The Marriage Game*

Marco Li Calzi and M. Cristina Molinari

The happiness of a married man depends on the people he has not married.

Oscar Wilde

The marriage game is one of the world's oldest games. One takes some men and women (preferably unmarried) and lets them freely interact, hoping that they will manage to find a partner and enter into a happy and lasting marriage.

The marriage game happens to be one of the most difficult games to play, because nobody knows its rules exactly. In spite of a training phase during adolescence, the players need many years of experience and an enourmous emotional investment before they can learn the fundamental rules of the game. The game can take a cruel twist, especially for those who do not reach the final goal and end up as "bachelors". Moreover, it is only at the end of their lives that the players who get married and leave the game find out whether they have won or not. Sometimes, chance or necessity may put them back into the game against their will.

The same reasons that make the game so difficult to play probably explain the fascination it exerts on those looking at it from the outside. Obviously, there are different ways of playing the role of the observer. Readers of illustrated magazines or fans of soap opera look at this game very differently from writers, artists or movie directors, who find in the marriage game one of their main sources of inspiration.

The eye of the mathematician is particularly sensitive towards formal structures. Thus, when mathematics observes the marriage game, it is more apt to remark the characteristics that can be formalised in a model. For example, Gottman et al. [1] show how to construct dynamical models to describe and predict the evolution of a conversation between husband and wife. And Bearman et al. [2] use a graph to describe the emotional and sexual relations between about 800 students at a secondary school in a small town in the American Midwest.

* *Translated by Sarah Wolf*

uncorrected proof

Let us try to look at the marriage game through the eyes of a mathematician: we will construct a model (or better, a stylised version) of the marriage game and see which properties emerge.

The Marriage Game

There are n men and m women. All of them have a preference ranking over the partners of the other sex, based on which they are never perfectly indifferent between two suitors. These people try to find their match within the group, or in a more romantic language, to get married. Only strictly monogamous and heterosexual matchings are admitted.

Let us consider an example with three men and four women. Brian, Charles and David are the three men. Ann, Emma, Ingrid and Olivia are the four women. For brevity, we often use simply the initial letters of the names to denote who we are talking about. Note that the initials of the men are the first three consonants (B, C, D) and those of the women are the first four vowels (A, E, I, O).

Suppose that Brian's preference ranking has Ann before Emma before Ingrid before Olivia. Brian would prefer to marry Ann, but he is also willing to take either Emma or Ingrid as a wife. However, rather than marrying Olivia, he would prefer to remain single. To represent this preference ranking, we use the notation:

Brian: $A > E > I > *$

where "$>$" indicates preference while the star "$*$" denotes that from this position onwards Brian prefers not to get married.

For our example, let us suppose that the preferences of the seven players are the following:

Brian: $A > E > I > *$
Charles: $I > E > A > *$
David: $A > I > E > O$
Ann: $C > D > B$
Emma: $B > C > D$
Ingrid: $D > C > B$
Olivia: $D > B > C$

We are interested in describing which matchings can be expected to emerge. Ann, for example, is the first choice for both Brian and David. We would like to know who will manage to marry her, assuming that Ann does not succeed instead in marrying Charles, her own first choice. Of course, in general the answer depends on the circumstances. In a society where older men prevail, Ann would probably end up as the wife of whoever is older between Brian and David. In a society where women prevail, Ann would perhaps manage to take Charles for a husband.

To be more specific, let us analyse a modern (and Western) version of the marriage game where two *fundamental* rights have to be respected. The first one is that everybody can elect to remain single: nobody can impose marriage upon anybody

uncorrected proof

else. Therefore, if Charles does not want to marry Olivia (as his preferences state), we are sure that the match between Charles and Olivia cannot take place. The second right is that a married person can file for divorce should (s)he find a partner whom (s)he likes better than the current one and who is willing to marry him (or her). For example, if Ann is married to David, she can divorce him and get married with Charles if the latter one wants to wed her. Of course, if Charles is already legally committed to Ingrid (whom he prefers to Ann), Charles himself is not available. Therefore, Ann stays with David only if Charles is already married off to Ingrid (or Emma).

Stable matchings are those pairings of people which respect these two fundamental rights of the individuals. In our example there are only two stable matchings:

1. *B*–*E*, *C*–*I*, *D*–*A*, *O* single
2. *B*–*E*, *C*–*A*, *D*–*I*, *O* single

Therefore, for this particular version of the marriage game, we expect one of these two configurations to emerge at the end.

In general, depending on the number of people involved and on their preference rankings, there can be few or many stable matchings. Our example had two; however, even situations with a very small number of people can have a rather large number of stable matchings. Example 2.17 in [3] reports a situation devised by Knuth where a marriage game with four men and four women admits ten different stable matchings.

Mathematicians have found several properties that hold in general for our version of the marriage game. The monograph by Roth and Sotomayor [3] collects many of these results and provides proofs. Here we choose to quote only four of these.

The first result ensures that there always exists a stable matching, whatever the number of players and their preference rankings. In other words, any marriage game has at least one solution. Stable matchings can be determined by examining the possible configurations one by one to check whether they respect the two fundamental rights. This can be very time-consuming. In the next section we describe a constructive algorithm by Gale and Shapley [4] that finds one stable matching in a very simple and direct way. The existence of a stable matching is thus a corollary of the proposition that the Gale-Shapley algorithm always leads to a stable matching.

The second theorem states that the stable matching need not be unique, as our example also shows. In very many cases all we can do is to restrict the outcome of the marriage game to the subset of configurations that are stable matchings, but we are not able to foresee exactly which one is going to emerge. Luckily, not everything is predetermined. In our example, we expect Charles to marry Ann or Ingrid but the final outcome of the game is left to chance and to the ability of the players.

The third result says that a player who happens to stay single in a stable matching must remain single in any other stable matching. In our example, Olivia stays single in any stable matching. This theorem leaves no margin for regret: those who find themselves "bachelors" in a stable matching would have remained single even if things had gone differently. Perhaps this can be a source of consolation: after all,

uncorrected proof

no matter how things would have gone, Olivia could not have succeeded in finding a husband as part of a stable matching. Thus she has nothing to reproach herself about when she remains single.

The fourth theorem requires some preparation. Let us go back to our example, where there are just two stable matchings: [B–E, C–I, D–A, O single] and [B–E, C–A, D–I, O single]. We call the former M and the latter F. Let us find out how each of the men judges these two matchings.

Brian is indifferent, because he marries the same woman in both cases. Charles prefers M to F, because in M he marries Ingrid who is his first choice. And David prefers M to F as well, because in M he marries Ann who is his first choice. Therefore, for each of the men, the matching M is better than (or at least indifferent to) the matching F. The men unanimously prefer M to F.

The contrary happens for the women. Emma and Olivia are indifferent between M and F: the former marries the same man and the latter does not marry anyone anyway. Ann prefers F to M, because she marries Charles who is her first choice and Ingrid prefers F to M because she marries David who is her first choice. In this case the women unanimously prefer F to M. (The initials F and M help to remember which of the two sexes prefers the corresponding matching.)

Looking at the choice between M and F from the point of view of the two sexes, there is an obvious conflict. The men collectively prefer M and, similarly, the women collectively prefer F. Given that we started out from an example, this divergence in the collective preferences might just be coincidence. However, the fourth theorem states that this is an inevitable characteristic of the marriage game. If one of the sexes collectively prefers a certain stable matching to another one, then the opposite sex has an exactly contrary preference. In this case, the adage that men and women are in perpetual conflict seems to have some foundation.

How to Find a Stable Matching

There are many ways of finding a stable matching. The most natural one, and the first that was proposed, is the algorithm of Gale and Shapley [4]. Its main characteristic is to assign different tasks to the two sexes. The representatives of one sex have the burden to propose matrimony and the representatives of the other sex have the honour to accept it. For convenience in the following description, suppose it is the men who propose and the women who accept. Later on, we will see what happens when the roles are reversed.

The algorithm proceeds in stages. At the first stage, each man asks for the hand of his first choice (if he has one). Every woman judges the proposals she may have received and chooses whether to accept one of them or to remain single. If she accepts, she will become "engaged" to the proposing man; however, the engagement is not definitive and can be broken in the subsequent stages.

At the next stage, every man who is not engaged asks for the hand of his next best choice (if he has one). Every woman chooses between the proposals she may

Uncorrected proof

have received, her current fiancé and staying single. If she accepts a new proposal, she breaks the previous engagement and announces a new one.

The algorithm is repeated until it reaches a state where all men are engaged or the men who are still single have asked for the hand of all women they are willing to marry. Since the maximum number of women that a man can propose to is m, the algorithm must terminate after at most m iterations. At this point all engagements are confirmed and the weddings will (possibly) be celebrated.

We can test the functioning of this algorithm in our example. At the first stage, Brian and David both propose to Ann (who prefers David and thus becomes engaged to him) while Charles and Ingrid get engaged. At the second stage Brian – who is the only man not yet engaged – proposes to Emma, who accepts. All men are engaged and the algorithm terminates, producing configuration M as a stable matching.

Now, let us switch roles and have the women propose while the men accept. At the first stage, Ann becomes engaged to Charles, Emma to Brian and Ingrid to David (who declines Olivia's proposal). In the two following stages, Olivia approaches first Brian and then David, but is always rejected. At this point Olivia has proposed to all men she is willing to marry, and the algorithm terminates, producing the configuration F that is again a stable assigment.

Recall that men and women have opposite collective preferences over the stable matchings. The men prefer M to F and the women vice versa. In our example, the version of the algorithm where the men propose leads to the stable matching M; symmetrically, the version where the women propose leads to the stable matching F. This holds in general: among all stable matchings, this algorithm always finds the one that is collectively preferred by the sex to which the task of proposing is entrusted. Of course, more complex algorithms that are capable of finding less extreme stable matchings exist.

It is worth noting that the right to accept or decline a marriage offer is in fact less advantageous than having to make the offer. The reason is intuitively simple. The one who makes the offer starts out from his first choice and proceeds downwards, worsening his situation only when forced to. His role is active. The one who accepts offers, on the other hand, has access only to the best partner knocking on his door, but is not allowed to actively begin courting another partner that he considers to be better. Thus, he has a passive role.

Let us make two observations. The first one is that if the two sexes were to debate which version of the algorithm is preferable, the men would defend the first and the women the second one. The conflict between the sexes would shift from collective preferences over stable matchings to the choice of the method used to find one.

The second remark is of a more speculative nature. The version of the algorithm where the men propose roughly recalls the prevailing way the marriage game used to be organised in the Western world of the nineteenth century. Considering that this way favours the men, might we say that part of the progress towards equality between the sexes has been to teach the women of our century not to leave all the initiative to men?

Uncorrected proof

A Pinch of Poligamy

One of the crucial assumptions of the marriage game is that the matchings are monogamous. However, thinking of a version of the marriage game with a more exotic flavour, we could imagine that each man might be allowed to marry up to four women. What happens if we modify the marriage game allowing one of the sexes to practice poligamy?

Things become quite a bit more interesting if we change subject and, instead of men and women, speak about firms and workers or about universities and students: usually a firm hires more than one worker, while each worker is an employee of only one firm; similarly, a university admits various students, while each student is enrolled at only one university. Thus we can view their relationships as a form of marriage game in which one of the sides has the right to practice poligamy while the other one does not.

Consider n universities and m students. Each university can accept the enrollment of several students, possibly up to the number of positions that are available. Each student can enrol at only one university. Universities need not accept all applications they receive and students are not obliged to enrol at a university. Every student has a preference ranking over universities, based on which he is never indifferent between two universities.

Universities have a preference ranking over students with no ties, that satisfies also another assumption: the acceptance of a student by the university does not depend on who has already been admitted. In formal terms, if a university considers student x to be better than student y in a direct comparison, this remains true when the university has already accepted the enrollment of some students and has to evaluate whether it now prefers taking x or y. This assumption excludes cases of affirmative action, where a student who is individually preferred ends up being rejected because the school prefers to admit a student who is ranked worse but represents an ethnic group or some other socially disadvantaged group. Our assumption imposes that the evaluation of a student depends only on his intrinsic merits and does not take into consideration also how many students of the different ethnic groups have already been admitted.

Students and universities try to combine matchings amongst themselves, or in a more bureaucratic language, to form the classes for the coming school year. This is a form of marriage game where unilateral poligamy is allowed for the universities. By means of an appropriate transformation, this problem can be reduced to a marriage game without poligamy. Therefore, many of the mathematical properties we have seen persist in this model.

In particular, the following four theorems, which are analogous to those presented for the marriage game in the monogamous regime, hold:

- there always exists a stable matching;
- the stable matching need not be unique;
- if a student is not admitted to any university in some stable matching, he will not be admitted to any university in any stable matching; furthermore, the number of

uncorrected proof

available positions that a university succeeds in filling is the same in any stable matching and, if the number of students is less than the number of available positions, even the set of students that are accepted will be the same;
– universities and students have opposite collective preferences over the stable matchings.

The algorithm of Gale and Shapley for finding a stable matching generalises in a natural way. As before, we need to assign to one side the task of proposing and to the other one that of accepting. The side which proposes is favoured in the sense that the algorithm generates its collectively preferred stable matching. Given that students and colleges have opposite collective preferences, this again poses the problem of choosing which side to entrust with the task of proposing. In this case, however, the asymmetry between the poligamous side and the monogamous side suggests to favour the latter, which appears weaker. In fact, traditionally, it is the students who apply at a university or the workers who search for employment.

It is Not Just a Game

As in many other countries, US medical graduates have to do an internship (or residency) with a hospital department. In the first decades of the twentieth century, the competition between hospitals for the best interns and the competition between graduates for the best internships had produced a climate in which the offers of internships were made too early. In the fourties, offers were often made at the beginning of the third year of study, when there was not yet enough information on the abilities and the preparation of the student. On their part, students had to accept or reject an offer without knowing which other offers they could have received subsequently. All in all, the selection process for interns was chaotic.

Between 1945 and 1951 a serious attempt was made to establish a unique deadline for accepting an offer, but the effort did not achieve lasting effects. The hospitals continued to impose very short deadlines constraining students to decide in the dark. Hospitals, on their part, had to embark on complicated and troubling investigations because, when a student declined an offer, it was often too late to make an offer to the second choice. In 1951 the chaos was tamed by adopting a centralised algorithm for appointments, known as the National Resident Matching Program (NRMP, in the following), as reported in Roth [5].

The NRMP, which every year establishes the appointments of residents to hospitals for the whole of the United States, is obviously of great importance. However, although the algorithm worked rather well, nobody had a clear explanation of its success.

In 1962 Gale and Shapley [4], without knowing the NRMP, described the marriage game in the *American Mathematical Monthly*. In 1984, Alvin Roth [6] became aware that the NRMP was substantially equivalent to the algorithm of Gale and Shapley and rapidly understood the reason for its success: the problem of assigning residents to hospitals is a version of the marriage game with unilateral poligamy.

uncorrected Proof

The centralised algorithm of the NRMP works very well because it produces stable matchings.

This intuition has allowed to update the NRMP so as to take into account changes in the characteristics of the participants. For example, one of the biggest problems recently faced is that the number of couples of residents who look for two positions in the same geographical area has increased in the course of time (cf. [7]). This is a natural consequence of the evolution of customs. While in the fifties the typical resident was single or married to a partner who was prepared to follow him, now it is rather customary that medical students meet their partner while studying. Thus couples are formed, who at the end of their studies wish to live together and therefore look for two resident positions at hospitals in the same vicinity.

Applications of the marriage game do not end here. All over the world other centralised algorithms, similar to the NRMP, have been proposed or are effectively in use, often inspired by the NRMP or the theoretical study of the marriage game and the algorithm of Gale and Shapley [4]. Among the applications in which these algorithms have been successful are: mechanisms analogous to the NRMP in Canada and Great Britain; general rules for entering the legal profession in some provinces in Canada; the National Panhellenic Conference in the US for assigning female students to social organisations known as *sororities*, including housing assignment; the market for the recruitment of football players in American university colleges.

Amongst the mechanisms recommended to improve upon the existing situation, we mention the proposal to modify the system of assigning students to universities in Turkey, that already takes place in a centralised way on a national basis [8]. Another recent and important proposal, however based on techniques that were created to solve assignment problems different from the marriage game, suggests the establishment of a centralised market for kidney exchange between patients on the waiting list for a transplant and donors [9]. This proposal has been approved in September 2004 by the Renal Transplant Oversight Committee of New England; it will most likely be possible to judge its effectiveness soon.

Bibliography

[1] J.M. Gottman, J.D. Murray, C. Swanson, R. Tyson, K.R. Swanson (2003) *The Mathematics of marriage: Dynamic nonlinear models*, The MIT Press, Cambridge (USA).

[2] P.S. Bearman, J. Moody, K. Stovel (2004) Chains of affection: The structure of adolescent romantic and sexual networks, *American Journal of Sociology* 100, pp. 44–91.

[3] A.E. Roth, M.A.O. Sotomayor (1990) *Two-sided Matching: A Study in Game-Theoretic Modeling Analysis*, Cambridge University Press, Cambridge (UK).

[4] D. Gale and L.S. Shapley (1962) College admission and the stability of marriage, *American Mathematical Monthly* 69, pp. 9–14.

[5] A.E. Roth (2003) The origins, history and design of the Resident Match, *Journal of the American Medical Association* 289, pp. 909–912.

[6] A.E. Roth (1984) The evolution of the labor market for medical interns and residents: A case study in game theory, *Journal of Political Economy* 92, pp. 991–1016.

uncorrected proof

[7] A.E. Roth (2003) The economist as engineer: Game theory, experimentation, and computation as tools for design economics, *Econometrica* 70, pp. 1341–1378.
[8] M. Balinski and T. Sönmez (1999) A tale of two mechanisms: Student placement, *Journal of Economic Theory* 84, pp. 73-94.
[9] A.E. Roth, T. Sönmez, M.U. Ünver (2004) Kidney exchange, *Quarterly Journal of Economics* 119, pp. 457–488.

uncorrected proof