COLLEGE ADMISSIONS AND THE STABILITY OF MARRIAGE

David Gale*
Consultant, The RAND Corporation

and

L. S. Shapley
Mathematics Department
The RAND Corporation

P-2240

March 3, 1961

*Professor, Mathematics Department
Brown University

COLLEGE ADMISSIONS AND THE STABILITY OF MARRIAGE

D. Gale and L. S. Shapley

1. Introduction

The problem which we shall be concerned with relates to the following typical situation: A college is considering a set of n applicants of which it can admit a quota of only q. Having evaluated all their qualifications, the admissions office must decide which ones to admit. The procedure of offering admission only to the q best-qualified applicants will not generally be satisfactory, for it cannot be assumed that all who are offered admission will accept. The reason is of course that many applicants will apply to and be admitted by more than one college and hence will accept only their first choice. Accordingly, in order for a college to receive q acceptances, it will generally have to offer to admit more than q applicants. The problem of determining how many and which ones to admit in order to achieve the desired quota requires some rather involved guesswork. It may not be known (a) whether a given applicant has also applied elsewhere; if this is known it may not be known (b) how the applicant ranks the colleges to which he has applied; even if this is known it will not be known (c) which of the other colleges will offer to admit the given applicant. A result of all this uncertainty is that colleges can expect only that the entering class will come reasonably close in numbers to the desired quota, and be reasonably close to the attainable optimum in quality. Another result is an annual increase in the number of grey hairs on deans of admissions.

The usual admissions procedure, which we shall call the AAA procedure, (Apply, Admit, Accept), presents problems for the applicants as well as the colleges. An applicant who is asked to list in his application all other colleges applied for in order of preference may feel, perhaps not without reason, that by telling a college it is, say, his third choice he will be hurting his chances of being admitted.

One elaboration of the AAA procedure is the introduction of the "waiting list," whereby an applicant can be informed that he is not admitted but may be admitted later if a vacancy occurs. This introduces new problems. Suppose an applicant is accepted by one college and placed on the waiting list of another that he prefers. Should he play safe by accepting the first or take a chance that the second will admit him later? Is it ethical to accept the first without informing the second and then withdraw his acceptance if the second later admits him?

We contend that all the difficulties here described can be avoided. We shall describe a procedure for assigning applicants to colleges which should be satisfactory to both groups, which removes all uncertainties and which, assuming there are enough applicants, assigns to each college precisely its quota.

2. The Assignment Criteria

A set of n applicants is to be assigned among m colleges where q_i is the quota of the ith college. Each applicant ranks the colleges in the order of his preference, assigning 1 to his first choice, 2 to his second, etc., omitting only those colleges which he would never accept under any circumstances. For convenience we assume there are no ties; thus, if an applicant is indifferent between two or more colleges he is nevertheless required to list them in some order. Each college now ranks the students who have applied to it in order of preference, having first eliminated from consideration any of these applicants whom it would not admit under any circumstances even if it meant not filling its quota. From these data, consisting of the quotas of the colleges and the preferences of the colleges and the applicants, we wish to determine an assignment of applicants to colleges in accordance with some agreed upon criterion of fairness.

Stated in this way and looked at superficially, the solution may at first appear obvious. One merely makes the assignments "in accordance with" the given preferences. A little reflection shows that complications may arise. An example is the simple case of two colleges, A and B, and two applicants, α and β, in which α prefers A and β prefers B, but A prefers β and B prefers α. Here, no assignment can satisfy all preferences. One must decide what to do about this sort of situation. On the philosophy that the colleges exist for the students rather than the other way around, it would be fitting to assign α to A and β to B. This suggests the following admittedly vague principle: other things being equal, students should receive consideration over colleges. This remark is of little help in itself, but we will return to it later after taking up another more explicit matter.

The key idea in what follows is the assertion that--whatever assignment is finally decided on--it is clearly desirable that the situation described in the following definition should <u>not</u> occur:

DEFINITION: An assignment of applicants to colleges will be called unstable if there are two applicants α and β who are assigned to colleges A and B respectively, although β prefers A and A prefers β.

The term "unstable" is suggestive, for suppose the situation described above did occur. Then applicant β could indicate to college A that he would like to transfer to it, and A could respond by admitting β, after letting α go so as to remain within its quota. Both A and β would consider the change an improvement. The original assignment is therefore unstable in the sense that it can be upset by a college and applicant acting together in a manner which benefits both.

Our first requirement on an assignment is that it shall not exhibit instability. This immediately raises the mathematical question: will it always be possible to find an assignment which is free of instability? An affirmative answer to this question will be given in the next section, and while the proof is not difficult, the result seems not entirely obvious, as some examples will indicate. Assuming for the moment that stable assignments do exist, we must still decide which among possibly many stable solutions is to be preferred. We now return to the philosophical principle mentioned earlier and give it a precise formulation.

DEFINITION: A stable assignment is called <u>optimal</u> if every applicant is just as well off under it as under any other stable assignment.

Even granting the existence of stable assignments it is far from clear that there are optimal assignments. However, one thing that is clear is that the optimal assignment, if it exists, is unique. Indeed, if there were two such assignments, then, since they are distinct, at least one applicant must be differently assigned, and hence (by our "no tie" rule) would be better off under one than under the other; hence one of the assignments would not be optimal after all. Thus the principles of stability and optimality will lead us to a unique "best" method of assignment.

3. Stable Assignments and a Marriage Problem

In trying to settle the question of the existence of stable assignments we were led to look first at a special case, in which there are the same number of applicants as colleges and all quotas are unity. This situation is, of course, highly unnatural in the context of college admissions, but there is another "story" into which the above problem fits quite readily.

A certain community consists of n men and n women. Each person ranks those of the opposite sex in accordance with his or her preferences for a marriage partner. We seek a satisfactory way of marrying off all members of the community. Imitating our earlier definition, we call a set of marriages unstable (and here the suitability of the term is quite clear) if under it there are a man and a woman who are not married to each other but prefer each other to their actual mates.

Question: For any pattern of preferences is it possible to find a stable set of marriages?

Before giving the answer let us look at some examples.

Example 1. The following is the "ranking matrix" of three men, α, β, and γ, and three women, A, B, and C.

	A	B	C
α	1,3	2,2	3,1
β	3,1	1,3	2,2
γ	2,2	3,1	1,3

The first number of each pair in the matrix gives the ranking of women by the men, the second number is the ranking of the men by the women. Thus, α ranks A first, B second, C third, while A ranks β first, γ second, and α third, etc.

There are six possible sets of marriages; of these, three are stable. One of these is realized by giving each man his first choice, thus α marries A, β marries B, and γ marries C. Note that although each woman gets her last choice, the arrangement is nevertheless stable. Alternatively one may let the woman have their first choices and marry α to C,

β to A, and γ to B. The third stable arrangement is to give everyone his or her second choice and have α marry B, β marry C, and γ marry A. The reader will easily verify that all other arrangements are unstable.

Example 2. The ranking matrix is the following.

	A	B	C	D
α	1,3	2,3	3,2	4,3
β	1,4	4,1	3,3	2,2
γ	2,2	1,4	3,4	4,1
δ	4,1	2,2	3,1	1,4

There is only the one stable set of marriages indicated by the circled entries in the matrix. Note that in this situation no one can get his or her first choice if stability is to be achieved.

Example 3. A problem similar to the marriage problem is the "problem of the roommates." An even number of boys wish to divide up into pairs of roommates. A set of pairings is called <u>stable</u> if under it there are no two boys who are not roommates and who prefer each other to their actual roommates. An easy example shows that there can be situations in which there exists no stable pairing. Namely, consider boys α, β, γ and δ, where α ranks β first, β ranks γ first, γ

ranks α first, and α, β and γ all rank δ last. Then regardless of δ's preferences there can be no stable pairing, for whoever has to room with δ will want to move out, and one of the other two will be willing to take him in.

The above examples would indicate that the solution to the stability problem is not entirely obvious. Nevertheless,

Theorem 1. <u>There always exists a stable set of marriages</u>.

<u>Proof</u>. We shall prove existence by giving an iterative procedure for actually finding a stable set of marriages.

To start, let each boy propose to his favorite girl. Each girl who receives more than one proposal rejects all but her favorite from among those who have proposed to her. However, she does not accept him yet, but keeps him on a string to allow for the possibility that something better may come along.

We are now ready for the second stage. Those boys who have been rejected now propose to their second choice. Each girl receiving proposals chooses her favorite from the group consisting of the new proposals and the boy on her string, if any. She rejects all the rest and again keeps the favorite in suspense.

We proceed in the same manner. Those who are rejected at the second stage propose to their next choices, and the girls again reject all but the best proposal they have had so far.

Now if we continue in this way, eventually (in fact, in at most $n^2 - 2n + 2$ stages) every girl will have received a proposal, for as long as any girl has not been proposed to there will be rejections and as long as there are rejections there will be new proposals, but since no boy can propose to the same girl more than once, every girl is sure to get a proposal in due time. As soon as the last girl gets her proposal the courtship is over, and each girl is now required to accept the boy on her string.

We claim that this set of marriages is stable. Namely, suppose John and Mary are not married to each other but John prefers Mary to his wife Jane. Then John must have proposed to Mary at some stage and subsequently been rejected in favor of someone that Mary liked better. It is now clear that Mary must prefer her husband to John. Hence there is no instability.

The reader may amuse himself by applying the procedure of the proof to solve the problems of Examples 1 and 2, or the following example which requires ten iterations:

	A	B	C	D
α	1,3	2,2	3,1	4,3
β	1,4	2,3	3,2	4,4
γ	3,1	1,4	2,3	4,2
δ	2,2	3,1	1,4	4,1

The condition that there be the same number of boys and girls is not essential. If there are b boys and g girls with $b < g$, then the procedure terminates as soon as b girls have been proposed to. If $b > g$ the procedure ends when every boy is either on some girl's string or has been rejected by all of the girls. In either case the set of marriages that results is stable.

It is clear that there is an entirely symmetrical procedure, with girls proposing to boys, which must also lead to a stable set of marriages. The two solutions are not generally the same; indeed, we shall see in a moment that when the boys propose, the result is optimal for the boys, and when the girls propose it is optimal for the girls. Thus the solutions by the two procedures will be the same only when there is a unique stable set of marriages.

4. Stable Assignments and the Admissions Problem

The extension of our "deferred-acceptance" procedure to the problem of college admissions is straightforward. For convenience we will assume that if a college is not willing to accept a student under any circumstances, as described in Section 2, then that student will not even be permitted to apply to the college. With this understanding the procedure follows: First, all students apply to the college of their first choice. A college with a quota of q then places on its waiting list the q applicants who have applied to it who rank highest, or all who have applied if there are fewer than q, and rejects the rest. Rejected students then apply to their next choice and again each college selects the top q from among the group consisting of the new applicants and those on its waiting list from the previous round, and puts these on its new waiting list, rejecting the rest. The procedure terminates when every applicant is either on a waiting list or has been rejected by every college to which he is willing or permitted to apply. At this point each college admits everyone on its waiting list and the stable assignment has been achieved. The proof that the assignment is stable is entirely analogous to the proof given for the marriage problem and is left to the reader.

5. Optimality

We now show that the "deferred acceptance" procedure described in the previous section yields not only a stable but an optimal assignment of applicants. That is,

Theorem 2. *Every applicant is at least as well off under the assignment given by the procedure just described as he would be under any other stable assignment.*

<u>Proof</u>. Let us call a college "possible" for a particular applicant if there is a stable assignment that sends him there. The proof is by induction. Assume that up to a given point in the procedure no applicant has yet been turned away from a college that is possible for him. At this point suppose that college A, having received applications from a full quota of better-qualified applicants $\beta_1, \ldots, \beta_q$, rejects applicant α. We must show that A is impossible for α. We know that each β_i prefers college A to all the others, except for those that have previously rejected him, and hence (by assumption) are impossible for him. Consider a hypothetical assignment that sends α to A and everyone else to colleges that are possible for them. At least one of the β_i will have to go to a less desirable place than A. But this arrangement is unstable, since β_i and A could upset it to the benefit of both. Hence the hypothetical assignment is unstable and A is impossible for α. The conclusion is that our procedure only rejects applicants from colleges which they could not possibly be admitted to in any stable assignment. The resulting assignment is therefore optimal.

Parenthetically we may remark that even though we no longer have the symmetry of the marriage problem, we can still invert our admissions procedure to obtain the unique "college optimal" stable assignment. The inverted method bears some resemblance to a fraternity "rush week"; it starts with each college making bids to those applicants it considers most desirable, up to its quota limit, and then the bid-for students reject all but the most attractive offer, and so on.

6. Some Practical Questions

The theorems we have presented are based on a mathematically "constructive" procedure for arriving at an optimal assignment. It may well be asked how close this procedure is to a truly practicable method for actually assigning real students to real colleges. We shall mention briefly some of the difficulties that would have to be overcome.

On the face of it, our "deferred acceptance" method involves a great deal of communication back and forth if the description we have given is to be followed literally, on a nationwide scale. The time required, if not the cost, would be prohibitive. An obvious modification would be to collect all relevant information at some central clearing house, and have an electronic computer run through the motions of "application," rejection," "placement on a waiting list," etc. The only drawback to this solution is the enormous amount of effort that would be expended by the participants in making up their complete preference orderings, most of it completely wasted. There would be no time for careful consideration of alternatives, by either the colleges or the applicants. This indicates that some sort of compromise between complete mechanization and no mechanization would have the best chance of success, with the candidates submitting to the clearing house only their top four or five choices (say) at the start.

In this connection the following observation is of interest: It is not essential in our procedure for a rejected applicant to apply _immediately_ to the next college of his choice. The final outcome is not changed if he "sits out" one or more rounds while other rejectees are making new applications. This means that the whole process would not have to grind to a halt whenever a single individual's partial list of preferences happens to be exhausted through repeated rejections.

There are several other complicating factors aside from the sheer magnitude of the computation. One of them is that admission, or the acceptance of admission, is often tied in with other things, like the award of financial aid, advanced academic standing, or the like. The college's scholarship budget constitutes a new constraint, comparable to the quota q, and a more elaborate mathematical model would clearly be called for. An adequate theory would have to cope with such details as scholarship awards of different sizes, or students who would rather go to college A without a scholarship than to college B with one. Also, some colleges might prefer to give a poor boy a scholarship rather than admit a better-qualified rich boy and have money left over.

Still another problem is presented by colleges that pay attention to the overall composition of the entering class. Their preference scales are obviously more complex than a simple ranking of individual applicants. To meet this situation, a further generalization of our theory is called for, particularly in the definition of stability.

While it is not clear that either of the generalizations indicated above can be formulated in such a way as to preserve the content of Theorems 1 and 2, it is the authors' opinion that the "deferred acceptance" idea, insofar as it can be made practicable, is capable of producing assignments that are markedly superior to those obtained under the present system, or lack of system. The two important points at which "deferred acceptance" differs from current practice are the following: (1) Applicants are assured that they run no risk of losing out on their lower choices by default while their higher choices are being considered, and (2) colleges are assured that the candidates on their waiting lists have no other applications pending: if offered admission they will accept it.

7. Addendum on the Nature of Mathematics

The problem which we have discussed here seems to be of some interest both from an abstract mathematical point of view and from that of practical application. There is another aspect of the problem which may be worth mentioning. As an exercise in mathematical reasoning, it provides a counter-example to some of the stereotypes that non-mathematicians believe mathematics to be concerned with.

Most mathematicians at one time or another have probably found themselves in the position of trying to refute the notion that they are people with "a head for figures," or that they "know a lot of formulas." At such times, it may be convenient to have an illustration at hand to show that mathematics need not be concerned with figures, either numerical or geometrical. For this purpose, we recommend the statement and proof of our Theorem 1. The argument is carried out not in mathematical symbols but in ordinary English, there are no obscure or technical terms. Knowledge of calculus is not presupposed. In fact, one hardly needs to know how to count. Yet any mathematician will immediately recognize the argument as mathematical.

What then, to raise the old question once more, is mathematics? The answer, it appears, is that any argument which is carried out with sufficient precision is mathematical, or, as has been remarked not entirely facetiously, the difference between mathematicians and other people is that a mathematician is able to conceive of an argument requiring more than two steps.