



WP/04/190

IMF Working Paper

The Internal Job Market of the IMF's Economist Program

Greg Barron and Felix Várdy

IMF Working Paper

IMF Institute

The Internal Job Market of the IMF's Economist Program

Prepared by Greg Barron and Felix Várdy¹

Authorized for distribution by Sunil Sharma

October 2004

Abstract

This Working Paper should not be reported as representing the views of the IMF.

The views expressed in this Working Paper are those of the author(s) and do not necessarily represent those of the IMF or IMF policy. Working Papers describe research in progress by the author(s) and are published to elicit comments and to further debate.

This paper shows how the internal job market for participants in the IMF Economist Program (EPs) could be redesigned to eliminate most of the shortcomings of the current system. The new design is based on Gale and Shapley's (1962) Deferred Acceptance Algorithm (DAA) and generates an efficient and stable outcome. An Excel-based computer program, *EP-Match*, implements the algorithm and applies it to the internal job market for EPs. The program can be downloaded from http://www.people.hbs.edu/gbarron/EP-Match_for_Excel.html.

JEL Classification Numbers: C78, D73

Keywords: Matching, stability, marriage problem, deferred acceptance algorithm, Gale-Shapley, entry-level job market

Author(s) E-Mail Address: gbarron@hbs.edu and fvardy@imf.org

¹ Greg Barron is an Assistant Professor at Harvard Business School and Felix Várdy is a second-year participant in the Economist Program, currently in the Asian Division of the IMF Institute. The authors would like to thank David Coe for encouragement to write this paper, and Abdul Abiad, Chris Clarke, Andrew Feltenstein, Mohsin Khan, Jorge Marquez-Ruarte, Dayalini Mendez, Benoit Mercereau, Miguel Messmacher, Jacques Miniane, Saleh Nsouli, Al Roth, Martin Schindler, Sunil Sharma, Clinton Shiells, Jay Surti, and Zhiwei Zhang for comments and suggestions. All errors are our own.

Contents	Page
I. Introduction	3
II. The Current Situation	3
A. Decentralized Matching of Transferring and Graduating EPs	3
B. Centralized Matching of Incoming EPs	5
III. The Deferred Acceptance Algorithm	5
A. Preliminaries	5
B. The Algorithm	6
C. Strategic Issues	8
Strategy-proofness	8
The DAA and the Pressure to Pre-commit	9
Departments or Positions?	9
D. Summary	10
IV. The Computer Program	10
A. Description	10
B. Examples	11
Example 1	11
Example 2	14
V. Conclusion	16
Appendix: Matching of Incoming EPs	17
A. Introduction	17
B. Methodology	17
C. Stability	18
D. Pareto Efficiency	18
E. Conclusion	18
References	26
Figures:	
1. <i>EP Match</i> : Instructions	20
2. <i>EP Match</i> : Departments' Preferences	21
3. <i>EP Match</i> : EPs' Preferences	22
4. <i>EP Match</i> : The Match	23
Tables:	
1. Stable and Unstable Centralized Mechanisms	19
2. Participants' Preferences	24
3. HRD's Match	25

I. INTRODUCTION

The Economist Program is the gate of entry into the IMF for young economists joining the Fund soon after graduate school. Each year, 35 to 45 participants, usually referred to as “EPs,” join the Program for a two-year period. During this time, EPs work in two different departments of the Fund. After graduating from the Program, EPs are considered for permanent staff positions.

In the context of the Economist Program, each EP is matched to a department on three occasions. First, when he joins the Fund in June or October;² then, when he transfers to a second-year position one year later; and, finally, when the EP is assigned to a permanent position at the end of the Program.

Matching of EPs to second-year assignments and permanent positions takes place through a decentralized, internal job market system. Matching of incoming EPs to first-year positions is done centrally, by the IMF’s Human Resources Department (HRD).

This paper argues that the procedures currently in place exhibit structural shortcomings. It then proposes an alternative centralized matching mechanism, based on Gale and Shapley’s (1962) Deferred Acceptance Algorithm (DAA), that resolves most of these shortcomings. Finally, it develops *EP-Match*, an Excel-based computer program that implements the DAA, and applies it to the internal job market for EPs.

II. THE CURRENT SITUATION

A. Decentralized Matching of Transferring and Graduating EPs

Transferring and graduating EPs are matched to second-year assignments and permanent positions through a decentralized job market system. However, to promote equality of opportunities, and to prevent unraveling of the market,³ EPs and departments are supposed to abide by strict rules as to when they can first contact each other, and when they can start making and accepting job offers.

² In this paper, the pronoun “he” refers to a person of unspecified gender.

³ By “unraveling” we mean the phenomenon, common in entry-level labor markets, that employers make job offers earlier from year to year. This is done in an effort to preempt the competition and recruit the best candidates. Once unraveling has started, it is hard to stop and often spins out of control. In the market for federal appellate law clerks, for example, job offers have recently been made almost two years in advance of employment. (See Avery and others, 2001.) The same holds for entry-level gastroenterology positions for medical doctors. Interviews for these positions now take place even before students have had the opportunity to explore other subspecialties (see Niederle and Roth, 2003). In markets that experience unraveling, applicants typically receive “exploding” offers that must be accepted or rejected on short notice.

In terms of Roth and Xing's (1994) taxonomy of entry-level labor markets, the internal job market for transferring and graduating EPs is in "Stage 2." Here, "Stage 1" refers to an entirely unregulated job market, while "Stage 3" refers to a fully centralized market. Their detailed description of the symptoms of a typical "Stage 2" market is so recognizable that we quote it here at some length.

"In stage 2, a (...) market organization (...) attempts to establish a uniform date before which offers should not be made, and often an earlier date before which interviews should not be conducted, and a later date (or time) before which candidates who have received offers should not be required to respond. Sometimes this is hardly successful at all, with many market participants ignoring or circumventing the rules, and those who obey them quickly finding that this puts them at a disadvantage. And even when uniform dates are successfully established and maintained, the market often experiences a great deal of congestion and chaotic behavior, as the deadlines for accepting or rejecting offers grows near. A firm is eager to know if its offers will be accepted, in time so that if it has unfilled positions it may approach its most preferred alternative candidates before they have had to accept any offers they have received. And candidates who have received offers, but not from their first choice firm, are intent upon waiting until the last allowable moment before accepting any offer, in the hope of receiving a better one. Particularly if, as often seems to be the case, some fraction of candidates holds on to multiple offers as the final deadline approaches, this means that just before the deadline expires many transactions still remain to be made. Firms whose first choice candidates reject them may now find that their next dozen candidates have already accepted offers, and candidates may receive preferred offers moments after making a verbal commitment to accept an earlier offer. In some markets such verbal commitments are virtually always honored, and in others they are sometimes reneged on. In either event, in the aftermath, many firms and candidates have just missed making connections they would have preferred. The result is that the following year witnesses a resurgence of strategic behaviors designed to avoid being caught short at the end of the market. Often new rules are formulated to prohibit the more brazen of these, and new adaptations are made. While some markets have persisted for many seasons in this fashion, systems of formalized dates are often abandoned, with the market either reverting to stage 1 [an unregulated market], or moving on to stage 3 [a centralized market]." [Roth and Xing, 1994, p. 996]

From personal experience, one of the authors can attest that the internal EP job market for transferring and graduating EPs exhibits virtually all of the symptoms described by Roth and Xing. Specifically:

- The IMF's Human Resources Department (HRD) tries to prevent unraveling by enforcing rules on market timing, e.g., timing of first contact, interviewing, making and accepting offers.
- Departments and EPs routinely try to circumvent these rules, while those participants who do obey the rules often end up with rather unfavorable outcomes. A prime example is the rule that EPs and departments should not contact each other before a certain date. This rule is clearly unenforceable.
- When a department makes a job offer to an EP, it often pressures the EP to accept or reject the offer on the spot, or on very short notice.
- Conversely, if an EP can get away with it, he often tries to postpone a decision on a particular offer until the very last moment, in the hope of receiving a better offer.

- Sometimes, the hoarding of multiple offers by a few sought after EPs leads to a virtual standstill in the market. And once these EPs have made up their minds, a flurry of offers and acceptances takes place in a very short time period. (This kind of “congestion” is studied extensively in Roth and Xing, 1997.)
- In response to continuous complaints, HRD has modified the rules governing the market numerous times, apparently without much success.

In addition to the symptoms mentioned by Roth and Xing, it has been noted that:

- The current system hurts EPs who are traveling - or prevents departments from sending EPs on a business trip during the job market period - because of the limited time-frame in which EPs and departments are allowed to interact.
- The current system creates a lot of stress for the participants, as they consider the process to be rather unpredictable and not particularly fair.

B. Centralized Matching of Incoming EPs

HRD matches incoming EPs to departments on the basis of preferences of EPs and departments over each other. These preferences are solicited in the form of rank-order lists (ROLs). Participants who do not respond are assumed to be indifferent.

The process by which HRD translates the participants’ preferences into a match of EPs and departments is heuristic. This means that HRD tries to find the “best fit” without formulating an explicit objective function or a procedure to resolve conflicting preferences.

The drawback of a heuristic approach is that the resulting match is not guaranteed to be Pareto efficient and/or stable.⁴ Indeed, we show in the Appendix that, for a recent cohort of incoming EPs for which we were able to obtain the necessary data, HRD’s match was in fact neither.

Motivated by the shortcomings of the matching systems currently in place, in the next section, we discuss Gale and Shapley’s Deferred Acceptance Algorithm.

III. THE DEFERRED ACCEPTANCE ALGORITHM

A. Preliminaries

The internal job market for EPs can be viewed as a “College Admissions Problem” (CAP). The CAP is concerned with two disjoint groups, which we shall refer to as EPs and departments, that have to be matched to each other while taking into account their mutual

⁴ A formal definition of stability is given in the next section.

preferences. The CAP was first formulated by David Gale and Lloyd Shapley in a path-breaking paper in the *American Mathematical Monthly* (Gale and Shapley, 1962).

Gale and Shapley (GS) argue that a match of EPs and departments, denoted by m , solves the CAP only if it is *stable*. That is, only if there is no (EP, department)-pair that would prefer to be matched to each other instead of sticking to their matches prescribed by m . Obviously, a matching process that violates this property can easily be upset by an EP and a department getting together in a manner that benefits both. Such an (EP, department)-pair is called a “blocking pair.” The empirical literature indeed confirms that stability is a very important property when designing a matching market to fix unraveling and associated “Stage 2” market failures (see Table 1).

GS prove that, for the CAP, a stable match always exists.⁵ Although there may be multiple stable matches for a given preference profile, when preferences are strict, the set of positions filled and EPs employed is the same in all stable matches (Roth and Sotomayor, 1990). Another attractive feature of stability in the CAP is that it implies Pareto efficiency. (This follows from the fact that all stable matches are in the “Core.” See, e.g., Roth and Sotomayor, 1990.) The reverse, however, is not true. That is, typically, there are many Pareto efficient matches that are unstable. Thus, stability is a strictly stronger requirement than Pareto efficiency.

In their original paper, GS also provide an iterative procedure for finding a stable match with respect to the participants’ stated preferences. This procedure is called the Deferred Acceptance Algorithm. Currently, (a variant of) the DAA is used to match applicants and employers in virtually all medical and dental residency programs, pharmacy practices, and psychology internships, both in the United States and Canada.⁶ (See, e.g., www.nrmp.org.)

B. The Algorithm

The DAA is best illustrated in a simplified environment in which each department has exactly one vacancy and the total number of vacancies is equal to the total number of EPs. This simple set-up is known in the literature as the “Marriage Problem.” Subsequently, we

⁵ That this result is far from trivial is illustrated by the closely related “Roommates Problem.” In the Roommates Problem, an even number of boys has to be paired-up amongst each other. GS show that, in this case, a stable matching may not exist. Instead, cycles of blocking pairs may cause the set of stable matchings to be empty: Consider four boys, A, B, C, and D. Let A rank B first (i.e., B is A’s most preferred roommate), while B ranks C first, and C ranks A first. At the same time, A, B, and C all rank D last. Then, regardless of D’s preferences, there exists no stable matching; whoever is matched to D will want to move out, and one of the other two will be willing to take him in.

⁶ The variant used is the Roth and Peranson algorithm (1999). Its advantage over the standard DAA is that it can handle couples applying together.

can modify the DAA to handle the general case with an arbitrary number of departments, vacancies per department, and EPs.

It should be noted that, under the DAA, EPs and departments continue to make initial contact and arrange interviews in a decentralized manner. But in the end, instead of making offers directly, participants submit rank-order lists (ROLs) of each other to the market maker, i.e., HRD, who then runs the DAA. The output of the DAA is a stable match. That is, a stable allocation of EPs over departments.

For ease of exposition, we describe the algorithm as if offers and rejections in the DAA are made by the departments and EPs. In reality, it is the market maker that makes these decisions for them, based on the ROLs submitted by the participants.

For the simple case (“Marriage Problem”), the DAA works as follows:

Round 0: EPs and Departments rank each other in order of attractiveness.

Round 1: Each department makes a job offer to its top-ranked EP. An EP who receives one or more job offers rejects all but the most preferred among the offers received. The relatively most preferred job offer is kept on hold.

Round i ($i \geq 2$): A department that was rejected in the previous round makes an offer to its top-ranked EP among those who have not yet rejected it. An EP who receives new offers rejects all but the most preferred among the new offers received and the offer kept on hold from the previous round. Again, the relatively most preferred offer is kept on hold.

Stop: The algorithm terminates when no new job offers are made. At that point each EP has exactly one offer which he accepts.

Note that the roles of EPs and departments in the DAA can be reversed, such that it is the EPs who make the offers to the departments. While both procedures lead to a stable match, the outcomes are not necessarily identical. In fact, Gale and Shapley show that the DAA-generated stable match is (weakly) preferred over all other stable matches by the side making the offers, provided that preferences are strict.

The extension of the DAA to the general case is now straightforward. With k departments each having a quota of q_i positions, n EPs, and *EPs making the offers*, the algorithm proceeds as follows. First, all EPs apply to the department of their first choice. Department i with q_i positions then places on its wait list the q_i EPs whom it likes best, or all applying EPs if there are fewer than q_i . It rejects the rest. Rejected EPs then apply to the department of their second choice. Again, each department selects the top q_i among the new applicants and those on the waitlist, puts these on its new wait list and rejects the rest. The procedure terminates when every EP is either on a wait list or has been rejected by every department to which he is willing to apply. At this point, each department accepts the offers from everyone on its wait list and the EP-optimal stable match has been achieved.

Though we have implicitly assumed that all departments rank all EPs (and vice versa), in practice this is not really necessary. Instead, for each of its positions, a department should rank sufficiently many EPs such that it is reasonably assured of at least one EP accepting the job.⁷ This implies that the effort spent on interviewing and ranking EPs under the DAA will not differ much from the effort expended under the current system. Also in the current system, a department will want to avoid a situation in which, as a result of too many rejections, it has to start interviewing additional EPs after the market has opened and offers have been made. Because, by then, very few EPs are left to choose from.

Finally, rankings do not have to be strict. That is, being indifferent between some or all counterparts is allowed and resolved by randomization.⁸ Also, departments can classify certain EPs as unacceptable, and EPs can do the same for departments. We will get back to these issues later.

C. Strategic Issues

Strategy-proofness

Roth (1982) proves that there exists no stable matching mechanism, including the DAA, that is one-hundred-percent strategy-proof. That is, even for the DAA, one can “cook-up” a constellation of preferences such that at least one participant would gain by distorting his true preferences.

The catch is that strategic misrepresentation tends to require much more information about other participants’ preferences than is usually available. When participants in the DAA are sufficiently uncertain about the preferences of others, Roth and Rothblum (1999) show that they cannot gain by reversing the order of their true preferences. However, if towards the bottom of his ROL a participant is close to indifferent between being matched or not, and the

⁷ For example, in the National Resident Matching Program, each employer interviews and ranks only a tiny fraction of all potential candidates. This does not interfere with the workings of the DAA.

⁸ When a large number of participants on *both* sides of the market are indifferent between many counterparts, a complication may arise. For example, suppose that some EPs and some departments do not care who they are matched to, or have not submitted a ROL. In that case, the DAA proceeds with randomly generated (strict) preferences. Now, what if on the basis of these randomly generated preferences, an “indifferent” EP and an “indifferent” department end up matched to each other? Obviously, neither of them can or will object to that. However, there may be other, “non-indifferent” EPs or departments that would very much like to be matched to one of these “indifferent” participants. In that case, we might be able to accommodate one of them without hurting anybody else. In other words, when the randomization is such that “indifferent” participants end up matched to each other, the matching may not be strictly Pareto optimal. In that case, we should either redo the DAA with newly generated random preferences, or “manually” modify the matching to restore strict Pareto optimality. See the Appendix for a real-world example.

other side makes the offers, a participant might gain by cutting his ROL short, falsely claiming that the least preferred options are “unacceptable.” But in reality such a strategy is unlikely to be used much, because of its high risk: it may cause the EP or department to remain unmatched involuntarily.

In any case, truncation clearly is not an option for EPs, unless they are prepared to quit and leave the Fund. For departments, truncation is only possible in the market for permanent positions, since, at the end of the day, all transferring EPs have to be matched to a department.

Roth and Rothblum’s result of “almost-strategy-proofness in practice” goes a long way in explaining the popularity and robustness of the DAA in real-world applications such as the National Residency Matching Program for physicians. Indeed, Roth (1984, 1990, 1991) and Roth and Xing (1994) show that, with very few exceptions, centralized matching mechanisms survive if, and only if, they produce stable matches with respect to the participants (stated) preferences. This is illustrated in Table 1, taken from Roth and Rothblum, 1999.

The DAA and the Pressure to Pre-commit

In the current system, departments and EPs routinely flout the rules by committing to each other before the official opening of the market. This raises the question why the situation would be any different under the DAA. Specifically, what prevents departments and EPs from reducing the DAA to a mere formality by pre-committing to top-rank each other?

In the current system, having an offer rejected in the official market is potentially very costly for a department. The reason is that a rejected department may find that all other attractive EPs have already accepted offers from other departments. To mitigate this risk, departments routinely put extreme pressure on EPs to pre-commit to accepting a forthcoming job offer.

By contrast, rejections are costless in the DAA (at least, from a strategic perspective). The reason is that all acceptances and rejections occur simultaneously. Therefore, a department has no particular interest in making sure that its job offer is accepted by the first EP it is offered to. Indeed, the zero rejection cost in the DAA implies that a department can and should top-rank its favorite EP even if the probability of the EP accepting the offer is very small. In turn, this undermines the credibility of the exploding nature of a pre-market offer, as the department will top-rank its favorite EP, even if he refuses to pre-commit.

These considerations explain why, over time, unraveling and pre-commitment tend to occur less under the DAA than in decentralized markets. Nevertheless, it is true that pre-commitment might still occur on a smaller scale, particularly with repeated interaction. In this context, Niederle and Roth (2004) stress the importance of “market culture,” such as norms governing exploding offers.

Departments or Positions?

So far, we have assumed that EPs rank, and are matched to, departments. Usually, however, EPs do not have preferences over departments as such, but over specific positions within

departments. For example, while an EP's first choice might be to work on the China desk in the Asia and Pacific Department (APD), he might very well prefer to work on South Africa in the African Department over working on Tuvalu in APD. In this case, APD can do two things: either list the China and Tuvalu jobs separately, thereby allowing EPs to express interest specifically in China or specifically in Tuvalu; or lump both jobs together, only allowing for expressions of interest in working for APD without formal assurance of being assigned to any job in particular.

When EPs have preferences over particular positions instead of over departments, there is a trade-off between ex-ante clarity for EPs and ex-post flexibility for departments. If a department chooses to specify its vacancies in great detail, it lowers the risk and/or ambiguity for the applicants. In turn, this may help the department attract better EPs. On the other hand, if a department chooses to be less specific, it retains full flexibility to assign EPs to one job or another, depending on the need of the moment.

Presumably, competition will force all but the most popular departments to list their vacancies separately. From a transparency perspective, separate listings are indeed preferable.

D. Summary

The DAA solves many of the problems commonly encountered in matching markets. In the DAA, all final acceptances occur simultaneously. This eliminates premature decisions based on incomplete information. Also, undesirable phenomena such as unraveling, hoarding of multiple offers, and renegeing on a prior acceptance will no longer occur. At the same time, it is no longer necessary to impose rules on the participants that are hard to enforce. In particular, EPs and departments can communicate with each other as early and as often as they wish. The only hard constraint is that participants hand in their rankings of each other on a pre-announced date. The submitted rank-order lists (ROLs) then determine the matching of EPs to departments (or positions) through application of the DAA. Operationally, a market maker running the algorithm can execute all steps of the DAA in "real-time," one after the other, provided that EPs and departments have handed in their ROLs. This eliminates the need for repeated communication between EPs and departments, and thereby the need for EPs to be in Washington during any particular time period.

IV. THE COMPUTER PROGRAM

A. Description

To facilitate implementation of the DAA at the Fund, we have developed an Excel-based computer program, called *EP-Match*, that executes the algorithm and applies it to the Fund's internal job market for EPs. The user-interface of the program consists of four spreadsheets:

1. "Instructions" (see Figure 1).
2. "Departments' Preferences" (see Figure 2).

3. “EPs’ Preferences” (see Figure 3).
4. “Match” (see Figure 4).

To find the stable match generated by the DAA, the user, i.e., the market maker, sequentially works his way through the four spreadsheets. The first spreadsheet gives instructions on how to use the program and requires the market maker to enter the number of EPs and the number of departments (or positions) participating in the matching process. Here, the market maker must also indicate whether he wants to run the EP-proposing or department-proposing variant of the DAA. In the second spreadsheet, the market maker specifies the number of vacancies in each department and enters the departments’ rank-order lists (ROLs) in a matrix. The EPs’ ROLs go into the third spreadsheet. Finally, the fourth and last spreadsheet presents the resulting stable match.

The program is quite versatile. First, it allows for an arbitrary number of EPs, departments, and positions per department. Second, departments are allowed to be indifferent between certain EPs, and EPs between positions. Third, a department can be assured it will not be matched to an EP it finds unacceptable by not ranking that EP at all. Of course, it should only do this if it prefers the position to remain unfilled over hiring that EP. In principle, EPs can do the same with departments they do not want to join under any circumstance.

Finally, the program can also handle the following, slightly more complex situation. Suppose department Dx has multiple positions on offer, but its ranking of EPs is not the same for all positions. This might happen if a certain position requires particular skills, such as fluency in a foreign language or financial expertise. In that case, we partition the set of Dx ’s positions into r disjoint sub-sets S_1, \dots, S_r , such that, for all positions in a particular sub-set, Dx ’s ranking of EPs is in fact the same. Then, department Dx is replaced by multiple “virtual departments” Dx_1, \dots, Dx_r , such that Dx_i offers positions S_i and ranks the EPs accordingly. Finally, the program is run on the modified matching problem.

B. Examples

To illustrate how *EP-Match* works in practice, we now go through two examples.

Example 1

There are 3 departments, D1, D2, D3, and 6 EPs, EP1, ..., EP6. Each department has 2 vacancies. Let the departments’ ROLs be as follows:

Departments' Rankings of EPs

	D1	D2	D3
1 st	EP3	EP1	EP3
2 nd	EP1	EP3	EP1
3 rd	EP4	EP4	EP4
4 th	EP6	EP5	EP2
5 th	EP2	EP2	EP6
6 th	EP5	EP6	EP5

This means, for example, that D1 likes EP3 best, EP1 second-best, and so on.

Similarly, the EPs' ROLs are:

EPs' Rankings of Departments

	EP1	EP2	EP3	EP4	EP5	EP6
1 st	D1	D1	D2	D3	D1	D1
2 nd	D2	D3	D1	D1	D3	D3
3 rd	D3	D2	D3	D2	D2	D2

This means, for example, that EP4 has a preference ordering of departments given by D3, D1, D2.

To solve for the associated department-optimal stable match using *EP-Match*, the market maker first enters the number of departments (3) and EPs (6) in the appropriate boxes on the "Instructions" spreadsheet, and then checks that the option "Dept. Proposing" is selected. On the "Departments' Preferences" spreadsheet, the departments' rankings of the EPs is coded by assigning rank-order "1" to the most-preferred EP, rank-order "2" to the next-most-preferred EP, and so on. On the same spreadsheet, the number of vacancies ("positions") in each department (i.e., 2, 2, 2) is entered. The departments' preference matrix then looks like:

Clear All	D1	D2	D3
Positions	2	2	2
EP1	2	1	2
EP2	5	5	4
EP3	1	2	1
EP4	3	3	3
EP5	6	4	6
EP6	4	6	5

Similarly, on the “EPs’ Preferences” spreadsheet, the market maker codes the EPs’ rankings of the departments by assigning rank-order “1” to the most preferred department, rank-order “2” to the next-most preferred, and rank-order “3” to the least preferred department. The EPs’ preference matrix then looks like:

Clear All	EP1	EP2	EP3	EP4	EP5	EP6
D1	1	1	2	2	1	1
D2	2	3	1	3	3	3
D3	3	2	3	1	2	2

Finally, the “Match” button on the “Match” spreadsheet can be pressed, and the department-optimal stable match appears on screen as follows:

Match			
<u>EP</u>	<u>Department</u>	<u>Department</u>	<u>EP</u>
EP1	D1	D1	EP1
EP2	D3	D1	EP6
EP3	D2	D2	EP3
EP4	D3	D2	EP5
EP5	D2	D3	EP2
EP6	D1	D3	EP4

Here, the list on the left is ordered by EP, while the list on the right is ordered by department. Note that stability implies that no department will be able to lure a more preferred EP from another department. Neither can any EP succeed in “stealing” a more preferred position from another EP.

To calculate the EP-optimal stable match, select the option “EP Proposing” on the “Instructions” spreadsheet and press the “Match” button on the “Match” spreadsheet. It is easily verified that, in this example, the EP-optimal and departmental-optimal stable matches are identical.

Example 2

In our second example we illustrate how *EP-Match* can be used to solve more complicated scenarios, when not all EPs are acceptable to all departments, when there are more (or less) positions than EPs, and when departments are indifferent between some EPs, and EPs are indifferent between some departments.

Let us take the set-up of Example 1 as our point of departure but assume that department D1 is indifferent between EP1 and EP4, and also between EP2 and EP5. Moreover, department D2 now considers EP2, EP5, and EP6 to be unacceptable, while D3 has two additional vacancies (i.e., 4 instead of 2). On the EP side, EP2, EP4 and EP6 are assumed to be indifferent between the department they have ranked second and the department they have ranked third. For the rest, the situation is the same as in Example 1.

In *EP-Match*, on-screen instructions explain how indifference and unacceptability can be coded into the preference matrices. To express that department D_i is in fact indifferent between EP_x and the EP ranked one lower than EP_x , add a "*" to the rank-order number of EP_x on the "Departments' Preferences" spreadsheet. In our example, this means that D1's rank-order numbers of EP1 and EP2, i.e., "2" and "5," respectively, get a "*". To express that department D_i considers EP_y unacceptable, leave cell (EP_y, D_i) empty. In our example, this means that cells $(EP2, D2)$, $(EP5, D2)$, and $(EP6, D2)$ remain empty.

The departments' preference matrix then looks like:

Clear All	D1	D2	D3
Positions	2	2	4
EP1	2*	1	2
EP2	5*		4
EP3	1	2	1
EP4	3	3	3
EP5	6		6
EP6	4		5

Incorporating EP2, EP4, and EP6's indifference between their second and third-ranked departments, we get for the EPs' preference matrix:

Clear All	EP1	EP2	EP3	EP4	EP5	EP6
D1	1	1	2	2*	1	1
D2	2	3	1	3	3	3
D3	3	2*	3	1	2	2*

Now, the department proposing stable match is:

Match			
<u>EP</u>	<u>Department</u>	<u>Department</u>	<u>EP</u>
EP1	D1	D1	EP1
EP2	D3	D1	EP6
EP3	D2	D2	EP3
EP4	D3	D2	--- not matched ---
EP5	D3	D3	EP2
EP6	D1	D3	EP4
		D3	EP5
		D3	--- not matched ---

Again, the list on the left is ordered by EP, and the list on the right is ordered by department.

This stable match differs little from the stable match of Example 1. EP5 has moved from D2 to D3, filling one of D3's two additional vacancies, while D2's second and D3's fourth vacancy go unfilled. The latter reflects the fact that, in our example, there are more positions than EPs. The EP-proposing stable match is again identical to the department-proposing stable match.

V. CONCLUSION

In this note we have argued that adopting the Deferred Acceptance Algorithm to match EPs and departments can enhance the well-being of all parties involved. Unlike the current system, the DAA generates an efficient and stable outcome by optimally using all available information on the preferences of the participants as contained in their rank-order lists. In addition to its attractive theoretical properties, the DAA has been extensively tried-and-tested in practice, and it was shown to produce excellent results. To facilitate implementation of the DAA at the Fund, we have developed an Excel-based computer program that executes the algorithm on the basis of EPs' and departments' rankings of each other.

MATCHING OF INCOMING EPS

A. Introduction

As described in Section II, HRD matches incoming EPs to first-year positions on the basis of EPs' and departments' preferences over each other. The question that arises is whether these matches are stable and/or Pareto efficient. Because HRD uses rules of thumb to do the matching (instead of an algorithm), this question can only be addressed by looking at actual data.

In this appendix we show that HRD's match was neither stable nor Pareto efficient in the case of a cohort of incoming EPs for which we were able to retrieve the necessary data. These data consist of the participants' preferences, given by their rank-order lists (ROLs) (Table 2), and the match implemented by HRD (Table 3).⁹

B. Methodology

Participants (i.e., departments and EPs) who chose not to submit a ROL are assumed to be indifferent as to who they are matched to. Similarly, participants who ranked only a subset of counterparts are assumed to be indifferent between those left unranked, while strictly preferring ranked over unranked counterparts.

To check the stability of HRD's match, one might be tempted to run *EP-Match* on the participants' preferences and compare the outcomes. However, this would be misguided, because the set of stable matches usually contains many elements while, in principle, *EP-Match* only finds two of them; namely, the department-proposing and the EP-proposing stable matches. Thus, discrepancies between HRD's match and those calculated by *EP-Match* are not necessarily informative.

A second complication arises from the large number of declared and imputed "indifferences" in participants' preferences. Only when preferences are strict does stability imply Pareto efficiency. Else, stable matches may contain "suspect pairs." A suspect pair is a match of an EP and a department, such that both the EP and the department are indifferent between remaining matched to each other, or being matched to someone else. If the EP and/or the department making up the suspect pair are very sought after among other departments or EPs, it may be possible to make one or more of these other departments or EPs better off by breaking up the suspect pair, without hurting anybody else. This implies that, with a lot of indifference on both sides of the market, a stable match is only a *candidate* solution. To be a full-fledged solution, it must be checked that none of the suspect pairs, if any, give rise to Pareto inefficiencies.

⁹ To protect confidentiality we have made certain changes to the data. In particular, "dummy" EPs, departments, and preferences have been added. However, the gist of the example and the validity of the conclusions are fully retained.

Therefore, a two-step procedure is followed to assess the optimality of HRD’s match with respect to the participants’ preferences:

- Check the stability of the match by looking for *blocking pairs*.
- Check whether suspect pairs, if any, give rise to Pareto inefficiencies.

C. Stability

Is HRD’s match stable with respect to the participants’ stated preferences? The answer is no. To see why, note that EP7 and department D1 form a blocking pair: EP7 strictly prefers D1 over his current match D5, and D1 strictly prefers EP7 over its current match EP2.

In this case, a simple exchange is sufficient to transform HRD’s match into a stable one. That is, we reassign EP7 to D1 and EP2 to D5, while keeping everything else the same. Both EP7 and D5 go from being matched to a counterpart who was not even on their submitted preference list to being matched to their 1st choice. On the other hand, EP2 goes from his 1st choice to his 2nd choice, while D5 goes from its 3rd choice to its 6th choice.

D. Pareto Efficiency

It is easily checked that HRD’s match as given in Table 3 contains four suspect pairs, namely, (EP3, D12), (EP10, D13), (EP13, D8), and (EP20, D8). Not all of these suspect pairs give rise to Pareto inefficiencies, however. Only the first two suspect pairs do. By breaking up (EP3, D12) we can make EP5 strictly better off without hurting anyone else, while breaking up (EP10, D13) allows us to make EP18 strictly better off.

Specifically, we perform the following Pareto improving reassignments:

1. $\begin{array}{l} (EP3, D12) \\ (EP5, D11) \end{array} \longrightarrow \begin{array}{l} (EP3, D11) \\ (EP5, D12) \end{array}$
2. $\begin{array}{l} (EP10, D13) \\ (EP18, D11) \end{array} \longrightarrow \begin{array}{l} (EP10, D11) \\ (EP18, D13) \end{array}$

In this case, EP5 achieves his 1st choice instead of his 2nd choice, while EP18 achieves his 1st choice instead of being assigned to a department not on his preference list. At the same time, none of the others is worse off than before.

E. Conclusion

We have shown that HRD’s match of incoming EPs to departments was neither stable nor Pareto efficient. Application of *EP-match* followed by a critical assessment of all suspect pairs would have resulted in a better outcome.

Table 1

STABLE AND UNSTABLE CENTRALIZED MECHANISMS

Market	Stable/Based on DA Algorithm	Still in Use
Entry level medical markets:		
U.S. (NRMP)	yes/yes	yes
Edinburgh ('69)	yes/yes	yes
Cardiff	yes/yes	yes
Canada	yes/yes	yes
Cambridge	no	yes
London Hospital	no	yes
Birmingham	no	no
Edinburgh ('67)	no	no
Newcastle	no	no
Sheffield	no	no
Other markets:		
Medical Specialties		
(approximately 30 markets)	yes/yes	yes
Canadian Lawyers (articling positions)		
Toronto	yes/yes	yes
Vancouver	yes/yes	no (abandoned in '96)
Calgary & Edmonton	yes/yes	yes
Dental Residencies (5 specialties, 2 general programs)		
Osteopaths (< '94)	no	no
Osteopaths (> '94)	yes/yes	yes
Pharmacists	yes/yes	yes
Sororities	yes (at equilibrium)/no	yes

(Source: Roth and Rothblum, 1999.)

Figure 1

EP-Match for Excel

Instructions:

1. Enter the number of Departments and EPs in the box below.

Number of Departments:	3
Number of EPs:	6
Dept. Proposing	<input type="checkbox"/>
EP Proposing	<input type="checkbox"/>

2. Fill out the "Departments' preferences" worksheet according to the instructions on the sheet.
You can customise the Departments' and EPs' labels.
3. Indicate the number of open positions for each firm on the "Departments' preferences" worksheet.
4. Fill out the "EPs' preferences" worksheets according to the instructions on the sheet.
Here too, you can customise the Departments' and EPs' labels.
5. Click the [Match] button.
The Department-optimal stable match will appear in the "Match" spreadsheet. Save or print your results!

Do not add/remove rows or columns from these 4 worksheets or change their order.

Credits:

Created by: Greg Barron - 11/2003

CLER Research fellow - Research Computing Services

Harvard Business School

Thanks to Prof. Al Roth for valuable advice. Any remaining errors are the author's.
The deferred acceptance algorithm: David Gale and Lloyd Shapley (1962).

Figure 2

Departments' preferences:

- Each column contains the preference list of one department.
- You can change the names of departments and EPs.
- For each column, enter "1" in the row of the most preferred EP, "2" in the row of the 2nd-most preferred EP and so on.
- Add a "*" to EP x's rank-order to express a department's indifference between EP x and the EP y ranked one lower than EP x. E.g., "3*" and "4*" in the same column means that the department is indifferent between the EPs (s)he has ranked 3rd, 4th and 5th.
- Empty cells represent unacceptable EPs.
- Enter the number of positions being offered by each department in the "positions"-row.
- Do not add/delete rows or columns from this sheet.

Clear All	Department 1 Department 2 Department 3			Department 4
Positions	2	2	2	
EP 1	2*	1	2	
EP 2	5*	5	4	
EP 3	1	2	1	
EP 4	3	3	3	
EP 5	6	4	6	
EP 6	4	6	5	
EP 7				
EP 8				
EP 9				

Figure 4

The Match:

- Click the [Match] button below.
- The left list is ordered by EP, and the right list is ordered by Department

Match

<u>EP</u>	<u>Department</u>	<u>Department</u>	<u>EP</u>
EP 1	Department 1	Department 1	EP 1
EP 2	Department 3	Department 1	EP 6
EP 3	Department 2	Department 2	EP 3
EP 4	Department 3	Department 2	EP 5
EP 5	Department 2	Department 3	EP 2
EP 6	Department 1	Department 3	EP 4

Table 2

EPs' Preferences over Departments		Departments' Preferences over EPs	
EP Participant	Preferences 1/	Department	Preferences 1/
EP1	D12, D11*, D13*, D5*, D7*, D4	D1	EP7, EP15, EP17
EP2	D1, D11*, D13*, D5*, D7*, D12*, D4	D2	EP12, EP7, EP25
EP3	No preference indicated	D3	EP19, EP22*, EP11*, EP16
EP4	D1, D6, D13, D12	D4	EP22, EP19, EP6, EP11, EP12
EP5	D12, D11	D5	EP15, EP21, EP7, EP6, EP19, EP2
EP6	No preference indicated	D6	EP15, EP1, EP26
EP7	D1	D7	EP17, EP14, EP25
EP8	No preference indicated	D8	No preference indicated
EP9	No preference indicated	D9	No preference indicated
EP10	No preference indicated	D10	EP9*, EP19
EP11	No preference indicated	D11	EP22, EP6, EP19, EP21, EP23, EP7
EP12	D1, D5, D13	D12	No preference indicated
EP13	No preference indicated	D13	No preference indicated
EP14	No preference indicated	D14	E15, EP24
EP15	D1, D2, D12		
EP16	D1, D3, D10, D13, D5		1/ Preferences listed in ranked order
EP17	No preference indicated		* Indicates indifference between EP _x and the EP ranked one lower than EP _x .
EP18	D13, D1		
EP19	No preference indicated		
EP20	No preference indicated		
EP21	D1, D5, D12, D13		
EP22	No preference indicated		
EP23	D5		
EP24	D14, D1		
EP25	No preference indicated		
EP26	D1, D6, D2		

Table 3

EPs' Departmental Assignments		Departments' EP Assignments	
EP	Assigned Department	Department	Assigned EP
EP1	D6	D1	EP2
EP2	D1	D1	EP15
EP3	D12	D1	EP17
EP4	D4	D2	EP12
EP5	D11	D3	EP16
EP6	D11	D3	EP19
EP7	D5	D4	EP4
EP8	D11	D4	EP11
EP9	D10	D4	EP22
EP10	D13	D5	EP7
EP11	D4	D5	EP21
EP12	D2	D6	EP1
EP13	D8	D6	EP26
EP14	D7	D7	EP14
EP15	D1	D7	EP25
EP16	D3	D8	EP13
EP17	D1	D8	EP20
EP18	D11	D10	EP9
EP19	D3	D11	EP5
EP20	D8	D11	EP6
EP21	D5	D11	EP8
EP22	D4	D11	EP18
EP23	D11	D11	EP23
EP24	D14	D12	EP3
EP25	D7	D13	EP10
EP26	D6	D14	EP24

References

- Avery, C., C. Jolls, R. Posner, and A.E. Roth, 2001, "The Market for Federal Judicial Law Clerks," *University of Chicago Law Review* Vol. 68, pp. 793–902.
- Blum, Y., A.E. Roth, and U.G. Rothblum, 1997, "Vacancy Chains and Equilibration in Senior-Level Labor Markets," *Journal of Economic Theory* Vol. 76, pp. 362–411.
- Gale, D., 2001, "The Two-Sided Matching Problem: Origin, Development, and Current Issues," *International Game Theory Review* Vol. 3, pp. 237–52.
- , and L. Shapley, 1962, "College Admissions and the Stability of Marriage," *American Mathematical Monthly* Vol. 69, pp. 9–15.
- Niederle, M., and A.E. Roth, 2003, "Unraveling Reduces Mobility in a Labor Market: Gastroenterology With and Without a Centralized Match," *Journal of Political Economy* Vol. 111, pp. 1,342–52.
- , 2004a, "Market Culture: How Norms Governing Exploding Offers Affect Market Performance," *NBER Working Paper* w10256, Cambridge, Massachusetts: NBER.
- , 2004b, "The Gastroenterology Fellowship Match: How It Failed and Why It Could Succeed Again," *Gastroenterology* Vol. 127, pp. 658–66.
- Kagel, J.H., and A.E. Roth, 2000, "The Dynamics of Reorganization in Matching Markets: A Laboratory Experiment Motivated by a Natural Experiment," *Quarterly Journal of Economics*, Vol. 115, pp. 201–35.
- Roth, A.E., 1984a, "The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory," *Journal of Political Economy* Vol. 92, pp. 991–1,016.
- , 1984b, "Misrepresentation and Stability in the Marriage Problem," *Journal of Economic Theory* Vol. 34, pp. 383–87.
- , 1984c, "Stability and Polarization of Interests in Job Matching," *Econometrica* Vol. 52, pp. 47–57.
- , 1985, "The College Admissions Problem is not Equivalent to the Marriage Problem," *Journal of Economic Theory* Vol. 36, pp. 277–88.
- , 1986, "On the Allocation of Residents to Rural Hospitals: A General Property of Two-Sided Matching Markets," *Econometrica* Vol. 54, pp. 425–27.
- , 1989, "Two Sided Matching with Incomplete Information about Others' Preferences," *Games and Economic Behavior* Vol. 1, pp. 191–209.
- , 1990, "New Physicians: A Natural Experiment in Market Organization," *Science* Vol. 250, pp. 1,524–28.

- , 1991, “A Natural Experiment in the Organization of Entry Level Labor Markets: Regional Markets for New Physicians and Surgeons in the U.K.,” *American Economic Review* Vol. 81, pp. 415–40.
- , 2002, “The Economist as Engineer: Game Theory, Experimental Economics and Computation as Tools of Design Economics,” Fisher Schultz lecture, *Econometrica* Vol. 70, pp. 1,341–78.
- , 2003, “The Origins, History, and Design of the Resident Match,” *Journal of the American Medical Association* Vol. 289, pp. 909–12.
- Roth, A.E., and E. Peranson, 1999, “The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design,” *American Economic Review* Vol. 89, pp. 748–80.
- Roth, A.E., and U.G. Rothblum, 1999, “Truncation Strategies in Matching Markets: In Search of Advice for Participants,” *Econometrica* Vol. 67, pp. 21–43.
- Roth, A.E., and M. Sotomayor, 1989, “The College Admissions Problem Revisited,” *Econometrica* Vol. 57, pp. 559–70.
- , 1990, “Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis,” *Econometric Society Monograph Series*, Cambridge: Cambridge University Press.
- , 1992, “Two-Sided Matching,” in *Handbook of Game Theory* 1, ed. by R.J. Aumann and S. Hart, Amsterdam: Elsevier, pp. 485–541.
- Roth, A.E., and X. Xing, 1994, “Jumping the Gun: Imperfections and Institutions Related to the Timing of Market Transactions,” *American Economic Review* Vol. 84, pp. 992–1,044.
- , 1997, “Turnaround Time and Bottlenecks in Market Clearing: Decentralized Matching in the Market for Clinical Psychologists,” *Journal of Political Economy* Vol. 105, pp. 284–329.
- Zhou, L., 1991, “Stable Matchings and Equilibrium Outcomes of the Gale-Shapley’s Algorithm for the Marriage Problem,” *Economics Letters* Vol. 36, pp. 25–29.