, it is quite possible that all "nominees" for the role of \tilde{w}' have already rejected quite a few men during the previous rejection cycle (that which was triggered by \tilde{w}). In this case, adjusting the preference list of \tilde{w}' to reject the single man serenading under her window entails having her blacklist not only that man, but also every man she rejected in the previous rejection cycle in favour of this man, which would not yield the required low combined blacklist size (nor necessarily yield disjoint blacklists). In this case, we take a step back and "merge" (see Lemma 8) what would have been the rejection cycle triggered by \tilde{w}' into the previous rejection cycle triggered by \tilde{w} , i.e. modify the preferences of W without blacklisting an "excessive" amount of men, so that the "chain-reaction" triggered by the rejection of m by \tilde{w} , would cause not only all rejections from the cycle triggered by \tilde{w} (as originally defined), but also the rejections from the cycle that would have been triggered by \tilde{w}' . Luckily, such successive "merging" can be done in an efficient manner, without the need to resimulate rejection cycles after each step. Only when "merging" of additional rejection cycles is no longer possible (see Lemma 9), do we start another rejection cycle (see induction step in the proof of Theorem 11), knowing that any woman that we pick to trigger this new rejection cycle has not yet rejected any man in previous rejection cycles.

In the general case, if we let the Gale-Shapley algorithm run its course (using arbitrary preferences s.t. each woman w prefers $\mu(w)$ over any other man), then by the time the algorithm halts, it may already be the case that every woman has rejected quite a few men, and therefore, as in the previous case, cannot be used to trigger a rejection cycle without blacklisting an "excessive" amount of men. In this case, we show (see the induction step in the proof of Theorem 4) that there exists at least one woman s.t. the rejection cycle that would have been triggered by her can be "merged" into the run of the algorithm before it has halted, without blacklisting an "excessive" amount of men. The analysis, and the corresponding modifications to the preferences of W , are considerably more delicate in the case. A side effect is that the time complexity is also somewhat higher (a worst case of $O(n^3)$ instead of $O(n^2)$), however certain properties of random permutations are used to show that the increment in the average-case time complexity is considerably less significant (from $O(n^2)$ to $O(n^2 \cdot \log n)$).

3.2 Partial Matchings

It turns out that the case in which $|W| = |M|$ is a singular case at which a phase change occurs. If there are more women than men (and therefore not all women are matched in μ), then each unmatched woman may be required to blacklist as many as all men; in contrast, if there are more men than women (and therefore not all men are matched in μ), and if the unmatched men do not blacklist any woman,⁶ then no blacklists are whatsoever required, turning the dichotomy between Examples 1 and 2 on its head. It is interesting to note that Ashlagi et al. (2013) have shown that a somewhat similar phase change occurs w.r.t. the expected ranking of each participant's partner in the participant's preference list. Both that paper and this one show, in different senses, that the preferences of the smaller side of the market, even if it is only slightly smaller, play a far more significant role than may be expected in determining the stable matchings, and those of the larger side — a considerably insignificant one. Theorem 6 formalizes these results, as well as their generalizations for partial matchings with both unmatched women and unmatched men. Theorem 7 shows that these results are also tight in the same strong sense in which Theorem 5 shows that Theorem 4 is tight. Before we formulate these theorems, we first define some notation.

Definition 7. Let W and M be sets of women and men, respectively. For every $\tilde{W} \subseteq W$ (resp. $\tilde{M} \subseteq M$), we define $\tilde{W}^c = W \setminus \tilde{W}$ (resp. $\tilde{M}^c = M \setminus \tilde{M}$). (The set W (resp. M) will be clear from context.)

Theorem 6 (Manipulation with Minimal Blacklists to obtain a Partial Matching). *Let W and M be sets of women and men of arbitrary sizes, let \mathcal{P}_M be a profile of preference lists for M , and let μ be a (possibly-partial) M -rational matching. Define $n_h \triangleq |\{w \in W_\mu \mid \exists m \in M_\mu^c : w \notin B_m(\mathcal{P}_M)\}|$. There exists a profile \mathcal{P}_W of preference lists for W , s.t. all of the following hold.*

1. *The only stable matching, given \mathcal{P}_W and \mathcal{P}_M , is μ .*
 2. *The blacklists of W_μ in \mathcal{P}_W are pairwise disjoint, and contain only members of M_μ .*
 3. *n_b , the number of women in W_μ who have nonempty blacklists in \mathcal{P}_W , is at most $\lfloor \frac{n_\mu - n_h}{2} \rfloor$.*
 4. *The combined size of the blacklists of W_μ in \mathcal{P}_W is at most $n_\mu - n_h - n_b$.*
- Furthermore, \mathcal{P}_W can be computed in worst-case $O(\max\{|W| \cdot |M|, (n_\mu - n_h) \cdot n_\mu^2\})$ time, best-case $O(|W| \cdot |M|)$ time and average-case $O(\max\{|W| \cdot |M|, n_\mu^2 \log \frac{n_\mu + 1}{n_h + 1}\})$ time.*

Corollary 2. *Under the conditions of Theorem 4,*

1. *The combined size of all blacklists of W_μ in \mathcal{P}_W is at most $n_\mu - n_h - 1 \leq n - 1$, i.e. the average blacklist size of W_μ is less than 1.*
2. *If all women are matched in μ (i.e. $W_\mu = W$), and if for each $w \in W$ except for perhaps one, there exists $m \in M_\mu^c$ who does not blacklist w (e.g. this condition holds if at least one $m \in M_\mu^c$ blacklists at most one woman), then all blacklists in \mathcal{P}_W are empty, and \mathcal{P}_W can be computed in $O(|W| \cdot |M|)$ time, which constitutes a tight bound.*

Theorem 7 (Tightness of Theorem 6 and Corollary 2). *Let W and M be sets of women and men of arbitrary sizes, and let $\tilde{W} \subseteq W$ and $\tilde{M} \subseteq M$ s.t. $|\tilde{W}| = |\tilde{M}|$. For each $w \in \tilde{W}^c$, let $B_w \subseteq M$, and for each $m \in \tilde{M}^c$, let $B_m \subseteq W$. Define $n_h \triangleq |\{w \in \tilde{W} \mid \exists m \in \tilde{M}^c : w \notin B_m\}|$. For every $0 \leq n_b \leq \lfloor \frac{n_\mu - n_h}{2} \rfloor$ and for every $l_1, \dots, l_{n_b} > 0$ s.t. $l_1 + \dots + l_{n_b} \leq n_\mu - n_h - n_b$, there exist a profile \mathcal{P}_M of preference lists for M s.t. $B_m(\mathcal{P}_M) = B_m$ for each $m \in \tilde{M}^c$, and a matching μ s.t. $W_\mu = \tilde{W}$ and $M_\mu = \tilde{M}$, s.t. both of the following hold.*

1. *There exists a profile \mathcal{P}_W of preference lists for W , in which $B_w(\mathcal{P}_W) \supseteq B_w$ for each $w \in \tilde{W}^c$, and in which \tilde{W} have precisely n_b nonempty blacklists, all pairwise disjoint and containing only members of \tilde{M} , and of sizes l_1, \dots, l_{n_b} , s.t. the only stable matching, given \mathcal{P}_W and \mathcal{P}_M , is μ .*

⁶A much weaker condition suffices. See e.g. Corollary 2(2) below.

2. For every profile \mathcal{P}'_W of preference lists for W s.t. the only stable matching, given \mathcal{P}'_W and \mathcal{P}_M , is μ , it holds that $B_w(\mathcal{P}'_W) \supseteq \{m \in B_w(\mathcal{P}_W) \mid w \notin B_m\}$ for each $w \in W^c$, and in addition, there exist n_b distinct women $w_1, \dots, w_{n_b} \in \tilde{W}$ s.t. $|B_w(\mathcal{P}'_W) \cap \tilde{M}| \geq l_i$ for every $0 \leq i \leq n_b$.

Remark 2. As in Remark 1, we note the following also regarding Theorem 7(2).

1. As shown in our proof below, the requirement that the only stable matching, given \mathcal{P}'_W and \mathcal{P}_M , is μ , may be replaced by the weaker assumption that the M -optimal stable matching, given \mathcal{P}'_W and \mathcal{P}_M , be μ .
2. No limitations are imposed regarding the number of blacklists in \mathcal{P}'_W in which each man appears; in particular, not even the blacklists of \tilde{W} in \mathcal{P}'_W need necessarily be pairwise disjoint.

4 Divorces

In Section 3, we showed that women can force any (M -rational) matching as the outcome of the Gale-Shapley algorithm, via a small amount of nonempty blacklists, and average blacklist size less than one; nonetheless, as shown in that section, some of these blacklists may be quite long, having length in the order of magnitude of the number of men. We now show that, turning blacklisting into a dynamic process in the Gale-Shapley algorithm, women can force any M -rational matching as the outcome by means of at most one blacklist-driven rejection per woman (*without averaging*), but by a possibly-larger number of women than in Section 3.

Definition 8 (Gale-Shapley algorithm with Divorces). We extend the Gale-Shapley algorithm to allow for *divorces* as follows. The extended algorithm is divided into stages to which we refer as *seasons*. In the first season, the *vanilla* (i.e. without divorces) Gale-Shapley algorithm runs until it converges. At the end of each season s , if any woman asks to divorce her then-current partner, then one of these women, denoted w , is chosen (arbitrarily, or according to some predefined rule); in season $s + 1$, the vanilla Gale-Shapley algorithm runs once more, starting at the state concluding season s , with w rejecting her then-current partner on the first night (e.g. on the second night, this partner serenades under the window of his next choice after w), and running until it converges once more. When a season concludes with no woman asking for a divorce, the Gale-Shapley algorithm with divorces concludes.

Remark 3. If all women have a divorce strategy of “never-divorce”, then the Gale-Shapley algorithm with divorces concludes after one season, and is thus equivalent to the vanilla Gale-Shapley algorithm.

We note that we define the Gale-Shapley algorithm with divorces in terms of seasons for ease of presentation; indeed, all the results in this section hold verbatim, via conceptually-similar proofs, if we allow every woman to divorce her partner instantaneously whenever she so desires.

Before we present our results regarding divorces, we first exhibit the differences between blacklists and divorces on one hand, and the similarities between these on the other hand. To exhibit the qualitative distinction between these two concepts, consider for a moment a woman w who wishes to blacklist all men who approach her on the first night of the vanilla Gale-Shapley algorithm; such a (potentially-long) blacklist can be “replaced” by a single divorce in the following manner: w places all such men at the end of her preference list in arbitrary order, and thus rejects all such men but one of them, denoted m , on the first night; if w does not reject m by the end of the season, then she asks to divorce him (possibly repeatedly, until she is chosen to do so).⁷ Indeed, as shown in Theorems 9 and 10, no woman

⁷We note that if men were allowed to divorce as well, then they could also use divorces to replace blacklists, but to achieve a dual effect, i.e. instead of one divorce replacing the entire blacklist of one man, one divorce could allow

need divorce more than one man in order to achieve the lower bound on the required number of divorces; nonetheless, as we show in Theorem 10, the tight bound of Corollary 1 on the combined size of all blacklists cannot be improved upon by replacing blacklists with divorces. Proposition 8, which is an immediate consequence of the invariance of the outcome of the vanilla Gale-Shapley algorithm under timing changes (Dubins and Freedman, 1981), exhibits an equivalence, in a sense, of the strengths of blacklists and divorces.

Proposition 8 (Equivalent Strength of Divorces and Blacklists). *Let W and M be equal-sized sets of women and men, respectively, let \mathcal{P}_W be a profile of preference lists and divorce strategies for W and let \mathcal{P}_M be a profile of preference lists for M . Let $w \in W$.*

1. *If w 's divorce strategy in \mathcal{P}_W is that of "never-divorce", then removing w 's blacklist, placing its members in the end of w 's preference list in arbitrary order, and replacing w 's divorce strategy with "always ask for a divorce if w 's current partner is in $B_w(\mathcal{P}_W)$ ", does not alter the outcome of the Gale-Shapley algorithm with divorces.*
2. *If w has an empty blacklist in \mathcal{P}_W , then replacing her divorce strategy with a "never-divorce" strategy, and blacklisting the minimal set B s.t. B contains all the men w divorces during the run of the Gale-Shapley algorithm with divorces given \mathcal{P}_W and \mathcal{P}_M and s.t. every man that w rejects during this run in favour of a man in B is also in B ,⁸ does not alter the outcome of the Gale-Shapley algorithm with divorces.*

We are now ready to formulate our results regarding the Gale-Shapley algorithm with divorces.

Theorem 9 (One Divorce per Woman). *Let W and M be equal-sized sets of women and men, respectively. Define $n \triangleq |W| = |M|$. Let \mathcal{P}_M be a profile of preference lists for M . For every M -rational full matching μ , there exist a profile \mathcal{P}_W of preference lists and divorce strategies for W , s.t. all of the following hold.*

1. *The run of the Gale-Shapley algorithm with divorces according to \mathcal{P}_W and \mathcal{P}_M , which we henceforth denote by R , yields μ .*
2. *All blacklists in \mathcal{P}_W are empty.*
3. *At most $n - 1$ divorces occur during R , each of them by a distinct woman. Moreover, if a season of R commences with the divorce of a woman w , then this season concludes with w becoming matched with $\mu(w)$.*

Theorem 10 (Tightness of Theorem 9). *Let $n \in \mathbb{N}$ and let W and M be sets of women and men, respectively, s.t. $|W| = |M| = n$. There exist a profile \mathcal{P}_M of preference lists for M , and an M -rational full matching μ , s.t. for every profile \mathcal{P}'_W of preference lists and divorce strategies for W in which all blacklists are empty s.t. the run of the Gale-Shapley algorithm with divorces, according to \mathcal{P}'_W and \mathcal{P}_M , yields μ , at least $n - 1$ divorces occur in this run.*

Similarly to Theorems 6 and 7, it is possible to show that in the case of partial matchings, no divorces (nor blacklists) are required by any woman who is not blacklisted by at least one unmatched man. While any unmatched woman, in the context of Theorem 6, may be required to blacklist as many as all men, in many cases when divorces are allowed, such blacklists of all men can each be replaced by a single divorce.

the blacklisting of a woman by a single one man instead of by a set of men. Indeed, consider e.g. the scenario in which all men wish to blacklist all women, forcing the empty matching. In this case, the men could all set the same order of preference for themselves; they would thus all serenade under the window of the same woman w on the first night and w would thus reject all of them but a single man m , who would later divorce her. Similarly, every woman would have to be divorced only by her most-favourite man, thus replacing blacklists with combined size n^2 , with n divorces, each of a distinct woman but not necessarily by a distinct man. We note that in this example, due to the dynamic nature of divorces, this divorce strategy for the men would force the empty matching as the outcome against every profile \mathcal{P}_W of preference lists and divorces for W , just as blacklisting all women would.

⁸ B is the set of men reachable by means of the transitive closure of the operation "all men who are rejected by w during this run in favour of ...", starting from a man divorced by w during this run.

5 Acknowledgements

This work was supported in part by an ISF grant, by the Google Inter-university center for Electronic Markets and Auctions, and by the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no. [249159]. The author would like to thank Sergiu Hart and Noam Nisan, his Ph.D. advisors, for useful discussions and comments; Assaf Romm for providing him with useful references to the literature; and Jacob Leshno for suggesting the implication in goods allocation.

References

- A. Abdulkadiroğlu and T. Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–710, 1998.
- A. Abdulkadiroğlu and T. Sönmez. School choice: A mechanism design approach. *American Economic Review*, 93(3):729–747, 2003.
- A. Abdulkadiroğlu, P. A. Pathak, and A. E. Roth. The new york city high school match. *American Economic Review*, 95(2):364–367, 2005a.
- A. Abdulkadiroğlu, P. A. Pathak, A. E. Roth, and T. Sönmez. The boston public school match. *American Economic Review*, 95(2):368–371, 2005b.
- A. Abdulkadiroğlu, P. A. Pathak, and A. E. Roth. Strategy-proofness versus efficiency in matching with indifferences: Redesigning the nyc high school match. *American Economic Review*, 99(5):1954–1978, 2009.
- R. Arratia, A. Barbour, and S. Tavaré. *Logarithmic Combinatorial Structures: a Probabilistic Approach*. EMS Monographs in Mathematics. European Mathematical Society, Zurich, Switzerland, 2003.
- I. Ashlagi, Y. Kanoria, and J. D. Leshno. Unbalanced random matching markets. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC)*, pages 27–28, 2013.
- G. Demange, D. Gale, and M. Sotomayor. A further note on the stable matching problem. *Discrete Applied Mathematics*, 16(3):217–222, 1987.
- L. E. Dubins and D. Freedman. Machiavelli and the Gale-Shapley algorithm. *American Mathematical Monthly*, 88(7):485–494, 1981.
- D. Gale and L. S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69(1):9–15, 1962.
- D. Gale and M. Sotomayor. Ms. Machiavelli and the stable matching problem. *American Mathematical Monthly*, 92(4):261–268, 1985.
- Y. A. Gonczarowski. The distribution of the combined length of spanned cycles in a random permutation. Discussion Paper 650, Center for the Study of Rationality, Hebrew University of Jerusalem, 2013.
- Y. A. Gonczarowski and E. Friedgut. Sisterhood in the Gale-Shapley matching algorithm. *The Electronic Journal of Combinatorics*, 20(2):#P12 (18pp), 2013.
- D. Gusfield and R. W. Irving. *The Stable Marriage Problem: Structure and Algorithms*. The MIT Press, Cambridge, MA, USA, 1989.

- C.-C. Huang. Cheating by men in the gale-shapley stable matching algorithm. In *Proceedings of the 14th Annual European Symposium on Algorithms (ESA)*, pages 418–341, 2006.
- J. L. W. V. Jensen. Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta Mathematica*, 30(1):175–193, 1906.
- D. G. McVitie and L. B. Wilson. Stable marriage assignment for unequal sets. *BIT*, 10(3):295–309, 1970.
- D. G. McVitie and L. B. Wilson. The stable marriage problem. *Communications of the ACM*, 14(7):486–490, 1971.
- B. Pittel. The average number of stable matchings. *SIAM Journal on Discrete Mathematics*, 2(4): 530–549, 1989.
- A. E. Roth. The evolution of the labor market for medical interns and residents: a case study in game theory. *Journal of Political Economy*, 92(6):991–1016, 1984.
- A. E. Roth. The economist as engineer: Game theory, experimentation and computation as tools for design economics. *Econometrica*, 70(4):1341–1378, 2002.
- A. E. Roth, T. Sönmez, and M. U. Ünver. A kidney exchange clearinghouse in new england. *American Economic Review*, 95(2):376–380, 2005.

A Proofs

A.1 Proof of Theorem 4

Let W , M and n be as in Theorem 4. We begin by defining some notation.

Definition 9. Let \mathcal{P}_W and \mathcal{P}_M be profiles of preferences lists for W and for M , respectively.

- We denote by $R^{\mathcal{P}_W, \mathcal{P}_M}$ the run of the Gale-Shapley algorithm according to \mathcal{P}_W and \mathcal{P}_M , and denote the matching that $R^{\mathcal{P}_W, \mathcal{P}_M}$ yields (i.e. the M -optimal stable matching for \mathcal{P}_W and \mathcal{P}_M) by $MenOpt(\mathcal{P}_W, \mathcal{P}_M)$.
- Let μ' be a matching. We denote by $R_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$ the run of the Gale-Shapley algorithm according to \mathcal{P}_W and \mathcal{P}_M , starting with μ' as the initial state, and denote the matching that $R_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$ yields by $\mu'_{\mathcal{P}_W, \mathcal{P}_M}$.

At the heart of our proof of Theorem 4 lies a combinatorial structure that we call a *cycle*, and which we now define.

Definition 10.

1. $(w_1 \xrightarrow{m_1} w_2 \xrightarrow{m_2} \dots w_{d-1} \xrightarrow{m_{d-1}} w_d)$, for $d \in \mathbb{N}$, is called a *cycle* if for every $i < d$, $w_i \in W$ and $m_i \in M$, and if $w_d = w_1$.
2. We say that a cycle $C \triangleq (w_1 \xrightarrow{m_1} w_2 \xrightarrow{m_2} \dots w_{d-1} \xrightarrow{m_{d-1}} w_d)$ is *simple* if $w_1, \dots, w_{d-1}, m_1, \dots, m_{d-1}$ are all distinct, i.e. if the only participant appearing in C more than once is w_1 , appearing both as w_1 and as w_d .
3. We say that two cycles C, \tilde{C} are *disjoint* if no participant appears in both.
4. We say that a cycle \tilde{C} is a *cyclic shift* of a cycle $C \triangleq (w_1 \xrightarrow{m_1} w_2 \xrightarrow{m_2} \dots w_{d-1} \xrightarrow{m_{d-1}} w_d)$ if there exists $1 \leq \ell \leq d$ s.t. $\tilde{C} = (w_\ell \xrightarrow{m_\ell} w_{\ell+1} \dots \xrightarrow{m_{d-1}} w_d = w_1 \xrightarrow{m_1} w_2 \xrightarrow{m_2} \dots w_\ell)$.

Cycles can be used to naturally describe the dynamics following a rejection in the Gale-Shapley algorithm, as we now show.

Definition 11. Let μ' be a matching. We say that profiles \mathcal{P}_W and \mathcal{P}_M of preferences lists for W and M , respectively, are μ' -*cycle generating* if all of the following hold.

1. For every $m \in M$, there exists a woman w_m s.t. all of the following hold.
 - i. m is w_m 's top choice.
 - ii. m does not blacklist w_m .
 - iii. m weakly prefers $\mu'(m)$ over w_m .
2. There exists a unique woman $\tilde{w} \in W$ s.t. \tilde{w} blacklists $\mu'(\tilde{w})$. We call \tilde{w} the *cycle trigger*.

Lemma 1. Let μ' be a matching. For every μ' -cycle-generating profiles \mathcal{P}_W and \mathcal{P}_M of preferences lists for W and M , respectively, all of W and M are matched according to $\mu'_{\mathcal{P}_W, \mathcal{P}_M}$.

Proof. As $|W| = |M|$, it is enough to show that every $m \in M$ is matched according to $\mu'_{\mathcal{P}_W, \mathcal{P}_M}$. Assume for contradiction that $m \in M$ is unmatched according to $\mu'_{\mathcal{P}_W, \mathcal{P}_M}$. As w_m is not blacklisted by m , and as m weakly prefers $\mu'(m)$ over w_m , we have, by m being unmatched in $\mu'_{\mathcal{P}_W, \mathcal{P}_M}$, that m serenades under w_m 's window during $R_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$, and that w_m rejects w during this run; but m is w_m 's top choice — a contradiction. \square

Definition 12. Let μ' be a matching. Let \mathcal{P}_W and \mathcal{P}_M be μ' -cycle-generating profiles of preferences lists for W and M , respectively, with cycle trigger \tilde{w} . We define the *rejection cycle* $(\mathcal{P}_W, \mathcal{P}_M)$ -generated from μ' , denoted by $C_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$, as $(w_1 \xrightarrow{m_1} w_2 \xrightarrow{m_2} \dots w_{d-1} \xrightarrow{m_{d-1}} w_d)$, where d , $(w_i)_{i=1}^d$ and $(m_i)_{i=1}^d$ are defined by following $R_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$:

1. $w_1 = \tilde{w}$, $m_1 = \mu'(\tilde{w})$. (w_1 rejects m_1 on the first night.)

2. On the nights following the rejection of m_i by w_i , m_i serenades under the windows of the women following w_i in his preferences list. Denote the first of these women to provisionally accept him by w_{i+1} . If w_{i+1} does not reject any man on the night on which she first provisionally accepted m_i , then denote $d = i + 1$ and the algorithm stops. Otherwise, denote the unique man rejected by w_{i+1} on that night by m_{i+1} .

Lemma 2. *Under the conditions of Definition 12, $C_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$ is a well-defined cycle.*

Proof. We show by induction that on every night of the algorithm, one of the following holds: (The claim follows directly from this proof.)

1. Each man serenades under a distinct woman's window. In this case, either the algorithm stops or only \tilde{w} rejects the man serenading under her window.
2. No man serenades under \tilde{w} 's window, two men serenade under the window of some other woman w , and each other woman is serenaded-to by one man. In this case, one of the men serenading under w 's window is the only man rejected on this night.

Base: By definition, on the first night each man $m \in M$ serenades under $\mu'(m)$'s window, and all of these women are distinct. By definition of \tilde{w} , indeed the only rejection on the first night is of m_1 by w_1 . Therefore, the first night is of Type 1.

Step: Assume that the algorithm does not stop at the end of the i 'th night. Regardless of the type of the i 'th night, by the induction hypothesis exactly one man m is rejected during it, while $n - 1$ men are provisionally accepted during it, each by a distinct woman from $W \setminus \{\tilde{w}\}$. By Lemma 1, on the $i + 1$ 'th night m serenades under some woman w 's window. If $w = \tilde{w}$, we show that the $i + 1$ 'th night is of Type 1. Indeed, in this case each woman is serenaded-to by exactly one man. Furthermore, each woman $w' \in W \setminus \{\tilde{w}\}$ is serenaded-to by the man she provisionally accepts on the i 'th night; therefore, this man is not in w 's blacklist, and being the only man serenading under w 's window on the $i + 1$ 'th night, he is not rejected on this night. Thus, no man but m is rejected the $i + 1$ 'th night. If m is not rejected on this night, then the algorithm stops; otherwise, m is rejected by the woman under whose window he serenades, that is, by \tilde{w} .

Otherwise, $w \neq \tilde{w}$ and in this case we show that the $i + 1$ 'th night is of Type 2. Indeed, in this case, on the $i + 1$ 'th night no one serenades under \tilde{w} 's window, two men serenades under w 's window (m and the man w provisionally accepts on the i 'th night, denoted henceforth as m'), and each woman $W \setminus \{\tilde{w}, w\}$ is serenaded-to by exactly one man. Similarly to the previous case, each woman in $W \setminus \{w, \tilde{w}\}$ is serenaded-to by the man she provisionally accepts on the i 'th night; therefore, this man is not in this woman's blacklist, and being the only man serenading under her window on the $i + 1$ 'th night, he is not rejected on this night. Thus, no woman but w rejects any man during this night. As m' is provisionally accepted by w on the i 'th night, he is not blacklisted by her, and thus she does not reject both m and m' on the $i + 1$ 'th night. Thus, she rejects exactly one of them on this night and the proof is complete. \square

Lemma 3. *Under the conditions of Definition 12, the people for whom $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(p) \neq \mu'(p)$ are exactly the people in $C_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$.*

Proof. By Definition 12, the people in $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(p)$ are exactly the people who reject or are rejected during $R_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$. \square

We now move on to define, given two matchings, a cycle associated with both. While this definition is syntactic, and *a priori* not related to any run, in Lemmas 7 and 8 we show its relation to Definition 12.

Definition 13. Let μ' and μ be two matchings and let $w \in W$. We define a sequence $S_{\mu'}^{\mu}(w) \triangleq ((w_i, m_i))_{i=1}^{\infty} \in (W \times M)^{\mathbb{N}}$ as follows:

1. $w_1 = w$.

2. $m_i = \mu'(w_i)$.
3. $w_{i+1} = \mu(m_i)$.

Let $d > 1$ be minimal s.t. $w_d = w_1$. If $d > 2$, then we define the $(\mu' \rightarrow \mu)$ -cycle of w , denoted by $C_{\mu'}^\mu(w)$, as $(w_1 \xrightarrow{m_1} w_2 \xrightarrow{m_2} \dots \rightarrow w_{d-1} \xrightarrow{m_{d-1}} w_d)$. Otherwise (i.e. if $\mu(w) = \mu'(w)$), we define $C_{\mu'}^\mu(w) = (w)$.

Lemma 4.

1. $C_{\mu'}^\mu(w)$ is a well-defined simple cycle, for every $w \in W$.
2. For every $w, w' \in W$, $C_{\mu'}^\mu(w)$ and $C_{\mu'}^\mu(w')$ are either disjoint, or cyclic shifts of one another.
3. W is partitioned into (the women from) $(\mu' \rightarrow \mu)$ -cycles.

Proof. Let $w \in W$ and denote $((w_n, m_n))_{n=1}^\infty \triangleq S_{\mu'}^\mu(w)$. By definitions of μ and μ' , for each i , w_i uniquely determines m_i and (if $i \neq 1$) w_{i-1} ; similarly, m_i uniquely determines w_i and w_{i+1} . Thus, $S_{\mu'}^\mu(w)$ is periodic, and a single period of this sequence contains no participant more than once; therefore, Part 1 follows. We furthermore conclude that for every $w, w' \in W$, $S_{\mu'}^\mu(w)$ and $S_{\mu'}^\mu(w')$ are either disjoint, or each a suffix of the other, and thus Part 2 follows. Finally, Part 3 is an immediate consequence of Part 2. \square

Remark 4. There is a natural isomorphisms between $(\mu' \rightarrow \mu)$ -cycles modulo cyclic shifts, and cycles of the permutation $\mu \circ \mu'^{-1}$ (resp. of the permutation $\mu' \circ \mu^{-1}$), given by taking all women (resp. all men; or $\mu(w)$, if $d = 1$) of a $(\mu' \rightarrow \mu)$ -cycle in order.

Our proof strategy for Theorem 4, as seen below, is to iteratively modify the preferences of W , so that the M -optimal matching they yield becomes “closer and closer”, in a sense, to the target matching μ . We now define the property required of these intermediate M -optimal matchings μ' w.r.t. the preferences of M , for this process to be able to continue until μ is achieved.

Definition 14. Let μ' and μ be matchings. We say that a profile of preference lists for M is $(\mu' \rightarrow \mu)$ -compatible if $m \in M$ weakly prefers $\mu'(m)$ over $\mu(m)$, but does not blacklist $\mu(m)$.

Lemma 5. Let μ' and μ be matchings and let \mathcal{P}_M be an $(\mu' \rightarrow \mu)$ -compatible profile of preference lists for M . For every profile \mathcal{P}_W of preference lists for W according to which each woman w 's top choice is $\mu(w)$, all of W and M are matched according to $\mu'_{\mathcal{P}_W, \mathcal{P}_M}$, and \mathcal{P}_M is also $(\mu'_{\mathcal{P}_W, \mathcal{P}_M} \rightarrow \mu)$ -compatible.

Proof. The proof is similar to that of Lemma 1. Let $m \in M$. We need only show that m is matched in $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(m)$ and weakly prefers $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(m)$ over $\mu(m)$. Since m weakly prefers $\mu'(m)$ over $\mu(m)$, whom he does not blacklist, we have that m would serenade under $\mu(m)$'s window during $R_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$ before serenading under the window of any woman he prefers less, and before becoming unmatched. As $\mu(m)$'s top choice is m , she would never reject him, and the proof is complete. \square

We now begin the journey toward developing Lemma 9, which is the main technical tool underlying our proof of Theorem 4. Lemma 9 is proven inductively, by continuously tweaking the preferences of W until they meet a certain property. Definition 15 defines the induction-invariant properties of the preferences of W during this process, while Lemmas 7 and 8 define the induction base and step of this process, respectively.

Definition 15. Let μ' and μ be matchings and let \mathcal{P}_M be an $(\mu' \rightarrow \mu)$ -compatible profile of preference lists for M . We say that a profile \mathcal{P}_W of preference lists for W is $(\mu' \xrightarrow{\mathcal{P}_M} \mu)$ -compatible if all the following hold.

1. \mathcal{P}_W and \mathcal{P}_M are μ' -cycle generating.
2. Each woman w 's top choice is $\mu(w)$.
3. The cycle trigger for $C_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$ blacklists (at least) every $m \in M$ s.t. $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(m) \neq \mu(m)$.

4. For every woman w other than the cycle trigger, every $m \in M$ that w prefers over $\mu'(w)$, except perhaps for $\mu(w)$, satisfies $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(m) = \mu(m)$.
5. For every $w \in W$, $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(w) \in \{\mu'(w), \mu(w)\}$.

Lemma 6. *Under the conditions of Definition 15, all of the following hold.*

1. Let $m \in M$ s.t. $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(m) \neq \mu(m)$. The only $w \in W$ who prefers m over $\mu'(w)$ is $\mu(m)$.
2. For every $p \in W \cup M$, $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(p) \in \{\mu'(p), \mu(p)\}$.
3. $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(p) = \mu(p) \neq \mu'(p)$ for every participant p in $C_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$.
4. $\{p \in W \cup M \mid \mu'_{\mathcal{P}_W, \mathcal{P}_M}(p) = \mu(p)\}$ is a (disjoint) union of $(\mu' \rightarrow \mu)$ -cycles.

Proof.

1. Follows from Definition 15(4).
2. Follows from Definition 15(5).
3. By Lemma 3, every participant p in $C_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$ satisfies $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(p) \neq \mu'(p)$. Thus, by Part 2, we also have $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(p) = \mu(p)$.
4. Let w be a woman s.t. $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(w) = \mu(w)$. If $\mu'(w) = \mu(w)$, then $C_{\mu'}^{\mu}(w) = (w)$. Otherwise, denote $C_{\mu'}^{\mu}(w) = (w_1 \xrightarrow{m_1} w_2 \xrightarrow{m_2} \dots w_{d-1} \xrightarrow{m_{d-1}} w_d)$, and we show by induction that for all people p in $C_{\mu'}^{\mu}(w)$, $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(p) = \mu(p) \neq \mu'(p)$.

Base: By definition, $w_1 = w$, and thus $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(w_1) = \mu(w_1)$ and by assumption, $\mu(w_1) \neq \mu'(w_1)$.

Step 1: Assume that $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(w_i) = \mu(w_i) \neq \mu'(w_i)$ for some $1 \leq i < d$. Recall that $m_i = \mu'(w_i)$. Since $m_i = \mu'(w_i) \neq \mu(w_i)$, we have $\mu'(m_i) = w_i \neq \mu(m_i)$. Since $m_i = \mu'(w_i) \neq \mu'_{\mathcal{P}_W, \mathcal{P}_M}(w_i)$, we have $w_i = \mu'(m_i) \neq \mu'_{\mathcal{P}_W, \mathcal{P}_M}(m_i)$. Thus, by Definition 15(5), $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(m_i) = \mu(m_i)$.

Step 2: Assume that $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(m_i) = \mu(m_i) \neq \mu'(m_i)$ for some $1 \leq i < d$. Recall that $w_{i+1} = \mu(m_i)$. Since $w_{i+1} = \mu(m_i) = \mu'_{\mathcal{P}_W, \mathcal{P}_M}(m_i)$, we have $\mu(w_{i+1}) = m_i = \mu'_{\mathcal{P}_W, \mathcal{P}_M}(w_{i+1})$. Since $w_{i+1} = \mu(m_i) \neq \mu'(m_i)$, we have $\mu(w_{i+1}) = m_i \neq \mu'(w_{i+1})$.

The proof that for every $m \in M$ s.t. $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(m) = \mu(m) \neq \mu'(m)$, we have $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(p) = \mu(p) \neq \mu'(p)$ for all people p in $C_{\mu'}^{\mu}(\mu(m))$ (this cycle includes m as m_{d-1}), uses Step 2 as the induction base and Steps 1 and 2 as the induction steps. \square

Lemma 7. *Let μ' and μ be matchings, let $\tilde{w} \in W$ s.t. $\mu(\tilde{w}) \neq \mu'(\tilde{w})$, and let \mathcal{P}_W be a profile of preference lists for W according to which each woman $w \neq \tilde{w}$ has preference list starting with $\mu(w)$, followed immediately by $\mu'(w)$ (if $\mu'(w) \neq \mu(w)$), followed by some or all other men in arbitrary order, and in which \tilde{w} has preference list $\mu(\tilde{w})$, with all other men blacklisted. For every $(\mu' \rightarrow \mu)$ -compatible profile \mathcal{P}_M of preference lists for M , both of the following hold.*

1. \mathcal{P}_W is $(\mu' \xrightarrow{\mathcal{P}_M} \mu)$ -compatible.
2. $C_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M} = C_{\mu'}^{\mu}(\tilde{w})$.

Proof. It is straightforward to check, immediately from the definition of \mathcal{P}_W , that \mathcal{P}_W and \mathcal{P}_M are μ' -cycle generating with cycle trigger \tilde{w} and with $w_m = \mu(m)$ for every $m \in M$. Thus, in particular Part 1 of Definition 15 holds. By definition of \mathcal{P}_W , we trivially have that Part 2 of Definition 15 holds.

By Lemma 1 and by definition of \mathcal{P}_W , we have $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(\tilde{w}) = \mu(\tilde{w})$. Thus, $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(\mu(\tilde{w})) = \mu(\mu(\tilde{w}))$, and so, by definition of \mathcal{P}_W we have that \tilde{w} blacklists every $m \in M$ s.t. $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(m) \neq \mu(m)$ and Part 3 of Definition 15 holds.

Let $w \in W \setminus \{\tilde{w}\}$. By definition of \mathcal{P}_W , the only man that w possibly prefers over $\mu'(w)$ is $\mu(w)$ (and this only happens if $\mu(w) \neq \mu'(w)$), and thus Part 4 of Definition 15 vacuously holds.

For every woman $w \in W$, by a well-known property of the Gale-Shapley algorithm, we have that $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(w)$ is w 's most-preferred choice out of all the men that serenade under her window during $R_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$. Thus, since on the first night of this run, each woman $w \in W \setminus \{\tilde{w}\}$

is serenaded-to by $\mu'(w)$, she weakly prefers $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(w)$ over $\mu'(w)$; therefore, by definition of \mathcal{P}_W , we obtain that $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(w) \in \{\mu'(w), \mu(\tilde{w})\}$ for every $w \in W \setminus \{\tilde{w}\}$. Recall that $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(\tilde{w}) = \mu(\tilde{w})$. By both of these, Part 5 of Definition 15 holds and the proof of Part 1 is complete.

We now prove Part 2. Define $C_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M} = (w_1 \xrightarrow{m_1} w_2 \xrightarrow{m_2} \dots w_{d-1} \xrightarrow{m_{d-1}} w_d)$. We show by induction that $(w_1 \xrightarrow{m_1} w_2 \xrightarrow{m_2} \dots w_{d-1} \xrightarrow{m_{d-1}} w_d) = C_{\mu'}^{\mu}(\tilde{w})$.

Base: As \tilde{w} is the cycle trigger for $C_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$, we have, by Definition 12 that $w_1 = \tilde{w}$ and $m_1 = \mu'(\tilde{w})$, agreeing with the definition of $C_{\mu'}^{\mu}(\tilde{w})$.

Step: Let $1 < i \leq d$ and assume that w_j and m_j agree with the definition of $C_{\mu'}^{\mu}(\tilde{w})$ for every $j \leq i$. We show that w_i and, if $i < d$, also m_i , agrees with the definition of $C_{\mu'}^{\mu}(\tilde{w})$ as well. Let t be the night on which w_i rejects m_i . On the first night, every $w \in W \setminus \{\tilde{w}\}$ is serenaded-to by $\mu'(w)$, whom she does not blacklist; therefore, once more by the same well-known property of the Gale-Shapley algorithm as above, on every night there is a man serenading under w 's window that she weakly prefers over $\mu'(w)$. Thus, by definition of \mathcal{P}_W , every $w \in W \setminus \{\tilde{w}, \mu(m_i), \mu'(m_i)\}$ is serenaded-to on every night by a man she strictly prefers over m_i , and would thus reject m_i if he ever serenaded under w 's window. Furthermore, as, by the induction hypothesis, $w_i = \mu'(m_i)$ rejects m_i on night t , we have that he does not serenade under her window on any later night. Thus, $w_{i+1} \in \{\tilde{w}, \mu(m_i)\}$. If $\tilde{w} = \mu(m_i)$, then $w_{i+1} = \tilde{w}$ and by definition of \mathcal{P}_W , she does not blacklist m_i . Thus, the algorithm stops, yielding $d = i + 1$. Thus, w_{i+1} and d agree with the definition of $C_{\mu'}^{\mu}(\tilde{w})$. Otherwise, $\tilde{w} \neq \mu(m_i)$, and thus, by definition of \mathcal{P}_W , \tilde{w} blacklists w , and thus $w_{i+1} \neq \tilde{w}$, and we have $w_{i+1} = \mu(m_i)$. Let t' be the night on which m_i serenades under w_{i+1} 's window for the first time. By the induction hypothesis and by Lemma 4(1), we have that w_{i+1} has yet to reject anyone by night t , and thus by night $t' - 1$, and thus on night t' she rejects the man serenading under her window continuously from the first night, namely $\mu'(w_{i+1})$, and thus $m_{i+1} = \mu'(w_{i+1})$ and the proof by induction is complete. \square

Corollary 3. *Under the conditions of Lemma 7, $\mu'_{\tilde{\mathcal{P}}_W, \mathcal{P}_M}(w) = \mu(w)$ for all women w in $C_{\mu'}^{\mu}(\tilde{w})$ (and thus in particular also for $w = \tilde{w}$), as well as for all w for whom $\mu'(w) = \mu(w)$.*

Proof. By the proof of Lemma 7, or immediately by Lemma 6(3) and by Definition 15(5). \square

Lemma 8. *Let μ' and μ be matchings and let \mathcal{P}_M be a $(\mu' \rightarrow \mu)$ -compatible profile of preference lists for M . Let \mathcal{P}_W be a $(\mu' \xrightarrow{\mathcal{P}_M} \mu)$ -compatible profile of preference lists for W . Let $\tilde{w} \in W$ be a woman who rejects a man \tilde{m} during $R_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$, but s.t. $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(\tilde{w}) \neq \mu(\tilde{w})$. Let $\tilde{\mathcal{P}}_W$ be a profile of preference lists for W , obtained from \mathcal{P}_W by modifying the preference list of \tilde{w} to start with $\mu(\tilde{w})$, followed immediately by \tilde{m} , followed immediately by $\mu'(\tilde{w})$, followed by some or all other men in arbitrary order. $\tilde{\mathcal{P}}_W$ satisfies both of the following.*

1. $\tilde{\mathcal{P}}_W$ is $(\mu' \xrightarrow{\mathcal{P}_M} \mu)$ -compatible.
2. $C_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M} = (w_1 \xrightarrow{m_1} w_2 \dots \xrightarrow{m_{\ell-1}} w_{\ell} \xrightarrow{m_{\ell}} C_{\mu'}^{\mu}(\tilde{w}) \xrightarrow{m_{\ell}} w_{\ell+1} \xrightarrow{m_{\ell+1}} w_{\ell+2} \xrightarrow{m_{\ell+2}} \dots w_d)$, where $(w_1 \xrightarrow{m_1} w_2 \xrightarrow{m_2} \dots w_{d-1} \xrightarrow{m_{d-1}} w_d) \triangleq C_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$ and where $m_{\ell} = \tilde{m}$ is rejected by \tilde{w} during $R_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$ between his rejection by w_{ℓ} and his provisional acceptance by $w_{\ell+1}$.

Proof. It is straightforward to check, immediately from the definition of $\tilde{\mathcal{P}}_W$, that since \mathcal{P}_W and \mathcal{P}_M are μ' -cycle generating, so are $\tilde{\mathcal{P}}_W$ and \mathcal{P}_M (and with the same cycle trigger and w_m 's). Before continuing to prove Part 1, we now prove Part 2.

By Lemma 6(3), $\tilde{w} \notin \{w_i\}_{i=1}^d$; thus, \tilde{w} rejects \tilde{m} during $R_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$ "on his way" between two women w_{ℓ} and $w_{\ell+1}$ for some ℓ , and thus $m_{\ell} = \tilde{m}$. As \tilde{m} does not serenade under \tilde{w} 's window after she rejects him, ℓ is well-defined. We note that by definition of $\tilde{\mathcal{P}}_W$, $R_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ and $R_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$

coincide as long as \tilde{m} does not serenade under \tilde{w} 's window, and thus $C_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ has prefix $(w_1 \xrightarrow{m_1} w_2 \cdots \xrightarrow{m_{\ell-1}} w_\ell \xrightarrow{m_\ell})$.

By Lemma 6(4), and by Lemma 4(2), we have that for every m in $C_{\mu'}^\mu(\tilde{w})$, m is not in $C_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$; thus, such m is never rejected during $R_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$, and thus serenades throughout all nights of $R_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$ under the window of the same woman — the one under whose window m serenades on the first night of $R_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$, namely $\mu'(m)$, by definition.

We now show by induction that $C_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ has prefix $(w_1 \xrightarrow{m_1} w_2 \cdots \xrightarrow{m_{\ell-1}} w_\ell \xrightarrow{m_\ell} C_{\mu'}^\mu(\tilde{w}))$.

Denote $C_{\mu'}^\mu(\tilde{w}) = (w'_1 \xrightarrow{m'_1} w'_2 \xrightarrow{m'_2} \cdots w'_{d'-1} \xrightarrow{m'_{d'-1}} w'_{d'})$.

Base: As explained above, up to the night on which \tilde{m} serenades under \tilde{w} 's window, we have that $C_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ has prefix $(w_1 \xrightarrow{m_1} w_2 \cdots \xrightarrow{m_{\ell-1}} w_\ell \xrightarrow{m_\ell})$. As explained above, the other man serenading under \tilde{w} 's window on that night is $\mu'(\tilde{w})$, and by definition of $\tilde{\mathcal{P}}_W$, she thus rejects him in favour of \tilde{m} during that night of $R_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$. Thus, the next woman in $C_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ is $\tilde{w} = w'_1$, and the next man — $\tilde{m} = m'_1$.

Step: Let $1 < i \leq d'$ and assume that we have shown that $C_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ has prefix $(w_1 \xrightarrow{m_1} w_2 \cdots \xrightarrow{m_{\ell-1}} w_\ell \xrightarrow{m_\ell} w'_1 \xrightarrow{m'_1} \cdots w'_i \xrightarrow{m'_i})$. As explained above, we have $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(m'_i) \neq \mu(m'_i)$, and thus, by Definition 15(3) (since $w_1 \neq \tilde{w}$), we have that m'_i is blacklisted by w_1 . Thus, the next woman in $C_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ is not w_1 . Once more by the same well-known property of the Gale-Shapley algorithm as above, on every night each woman $w \in W \setminus \{w_1\}$ is serenaded-to by a man she weakly prefers over $\mu'(w)$. Thus, by Lemma 6(1), and by Lemma 1, we have that the next woman in $C_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ is $\mu(m'_i) = w'_{i+1}$. If $i = d'$, the proof by induction is complete. Otherwise, let t be the night on which m'_i first serenades under w'_{i+1} 's window during $R_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$. As $m'_{i+1} = \mu'(w'_{i+1})$ is not in $C_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$ and by Lemma 4(1), we have that m'_{i+1} has not been rejected before night t , and thus he is the other man (in addition to m'_i) serenading under $\mu'(w'_{i+1})$'s window on night t . Therefore, by Definition 15(2) (and since $w'_{i+1} \neq \tilde{w}$, by Lemma 4(1)), the next man in $C_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ is m'_{i+1} , and the proof by induction is complete.

Denote by t the night on which \tilde{m} first serenades under \tilde{w} 's window and by t' — the night on which $m'_{d'-1}$ first serenades under her window. As $C_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$ and $C_{\mu'}^\mu(\tilde{w})$ are disjoint, and as $\tilde{m} = m_\ell$ is in $C_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$, it follows that \tilde{m} is not rejected between nights t and t' , exclusive of t' . Thus, and as $w'_{d'} = w'_1 = \tilde{w}$ provisionally accepts \tilde{m} at time t , it follows that he is the man rejected by $w'_{d'}$ in favour of $m'_{d'-1}$ during $R_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$.

So far, we have that $C_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ has prefix $(w_1 \xrightarrow{m_1} w_2 \cdots \xrightarrow{m_{\ell-1}} w_\ell \xrightarrow{m_\ell} C_{\mu'}^\mu(\tilde{w}) \xrightarrow{m_\ell})$. We conclude the proof inductively.

Assume that for some $\ell \leq i < d'$, $C_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ has prefix $(w_1 \xrightarrow{m_1} w_2 \cdots \xrightarrow{m_{\ell-1}} w_\ell \xrightarrow{m_\ell} C_{\mu'}^\mu(\tilde{w}) \xrightarrow{m_\ell} w_\ell \xrightarrow{m_{\ell+1}} \cdots w_i \xrightarrow{m_i})$.

If $i = \ell$, denote $r = \tilde{w} = w'_{d'}$; otherwise, denote $r = w_i$. Let t be the night on which r rejects m_i during $R_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$ and let t' be the night on which r rejects m_i during $R_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$. In order to show that the next woman in $C_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ is w_{i+1} , we have to show that on night t' in $R_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$, every woman that m_i prefers over w_{i+1} but less than r , prefers her provisional match over m_i , while w_{i+1} prefers m_i over her provisional match. Let $w \in W$ s.t. m_i prefers r over w and weakly prefers w over w_{i+1} . If w is in $C_{\mu'}^\mu(\tilde{w})$, then as $C_{\mu'}^\mu(\tilde{w})$ and $C_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$ are distinct, $w \neq w_{i+1}$. Furthermore, as w is in $C_{\mu'}^\mu(\tilde{w})$, then by the induction hypothesis she is provisionally matched on some night before night t' in $R_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ with $\mu(w)$, and by Definition 15(2) and by definition of $\tilde{\mathcal{P}}_W$ prefers him over any other man (including m_i), and is thus, in particular, also provisionally

matched with $\mu(w)$ on night t' . Otherwise, w is not in $C_{\mu'}^\mu(\tilde{w})$ and thus in particular $w \neq \tilde{w}$. As $w \neq \tilde{w}$, she has the same preferences according to $\tilde{\mathcal{P}}_W$ as she does according to \mathcal{P}_W , and in particular m_i serenades under w 's window "on his way" to w_{i+1} (inclusive) during $R_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$, and thus it is enough to show that she has the same provisional match on night t in $R_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$ as she does on night t' in $R_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$. Indeed, if there does not exist $j < i$ s.t. $w = w_j$, then by the induction hypothesis this provisional match on both nights is $\mu'(w)$; otherwise, let $j < i$ be maximal s.t. $w = w_j$ — if $j > 1$ then this provisional match on both nights in both runs is m_{j-1} and otherwise she is not provisionally matched to anyone in either night in either run. Thus, we conclude that the next woman in $C_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ is indeed w_{i+1} .

If $w_{i+1} = \tilde{w}$, then both $R_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$ and $R_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ stop, and thus $i + 1 = d$ and the proof is complete; assume, therefore, that $w_{i+1} \neq \tilde{w}$. Recall that $C_{\mu'}^\mu(\tilde{w})$ and $C_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$ are disjoint; thus, w_{i+1} is not in $C_{\mu'}^\mu(\tilde{w})$ and therefore, as explained above, her provisional match in $R_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ at night t' is the same as in $R_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$ at night t , namely m_{i+1} , and thus he is the next man in $C_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ and the proof by induction is complete.

We now turn to finish proving Part 1. Part 2 of Definition 15 holds for $\tilde{\mathcal{P}}_W$ since it holds for \mathcal{P}_W and by definition of $\tilde{\mathcal{P}}_W$.

To show that Parts 3 to 5 of Definition 15 hold, it is enough to show that for every man $m \in M$ s.t. $\mu'_{\tilde{\mathcal{P}}_W, \mathcal{P}_M}(m) \neq \mu'_{\mathcal{P}_W, \mathcal{P}_M}(m)$, we have $\mu'_{\tilde{\mathcal{P}}_W, \mathcal{P}_M}(m) = \mu(m)$. (To deduce Part 4 of Definition 15, we note that $\mu'_{\tilde{\mathcal{P}}_W, \mathcal{P}_M}(\tilde{m}) = \mu(m)$, as seen above.) Let $m \in M$. If m is in $C_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$, then let i be maximal s.t. $m_i = m$. As explained above, $\mu'_{\tilde{\mathcal{P}}_W, \mathcal{P}_M}(m) = w_i = \mu'_{\mathcal{P}_W, \mathcal{P}_M}(m)$. (And by Lemma 6(3), $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(m) = \mu(m)$ anyway.) Otherwise, if m is in $C_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$, then m is in $C_{\mu'}^\mu(\tilde{w})$, and as shown above, $\mu'_{\tilde{\mathcal{P}}_W, \mathcal{P}_M}(m) = \mu(m)$. Finally, if m is not in $C_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$, then m is not in $C_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$ either, and by Lemma 3, we have $\mu'_{\tilde{\mathcal{P}}_W, \mathcal{P}_M}(m) = \mu'(m) = \mu'_{\mathcal{P}_W, \mathcal{P}_M}(m)$ and the proof is complete. \square

Corollary 4. *Under the conditions of Lemma 8, $\mu'_{\tilde{\mathcal{P}}_W, \mathcal{P}_M}(w) = \mu(w)$ for all women w in $C_{\mu'}^\mu(\tilde{w})$ (and thus in particular also for $w = \tilde{w}$), as well as for all w for whom $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(w) = \mu(w)$.*

Proof. By the proof of Lemma 8, or immediately by Lemma 6(3) and by Definition 15(5). \square

We are now ready to define the inductive result stemming from Lemma 7 as the induction base and from Lemma 8 as the induction step.

Lemma 9. *Let μ' and μ be matchings, let \mathcal{P}_M be a $(\mu' \rightarrow \mu)$ -compatible profile of preference lists for M and let \mathcal{P}_W be a profile of preference lists for W according to which each $w \in W$ has preference list starting with $\mu(w)$, followed immediately by $\mu'(w)$ (if $\mu'(w) \neq \mu(w)$), followed by some or all other men in arbitrary order. For every $\tilde{w} \in W$ s.t. $\mu(\tilde{w}) \neq \mu'(\tilde{w})$, there exists a profile $\tilde{\mathcal{P}}_W$ of preference lists for W s.t. all of the following hold.*

1. $\mu'_{\tilde{\mathcal{P}}_W, \mathcal{P}_M}(\tilde{w}) = \mu(\tilde{w})$, and \tilde{w} still has $\mu(\tilde{w})$ at the top of her preference list, followed only by men who do not serenade under \tilde{w} 's window during $R_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$.
2. Every $w \in W \setminus \{\tilde{w}\}$ whose preference lists according to \mathcal{P}_W and to $\tilde{\mathcal{P}}_W$ differ satisfies all of the following.
 - i. w still has $\mu(w)$ at the top of her preference list.
 - ii. The difference in w 's preference list is only in the promotion of a man m who strictly prefers $\mu'(m)$ over w .
 - iii. $\mu'_{\tilde{\mathcal{P}}_W, \mathcal{P}_M}(w) = \mu(w) \neq \mu'(w)$
3. $\{w \in W \mid \mu'(w) = \mu(w)\} \subsetneq \{w \in W \mid \mu'_{\tilde{\mathcal{P}}_W, \mathcal{P}_M}(w) = \mu(w)\}$.

4. Every $w \in W$ who rejects any man during $R_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ satisfies $\mu'_{\tilde{\mathcal{P}}_W, \mathcal{P}_M}(w) = \mu(w)$.
5. No woman's blacklist differs between \mathcal{P}_W and $\tilde{\mathcal{P}}_W$, except perhaps for that of \tilde{w} , each of the men m she blacklists according to $\tilde{\mathcal{P}}_W$ satisfying $\mu'_{\tilde{\mathcal{P}}_W, \mathcal{P}_M}(m) = \mu(m) \neq \mu'(m)$.

Proof. Let $\tilde{\mathcal{P}}_W$ be the profile of preferences obtained by \mathcal{P}_W by modifying the preference list of \tilde{w} to consist solely of $\mu(\tilde{w})$, with all other men blacklisted. By Lemma 7, $\tilde{\mathcal{P}}_W$ is $(\mu' \xrightarrow{\mathcal{P}_M} \mu)$ -compatible. As long as there exists a woman $w \in W$ who rejects a man during $R_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ but s.t. $\mu'_{\tilde{\mathcal{P}}_W, \mathcal{P}_M}(w) \neq \mu(w)$, we may modify $\tilde{\mathcal{P}}_W$ using Lemma 8 (without modifying w 's blacklist), keeping it $(\mu' \xrightarrow{\mathcal{P}_M} \mu)$ -compatible and only enlarging $C_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$. We thus repeatedly apply Lemma 8 to modify $\tilde{\mathcal{P}}_W$ until no such woman exists. We note that this process stops after at most $\lfloor \frac{n-1}{2} \rfloor$ applications of Lemma 8, as by Corollary 4 the number of women $w \in W$ for whom $\mu'_{\tilde{\mathcal{P}}_W, \mathcal{P}_M}(w) = \mu(w)$ increases by at least 2 following each application of Lemma 8 (and by Corollary 3 equals at least 2 immediately after applying Lemma 7). Finally, we shorten the blacklist of \tilde{w} in $\tilde{\mathcal{P}}_W$ to contain only the men that she actually rejects during $R_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ ($R_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ is thus completely unaffected by this).

By Corollaries 3 and 4, we have that that Parts 1 and 3 hold. (The second part of Part 1 follows from the final stage of the construction of $\tilde{\mathcal{P}}_W$.) We now show that Part 2 holds. Let $w \in W \setminus \{\tilde{w}\}$ whose preference list differs between \mathcal{P}_W and $\tilde{\mathcal{P}}_W$. By our construction, w 's preference list was changed by an application of Lemma 8 with w in the role of \tilde{w} (as defined there), and therefore $\mu'(w) \neq \mu(w)$, and by Corollary 4, also $\mu'_{\tilde{\mathcal{P}}_W, \mathcal{P}_M}(w) = \mu(w)$. Furthermore, the change of w 's preference list is in the promotion to second place of a man \tilde{m} who, by Lemma 8(2), serenades under w 's window during $R_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ following his rejection by another woman; thus, \tilde{m} strictly prefers $\mu'(m)$, under whose window he serenades on the first night of $R_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$, over w , as required.

Part 4 follows by construction, as it was no longer possible to apply Lemma 8 when the construction stopped.

It thus remains to prove Part 5. By construction, the only blacklist that was changed during construction of $\tilde{\mathcal{P}}_W$ from \mathcal{P}_W is indeed that of \tilde{w} . Furthermore, by the last step of the construction, every man blacklisted by \tilde{w} is rejected by her during $R_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$. Thus, every such man is in $C_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ and by Lemma 6(3) the proof is complete. \square

Corollary 5. *Under the conditions of Lemma 9, denote the size of \tilde{w} 's blacklist, according to $\tilde{\mathcal{P}}_W$, by b . At least $b + 1$ men m satisfy $\mu'_{\tilde{\mathcal{P}}_W, \mathcal{P}_M}(m) = \mu(m) \neq \mu'(m)$.*

Proof. By Lemma 9(5), all b distinct men in \tilde{w} 's blacklist (according to $\tilde{\mathcal{P}}_W$) satisfy the required condition. Furthermore, by definition of \tilde{w} and by Lemma 9(1), we have that $\mu'_{\tilde{\mathcal{P}}_W, \mathcal{P}_M}(\tilde{w}) = \mu(\tilde{w}) \neq \mu'(\tilde{w})$, and thus $\mu'_{\tilde{\mathcal{P}}_W, \mathcal{P}_M}(\tilde{w})$, who is by definition not blacklisted by \tilde{w} and is thus distinct from the b men in her blacklist, satisfies the required condition as well. \square

While a considerable amount of reasoning went into proving Lemma 9 (and the lemmas supporting it), and while \mathcal{P}_W is constructed in its proof via an inductive process naively requiring a simulation of a run of the Gale-Shapley algorithm on every iteration, we now show that \mathcal{P}_W can be calculated quite efficiently, by leveraging the theory of cycles developed above. Indeed, without the aid of this theory, Algorithm 1 may seem quite obscure.

Lemma 10. *$\tilde{\mathcal{P}}_W$ from Lemma 9 can be computed in $O(n \cdot k)$ time, where k is the number of men m (alternatively, women) satisfying $\mu'_{\tilde{\mathcal{P}}_W, \mathcal{P}_M}(m) = \mu(m) \neq \mu'(m)$.*

Proof. Algorithm 1 computes $\tilde{\mathcal{P}}_W$, as it is equivalent to the construction of Lemma 9, each time applying Lemma 8 to the woman who performs the earliest rejection out of the women w for who are not eventually-matched with $\mu(w)$.

Algorithm 1 In-place computation of $\tilde{\mathcal{P}}_W$ from Lemma 9, using $O(n)$ additional memory, in $O(n \cdot k)$ time. (All lines run in $O(1)$ time, unless otherwise noted.)

```

1: // See Algorithm 2 for the reason for separation of the initialization.
2: procedure INITIALIZE( $\mu', \mu$ )
3:   // Initialize todo.
4:   todo  $\leftarrow \{w \in W \mid \mu'(w) \neq \mu(w)\}$  // implemented as hash set;  $O(n)$ 
5: end procedure
6: procedure COMPUTELEMMA9( $\mu', \mu, \mathcal{P}_M, \tilde{w}, \tilde{\mathcal{P}}_W$ ) //  $\tilde{\mathcal{P}}_W = \mathcal{P}_W$  when called
7:   // Initialize  $R_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$  simulation.
8:   provisionalMatching  $\leftarrow \mu'$  //  $O(n)$ 
9:   shouldBlacklist  $\leftarrow$  empty hash set // the men to be blacklisted by  $\tilde{w}$ 
10:  currentPref  $\leftarrow$  array indexed by  $M$  // cache, initialized in lines 18 and 42
11:  // Initialize  $\tilde{\mathcal{P}}_W$  as in Lemma 7.
12:  set preference list of  $\tilde{w}$  in  $\tilde{\mathcal{P}}_W$  to  $(\mu(\tilde{w}))$  // not required for correctness
13:  // Simulate  $R_{\mu'}^{\tilde{\mathcal{P}}_W, \mathcal{P}_M}$ , adjusting  $\tilde{\mathcal{P}}_W$  as needed.
14:  insert  $m$  into shouldBlacklist
15:  provisionalMatching $[\tilde{w}] \leftarrow \emptyset$  // instead of  $m$ ; not required for correctness
16:   $m \leftarrow \mu'(\tilde{w})$  //  $m$  is the man currently "on the move"
17:  MARKASDONECYCLEOF( $\tilde{w}$ ) // see line 53
18:  currentPref $[m] \leftarrow$  pointer to  $m$ 's preference list right after  $\tilde{w}$  //  $O(n)$ 
19:  while  $m \neq \mu(\tilde{w})$  or deref currentPref $[m] \neq \tilde{w}$  do //  $< k \cdot (n - 1)$  iterations
20:    // Loop Invariant:
21:    // todo =  $\{w \in W \mid \mu'_{\tilde{\mathcal{P}}_W, \mathcal{P}_M}(w) \neq \mu(w)\}$ .
22:     $w \leftarrow$  deref currentPref $[m]$  // the woman currently approached by  $m$ 
23:    advance currentPref $[m]$ 
24:    let provisionallyAccept be a boolean // should  $w$  provisionally accept  $m$ ?
25:    if  $w = \tilde{w}$  then // by the loop condition,  $m \neq \mu(\tilde{w})$ 
26:      insert  $m$  into shouldBlacklist
27:      provisionallyAccept  $\leftarrow$  false
28:    else if  $m = \mu(w)$  then
29:      provisionallyAccept  $\leftarrow$  true
30:      discard currentPref $[m]$ 
31:    else if  $w \in$  todo then //  $\leq \lfloor \frac{k-1}{2} \rfloor$  times throughout the run
32:      // Update  $\tilde{\mathcal{P}}_W$  as in Lemma 8.
33:      promote  $m$  to be second on  $w$ 's preference list in  $\tilde{\mathcal{P}}_W$  //  $O(n)$ 
34:      MARKASDONECYCLEOF( $w$ ) // see line 53
35:      provisionallyAccept  $\leftarrow$  true
36:    else
37:      provisionallyAccept  $\leftarrow$  false
38:    end if
39:    if provisionallyAccept then
40:      swap values of  $m$  and provisionalMatching $[w]$ 
41:      if undefined currentPref $[m]$  then //  $< k$  times throughout the run
42:        currentPref $[m] \leftarrow$  pointer to  $m$ 's preference list right after  $w$  //  $O(n)$ 
43:      end if
44:    end if // else, do not provisionally accept (i.e.  $m$  is rejected by  $w$ )

```

```

45:   end while
46:   provisionalMatching[ $\tilde{w}$ ]  $\leftarrow m$ 
47:   // Update  $\tilde{w}$ 's blacklist per the final step of the construction from Lemma 9.
48:   set  $\tilde{w}$ 's blacklist to be shouldBlacklist (keeping  $\mu(\tilde{w})$  most-preferred) //  $O(n)$ 
49:   // At this point, provisionalMatching =  $\mu'_{\mathcal{P}_W, \mathcal{P}_M}$ .
50:   return provisionalMatching
51: end procedure
52: // Removes all women in  $C_{\mu'}^\mu(w)$  from todo.
53: procedure MARKASDONECYCLOF( $w$ )
54:    $w' \leftarrow w$ 
55:   repeat //  $\leq n$  iterations throughout the run of COMPUTELEMMA9
56:     remove  $w'$  from todo
57:      $w' \leftarrow \mu(\mu'(w'))$ ;
58:   until  $w' = w$ 
59: end procedure

```

We note that as there are at most k rejected men (by Lemma 6(3)), each newly-serenading at most $n - 1$ times, the loop starting at line 19 indeed iterates less than $k \cdot (n - 1)$ times. \square

Before proving Theorem 4, we first prove a simpler special case thereof, in which each man approaches a distinct woman on the first night, e.g. as in Example 1. (We reference and reuse significant portions of this proof in the more-intricate proof of the general case of Theorem 4 below.)

Theorem 11. *Theorem 4 holds for the special case in which each man approaches a distinct woman on the first night. (In this case, in the absence of nonempty blacklists, the algorithm would terminate at the end of the first night.) Furthermore, in this case \mathcal{P}_W can be computed online in $O(n^2)$ time.*

Proof. Let \mathcal{P}_M and μ be as in Theorem 4, but s.t. the top choices of all men are distinct. We inductively define a sequence of profiles $(\mathcal{P}_W^i)_{i=1}^d$ of preference lists for W and a sequence of matchings $(\mu^i)_{i=1}^d$, for d to be defined below, satisfying the following properties for every $i \geq 1$.

1. $\mu^i = \text{MenOpt}(\mathcal{P}_W^i, \mathcal{P}_M)$.
2. If $i > 1$, then $\{w \in W \mid \mu^{i-1}(w) = \mu(w)\} \subsetneq \{w \in W \mid \mu^i(w) = \mu(w)\}$.
3. \mathcal{P}_M is $(\mu^i \rightarrow \mu)$ -compatible.
4. Each $w \in W$ has $\mu(w)$ first on her preference list according to \mathcal{P}_W^i .
5. Each $w \in W$ for whom $\mu^i(w) \neq \mu(w)$ has $\mu^i(w)$ second on her preference list (after $\mu(w)$) according to \mathcal{P}_W^i .
6. Each $w \in W$ for whom $\mu^i(w) \neq \mu(w)$ does not reject any man during $R^{\mathcal{P}_W^i, \mathcal{P}_M}$.
7. No man appears in more than one blacklist, and every blacklisted man $m \in M$ satisfies $\mu^i(m) = \mu(m)$.
8. For each $w \in W$ who has a nonempty blacklist, $\mu(w)$ appears in no blacklist and satisfies $\mu(\mu(w)) = \mu^i(\mu(w))$.

Base: Denote by μ^1 the (provisional) matching describing the first night of the algorithm. As μ is M -rational, and as each $m \in M$ weakly prefers $\mu^1(m)$ over any $w \in W$, \mathcal{P}_M is $(\mu^1 \rightarrow \mu)$ -compatible. Denote by \mathcal{P}_W^1 a profile of preference lists for W according to which each $w \in W$ has preference list starting with $\mu(w)$, followed immediately by $\mu^1(w)$ (if $\mu^1(w) \neq \mu(w)$), followed by all other men in arbitrary order. We note that as all blacklists in \mathcal{P}_W are empty, no rejections occur in $R^{\mathcal{P}_W^1, \mathcal{P}_M}$. In particular, we thus have $\text{MenOpt}(\mathcal{P}_W^1, \mathcal{P}_M) = \mu^1$. Properties 7 and 8 follow immediately as all the blacklists in \mathcal{P}_W^1 are empty.

Step: Assume that μ^i and \mathcal{P}_W^i have been defined for some $i \geq 0$. If $\mu^i = \mu$, we conclude the inductive process and set $d \triangleq i$ and $\mathcal{P}_W \triangleq \mathcal{P}_W^d$. Otherwise, let \tilde{w}^{i+1} be a woman s.t. $\mu^i(\tilde{w}^{i+1}) \neq \mu(\tilde{w}^{i+1})$. We denote by \mathcal{P}_W^{i+1} the profile of preference lists for W obtained by use of Lemma 9 when applied to $\mu, \mathcal{P}_W^i, \mathcal{P}_M, \mu' \triangleq \mu^i$ and $\tilde{w} \triangleq \tilde{w}^{i+1}$. Furthermore, we set

$\mu^{i+1} \triangleq \mu_{\mathcal{P}_W^{i+1}, \mathcal{P}_M}^i$. Lemma 9 is indeed applicable due to Properties 3 to 5 for i . We now show that Properties 1 to 8 hold for $i + 1$.

Consider the following timing for $R^{\mathcal{P}_W^{i+1}, \mathcal{P}_M}$: In the first part of the run, denoted by $R_{1/2}^{i+1}$, the algorithm runs normally, except that the rejection of $\mu^i(\tilde{w}^{i+1})$ by \tilde{w}^{i+1} does not (yet) take place. In the second part of the run, denoted by $R_{2/2}^{i+1}$, after $R_{1/2}^{i+1}$ converges, \tilde{w}^{i+1} 's rejection of $\mu^i(\tilde{w}^{i+1})$ takes place and the algorithm runs until it converges once again. As the outcome of the Gale-Shapley algorithm is invariant under timing changes (Dubins and Freedman, 1981), this timing also yields $\text{MenOpt}(\mathcal{P}_W^{i+1}, \mathcal{P}_M)$. We claim that $R_{1/2}^{i+1}$ is indistinguishable from $R^{\mathcal{P}_W^i, \mathcal{P}_M}$, while $R_{2/2}^{i+1}$ is indistinguishable from $R_{\mu^i}^{\mathcal{P}_W^{i+1}, \mathcal{P}_M}$. By Property 6 for i , \tilde{w}^{i+1} does not reject any man during $R^{\mathcal{P}_W^i, \mathcal{P}_M}$, and thus any change to her preference list has no effect on $R^{\mathcal{P}_W^i, \mathcal{P}_M}$, as long as any blacklist-induced rejection of $\mu^i(\tilde{w}^{i+1})$ (the only man serenading under \tilde{w}^{i+1} 's window during $R^{\mathcal{P}_W^i, \mathcal{P}_M}$) is deferred. By Lemma 9(2), the only difference in the preference list of any other woman $w \neq \tilde{w}^{i+1}$ between \mathcal{P}_W^i (which yields $R^{\mathcal{P}_W^i, \mathcal{P}_M}$) and \mathcal{P}_W^{i+1} (which yields $R_{1/2}^{i+1}$), is in the possible promotion of men m who prefer $\mu^i(m)$ over w . As by Property 1 for i , such men never serenade under w 's window during $R^{\mathcal{P}_W^i, \mathcal{P}_M}$, such a change has no effect on $R^{\mathcal{P}_W^i, \mathcal{P}_M}$ either. We thus indeed have that $R_{1/2}^{i+1}$ is indistinguishable from $R^{\mathcal{P}_W^i, \mathcal{P}_M}$, and hence concludes with the provisional matching $\text{MenOpt}(\mathcal{P}_W^i, \mathcal{P}_M) = \mu^i$; thus, $R_{2/2}^{i+1}$ is indistinguishable from $R_{\mu^i}^{\mathcal{P}_W^{i+1}, \mathcal{P}_M}$, by definition of the latter. Therefore, $\text{MenOpt}(\mathcal{P}_W^{i+1}, \mathcal{P}_M) = \mu_{\mathcal{P}_W^{i+1}, \mathcal{P}_M}^i$, and Property 1 holds for $i + 1$.

Property 2 for $i+1$ follows directly from Lemma 9(3). By Property 4 for i , and by Lemma 9(1 and 2), we have that Property 4 holds for $i + 1$ as well. Property 3 for $i + 1$ follows from Property 3 for i and from Property 4 for $i + 1$, by Lemma 5 and as $\mu^{i+1} = \mu_{\mathcal{P}_W^{i+1}, \mathcal{P}_M}^i$.

Let $w \in W$ s.t. $\mu^{i+1}(w) \neq \mu(w)$; thus, by Lemma 9(1), $w \neq \tilde{w}^{i+1}$. By Property 2 for $i + 1$, $\mu^i(w) \neq \mu(w)$ as well. Thus, by Lemma 9(2), we have that Property 5 for $i + 1$ follows from Property 5 for i . Property 6 for $i + 1$ follows from Property 6 for i (for $R_{1/2}^{i+1}$), and from Lemma 9(4) (for $R_{2/2}^{i+1}$).

To prove Property 7 for $i + 1$ from Property 7 for i , we must show that every newly-blacklisted man m is newly-blacklisted by exactly one woman, is not blacklisted in \mathcal{P}_W^i , and satisfies $\mu^{i+1}(m) = \mu(m)$. Indeed, by Lemma 9(5), any newly-blacklisted man m is newly-blacklisted only by \tilde{w}^{i+1} and satisfies $\mu^{i+1}(m) = \mu(m) \neq \mu^i(m)$. Thus, in particular, by Property 7 for i , any such m is not blacklisted in \mathcal{P}_W^i .

Finally, since by Lemma 9(5) no blacklist is changed but \tilde{w}^{i+1} 's, in order to prove Property 8 for $i + 1$ given Property 8 for i , it is enough to show that no woman blacklists $\mu(\tilde{w}^{i+1})$ in \mathcal{P}_W^i , and that there exists no woman w with a nonempty blacklist in \mathcal{P}_W^i s.t. $\mu(w)$ is blacklisted by \tilde{w}^{i+1} in \mathcal{P}_W^{i+1} . The former holds since by definition of \tilde{w}^{i+1} we have $\mu^i(\mu(\tilde{w}^{i+1})) \neq \tilde{w}^{i+1}$, and thus by Property 7 for i , $\mu(\tilde{w}^{i+1})$ is not blacklisted in \mathcal{P}_W^i . To show the latter, let w be a woman with a nonempty blacklist in \mathcal{P}_W^i . By Property 8 for i , $\mu^i(\mu(w)) = \mu(\mu(w))$, and thus, by Lemma 9(5), $\mu(w)$ is not blacklisted by \tilde{w}^{i+1} in \mathcal{P}_W^{i+1} and the proof of Property 8 for $i + 1$ is complete. Thus, the proof by induction is complete, as the process stops by Property 2, by finiteness of W .

We conclude the proof by showing that $\mathcal{P}_W = \mathcal{P}_W^d$ satisfies the conditions of Theorem 4. By definition of d , $\mu^d = \mu$; thus, by Property 1 for $i = d$, $\text{MenOpt}(\mathcal{P}_W, \mathcal{P}_M) = \mu$. Furthermore, the W -optimal stable matching is also μ , by Property 4 for $i = d$. As the W -optimal and M -optimal stable matchings coincide, and as the latter is also the W -worst stable matching, these constitute the unique stable matching.

By Property 7 for $i = d$, each man is blacklisted by at most one woman in \mathcal{P}_W . Let w_1, \dots, w_{n_b} be the women with nonempty blacklists in \mathcal{P}_W . By Property 8 for $i = d$, $m_1 =$

$\mu(w_1), \dots, m_{n_b} = \mu(w_{n_b})$ constitute n_b distinct men not blacklisted in \mathcal{P}_W . Hence, there exist at most $|M| - n_b = n - n_b$ men blacklisted in \mathcal{P}_W . Thus, and since every man is blacklisted by at most one woman, we have that the combined size of all blacklists is at most $n - n_b$. As each of w_1, \dots, w_{n_b} has a blacklist of size at least 1, we have that the combined size of all blacklists is at least n_b . Combining these, we obtain $n_b \leq n - n_b$, yielding $n_b \leq \lfloor \frac{n}{2} \rfloor$ and completing the proof.

Algorithm 2 follows the above construction for computing \mathcal{P}_W . This algorithm simulates the run $R^{\mathcal{P}_W, \mathcal{P}_M}$ (or more precisely, an equivalent run of the implementation of the Gale-Shapley algorithm proposed by Dubins and Freedman, 1981), while building \mathcal{P}_W online; i.e. the choice of who acts next and the decisions of whom each woman prefers or blacklists at any step depend solely on the history of the run (and not on the not-yet-acted-upon suffixes of the men's preference lists). Thus, if participants are not required to submit their preference lists in advance, but rather only to dynamically act upon them, then in Dubins and Freedman's implementation of the algorithm, if the women can control its scheduling, Algorithm 2 constitutes a strategy for the women that forces μ against *every* profile \mathcal{P}_M of preference lists for M .

Algorithm 2 On-line computation of \mathcal{P}_W from Theorem 4, for the special case of Theorem 11, using $O(n)$ additional memory, in $O(n^2)$ time. (All lines run in $O(1)$ time, unless otherwise noted.)

```

1: procedure COMPUTETHEOREM4SPECIALCASE( $\mathcal{P}_M, \mu$ )
2:   // Initialize  $\mathcal{P}_W$  to  $\mathcal{P}_W^1$  and  $\mu'$  to  $\mu^1$ .
3:    $\mu' \leftarrow$  matching between  $W$  and  $M$  // implemented e.g. using two arrays
4:   for  $m \in M$  do //  $n$  iterations
5:     match  $m$  in  $\mu'$  with its most-preferred woman according to  $\mathcal{P}_M$ 
6:   end for
7:    $\mathcal{P}_W \leftarrow$  profile of preference lists for  $W$ 
8:   for  $w \in W$  do //  $n$  iterations
9:     set preference list of  $w$  in  $\mathcal{P}_W$  to  $(\mu(w), \mu'(w), \text{all other men})$  //  $O(n)$ 
10:  end for
11:  // Initialize todo.
12:  INITIALIZE( $\mu', \mu$ ) // see Algorithm 1;  $O(n)$ 
13:  // Iterate until for every women  $w$  we have  $\mu(w) = \mu'(w)$ .
14:  while todo  $\neq \emptyset$  do
15:    let  $\tilde{w} \in \text{todo}$ 
16:    // Update  $\mathcal{P}_W$  from  $\mathcal{P}_W^i$  to  $\mathcal{P}_W^{i+1}$  and  $\mu'$  from  $\mu^i$  to  $\mu^{i+1}$ .
17:     $\mu' \leftarrow$  COMPUTELEMMA9( $\mu', \mu, \mathcal{P}_M, \tilde{w}, \mathcal{P}_W$ ) // see Algorithm 1
18:  end while
19:  return  $\mathcal{P}_W$ 
20: end procedure

```

By Lemma 10, each call to COMPUTELEMMA9 takes $O(n \cdot k)$ time, where k is the number of women w for whom $\mu'(w)$ is "fixed" from $\mu^1(w)$ to $\mu(w)$ by this call. Thus, all calls to COMPUTELEMMA9 take $O(n^2)$ time in total. \square

The basic ideas of the proof of Theorem 11 are used in the proof that we now give for Theorem 4 as well, however this proof is somewhat more delicate, as we must choose \tilde{w}^{i+1} carefully, and more intricately analyse the inductive step in certain cases.

Proof of Theorem 4. Let \mathcal{P}_M and μ be as in Theorem 4. We once again inductively define a sequence of profiles $(\mathcal{P}_W^i)_{i=1}^d$ of preference lists for W and a sequence of matchings $(\mu^i)_{i=1}^d$, for d to be defined below, satisfying the following, slightly modified, properties for every $i \geq 1$. (Note the omission of Property 6, which causes us significant hardship below, and the addition of Property 9.)

1. $\mu^i = \text{MenOpt}(\mathcal{P}_W^i, \mathcal{P}_M)$.
2. If $i > 1$, then $\{w \in W \mid \mu^{i-1}(w) = \mu(w)\} \subsetneq \{w \in W \mid \mu^i(w) = \mu(w)\}$.
3. \mathcal{P}_M is $(\mu^i \rightarrow \mu)$ -compatible.
4. Each $w \in W$ has $\mu(w)$ first on her preference list according to \mathcal{P}_W^i .
5. Each $w \in W$ for whom $\mu^i(w) \neq \mu(w)$ has $\mu^i(w)$ second on her preference list (after $\mu(w)$) according to \mathcal{P}_W^i .
7. No man appears in more than one blacklist, and every blacklisted man $m \in M$ satisfies $\mu^i(m) = \mu(m)$.
8. For each $w \in W$ who has a nonempty blacklist, $\mu(w)$ appears in no blacklist and satisfies $\mu(\mu(w)) = \mu^i(\mu(w))$.
9. For each $w \in W$ who is serenaded-to during $R^{\mathcal{P}_W^i, \mathcal{P}_M}$ solely by $\mu(w)$, $\mu(w)$ appears in no blacklist.

Base: Denote by \mathcal{P}_W^0 a profile of preference lists for W according to which each $w \in W$ has preference list starting with $\mu(w)$, followed by all other men in arbitrary order. Define $\mu^1 = \text{MenOpt}(\mathcal{P}_W^0, \mathcal{P}_M)$. We note that all men are matched in μ^1 . Indeed, let $m \in M$; similarly to the proof of Lemma 1, since by definition m does not blacklist $\mu(w)$, and since $\mu(w)$'s top choice is m , then m must be matched by μ^1 . Furthermore, we thus obtain that m weakly prefers $\mu^1(m)$ over $\mu(m)$, and as μ is by definition M -rational, we have that \mathcal{P}_M is $(\mu^1 \rightarrow \mu)$ -compatible. Denote by \mathcal{P}_W^1 the profile of preference lists for W obtained from \mathcal{P}_W^0 by promoting, for each woman w s.t. $\mu^1(w) \neq \mu(w)$, $\mu^1(w)$ to be second on w 's preference list (immediately following $\mu(w)$). Since by definition each such woman w never rejects $\mu^1(w)$ during $R^{\mathcal{P}_W^0, \mathcal{P}_M}$, we have that $R^{\mathcal{P}_W^0, \mathcal{P}_M}$ and $R^{\mathcal{P}_W^1, \mathcal{P}_M}$ are indistinguishable. Thus, $\mu^1 = \text{MenOpt}(\mathcal{P}_W^1, \mathcal{P}_M)$. Similarly to the induction base in the proof of Theorem 11, Properties 7 to 9 follow immediately as all blacklists exist in \mathcal{P}_W^1 are empty. We note that unlike the scenario analysed in the proof of Theorem 11, we generally have that Property 6 from that proof simply does not hold here, even during $R^{\mathcal{P}_W^0, \mathcal{P}_M}$.

Step: Assume that μ^i and \mathcal{P}_W^i have been defined for some $i \geq 0$. If $\mu^i = \mu$, then we conclude the inductive process and set $d \triangleq i$ and $\mathcal{P}_W \triangleq \mathcal{P}_W^d$. Otherwise, $T^i \triangleq \{w \in W \mid \mu^i(w) \neq \mu(w)\} \neq \emptyset$. For every woman w , denote by $t^i(w)$ the last night of $R^{\mathcal{P}_W^i, \mathcal{P}_M}$ on which any man newly-serenades under w 's window. Let $\tilde{w}^{i+1} \in T^i$ with largest $t^i(\tilde{w}^{i+1})$. (If more than one such woman exists, choose one of them arbitrarily.) We denote by \mathcal{P}'_W the profile of preference lists for W obtained by use of Lemma 9⁹ when applied to $\mu, \mathcal{P}_W^i, \mathcal{P}_M, \mu' \triangleq \mu^i$ and $\tilde{w} \triangleq \tilde{w}^{i+1}$. As in the proof of Theorem 11, Lemma 9 is applicable due to Properties 3 to 5 for i . We define $\mu^{i+1} \triangleq \mu'_{\mathcal{P}'_W, \mathcal{P}_M}$, and consider two separate cases when defining \mathcal{P}_W^{i+1} :

- If on night $t^i(\tilde{w}^{i+1})$ of $R^{\mathcal{P}_W^i, \mathcal{P}_M}$ only one man serenades under \tilde{w}^{i+1} 's window, then this night is the first on which any man serenades under her window during $R^{\mathcal{P}_W^i, \mathcal{P}_M}$; thus, \tilde{w}^{i+1} rejects no man before or on that night, and thus, by definition of \tilde{w}^{i+1} , she rejects no man during $R^{\mathcal{P}_W^i, \mathcal{P}_M}$.¹⁰ In this case, as in the inductive step of Theorem 11, we set $\mathcal{P}_W^{i+1} \triangleq \mathcal{P}'_W$ (and thus $\mu^{i+1} = \mu'_{\mathcal{P}'_W, \mathcal{P}_M}$). The proofs of Properties 1 to 5, 7 and 8 for

⁹We note that this proof, unlike that of Theorem 11, uses Lemma 9(4) solely to show that for every $w \in W$, $\mu^{i+1}(w) \in \{\mu^i(w), \mu(w)\}$. Thus, the use of Lemma 9 could be replaced by the use of a simple variant of Lemma 7, rendering Lemmas 8 and 9 unneeded. Nonetheless, we present these lemmas for several reasons. First, since they provide insight into the structure underlying our solution, and also allow for both a much cleaner and more structured proof of Theorem 11 and a more gradual introduction of the methods used in the second case of the current proof; second, since they yield an online algorithm for the case of Theorem 11, and since using Lemma 9 in the current proof provides for improved running-time of the resulting algorithm for many inputs, by turning iterations of the second case of this proof into (faster) applications of Lemma 8 as iterations in the proof of Lemma 9.

¹⁰An implementation note: in fact, any woman who never rejects any man during $R^{\mathcal{P}_W^i, \mathcal{P}_M}$ may be chosen as \tilde{w}^{i+1} and treated as detailed in this case, regardless of whether she has the largest t^i value or not (this condition is only needed in the second case, in which \tilde{w}^{i+1} rejects men during $R^{\mathcal{P}_W^i, \mathcal{P}_M}$), resulting in an asymptotically-faster algorithm for many inputs. The reason we nonetheless describe our proof by choosing \tilde{w}^{i+1} in the same manner on both cases above is for the sake of simplicity, in order to better reflect the analogies between these two cases.

$i + 1$ follow exactly as in the inductive step in the proof of Theorem 11 — we note that these use only Properties 1 to 5, 7 and 8 for i , and additionally Property 6 for i , but only to show that \tilde{w}^{i+1} rejects no man during $R^{\mathcal{P}_W^i, \mathcal{P}_M}$, which we now have by definition of this case, as explained above. It thus remains to show that Property 9 holds for $i + 1$. By equivalence of $R_{1/2}^{i+1}$ and $R^{\mathcal{P}_W^i, \mathcal{P}_M}$, any woman w serenaded-to during $R^{\mathcal{P}_W^{i+1}, \mathcal{P}_M}$ solely by $\mu(w)$ is also serenaded-to during $R^{\mathcal{P}_W^i, \mathcal{P}_M}$ solely by him. Thus, Property 9 for $i + 1$ follows from Property 9 for i and from Lemma 9(5) (as $\mu^i(\mu(w)) = \mu(\mu(w))$ for any such w).

- Otherwise, on night $t^i(\tilde{w}^{i+1})$ of $R^{\mathcal{P}_W^i, \mathcal{P}_M}$, by definition of $t^i(\cdot)$, more than one man serenades under \tilde{w}^{i+1} 's window. Denote by \tilde{m} a man rejected by her on that night. (If more than one such man exists, choose \tilde{m} arbitrarily among all such men.) Note that by definition of $t^i(\tilde{w}^{i+1})$, the man serenading under \tilde{w}^{i+1} 's window without being rejected by her on that night is $\mu^i(\tilde{w}^{i+1})$. Consider the following timings for $R^{\mathcal{P}_W^i, \mathcal{P}_M}$: In the first part of the run, denoted by $\tilde{R}_{1/2}^i$, the algorithm runs normally, except that the rejection of \tilde{m} by \tilde{w}^{i+1} in favour of $\mu^i(\tilde{w}^{i+1})$ does not take place (but on the other hand, neither does she reject $\mu^i(\tilde{w}^{i+1})$). Thus, by definition of \tilde{w}^{i+1} , when $\tilde{R}_{1/2}^i$ converges, two men (\tilde{m} and $\mu^i(\tilde{w}^{i+1})$) serenade under \tilde{w}^{i+1} 's window, while each of the remaining $n - 2$ men is the only man serenading under the window of some woman. Thus, at the end of $\tilde{R}_{1/2}^i$, there exists a unique woman \hat{w} , under whose window no man serenades. We note that by definition of $t^i(\tilde{w}^{i+1})$, we have that $\hat{w} \notin T^i$. In the second part of the run, denoted by $\tilde{R}_{2/2}^i$, after $\tilde{R}_{1/2}^i$ converges, \tilde{w}^{i+1} 's rejection of \tilde{m} in favour of $\mu^i(\tilde{w}^{i+1})$ takes place and the algorithm runs until it converges once more. Once more, as the outcome of the Gale-Shapley algorithm is invariant under timing changes (Dubins and Freedman, 1981), this timing also yields $MenOpt(\mathcal{P}_W^i, \mathcal{P}_M)$.

We denote by \mathcal{P}_W^{i+1} the profile of preference lists for W obtained by applying the following modifications to \mathcal{P}_W^i :

1. Along the lines of Lemma 8, we set the preference list of \tilde{w}^{i+1} to be as in \mathcal{P}_W^i , but with \tilde{m} promoted to the second place (immediately following $\mu(\tilde{w}^{i+1})$), maintaining the internal order of all other men.
2. For each $m \in M$ who satisfies $\mu(m) = \mu^{i+1}(m) \neq \mu^i(m)$, perform the following for each woman $w \notin T^i$ s.t. m (strictly) prefers $\mu^i(m)$ over w and w over $\mu(m)$:
 - i. If w has no provisional match at the end of $\tilde{R}_{1/2}^i$ (i.e. if $w = \hat{w}$), then, once again somewhat along the lines of Lemma 8, alter \hat{w} 's preferences to blacklist m .
 - ii. Otherwise, demote m on w 's preference list to somewhere below the man with whom w is provisionally matched at the end of $\tilde{R}_{1/2}^i$.

We note that Modification 2 is well-defined. Indeed, since for every such m , $\mu^i(m) \neq \mu(m)$, and thus $\mu^i(m) \in T^i$. Thus, by definition of $\tilde{R}_{1/2}^i$, m is provisionally matched with $\mu^i(m)$ at the end of $\tilde{R}_{1/2}^i$. Thus, m is not provisionally matched with any $w \notin T^i$ at the end of $\tilde{R}_{1/2}^i$, and thus Modification 2 does not specify the demotion of any man “below himself” in any preference list.

Consider the following timings, which we soon show to be well-defined, for $R^{\mathcal{P}_W^{i+1}, \mathcal{P}_M}$: In the first part of the run, denoted by $R_{1/3}^{i+1}$, the algorithm runs normally, except that the rejection of $\mu^i(\tilde{w}^{i+1})$ by \tilde{w}^{i+1} in favour of \tilde{m} does not take place (but on the other hand, neither does she reject \tilde{m}). In the second part of the run, denoted by $R_{2/3}^{i+1}$, after $R_{1/3}^{i+1}$ converges, \tilde{w}^{i+1} 's rejection of $\mu^i(\tilde{w}^{i+1})$ in favour of \tilde{m} takes place and the algorithm runs normally once more, except that \tilde{m} 's rejection by \tilde{w}^{i+1} in favour of $\mu(\tilde{w}^{i+1})$ does not take place (but on the either hand, neither does she reject $\mu(\tilde{w}^{i+1})$). In the third and final part of the run, denoted by $R_{3/3}^{i+1}$, after $R_{2/3}^{i+1}$ converges, the algorithm runs

until it converges. We now show that this timing is indeed well-defined, that $R_{1/3}^{i+1}$ is indistinguishable from $\tilde{R}_{1/2}^i$, that $R_{2/3}^{i+1}$ has the same rejections as $R_{\mu^i}^{\mathcal{P}'_W, \mathcal{P}_M}$, and that $R_{3/3}^{i+1}$ has the same rejections as $\tilde{R}_{2/2}^i$. (The scenario analysed above, in which on night $t^i(\tilde{w}^{i+1})$ only one man serenades under \tilde{w}^{i+1} 's window, may be viewed in a sense as a special case of the scenario discussed here, in which $\tilde{R}_{2/2}^i$ and $R_{3/3}^{i+1}$ are both empty.) From this timing being well-defined, we yet again have, as the outcome of the Gale-Shapley algorithm is invariant under timing changes (Dubins and Freedman, 1981), that this timing also yields $MenOpt(\mathcal{P}_W^{i+1}, \mathcal{P}_M)$.

We indeed start by showing that $R_{1/3}^{i+1}$ is identical to $\tilde{R}_{1/2}^i$. Modification 1 (when compared with the unmodified \mathcal{P}_W^i) has no effect on $R^{\mathcal{P}'_W, \mathcal{P}_M}$ before \tilde{m} is rejected by \tilde{w}^{i+1} in that run, and thus has no effect on $\tilde{R}_{1/2}^i$. As in the proof of Theorem 11, by Lemma 9(2), the only difference in the preference list of any other woman $w \neq \tilde{w}^{i+1}$ between \mathcal{P}_W^i and \mathcal{P}'_W is in the possible promotion of men m who prefer $\mu^i(m)$ over w . As by Property 1 such men never serenade under w 's window during $R^{\mathcal{P}'_W, \mathcal{P}_M}$, such a change has no effect on $R^{\mathcal{P}'_W, \mathcal{P}_M}$, and thus on $\tilde{R}_{1/2}^i$ in particular. Let m be a man satisfying the conditions of Modification 2; as Modification 2 only alters the preference, regarding m , of women that m prefers less than $\mu^i(m)$, who, as explained above, is his provisional match at the end of $\tilde{R}_{1/2}^i$, m does not serenade under any of their windows during $\tilde{R}_{1/2}^i$ and thus this modification has no effect on $\tilde{R}_{1/2}^i$ either. By all of the above, as long as \tilde{w}^{i+1} does not reject \tilde{m} nor $\mu^i(\tilde{w}^{i+1})$, there is no difference between $R^{\mathcal{P}'_W, \mathcal{P}_M}$ and $R^{\mathcal{P}_W^{i+1}, \mathcal{P}_M}$. By definition of \tilde{m} and $\mu^i(\tilde{w}^{i+1})$, deferring the rejection of \tilde{m} by \tilde{w}^{i+1} achieves this effect in $R^{\mathcal{P}'_W, \mathcal{P}_M}$; by Modification 1, deferring the rejection of $\mu^i(\tilde{w}^{i+1})$ by \tilde{w}^{i+1} achieves this effect in $R^{\mathcal{P}_W^{i+1}, \mathcal{P}_M}$. Thus, we have shown that $R_{1/3}^{i+1}$ is well-defined and indistinguishable from $\tilde{R}_{1/2}^i$.

We move on to showing that $R_{2/3}^{i+1}$ has the same rejections as $R_{\mu^i}^{\mathcal{P}'_W, \mathcal{P}_M}$. By definition of \tilde{w}^{i+1} , each $w \in T^i \setminus \{\tilde{w}^{i+1}\}$ is provisionally matched with $\mu^i(w)$ (with whom she is provisionally matched on the first night of $R_{\mu^i}^{\mathcal{P}'_W, \mathcal{P}_M}$) at the end of $\tilde{R}_{1/2}^i$, and thus also at the end of $R_{1/3}^{i+1}$. Furthermore, none of Modifications 1 and 2 apply to such w . By Lemma 9(1), \tilde{w}^{i+1} immediately rejects all men approaching her during during $R_{\mu^i}^{\mathcal{P}'_W, \mathcal{P}_M}$, except for $\mu(\tilde{w}^{i+1})$. By Modification 1, as \tilde{w}^{i+1} is provisionally matched with \tilde{m} on the first night of $R_{2/3}^{i+1}$, the only man she would not immediately reject after that night (except for \tilde{m}) is $\mu(\tilde{w}^{i+1})$. Furthermore, since during $R_{2/3}^{i+1}$ (as during $R_{\mu^i}^{\mathcal{P}'_W, \mathcal{P}_M}$), only one woman is newly-serenaded-to every night after the first, we therefore have that the $R_{2/3}^{i+1}$ would stop immediately if $\mu(\tilde{w}^{i+1})$ serenades under \tilde{w}^{i+1} 's window (as happens on the last night of $R_{\mu^i}^{\mathcal{P}'_W, \mathcal{P}_M}$). Finally, let m be a man rejected during $R_{\mu^i}^{\mathcal{P}'_W, \mathcal{P}_M}$, and let $w \notin T^i$ under whose window m serenades during $R_{\mu^i}^{\mathcal{P}'_W, \mathcal{P}_M}$. As $\mu^i(w) = \mu(w)$, and as by definition $\mu^i(m) \neq \mu(m)$, we have $\mu^i(m) \neq w$, and thus m serenades under w 's window during $R_{\mu^i}^{\mathcal{P}'_W, \mathcal{P}_M}$ following a rejection, and thus in particular, by Lemma 9(2) and as $w \notin T^i$, w rejects m as soon as he serenades under her window during $R_{\mu^i}^{\mathcal{P}'_W, \mathcal{P}_M}$. As m serenades under w 's window following a rejection in $R_{\mu^i}^{\mathcal{P}'_W, \mathcal{P}_M}$, m prefers $\mu^i(m)$ over w ; by Lemma 9(2), m also prefers w over $\mu(m)$ (as $w \neq \mu(m)$, since $\mu(m) \in T^i$). Furthermore, by Lemma 9(4), $\mu^{i+1}(m) = \mu(m)$. Thus, Modification 2 was applied to w 's preference list w.r.t. m . If $w = \tilde{w}$, then she thus blacklists m and would thus immediately reject him if he ever serenades under her window during $R^{\mathcal{P}_W^{i+1}, \mathcal{P}_M}$ (and thus in particular dur-

ing $R_{2/3}^{i+1}$); otherwise, $w \neq \hat{w}$, and as w 's provisional match only improves (according to \mathcal{P}_W^{i+1}) during $R_{W, \mathcal{P}_M}^{i+1}$, she would thus immediately reject m if he ever serenades under her window during $R_{2/3}^{i+1}$. We conclude that any woman who, during $R_{\mu^i}^{\mathcal{P}'_W, \mathcal{P}_M}$, immediately rejects all men serenading under her window for the first time after the first night, would reject all such men during $R_{2/3}^{i+1}$ as well, if any of them serenade under her window after the first night; we also conclude that any other woman has the same preference list and initial matching during $R_{\mu^i}^{\mathcal{P}'_W, \mathcal{P}_M}$ and during $R_{2/3}^{i+1}$, with the exception of \tilde{w}^{i+1} , who would immediately reject all men except $\mu(\tilde{w}^{i+1})$ in both $R_{\mu^i}^{\mathcal{P}'_W, \mathcal{P}_M}$ and $R_{2/3}^{i+1}$, and would accept $\mu(\tilde{w}^{i+1})$, in which case both $R_{\mu^i}^{\mathcal{P}'_W, \mathcal{P}_M}$ and $R_{2/3}^{i+1}$ stop. Thus, as the only rejection on the first night of both $R_{\mu^i}^{\mathcal{P}'_W, \mathcal{P}_M}$ and $R_{2/3}^{i+1}$ is of $\mu^i(\tilde{w}^{i+1})$ by \tilde{w}^{i+1} , we conclude that $R_{2/3}^{i+1}$ is well-defined and has the same rejections as $R_{\mu^i}^{\mathcal{P}'_W, \mathcal{P}_M}$.

Finally, we show that $R_{3/3}^{i+1}$ has the same rejections as $\tilde{R}_{2/2}^i$. Indeed, by definition of $t^i(\tilde{w}^{i+1})$, all rejections during $\tilde{R}_{2/2}^i$, in all nights but the first, are by women in $W \setminus T^i$. As explained above, the provisional matches of these women are unaltered during $R_{2/3}^{i+1}$, and thus are the same on the first nights of $R_{3/3}^{i+1}$ and $\tilde{R}_{2/2}^i$. As in both these first nights the only rejection is of \tilde{m} by \tilde{w}^{i+1} , it is enough to show that the preferences of all women in $W \setminus T^i$ agree w.r.t. all men provisionally matched to any of $W \setminus T^i$ in both of these first nights, as well as w.r.t. \tilde{m} . By definition of $t^i(\tilde{w}^{i+1})$, we have that all such men m (incl. \tilde{m}) satisfy $\mu^i(m) \notin T^i$, and thus $\mu^i(m) = \mu(m)$. Thus, Modification 2 does not involve demoting or blacklisting any such m . As $\tilde{w}^{i+1} \in T^i$, Modification 1 is irrelevant as well at this point. By Lemma 9(2), the preferences of $W \setminus T^i$ are the same in \mathcal{P}'_W as they are in \mathcal{P}_W^i . Thus, $R_{3/3}^{i+1}$ and $\tilde{R}_{2/2}^i$ have the same rejections. We conclude that the rejections occurring in $R_{W, \mathcal{P}_M}^{i+1}$ are exactly those occurring in either of R_{W, \mathcal{P}_M}^i or $R_{\mu^i}^{\mathcal{P}'_W, \mathcal{P}_M}$. Thus, as R_{W, \mathcal{P}_M}^i concludes with the initial matching of $R_{\mu^i}^{\mathcal{P}'_W, \mathcal{P}_M}$, and as each man is eventually matched with the woman he prefers most out of those who have not rejected him, we obtain that $\text{MenOpt}(\mathcal{P}_W^{i+1}, \mathcal{P}_M)$ is the matching obtained by first running R_{W, \mathcal{P}_M}^i and then $R_{\mu^i}^{\mathcal{P}'_W, \mathcal{P}_M}$, namely $\mu_{\mathcal{P}'_W, \mathcal{P}_M}^i = \mu^i$ and Property 1 holds for $i + 1$.

Property 2 for $i + 1$ follows once more directly from Lemma 9(3).

By Properties 4 and 5 for i , and by Lemma 9(1 and 2), in order to show that Properties 4 and 5 hold for $i + 1$, it is enough to show that for every $w \in W$, Modifications 1 and 2 do not affect the ranking of $\mu(w)$ and of $\mu^{i+1}(w)$ (if $\mu(w) \neq \mu^{i+1}(w)$). Indeed, Modification 1 affects \tilde{w}^{i+1} only, but leaves $\mu(\tilde{w}^{i+1})$ at the top of her preference list; by Lemma 9(1), this suffices. Modification 2 only affect the preference list of women w by demoting men who are not $\mu(w)$ or $\mu^i(w)$ on w 's preference list, and by Lemma 9(4), $\mu^{i+1}(w)$ is one of these two. Thus, Properties 4 and 5 hold for $i + 1$. Property 3 for $i + 1$ follows again from Property 3 for i and from Property 4 for $i + 1$, by Lemma 5 and as $\mu^{i+1} = \mu_{\mathcal{P}'_W, \mathcal{P}_M}^i$.

To prove Property 7 for $i + 1$ from Property 7 for i , we must show once more that every newly-blacklisted man m is newly-blacklisted by exactly one woman, is not blacklisted in \mathcal{P}'_W , and satisfies $\mu^{i+1}(m) = \mu(m)$. Indeed, by Lemma 9(5) and by Modifications 1 and 2, any man newly-blacklisted is newly-blacklisted only by \hat{w} and satisfies $\mu^{i+1}(m) = \mu(m) \neq \mu^i(m)$. Thus, in particular, by Property 7 for i , he is not blacklisted in \mathcal{P}_W^i .

Since by Lemma 9(5) and by Modifications 1 and 2, no blacklist is changed but perhaps \hat{w} 's, in order to prove Property 8 for $i + 1$ given Property 8 for i , it is enough to yet again show that no woman blacklists $\mu(\hat{w})$ in \mathcal{P}'_W , and that there exists no woman w with a nonempty blacklist in \mathcal{P}_W^i s.t. $\mu(w)$ is blacklisted by \hat{w} in \mathcal{P}_W^{i+1} . By Property 8 for i , it is enough to prove the former for the case in which \hat{w} has an empty blacklist in \mathcal{P}_W^i . In this

case, by definition of \hat{w} , no man serenades under her window during $\tilde{R}_{1/2}^i$, and $\tilde{R}_{2/2}^i$ stops as soon as any man serenades under her window. Thus, only one man serenades under \hat{w} 's window during $R^{\mathcal{P}_W^i, \mathcal{P}_M}$, and since $\hat{w} \notin T^i$, this man is $\mu(\hat{w})$. Thus, by Property 9 for i , $\mu(\hat{w})$ is not blacklisted in \mathcal{P}_W^i . To show the latter, let w be a woman with a nonempty blacklist in \mathcal{P}_W^i . By Property 8 for i , $\mu^i(\mu(w)) = \mu(\mu(w))$, and thus, by Modification 2, $\mu(w)$ is not blacklisted by \hat{w} in \mathcal{P}_W^{i+1} and the proof of Property 8 for $i+1$ is complete. Finally, by equivalence of $R_{1/3}^{i+1}$ and $\tilde{R}_{1/2}^i$, and of $R_{3/3}^{i+1}$ and $\tilde{R}_{2/2}^i$, any woman w serenaded-to during $R^{\mathcal{P}_W^{i+1}, \mathcal{P}_M}$ solely by $\mu(w)$ is also serenaded-to during $R^{\mathcal{P}_W^i, \mathcal{P}_M}$ solely by him. Thus, Property 9 for $i+1$ follows from Property 9 for i and from Modification 2 (as $\mu^i(\mu(w)) = \mu(\mu(w))$ for any such w).

Thus, the proof by induction is complete, as the process stops by Property 2 and by finiteness of W . The remainder of the proof follows verbatim as in the proof of Theorem 11. The $O(n^3)$ time complexity follows by naïvely implementing the above inductive process. We note that the only obstacle to obtaining $O(n^2)$ time complexity is the need to rerun the Gale-Shapley algorithm in order to identify \hat{w} on every iteration corresponding to the second case of the induction step; more formally, the algorithm runs in $O(n^2 \cdot (1+k))$ time, where k is the number of iterations of the second case of the induction step (thus yielding a best-case time complexity of $O(n^2)$, i.e. under the conditions of Theorem 11). Recall that by Lemmas 7 and 9, the number of all iterations is at most the number of $(\mu^1 \rightarrow \mu)$ -cycles. Since μ is uniformly distributed given \mathcal{P}_M (and thus given μ^1), we have that $\mu \circ \mu^{1-1}$ is uniformly distributed in S_n ; therefore, the expected number of cycles in $\mu \circ \mu^{1-1}$ is $H_n \triangleq \sum_{j=1}^n \frac{1}{j}$ (Arratia et al., 2003, p. 19). Thus, by Remark 4, we have that $\mathbb{E}[k] \leq H_n = O(\log n)$, and hence the average-case time complexity of the above algorithm is $O(n^2 \log n)$, as required.

A final implementation note: we observe that as long as \tilde{w}^{i+1} is chosen as detailed in the proof above, then for any i , until time $t^i(\tilde{w}^{i+1})$, $R^{\mathcal{P}_W^i, \mathcal{P}_M}$ is indistinguishable from $R^{\mathcal{P}_W^{i+1}, \mathcal{P}_M}$, and that after this time, in neither run is any woman in T^i newly-serenaded to. This follows since $(t^i(\tilde{w}^{i+1}))_{i=1}^d$ is weakly monotone-decreasing by its definition and by Lemma 9(4). \square

A.2 Proofs of the remaining Theorems from Section 3

Proof of Theorem 5. We start by showing that Theorem 5 holds for $n_b = 1$ and $l_1 = n - 1$. Denote the members of W by w_0, \dots, w_{n-1} , and of M — by m_0, \dots, m_{n-1} . Let μ be the matching and \mathcal{P}_M be the profile of preference lists for M s.t. for every $0 \leq j \leq n - 1$, both $\mu(w_j) \triangleq m_j$ and the preference list of m_j (from most-preferred to least-preferred) is: $w_{j+1}, w_{j+2}, \dots, w_{n-1}, w_0, w_1, \dots, w_j$. To show Part 1, let \mathcal{P}_W be the profile of preference lists by which, as in Lemma 7, the preference list of w_0 is m_0 (with all other men blacklisted), and for every $1 \leq j \leq n - 1$, the preference list of w_j is first m_j , followed immediately by m_{j-1} , followed by all other men in arbitrary order. Let μ' be the matching describing the first night of $R^{\mathcal{P}_W, \mathcal{P}_M}$, i.e. $\mu'(m_j) \triangleq w_{(j+1) \bmod n}$ for all $0 \leq j \leq n - 1$. It is straightforward to check that $C_{\mu'}^{\mu}(w_0) = (w_0 \xrightarrow{m_{n-1}} w_{n-1} \xrightarrow{m_{n-2}} w_{n-2} \xrightarrow{m_{n-3}} \dots w_1 \xrightarrow{m_1} w_0)$. Thus, by Lemma 7 and by Lemma 6(3), $\mu'_{\mathcal{P}_W, \mathcal{P}_M} = \mu$. (The interested reader may verify that \mathcal{P}_W is the profile of preference lists obtained by applying the above proofs/algorithms.)

We now move on to proving Part 2. Let \mathcal{P}'_W be a profile of preference lists for W s.t. $\text{MenOpt}(\mathcal{P}'_W, \mathcal{P}_M) = \mu$. We must show that there exists a woman whose preference list according to \mathcal{P}'_W consists of a single man. We consider the timing for $R^{\mathcal{P}'_W, \mathcal{P}_M} = R_{\mu'}^{\mathcal{P}'_W, \mathcal{P}_M}$ induced by the recursive implementation (McVitie and Wilson, 1971) of the Gale-Shapley algorithm, i.e. as long as any blacklist-induced rejection remains, one such rejection occurs, its aftermath runs until the algorithm converges once more (we call this an *iteration*), at which point, if any other blacklist-induced rejection remain, the process repeats. If, at any time, a woman is approached by two men she blacklists, she arbitrarily rejects one of them in favour of the other.

We first note that on each iteration, every woman rejects precisely one man. Indeed, assume that an iteration starts with a woman w_j rejecting her provisional match. By definition of \mathcal{P}_M , this provisional match then serenades under $w_{(j+1 \bmod n)}$'s window, who rejects either him or her provisional match, the rejected man then serenades under $w_{(j+2 \bmod n)}$'s window, who rejects either him or her provisional match, and so forth until the man rejected by $w_{(j+n-1 \bmod n)}$ serenades under the window of $w_{(j+n \bmod n)} = w_j$, who has no provisional match at that point, and thus the iteration concludes. So, exactly n rejections take place during each iteration. As $R^{\mathcal{P}'_W, \mathcal{P}_M}$ yields μ , each man is rejected by $n - 1$ women during this run, and thus $n \cdot (n - 1)$ rejections in total take place during it, and so the above-described timing for it consists of $n - 1$ iterations. Let w be the woman whose blacklist-induced rejection of her then-provisional match m triggers the last iteration of $R^{\mathcal{P}'_W, \mathcal{P}_M}$. In the previous iterations, w has already performed $n - 2$ rejections (one in each iteration), and since in the beginning of the last iteration she is provisionally matched with m , whom she blacklists, we conclude that the $n - 2$ men she has already rejected in prior iterations are also blacklisted by her, for a total of $n - 1$ men blacklisted by w and the proof is complete.

The general case of Theorem 5 follows in a similar way. Assume w.l.o.g. that $l_1 + \dots + l_{n_b} = n - n_b$, by adding l_i s equal to zero (and thus increasing the value of n_b) if necessary. Denote the members of W by w_j^i , for $1 \leq i \leq n_b$ and for each such i , for $0 \leq j \leq l_i$; similarly denote the members of M by m_j^i , for the same values of i and j . For every such i and j , we set $\mu(w_j^i) \triangleq m_j^i$ and set the preference list of m_j^i to start with $w_{j+1}^i, w_{j+2}^i, \dots, w_{l_i}^i, w_0^i, w_1^i, \dots, w_j^i$, in this order, followed by all other women in arbitrary order. The proof of Part 1 is similar to that of the special case, with w_0^i , for every $1 \leq i \leq n_b$ preferring m_0^i most, and blacklisting m_j^i for all $0 \leq j \leq l_i$; and with w_j^i , for every such i and for every $0 < j \leq l_i$, preferring m_j^i most, followed immediately by m_{j-1}^i , followed by all other men in arbitrary order. Finally, Part 2 follows by performing the analysis of the special case for each i separately, and by observing that since $\text{MenOpt}(\mathcal{P}'_W, \mathcal{P}_M) = \mu$, no man m_j^i is ever rejected by w_j^i during $R^{\mathcal{P}'_W, \mathcal{P}_M}$, and thus no such man ever serenades under the window of any woman $w_{k'}^i$, for any $i' \neq i$ and any k . \square

Proof of Theorem 6. The proof runs along the lines of the proof of Theorem 4, with a few differences that we now survey.

We begin by focusing our attention on W_μ^c . By having each women $w \in W_\mu^c$ blacklist all men¹¹, we guarantee that these women turn out to be unmatched, and can *de facto* ignore them from this point on.

The base of the inductive argument is as in Theorem 4, except that we place M_μ^c last in all preference lists in \mathcal{P}_W^0 . (The resulting matching μ^1 is thus the matching that would have been obtained by considering only M_μ in the runs defining μ^1 .) This guarantees that M_μ^c are exactly the men unmatched in μ^1 . Before commencing with the inductive steps of the proof of Theorem 4, we perform the following inductive step, as many times as possible (with the same induction invariants): if there exists any woman $w \in T^i$ who is serenaded-to by some $\tilde{m} \in M_\mu^c$ during $R^{\mathcal{P}'_W, \mathcal{P}_M}$ (equivalently, who is not blacklisted by some such \tilde{m} — the set of all such women can be precomputed once at the beginning of the algorithm), then we denote $\tilde{w}^{i+1} \triangleq w$. We continue the construction and proof as in the second case of the induction step of the proof of Theorem 4 (however w.r.t. \tilde{w}^{i+1} and \tilde{m} as defined here), only noting that all women are matched at the end of $\tilde{R}_{1/2}^i$, and therefore Modification 2(i) never takes place. (The interested reader may verify that Modification 2(ii) is a no-op, and therefore Modification 2 is redundant in its whole in this case.) Therefore, no blacklist is changed during this induction step, and thus when it is no longer possible to conduct any more such steps, all blacklists in \mathcal{P}_W^{i+1} are empty. The rest of the results, except for those regarding the worst- and average-case time complexities, follow as in the proof of Theorem 4.

¹¹It actually suffices for each such w to blacklist all $m \in M_\mu^c$ who do not blacklist her, in addition to all $m \in M_\mu$ who prefer w over $\mu(m)$. The proof is left to the interested reader.

Let r be the number of women w for whom, when it is no longer possible to conduct any more induction steps as above, it still holds that $\mu^{i+1}(w) \neq \mu(w)$. By Lemma 9(1), $r \leq n_\mu - n_h$, yielding the worst-case time complexity as required, via the same arguments as in the proof of Theorem 4. For the average-case analysis, note that each induction step as above causes not only \tilde{w}^{i+1} to be matched with $\mu(\tilde{w}^{i+1})$, but actually her entire $(\mu^1 \rightarrow \mu)$ -cycle to be matched with their partners in μ . As the expected combined size of the cycles containing n_h elements in a random permutation uniformly distributed in S_{n_μ} is $\frac{n_h \cdot (n_\mu + 1)}{n_h + 1}$ (Gonczarowski, 2013), we have $\mathbb{E}[r] = n_\mu - \frac{n_h \cdot (n_\mu + 1)}{n_h + 1} < \frac{n_\mu + 1}{n_h + 1}$. By Jensen's inequality (Jensen, 1906) and by concavity of the logarithm function, we have $\mathbb{E}[\log r] \leq \log \mathbb{E}[r] = O(\log \frac{n_\mu + 1}{n_h + 1})$. The rest of the analysis is as in the proof of Theorem 4. \square

Proof of Theorem 7. The proof runs along the lines of that of Theorem 5. Assume w.l.o.g. that $l_1 + \dots + l_{n_b} = n_\mu - n_h - n_b$, by adding l_i s equal to zero (and thus increasing the value of n_b) if necessary. Define $H \triangleq \{w \in \tilde{W} \mid \exists m \in \tilde{M}^c : w \notin B_m\}$, and note that by definition, $n_h = |H|$. Let $\{m_w\}_{w \in H}$ be n_h distinct arbitrary men from \tilde{M} . For every $m \in M$, define $P(m) \triangleq \{w \in \tilde{W}^c \mid m \in B_w\}$. Denote the members of $\tilde{W} \setminus H$ by w_j^i , for $1 \leq i \leq n_b$ and for each such i , for $0 \leq j \leq l_i$; similarly denote the members of $\tilde{M} \setminus \{m_w\}_{w \in H}$ by m_j^i , for the same values of i and j . For every such i and j , we set $\mu(w_j^i) \triangleq m_j^i$ and set the preference list of m_j^i to start with $P(m_j^i)$ in arbitrary order, followed immediately by $w_{j+1}^i, w_{j+2}^i, \dots, w_{l_i}^i, w_0^i, w_1^i, \dots, w_j^i$, in this order, followed by all other women in arbitrary order. Additionally, for every $w \in H$, we set $\mu(w) = m_w$ and set the preference list of m_w to start with $P(m_w)$ in arbitrary order, followed immediately by w , followed by all other women in arbitrary order. Finally, we set the blacklist of each $m \in \tilde{M}^c$ to consist of B_m , with all other women appearing in m 's preference list in arbitrary order.

The proof of Part 1 is similar to that of Theorem 5(1), with w_0^i , for every $1 \leq i \leq n_b$ preferring m_0^i most, and blacklisting m_j^i for all $0 < j \leq l_i$; and with w_j^i , for every such i and for every $0 < j \leq l_i$, preferring m_j^i most, followed immediately by m_{j-1}^i , followed by all other men in arbitrary order. Set the preference list of every $w \in H$ to start with m_w , followed by all other men in arbitrary order. Finally, set the blacklist of each $w \in \tilde{W}^c$ to consist of $B_w \cup \{m \in \tilde{M}^c \mid w \notin B_m\}$, with all other men appearing in w 's preference list in arbitrary order.

For every $w \in H$, w 's top choice is m_w , and m_w 's top choice, out of all women who do not blacklist him, is w (as all of $P(m_w)$ blacklist him), and thus w and m_w are matched by $\text{MenOpt}(\mathcal{P}_W, \mathcal{P}_M)$ and no man is ever rejected in favour of m_w , except by w . We thus *de facto* ignore H and $\{m_w\}_{w \in H}$ henceforth. Let us consider the timing for $R^{\mathcal{P}_W, \mathcal{P}_M}$ obtained by deferring the participation of \tilde{M}^c until the algorithm converges (we denote this part of the run by $R_{1/2}^{\mathcal{P}_W, \mathcal{P}_M}$), and only then introducing \tilde{M}^c into the market, until the algorithm converges once again (we denote this part of the run by $R_{2/2}^{\mathcal{P}_W, \mathcal{P}_M}$). As in the proof of Theorem 5(1), if \tilde{W}^c were to not participate in $R_{1/2}^{\mathcal{P}_W, \mathcal{P}_M}$, then at its end each m_j^i would be matched with w_j^i . As every such m_j^i prefers w_j^i over all of \tilde{W}^c (except for $P(m_j^i)$, who all blacklist him), we have that participation of \tilde{W}^c in $R_{1/2}^{\mathcal{P}_W, \mathcal{P}_M}$ would not change the resulting matching. Finally, it is straightforward to verify that at the end of the $R_{1/2}^{\mathcal{P}_W, \mathcal{P}_M}$, every $w \in \tilde{W}$ prefers her provisional match over all of \tilde{M}^c , and every $w \in \tilde{W}^c$ blacklists all of \tilde{M}^c who do not blacklist her. Thus, the provisional matching does not change at any point during $R_{2/2}^{\mathcal{P}_W, \mathcal{P}_M}$, and the proof of Part 1 is complete.

To prove Part 2, let \mathcal{P}'_W be a profile of preference lists for W s.t. $\text{MenOpt}(\mathcal{P}'_W, \mathcal{P}_M) = \mu$. We commence by examining \tilde{W}^c . Let $w \in \tilde{W}^c$ and let m be a man blacklisted by w in \mathcal{P}_W . If $m \in \tilde{M}^c$, then by the statement of Part 2, we need only consider the case in which $w \notin B_m = B_m(\mathcal{P}_M)$; in this case, the fact that both w and m are unmatched by $\text{MenOpt}(\mathcal{P}'_W, \mathcal{P}_M)$ implies

that w blacklists m in \mathcal{P}'_W , as required. Otherwise, $m \in \tilde{M}$ and thus, by definition of \mathcal{P}_W , $m \in B_w$; therefore, $w \in P(m)$ and so m prefers w over his match $\mu(m)$. As w is unmatched in μ , we thus have that she blacklists m in \mathcal{P}'_W in this case as well. Finally, the analysis of the blacklists of $\tilde{W} \setminus H$ is as in the proof of Theorem 5(2), since all of \tilde{M}^c blacklist them, and all of $\{m_w\}_{w \in H}$ prefer their final matches over them; thus, only $\tilde{M} \setminus \{m_w\}_{w \in H}$ serenade during $R^{\mathcal{P}'_W, \mathcal{P}_M}$ under their windows (and these men are blacklisted by any woman in $\tilde{W}^c \cup H$ that they prefer over their final match). \square

A.3 Proofs of the Theorems from Section 4

Theorem 9 readily follows from the following lemma, phrased in terms of the vanilla Gale-Shapley algorithm.

Lemma 11. *Let μ' and μ be matchings and let \mathcal{P}_M be a $(\mu' \rightarrow \mu)$ -compatible profile of preference lists for M . Let W and M be equal-sized sets of women and men, respectively. Define $n \triangleq |W| = |M|$. If $\mu' \neq \mu$, then there exists a woman \tilde{w} and a profile \mathcal{P}_W of preference lists for W , s.t. the only woman with a nonempty blacklist according to \mathcal{P}_W is \tilde{w} , having a blacklist consisting solely of $\mu'(\tilde{w})$, s.t. $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(\tilde{w}) = \mu(\tilde{w})$, and s.t. \mathcal{P}_M is $(\mu'_{\mathcal{P}_W, \mathcal{P}_M} \rightarrow \mu)$ -compatible.*

Proof. Denote by $\tilde{\mathcal{P}}_W$ a profile of preference lists for W , according to which each $w \in W$ prefers $\mu(w)$ most, followed immediately by $\mu'(w)$ (if $\mu'(w) \neq \mu(w)$), followed by all other men in arbitrary order. For every $w \in W$ s.t. $\mu'(w) \neq \mu(w)$, we denote by $\tilde{\mathcal{P}}_W^w$ the profile of preference lists for W obtained from $\tilde{\mathcal{P}}_W$ by having w blacklist $\mu'(w)$. We note that $\tilde{\mathcal{P}}_W^w$ and \mathcal{P}_M are μ' -cycle generating with trigger w . Let \tilde{w} be a woman for whom $C_{\mu'}^{\mathcal{P}_{\tilde{w}}, \mathcal{P}_M}$ is of greatest length.

We claim that $\hat{w} \triangleq \mu'(\mu(\tilde{w}))$ rejects some man during $R_{\mu'}^{\mathcal{P}_{\tilde{w}}, \mathcal{P}_M}$. Indeed, assume for contradiction that this is not the case; we hence show that $C_{\mu'}^{\mathcal{P}_{\tilde{w}}, \mathcal{P}_M}$ is of greater length than $C_{\mu'}^{\mathcal{P}_{\hat{w}}, \mathcal{P}_M}$. Indeed, we claim that $C_{\mu'}^{\mathcal{P}_{\tilde{w}}, \mathcal{P}_M}$ has prefix $(\hat{w} \xrightarrow{\mu(\tilde{w})} \tilde{w} \xrightarrow{m_1} w_2 \xrightarrow{m_2} w_3 \xrightarrow{m_3} \dots w_{d-1} \xrightarrow{m_{d-1}} w_d)$, where $(w_1 \xrightarrow{m_1} w_2 \xrightarrow{m_2} \dots w_{d-1} \xrightarrow{m_{d-1}} w_d) \triangleq C_{\mu'}^{\mathcal{P}_{\hat{w}}, \mathcal{P}_M}$. By definition of $\tilde{\mathcal{P}}_W$, and since by definition $\mu'(\hat{w}) = \mu(\tilde{w})$, we directly have that $C_{\mu'}^{\mathcal{P}_{\hat{w}}, \mathcal{P}_M}$ has prefix $(\hat{w} \xrightarrow{\mu(\tilde{w})} \tilde{w} \xrightarrow{m_1} \dots)$. As \tilde{w} is not serenaded-to except on the first and last nights of $R_{\mu'}^{\mathcal{P}_{\tilde{w}}, \mathcal{P}_M}$, and as, by assumption, \hat{w} is serenaded-to solely by $\mu'(\hat{w})$ throughout $R_{\mu'}^{\mathcal{P}_{\tilde{w}}, \mathcal{P}_M}$, the following nights of $R_{\mu'}^{\mathcal{P}_{\tilde{w}}, \mathcal{P}_M}$ have the exact same rejections as the second to one-before-last nights of $R_{\mu'}^{\mathcal{P}_{\hat{w}}, \mathcal{P}_M}$, respectively, and thus indeed $C_{\mu'}^{\mathcal{P}_{\tilde{w}}, \mathcal{P}_M}$ has the above prefix — a contradiction.

If \hat{w} rejects $\mu'(\hat{w}) = \mu(\tilde{w})$ during $R_{\mu'}^{\mathcal{P}_{\tilde{w}}, \mathcal{P}_M}$, then by definition of $\tilde{\mathcal{P}}_W^{\tilde{w}}$ we have $\mu'_{\tilde{\mathcal{P}}_W^{\tilde{w}}, \mathcal{P}_M}(\mu(\tilde{w})) = \tilde{w}$, and the proof is complete by setting $\mathcal{P}_W \triangleq \tilde{\mathcal{P}}_W^{\tilde{w}}$. Otherwise, \hat{w} rejects some man \tilde{m} in favour of $\mu'(\hat{w})$ during $R_{\mu'}^{\mathcal{P}_{\tilde{w}}, \mathcal{P}_M}$. Let \mathcal{P}_W be the profile of preference lists obtained from $\tilde{\mathcal{P}}_W^{\tilde{w}}$ by promoting \tilde{m} to be second on \hat{w} 's preference list (immediately following $\mu(\hat{w})$ and immediately followed by $\mu'(\hat{w})$). Thus, in $R_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$, $\mu'(\hat{w}) = \mu(\tilde{w})$ is rejected by \hat{w} in favour of \tilde{m} and once again, by definition of \mathcal{P}_W , we have $\mu'_{\mathcal{P}_W, \mathcal{P}_M}(\mu(\tilde{w})) = \tilde{w}$ and the proof is complete. (Either way, as no woman w rejects $\mu(w)$ during $R_{\mu'}^{\mathcal{P}_W, \mathcal{P}_M}$, we have that \mathcal{P}_M is $(\mu'_{\mathcal{P}_W, \mathcal{P}_M} \rightarrow \mu)$ -compatible, as required.) \square

Proof of Theorem 10. The proof runs parallel to that of the special case of Theorem 5(2), by noting that the first season concludes as soon as it commences, and that each consecutive season precisely corresponds to one iteration, as defined there, and since, as explained there, $n - 1$ iterations are required. \square