The Flexible-Salary Match:
A Proposal to Increase the Salary Flexibility of the National Resident Matching Program

Vincent P. Crawford[1]
Department of Economics, University of California, San Diego
9 August 2004; revised 17 August 2004

Abstract

This paper proposes a generalization of the algorithm the National Resident Matching Program ("the Match") uses to match medical students to hospital residencies. Instead of hospitals setting salaries in advance, students and hospitals would have the option to report rankings that depend on salaries and the algorithm would be adapted to use such rankings. Such a Flexible-Salary Match would allow participants to compete via match-specific salaries within a centralized procedure like the current Match with a moderate increase in complexity, in a way that may attenuate a perceived bias toward lower salaries and address the concerns of a recent lawsuit.

1. Introduction

Most graduating medical students in the United States obtain hospital residencies through the National Resident Matching Program (NRMP; see http://www.nrmp.org/). The NRMP, or "Match" as it is often called, begins with hospitals defining residency positions, including the associated salaries.[2] Next, both students and hospitals submit rankings of their potential partners, taking the predetermined salaries into account. Finally, the Match, keeping participants' rankings secret, runs a computer algorithm that matches students and positions. Both students and hospitals commit themselves in advance to abide by the results. The Match is generally agreed to be a workable and effective way to match students and residency positions (Roth (1984b, 2003), Roth and Peranson (1999), Bulow and Levin (2003)).

The success of a procedure like the Match plainly depends on the method used to construct a matching from the submitted rankings. Once hospitals have defined their positions and salaries, the Match works like Gale and Shapley's (1962) "deferred-acceptance procedure,"

---

[1]Internet: http://weber.ucsd.edu/~vcrawfor/. I am grateful to Jeremy Bulow, Muriel Niederle, Sara Robinson, Alvin Roth, and Joel Sobel for useful comments.

[2]Actually only some programs post salaries in advance. In other programs students can only extrapolate from the program's salaries in the previous year, and learn their exact salaries after they take up their positions. This fact strengthens the case for a Flexible-Salary Match of the kind proposed here.

which was originally proposed as a model of processes like those by which people sort themselves into marriages or students sort themselves into colleges.[3]

Such "markets" differ from standard competitive markets in two important ways. The participants fall naturally into two sides, each of whose members can benefit only by dealing with a member or members of the other side. And the market is cleared not by adjusting prices (dowries, scholarships, or salaries) but by matching participants according to preferences that can differ across matched pairs. In the Match, each student and each residency may differ from all others; a student's idiosyncrasies may interact differently with the special features of different positions; and the salaries of positions are predetermined before the Match is run.

In the deferred-acceptance procedure for a marriage market, participants on one side of the market, say men, begin by proposing to their favorite women. Women then tentatively accept their favorite proposals, rejecting any others. Rejected men then propose to their next-favorite women, and the process continues as before until no more rejections are issued.

Gale and Shapley (1962) showed that if participants' proposals, acceptances, and rejections follow their true rankings, this algorithm converges to a matching of men to women that is *stable* in the sense that no man and woman prefer each other to their mates. Further, if participants have strict rankings of potential mates, the algorithm with men proposing converges to a matching at least as good for all men as any other stable matching; and the analogous algorithm in which women propose has the same properties except that it favors women in that sense.[4] Gale and Shapley also showed how to generalize their algorithm and these results to "college-admissions markets," in which participants on one side of the market wish to match with more than one participant on the other side, as for most hospital programs.

A third way in which the Match differs from a standard competitive market is that it is centralized, with a procedure like the deferred-acceptance procedure used to construct the matching "behind the scenes" from participants' submitted rankings. Until 1998 the Match worked like the deferred-acceptance procedure with hospitals proposing, and so favored hospitals to some extent, just as men were favored in the marriage market. In 1998, in response

[3]The NRMP evolved independently of Gale and Shapley's procedure, and the equivalence of the two was discovered only decades later, by Roth (1984b).

[4]The advantage of being on the side that makes proposals depends on the number of different stable matchings, which in turn depends on how closely participants' rankings are aligned. If all men have the same ranking of the women, and vice versa, there is only one stable matching (in which the highest-ranked man marries the highest-ranked woman, and so on), and it does not matter which side makes proposals. With less perfectly aligned rankings, there can be many stable matchings, and the men's favorite matching can be significantly better for the men than the women's favorite matching. (It is a surprising feature of these markets that men (or women), who may disagree a great deal about the best matching, agree completely on the best *stable* matching.)

to perceived problems, the NRMP switched to a version in which students propose, thus favoring them somewhat more than previously.[5] But because the centralized procedure mimics a decentralized auction-like process, in which participants on one side of the market make offers, starting with their highest-ranked partners and adjusting their offers until a stable matching is reached, the fact that the Match is centralized may not matter much in itself.[6]

A potentially more important difference between the Match and standard competitive markets is that in the Match, hospitals set the salaries of their positions in advance. This would make little difference in a market for a homogeneous commodity like wheat, where neither suppliers nor demanders care which of the bushels being traded is used to fulfill a given contract, and a single price for wheat is enough to ensure that the market allocates wheat to its highest-valued uses. But given the idiosyncrasies of both students and positions in the market for residents, an efficient allocation—roughly, one in which students are matched in a way that maximizes total productivity—may depend on match-specific salaries; and a decentralized market would normally be cleared by matched participants negotiating such salaries.

In the Match, however, hospitals normally set salaries for positions rather than individual students, because they cannot predict whom they will hire.[7] This severely limits salary flexibility, and as a result the Match can yield very different salaries and different, possibly inefficient, matchings than a decentralized market in which salaries are free to adjust to the idiosyncrasies of matches. Salaries for residents are in fact surprisingly low, averaging less than $40,000 for an 80-hour work week; and plaintiffs in a recent class-action antitrust suit charge that the current Match design enables hospitals to underpay residents (Jung et al. versus Association of American Medical Colleges et al., 02-CV-00873 (DDC 2002); see also Chae (2003), Miller and Greaney (2003), and http://www.savetheresidents.com/index.asp).

It may seem that in a market organized like the Match, hospitals will set salaries for positions that are competitive "on average" for candidates they hope to attract, so setting salaries in advance adds noise but otherwise does not matter much. But in an illuminating

[5]See Roth and Peranson (1999), whose simulations suggest that proposers' advantage was small before and after the change and that opportunities for participants to gain by misrepresenting their rankings were and are limited.
[6]Niederle and Roth (2003) have shown that there is little difference between Match and non-Match subspecialties for medical fellowships, either within hospitals or in average salaries across hospitals.
[7]As Bulow and Levin (2003) vividly put it, "A hospital will pay its offer regardless of the resident it actually matches with; it cannot offer 5x for the obstetrical Barry Bonds, but only x for the obstetrical Mario Mendoza… Most readers will recognize Bonds as baseball's greatest player over the past 15 years; Mendoza is best known for his struggles in keeping his batting average above his weight; a standard that has become known as the 'Mendoza Line'." In theory match-specific salaries are now possible by agreement conditioned on the Match's results (Sara Robinson, private communication) but they are unlikely to occur without a change in the Match.

analysis inspired in part by the suit, Bulow and Levin (2003) argue that the effect of setting salaries in advance can be larger and more systematic than this intuition suggests. Focusing on the case in which all participants on a given side of the market agree on the rankings of potential partners, they show that hospitals' uncertainty about the consequences of salaries can stifle salary competition and depress salaries below those that would emerge in a decentralized competitive market with salaries freely negotiated between matched participants. They also argue that the inefficiency of matching that may result is comparatively small, and that salaries are depressed most for the highest-ranked students and programs. The main effect of setting salaries in advance is then an across-the-board transfer from students to hospitals, with the highest-ranked students losing the most and the highest-ranked hospitals gaining the most.

As a remedy for the claimed bias in salaries, the plaintiffs in the antitrust suit have proposed alternative rules that would make it easier for hospitals to hire outside the Match; and some have advocated eliminating the Match (http://www.savetheresidents.com/index.asp). This paper proposes a simple alternative remedy, called the "Flexible-Salary Match," that would alleviate salary bias and inefficiency, while preserving the basic structure of the Match. The Flexible-Salary Match generalizes the current Match to give students and hospitals the option of reporting rankings of potential partners that depend on salaries, and adapts the algorithm to consider such rankings in constructing the matching. Although a centralized procedure that duplicates the effects of participants negotiating match-specific salaries may seem unwieldy, I will show that such a procedure can be constructed and implemented at the cost of a moderate, and individually voluntary, increase in the complexity of participants' submitted rankings.

Section 2 describes the proposed Flexible-Salary Match, which is based on algorithms that extend the Gale-Shapley deferred-acceptance procedure to allow salaries to be determined endogenously along with the matching introduced by Crawford and Knoer (1981) and Kelso and Crawford (1982); see also Roth (1984a), who further generalized the Crawford-Knoer-Kelso ("CKK") algorithm. Like the deferred-acceptance procedure, these algorithms were intended as models of market processes, but can also be viewed as centralized market-clearing procedures as in the Flexible-Salary Match. Section 3 gives a detailed example to illustrate how the Flexible-Salary Match would work, and Section 4 discusses its pros and cons.

2. The Flexible-Salary Match

The Flexible-Salary Match would start as before with hospitals defining their positions, including the associated salaries. But a hospital would now have the option of defining each position with a small number, say three, of alternative salaries of its own choosing. Students and hospitals would then submit rankings of potential partners as now, but each participant would have the option of ranking a given position differently at different salaries. (A hospital could still specify only one salary per position, and participants could choose to report rankings independent of salaries, as in the current Match.) Finally, the NRMP, keeping the rankings secret, would run the CKK algorithm, matching students to positions and determining salaries.

It may be best to run the CKK algorithm with students proposing to minimize differences from the current Match; although a Flexible-Salary Match with hospitals proposing, as in the pre-1998 Match, would be equally workable. Each student begins by proposing to his or her highest-ranked position at the highest of its alternative salaries. Hospitals then tentatively accept their highest-ranked proposals, taking salary into account, and rejecting any others.[8] Rejected students then propose to their next-highest-ranked position, taking salary into account, except that a student who proposes to a position for which he or she has previously been rejected may do so only at a lower salary than previously proposed (if one exists; otherwise no such proposal is allowed). Hospitals then tentatively accept their highest-ranked proposals, rejecting any others, and the process continues until no more rejections are issued.

This Flexible-Salary Match mimics a decentralized salary negotiation process, which resembles an English (ascending price) auction run upside-down, as students bid the salaries of the most desirable positions down until the market clears. This centralized process yields match-specific salaries that support efficient matching just as a decentralized competitive market would. Generalizing results of Crawford and Knoer (1981) and Kelso and Crawford (1982), Roth (1984a) showed that if participants' proposals, acceptances, and rejections follow their true rankings, under certain assumptions on preferences this algorithm converges to a matching with salaries that is efficient in the sense appropriate to such markets and *stable* in the sense that no student and hospital can find a salary at which they would prefer to be matched with each other to their current positions and salaries.[9]

---

[8]If a hospital's or student's reported ranking does not respond to the salaries for a given position, the algorithm implicitly ranks them in the "normal" direction for the purpose of interpreting this step.

[9]Strictly speaking, the Flexible-Salary Match only yields a matching that is efficient relative to the possibilities achievable with the specified list of salaries, and so only approximates full efficiency, which might require a

3. An Example

To illustrate how the Flexible-Salary Match would work, imagine that there are two students, 1 and 2; and two hospitals, A and B, each with one position. In the current Match, hospitals' uncertainty about salaries' consequences could lead them to set different salaries depending on their anticipations of market conditions. In the example I assume that the current Match would lead both hospitals to define their positions with salary M(edium), a moderate salary that is not necessarily the same for A and B. I assume that the Flexible-Salary Match would lead Hospital A still to define its position with salary M, yielding a single position called (A, M) (short for "hospital A at salary M"); but that it would lead Hospital B to define its position with three alternative salaries, H(igh), M(edium), and L(ow), yielding mutually exclusive potential positions (B, H), (B, M), and (B, L).[10]

I assume that student 1 prefers hospital A to B at any salary, and so reports ranking A > B (where ">" means "strictly prefers") or equivalently, (A, M) > (B, M), in the current Match with salaries fixed at M; and reports ranking (A, M) > (B, H) > (B, M) > (B, L) in the Flexible-Salary Match. Student 2 prefers A to B unless B's salary is H, and so reports A > B or equivalently, (A, M) > (B, M), in the current Match; and reports ranking (B, H) > (A, M) > (B, M) > (B, L) in the Flexible-Salary Match; note that this ranking cannot be expressed without salaries. Hospital A prefers student 2 at the only salary it has defined, and so reports ranking 2 > 1 or equivalently, (2, M) (short for "student 2 at salary M") > (1, M), in the current Match; and reports ranking (2, M) > (1, M) in the Flexible-Salary Match. Hospital B prefers student 2 at any salary, and so reports ranking 2 > 1 or equivalently, (2, M) > (1, M), in the current Match; and reports the more detailed ranking (2, L) > (1, L) > (2, M) > (2, H) > (1, M) > (1, H) in the Flexible-Salary Match.[11]

---

more flexible specification of salaries. In more general models, stability would also require that no matched pair could find a mutually beneficial adjustment in the terms of their agreement. That issue does not arise here because salary is the only thing being negotiated, and any change in it would hurt one member of the pair. The most important assumption on preferences is Kelso and Crawford's "gross substitutes" condition, now called "substitutability". This assumption rules out complementarities between students in a hospital's ranking of the kind that would make it prefer to hire student 1 if it could also hire student 2, but otherwise prefer not to hire student 1. Without substitutability, a hospital might become dissatisfied with an earlier proposal that has not been rejected in the algorithm, which can destroy the stability of the final matching. Failures of substitutability are a genuine problem, but arguably they just as problematic in the current Match as in a Flexible-Salary Match.

[10]A "soft" budget constraint is implicit in hospital B's ability to fill its position at any one of three salaries.

[11]Hospitals' and students' rankings in the example are consistent with each assigning a monetary equivalent value to each potential partner and ranking partners according to those values net of salaries; such preferences are separable across partners' characteristics and salaries. Hospital B's rankings use salary flexibility to avoid a version of the "Mendoza" problem in the current Match that Bulow and Levin described (fn. 7), by expressing more willingness to pay a premium for student 2 than for student 1. To economize on reporting costs, student 1 could be allowed to report (A, M) > (B, H) > (B, M) > (B, L) as A > B, as in the current Match. Note that a

With salaries fixed at M as in the current Match, there is a unique stable matching: 2 to (A, M), 1 to (B, M). It is stable because, although student 1 would prefer (A, M), hospital A prefers (2, M) to (1, M); and only (B, H) would be preferable for student 2, but hospital B cannot offer 2 salary H.[12] The deferred-acceptance procedure with students proposing as in the current Match reaches this stable matching almost immediately: Both students initially propose to hospital A. A then rejects 1, who then proposes to B and is accepted, and the market clears. The deferred-acceptance procedure with hospitals proposing as in the pre-1998 Match reaches the same outcome by a different route: Both hospitals initially propose to student 2. 2 then rejects B, which then proposes to 1 and is accepted, and the market clears. Gale and Shapley's (1962) results can be used to show that the fact that both versions reach the same outcome implies that it is the unique stable matching when salaries are fixed at M.

With the flexible salaries defined by hospital B in the example, B can offer student 2 salary H, and 2 prefers (B, H) to (A, M), so the matching that was stable with salaries fixed at M is no longer stable with flexible salaries.[13] In fact such an offer leads directly to a matching that is stable with flexible salaries: 2 to (B, H), 1 to (A, M). In this matching there is no way to make student 2 better off. There is also no way to make hospital B better off except by matching B with student 1, to which 1 would not agree; or by lowering 2's salary in position B, to which 2 would not agree. Both students are better off in this stable matching with flexible salaries than in the unique stable matching with salaries fixed at M of 2 to (A, M), 1 to (B, M).

The Flexible-Salary Match with students proposing reaches this stable matching immediately: Student 1 initially proposes to hospital A at salary M (the only salary for A's position, but ranked higher than B even at salary H), and student 2 initially proposes to hospital B at salary H (his or her highest-ranked possible salary in position B). Neither is rejected, and the market clears at the stable matching 2 to (B, H), 1 to (A, M). But the Flexible-Salary Match with hospitals proposing reaches a different outcome: Hospital A initially proposes to student 2 at salary M, and hospital B initially proposes to student 2 at salary L. B is rejected, and then proposes to student 1 at salary L and is accepted, and the market clears at the stable matching 2 to (A, M), 1 to (B, L). Both students do worse than in the Flexible-Salary Match with students

---

student (hospital) never ranks a position lower (higher) when its salary increases; it is reasonable to build this into the reported rankings. In a more realistic setting a hospital would report rankings for each position and salary, and the rankings might interact, as when it seeks to hire in only one of two related subspecialties.

[12]More precisely, hospital B cannot offer 2 salary H without risking ending up with student 1 at salary H, which the example assumes, following Bulow and Levin, is enough of a risk to deter B from offering salary H.

[13]This is a normal occurrence, because salary flexibility often gives participants new ways to upset a matching.

proposing, and student 1 does worse than in the current Match (whoever proposes). Both hospitals do better than in the Flexible-Salary Match with students proposing, and B does better than in the current Match.

The example was chosen to illustrate how the Flexible-Salary Match would work, not to make general claims about who would benefit (or suffer) the most if it replaced the current Match. It is easy to find other examples where with fixed salaries as in the current Match, there are multiple stable matchings, which bear different relationships to the stable matchings with flexible salaries; and where the Flexible-Salary Match (with students or hospitals proposing) yields matchings with different relationships to those generated by the current Match.

What is generally true is that the Flexible-Salary Match with students (hospitals) proposing yields a stable matching that all students (hospitals) rank at least as high as any other stable matching; and that is at least approximately efficient, normally with at least as high or higher total productivity than the current Match. The example illustrates the reason for the efficiency result by showing how the Flexible-Salary Match with students proposing allows hospital B to outbid A for student 2 when 2 is sufficiently more productive at B to make this worthwhile; the proof, however, depends on a detailed mathematical analysis (Crawford and Knoer (1981), Kelso and Crawford (1982), Roth (1984a), and Shapley and Shubik (1971)). It is noteworthy that the "outbidding" in the example occurs even though the Flexible-Salary Match is centralized, and even though (in this case) it is the students who make offers.

The example also does not bear directly on the issues discussed by Bulow and Levin (2003), whose analysis turns on hospitals' optimal choices of salaries in the current Match and how they relate to the competitive salaries that would emerge in the Flexible-Salary Match. The example departs somewhat from their assumption that all participants have the same ranking of potential partners, and the moderate salary M that the example's hospitals choose in the current Match need not relate to participants' preferences in the same way. Thus, it does not fully reflect the logic of Bulow and Levin's argument that the current Match depresses salaries below competitive levels, to an increasing extent for more highly-ranked students.

Finally, the example greatly understates the difficulty of using the CKK algorithm to find stable matchings, which in more realistic settings can take many more than one or two steps.

4. Pros and Cons

This section briefly discusses the pros and cons of the proposed Flexible-Salary Match.

Probably the most important issues are the complexity of reporting rankings involving salaries and of computing the matchings they imply. Computation is probably not a serious obstacle, as running the algorithms should take computer time that is approximately quadratic in the number of salaries per position, a slow rate of increase by computer-science standards. With three salaries per position the effect would be roughly like tripling the number of students and positions, which would require roughly nine times the computer time of the current Match.

With regard to participants' reporting costs, the effect would be at most to multiply the length of ranking lists by the number of salaries per position. But this may overstate it, because simplifications are possible as indicated in Section 3, and participants would not be required to report detailed salary rankings for potential partners they do not expect to be relevant (low-ranked hospitals for star students, etc.).[14] It seems likely that many participants would be willing to pay the additional reporting cost to reap the benefits of flexible salaries, and others would still have the option of reporting salary-independent rankings as in the current Match.

A related concern is that even if the additional flexibility were available, participants would not take advantage of it. Here there is a useful analogy with academic labor markets for new Ph.D.s, which even though they are not centralized have many similarities to the Match. Universities and hospitals have similar funding sources and budget constraints, the importance of institutional prestige is comparable, and the morale issues of having different salaries for different people in the same job are similar. Among elite universities, for many years in some disciplines there was a kind of "gentleman's agreement" not to compete vigorously on salary. The result was a decentralized, sequential process that, like the deferred-acceptance procedure, cleared the market without adjusting salaries much as the centralized Match does now.[15] The custom then eroded, and there was a gradual transition to more vigorous salary competition, making the process more like a decentralized version of the CKK algorithm, or the proposed Flexible-Salary Match. Salaries ended up higher, but universities found the money to pay them, and there were no other serious problems. This analogy suggests that under the Flexible-Salary Match there would also be a smooth transition to more vigorous salary competition, and that it would yield comparable benefits in efficiency and increased salaries for students.

[14]To the extent that such expectations are correct failure to report detailed rankings would not alter the outcome. The proposers' advantage and opportunities for beneficial misrepresentation of rankings are also likely to remain small, as suggested (but not established) by Roth and Peranson's (1999) results for the current Match.
[15]As Bulow and Levin's (2003) analysis suggests, this process may also have benefited elite institutions and depressed the salaries of their assistant professors most, without much affecting who ended up where.

References

Jeremy Bulow and Jonathan Levin, "Matching and Price Competition," manuscript, Stanford University, 2003.

Sanders H. Chae, "Is the Match Illegal?," *The New England Journal of Medicine* 348 (January 23, 2003), 352-356.

Vincent P. Crawford and Elsie Marie Knoer, "Job Matching with Heterogeneous Firms and Workers," *Econometrica* 49 (March 1981), 437-450. (http://weber.ucsd.edu/~Evcrawfor/CrawfordKnoer81EMT.pdf)

David Gale and Lloyd Shapley, "College Admissions and the Stability of Marriage," *American Mathematical Monthly* 69 (1962), 9-15.

Alexander S. Kelso, Jr. and Vincent P. Crawford, "Job Matching, Coalition Formation, and Gross Substitutes," *Econometrica* 50 (November 1982), 1483-1504. (http://weber.ucsd.edu/~vcrawfor/KelsoCrawford82EMT.pdf)

Frances H. Miller and Thomas L. Greaney, "The National Resident Matching Program and Antitrust Law," *Journal of the American Medical Association* 289 (February 19, 2003), 913-918.

Muriel Niederle and Alvin Roth, "Relationship between Wages and Presence of a Match in Medical Fellowships," *Journal of the American Medical Association* 290 (September 3, 2003), 1153-1154.

Alvin Roth, "Stability and Polarization of Interests in Job Matching," *Econometrica* 52 (January 1984), pp. 47-58. (a)

Alvin Roth, "The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory," *Journal of Political Economy*, 92 (1984), 991-1016. (b)

Alvin Roth and Elliott Peranson, "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design," *American Economic Review*, 89 (September, 1999), 748-780.

Alvin Roth, "The Origins, History, and Design of the Resident Match," *Journal of the American Medical Association* 289 (February 19, 2003), 909-912.

Lloyd Shapley and Martin Shubik, "The Assignment Game I: The Core," *International Journal of Game Theory* 1 (1972), 111-130.