

The Marriage Theorem

E. L. Lady

(April 21, 2000)

The Marriage Theorem is often stated in terms of a rather implausible situation. Namely, one assumes that we have n unmarried men and m unmarried women. We say that a man and woman are **compatible** if they would be willing to marry each other. The Marriage Problem is to match each woman up with a man with whom she is compatible, so that all the women will be married. Compatibility here is assumed to be mutual, i. e. a couple is compatible if and only if each is an acceptable mate for the other. Furthermore, it is apparently assumed that each individual will be satisfied if married to any one of the potential partners s/he is compatible with.

This probably seemed a little less absurd in the time of Jane Austen, although I suspect that even then it had a satirical flavor to it. (We know that mathematicians often like to sneak a little social criticism into their theorems.) In any case, one can make the problem more plausible for our era by assuming that the situation is that of a dance, and the problem is to match couples up so that every woman has a partner to dance with.

We say that the **Marriage Condition** holds if for each set $\{x_1, \dots, x_k\}$ of k women, there is a set A of at least k men such that each man in A is compatible with at least one of the women x_1, \dots, x_k . (If one imagines that the women x_1, \dots, x_k gather in one corner of the room, and all the men who are compatible with any one (or more) of these women gather around them, then the Marriage Condition states that anytime this happens, there will always be at least as many men in the group as there are women.)

In particular, the Marriage Condition requires that there be at least one compatible man for each woman. (This is the case $k = 1$.)

The **Marriage Theorem** states that it is possible for all the women to get married if and only if the Marriage Condition holds.

The “only if” part of the theorem is clear. If the Marriage Condition fails to hold, this means that when at least one group of women x_1, \dots, x_k gather in a corner of the room, there are fewer than k men who gather around them as potential partners. Clearly it is not possible for these k women in particular to all be married to compatible partners.

The proof that will be given here for the “if” part of the Marriage Theorem is essentially a restatement of the graph-theoretic proof which I think is fairly standard. (C.f. Brualdi, *Introductory Combinatorics*, 3rd Edition, Chapter 9.) Basically, the proof involves describing an algorithm by which a given matching of couples can be improved, i. e. another matching will be found which has a larger number of matched couples. We will see that by repeated

use of this algorithm we can reach a point where either all the women have been matched or where there is a clear contradiction to the Marriage Condition.


To start with, it's important to get the terminology straight. Assume that we have a (partial) matching, so that some of the women are paired with compatible men. We will say that a man or woman is **matched** if they are matched to someone under the current matching. Remember that a **matched** couple is always **compatible** but that a man and woman who are **compatible** are not necessarily **matched**.

Start with a preliminary matching. We may assume that there are no compatible pairs where both the man and the woman are **unmatched**, otherwise we could match these pairs and get a better matching. In the example below, compatible couples are always joined by a dashed or solid line segment or arrow. Double line segments are used to indicate **matched** couples.

The basic idea is as follows: We start with an **unmatched** woman. If the Marriage Condition holds, there must be men she is compatible with. If one of these men is **unmatched**, then of course we can match the two. If not, then we might match the woman in question with one of the **matched** men she is compatible with, but only by breaking apart the other couple. This would produce a new unmatched woman. Proceeding in this way, we move through a chain where each woman is succeeded by a man she is compatible with who has not previously occurred in the chain, and this man, if **matched**, is succeeded by his current partner. (In Brualdi, this is called an **alternating chain**.) If this chain finally arrives at an **unmatched** man, then we can rematch all the men and women in the chain and wind up with one more matched couple than before.

Schematically, this looks like this, for instance:

An alternating chain with three matches, and ending with an unmatched man



leads to a new matching with four matches.



However there is no guarantee that a particular alternating chain will eventually arrive at an unmatched man. Instead an alternating chain may terminate by reaching a woman who is not compatible with any man who has not previously been used in the chain.

An alternating chain ending with a woman who is only compatible with men already in the chain.



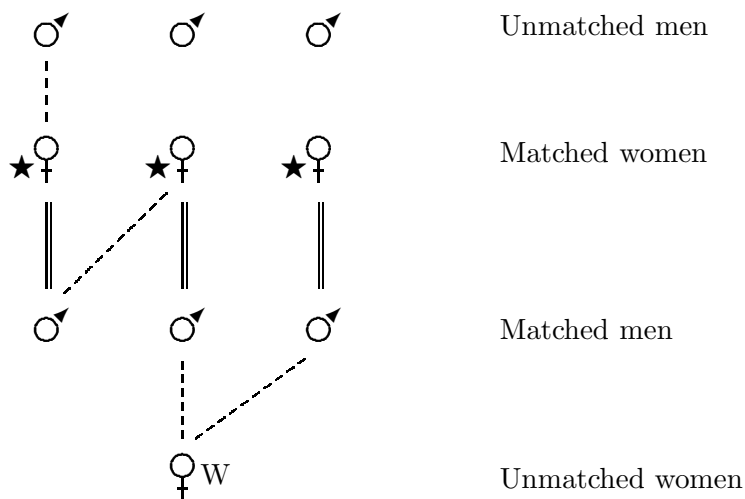
In this case, the chain does not enable us to produce a greater number of matched couples than already exists.

The idea of the proof, though, is that if we consider *all* the alternating chains starting with a particular **unmatched** woman, and if the Marriage Condition holds, then at least one of these chains will eventually reach an **unmatched** man.

We will now give an algorithm which will find such an alternating chain.

ALGORITHM

Start with a preliminary matching. Assign a star to each **matched** couple, initially pinning the star on the woman in the couple.




The algorithm will now improve this matching. We will start with an **unstarred** woman (who at this point, at the beginning of the iteration, must therefore also be unmatched) and look at the set of men she is compatible with. If one of these men is **unmatched**, then this is fine; go to **Breakthrough**, below. Otherwise, we will begin a process which will in effect construct longer and longer alternating chains until we find an alternating chain connecting an **unmatched** woman with an **unmatched** man.

Fortunately, we don't actually have to maintain a picture of all the alternating chains being constructed. The stars will enable us to keep track of the process and avoid cycling

endlessly. As we go through the re-matching process, we will transfer stars from women to men **within couples**. (There will always be one star for each matched couple.) A star on a woman means that she is not yet eligible to become part of one of the alternating chains being constructed, whereas a star on a man indicates that he is already part of an alternating chain (and thus also not eligible for further chains).

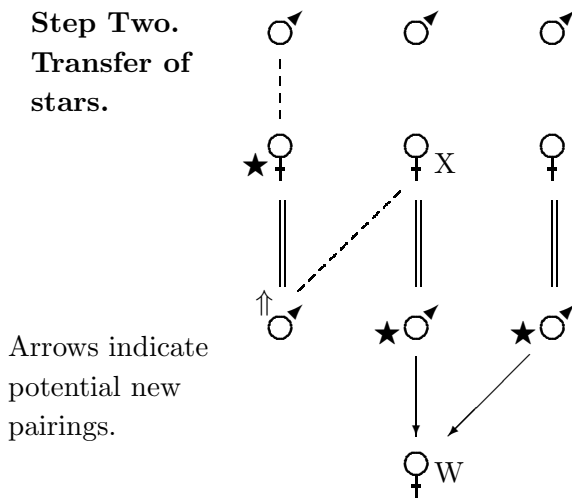
We will show below that if at any point there are no unstarred men who are compatible with unstarred women, then either all women have been matched or there is a set of k women such that there are fewer than k men compatible with any of these woman, so the Marriage Condition fails.

Begin Star Transfer. If no **unstarred** women are compatible with unstarred men, the algorithm terminates. Otherwise, choose an unstarred woman W who is compatible with at least one unstarred man. (In subsequent iterations, this applies even if this woman is **matched**. In the first iteration, of course, all matched women are also starred.) If one of the men is compatible with W is **unmatched**, then go to **Breakthrough** below. Otherwise we will say that those unstarred (but matched) men who are compatible with W have been **chosen** by W . The algorithm proceeds by transferring stars **within matched couples** to the men **chosen** by W . Each woman who is **matched** to an unstarred man compatible with W gives her star to her partner, who thus now becomes starred. (Couples where the male partner is not compatible with W are not affected.) The men receiving the stars need to remember that W is the woman who has chosen them. These starred men are now ineligible to be chosen in subsequent steps.



*W is compatible with the unstarred men M and N . We say that these men have been **chosen** by W . Their partners X and Z will now transfer their stars to M and N , who will no longer be eligible to be chosen. (Remember that stars are always transferred **within** matched couples.)*

There will now be new **unstarred but matched** women. There continues to be exactly one star in each **matched** couple, but now some of these stars have been transferred to men.



Stars are transferred to the men who are compatible with the unstarred woman W. A woman (X) who has given up her star is now compatible with an unstarred man at the beginning of the third row. Since this man is unstarred, he is a permissible new mate for X. But since he is currently **matched**, re-matching him with X will also require re-matching the woman he is currently paired with.

At the end of each iterative step, as just described, one of three things must happen:

1. An unstarred woman is compatible with an **unmatched** man. In this case, go to **Breakthrough** below.
2. There exists an unstarred woman X (who may in fact be **matched**) compatible with an unstarred but **matched** man. In this case, repeat the star transfer process, either going to **Breakthrough** if any of the men compatible with X are **unmatched**, or else transferring stars to the **matched** but unstarred men compatible with X.
3. No unstarred women are compatible with any unstarred men and so the algorithm terminates.

The stars are a mechanism for ensuring that a man who has been chosen will not be chosen again until **Breakthrough** and that a woman who once chooses will not choose again until **Breakthrough**.

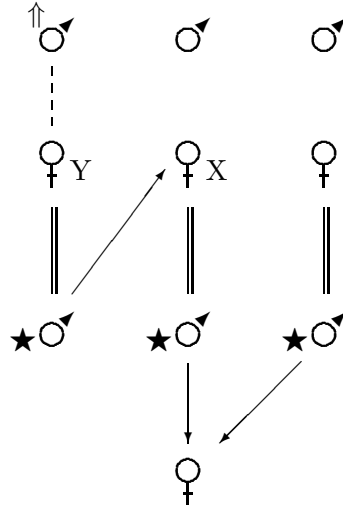
A star on a woman means that she is not currently available for re-pairing, because the man she is matched with has at the moment no other permissible partner.

A star on a man means that he has already been chosen by a possible new partner and thus is now longer available to be chosen.

Stars always move within existing **matched** couples from women to men, never vice-versa, except when stars are re-assigned after **Breakthrough**.

The number of stars is equal to the number of starred women plus the number of starred men and is also equal to the number of **matched** couples. If not all women have been **matched**, then the number of stars is less than the number of women. It follows that if not all women have been matched, then the number of **starred** men will be less than the number of **unstarred** women.

Step Three.



Stars are transferred to the men who are compatible with the unstarred woman X. A woman (Y) who has given up her star is now compatible with an **unmatched** man. This is called **breakthrough**.


Since there are only a finite number of stars, and at least one star is transferred from a woman to a man in each iteration, as we keep repeating this process either one of two things must happen: (1) We will find an unstarred woman compatible with an **unmatched** man, as in the diagram above. This is **breakthrough**. (2) We will reach a point where we are unable to repeat the process because no unstarred woman is compatible with any unstarred man.

Breakthrough. Breakthrough occurs when an **unstarred woman** Y is compatible with an unmatched man M. Starting with M, we can work backwards through a chain of re-matches, and finish with one more matched couple than previously.

In fact, we match M with the woman who has chosen him. If this woman is in an existing match, we break up this match and re-match her partner with the woman who has chosen him. We continue re-matching in this way until we finally get to an unmatched woman. (This must eventually occur, because there are only a finite number of women.)

Schematically, the process looks like this:

Replace this (three matches)




with this (four matches)



Looking across the top chain here from right to left, we recall that each woman in the chain chose the man on her left, at which point that man acquired his star. So each man in the chain acquired his star earlier than all the men to his left, and later than all the men to his right. Thus we can be sure that all the men in the chain are in fact distinct, and thus all the women are also distinct. Thus the chain does not cycle, and hence eventually must end, and this can only happen when it reaches an unmatched woman at the right end.

In terms of our earlier example, breakthrough looks like this.

**Step Four.
Breakthrough.**



On reaching breakthrough, we go through a chain reaction of re-matchings. There is now one more matched couple than before. (Stars will now be re-distributed, given initially to the woman in each new matched couple.)

After breakthrough, we have one more matched pair than before. If any women are still **unmatched**, then start the whole process again from the very beginning by assigning a star to the female member of each matched couple.


We can continue **until we come to with a situation where there are no unstarred women compatible with unstarred men** and so the algorithm terminates.

Proposition. The matching algorithm terminates only when a matching of maximum possible size has been found.

PROOF: The matching algorithm terminates when no **unstarred** woman is compatible with an **unstarred** men. In other words, in every possible **compatible** pair, at least one of the two people is starred. This shows that there in any conceivable matching, the number of matched couples cannot be greater than the current number of stars. But the number of stars equals the number of matched couples under the present matching. \square

Let m be the total number of women and k the number of unstarred women. By assumption, when the algorithm terminates the k unstarred women are only compatible with starred men. If any woman are **unmatched**, this means that there are fewer than

m matched pairs, and therefore fewer than m stars. Since there are $m - k$ starred women, it follows that there are fewer than k starred men.



Since the k unstarred women are compatible only with starred men, **we see that if not all women are matched when the algorithm terminates, then the Marriage Condition does not hold.**

Let σ be the largest possible shortfall between any set of women and the set of men who are compatible with any of these women. I. e. σ is the largest non-negative value $k - r$ such that there exists a some set of k women where there are only r men compatible with any of these women. (Thus the Marriage Condition holds if and only if $\sigma = 0$.)

Theorem. In the best possible matching, there will be σ women who are unmatched. Thus the maximum number of couples that can be matched in any possible matching is $m - \sigma$.

PROOF: For some ℓ there exists a group of ℓ women with only r compatible partners for any of these women, and with $\sigma = \ell - r$. This means that in any conceivable matching, there will be σ women at least who cannot be matched. Thus there can never be more than $m - \sigma$ matched couples.

Now apply the Marriage Algorithm above until we reach a point where there are k unstarred women and these are compatible with only starred men. There are $m - k$ starred women, so if there are r starred men at this point, then the number of matched couples equals the total number of stars, and this equals $m - k + r$. Thus the first paragraph shows that $m - k + r \leq m - \sigma$. On the other hand, by definition of σ , $k - r \leq \sigma$, so that $k - \sigma \leq r$ and $m - \sigma = (m - k) + k - \sigma \leq m - k + r =$ the number of matched couples. Thus there exists a matching such that $m - \sigma$ couples are matched. \square