

The Flexible-Salary Match:

A Proposal to Increase the Salary Flexibility of the National Resident Matching Program

Vincent P. Crawford¹

First version 9 August 2004; this revision 10 February 2005

Most graduating medical students in the United States obtain hospital residencies through the National Resident Matching Program ("NRMP"; <http://www.nrmp.org/>). The NRMP, or "Match" as it is usually called, is a centralized procedure that begins each year with hospitals defining residency positions, including the associated salaries. Both students and hospitals then submit rankings of potential partners, taking hospitals' predetermined positions and salaries into account.² A computer algorithm is then used to match students and positions, keeping the rankings secret. Students and hospitals commit themselves in advance to abide by the results.

Given hospitals' positions and salaries, the computer algorithm works like Gale and Shapley's (1962) *deferred-acceptance procedure*, explained below, which they proposed as a model of decentralized adjustment processes in *matching markets* like those by which men and women sort themselves into marriages or students sort themselves into colleges. In a striking example of independent discovery, the deferred-acceptance procedure and the Match evolved separately. Their equivalence was discovered only later, in exchanges between David Gale and then-NRMP administrator Elliott Peranson (Roth (1984b, footnote 18)). (Roth and Sotomayor (1990) give an overview of matching theory, including its application to the Match.)

The Match is generally agreed to be a workable and sensible way to match students and residency positions. It overcomes or mitigates some well-documented problems—unraveling of the timing of offers, misrepresentation of rankings, and *recontracting*, or the unwillingness of participants to abide by the results—that have arisen in decentralized professional labor markets, including the U.S. market for residencies before the Match was instituted (Roth (1984b, 2003), Roth and Peranson (1999), Niederle and Roth (2003b)). However, plaintiffs in a recent class-action antitrust suit (Jung et al. versus Association of American Medical Colleges

¹Vincent P. Crawford is Distinguished Professor of Economics at the University of California, San Diego. His email address is vcrawfor@weber.ucsd.edu. This paper describes the proposal discussed in Robinson (2004b), linked at <http://weber.ucsd.edu/~vcrawfor/24matc.htm> and <http://weber.ucsd.edu/~vcrawfor/match.jpg>, which gives additional background and commentary. The proposal was presented to the NRMP Board of Directors by one of its members, Daniel Leung, in October 2004. The Board referred it for study to its Executive Committee, to report back at the May 2005 meeting (Leung, private communication). I am grateful to Jeremy Bulow, Muriel Niederle, Alvin Roth, Joel Sobel, and especially Elliott Peranson and Sara Robinson for helpful comments.

²In fact only some hospital programs post salaries in advance. In other programs students must extrapolate from the program's salaries in the previous year, and learn their exact salaries only after they take up their positions.

et al., 02-CV-00873 (DDC 2002)) have charged, among other things, that by having hospitals set salaries and other aspects of positions in advance of the matching process, the current Match stifles salary competition and thereby keeps residents' salaries artificially low.³

Residents' salaries are in fact surprisingly low, given their duties and training, averaging less than \$40,000 for an 80-hour work week. And in an illuminating analysis inspired in part by the lawsuit, Bulow and Levin (2003) have given theoretical support to the plaintiffs' claim. In their model of markets organized like the Match, setting salaries in advance, together with the unpredictability of matching, effectively restricts hospitals to setting salaries for positions rather than students. Under plausible assumptions, this stifles competition and causes salary compression, making salaries lower than in a competitive market with match-specific salaries.

Bulow and Levin also show that markets like the Match can interfere with Pareto-efficiency by limiting the salary variation needed to support the most productive matching of students to positions.⁴ Here, Pareto-efficiency is defined in the usual way, taking hospitals' as well as students' preferences into account, and given the full possibilities of the market, including flexible match-specific salaries. As explained below, the Match may still produce outcomes that are efficient in a restricted sense, with salaries assumed fixed at predetermined levels; but those outcomes may be Pareto-inferior to some attainable with flexible salaries.

Although these arguments leave it less than certain that the current Match distorts salaries and causes inefficient allocation, they do suggest that alternatives that would allow more salary competition are worth considering. So far the plaintiffs in the lawsuit have focused on legal remedies that would make it easier for hospitals to hire outside the Match; and others have advocated eliminating the Match altogether (<http://www.savetheresidents.com/index.asp>). Both suggestions are motivated by a desire to move closer to the outcome of a decentralized market in which salaries (and other aspects of positions) are negotiated between hospitals and students along with the matching, coupled with a belief that no centralized procedure can mimic such an outcome. But given the well-documented problems in decentralized professional labor markets, it is unclear whether a decentralized residency market would be desirable, or even workable.

This paper takes an alternate route by proposing a centralized procedure designed to mimic the outcome of a decentralized market, which I will call the "Flexible-Salary Match."

³The lawsuit was recently dismissed following enactment of an amendment to antitrust law directly targeted at the case (Robinson (2004a)), but plaintiffs are in the process of appealing. Salary and other background information can be found at <http://www.savetheresidents.com/index.asp> and <http://www.savethematch.org/>.

⁴Bulow and Levin suggest that in their model the inefficiency is likely to be small, and caution readers not to apply their analysis too literally to the NRMP, for which not all of their assumptions are realistic.

The Flexible-Salary Match gives students and hospitals the option of reporting rankings of potential partners that depend on salaries and adapts the current algorithm to consider such rankings in determining the matching and salaries, otherwise preserving the current structure of the Match. The adaptation uses algorithms introduced by Crawford and Knoer (1981), Kelso and Crawford (1982), and Roth (1984a) (henceforth "CKKR"), which generalize Gale and Shapley's (1962) deferred-acceptance procedure to allow salaries to be determined endogenously along with the matching. I will argue that the Flexible-Salary Match could be implemented with a moderate, voluntary increase in the complexity of submitted rankings, and that it would allow the market to be cleared by flexible match-specific salaries, promoting Pareto-efficient allocation while avoiding the main problems of a decentralized market.

The next section describes the U.S. market for residencies and the current Match, and reviews the relevant theory of matching markets. The following section introduces the proposed Flexible-Salary Match and reviews the associated theory. The section after that gives a detailed example comparing the current Match and the Flexible-Salary Match. The final section concludes by summarizing the case for the Flexible-Salary Match.

The U.S. Market for Residencies, the Match, and the Theory of Matching Markets

This section describes the U.S. market for residencies and the current Match, and reviews the relevant theory of matching markets. The U.S. residency market differs from standard markets in several ways that are typical of professional labor markets and of matching markets:

- i. Its participants fall naturally into two sides—in this case students and hospital positions—such that students can only match with hospital positions, and vice versa;
 - ii. Matching is all-or-nothing, so that a student will normally work full-time in a single position and a position will normally be filled by a single student; and
 - iii. A given student's productivity may differ across positions, and s/he may have no perfect substitutes in a given position and may have imperfect substitutes of varying closeness.
- Under the current Match, the U.S. residency market also differs from standard markets in that:
- iv. The salaries of positions are normally predetermined before the Match is run, so that the market is cleared not by adjusting salaries but by adjusting the matching according to participants' fixed rankings of potential partners at the given salaries.

The Match and Gale and Shapley's Deferred-Acceptance Procedure

Once hospitals have defined their positions and salaries and participants have submitted their rankings, the current Match uses a computer algorithm that is equivalent to Gale and

Shapley's (1962) deferred-acceptance procedure. Gale and Shapley studied two classes of markets that share all four features just listed: *marriage markets*, in which a participant on one side (say, a man) can match with only one on the other side (a woman); and *college admissions markets*, in which a participant on one side (say, a college or hospital) can match with more than one on the other side (a student). They used the deferred-acceptance procedure to demonstrate the existence of a matching that is *stable* in the sense that there are no unmatched participants who prefer each other to their partners, a notion they introduced to characterize the implications of frictionless—though priceless—competition in a matching market. In the deferred-acceptance procedure for a marriage market, each participant on one side, say the men's, begins by proposing to his highest-ranked woman. Each woman then tentatively accepts her favorite proposal and rejects any others. Each rejected man then proposes to his next-favorite woman, and the process continues as before until no more rejections are issued. This procedure is readily adapted to the college-admissions markets that most closely resemble the market for residencies as organized by the Match. Before 1998 the Match used an algorithm with hospitals proposing, after 1998 one with students proposing (Roth and Peranson (1999)).

Gale and Shapley showed that in a marriage market, the deferred-acceptance procedure converges in a finite number of steps to a matching that is stable for the rankings participants express in their proposal and acceptance decisions. For, if the procedure converged to an unstable matching, then some man and some woman must prefer each other to their assigned partners. Such a man must have proposed at some time to the woman, and been rejected. But in the procedure, women reject men only in favor of offers from men they prefer, and they never lose offers involuntarily, a contradiction. If, in addition, the expressed rankings are strict, the procedure with men proposing converges to a stable matching ranked at least as highly for all men as any other stable matching; and the procedure with women proposing converges to a generally different stable matching that favors women in the same way. The extent of the proposers' advantage depends on how close participants' rankings are. If, for instance, all men have the same ranking of women and vice versa, then 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 proposes. If rankings differ, there can be many stable matchings, with the best stable matching for the men significantly better for them than the best one for the women.

These results extend to college-admissions markets, with qualifications regarding the ranking of the stable matching to which the procedure converges (Roth (1985)) and colleges' or

hospitals' preferences. With regard to preferences, it is necessary to rule out externalities across matches and to assume separability (Roth's "responsiveness") or substitutability (Kelso and Crawford's (1982) "gross substitutes" condition, which includes separability as a special case) of hospitals' preferences over sets of students. Substitutability rules out complementarities between students in a hospital's ranking that would make it prefer to hire student 1 if it could also hire student 2, but not otherwise. Without it, a hospital might regret its offer to student 1 if student 2 gets a better offer, even though student 1 has not rejected the hospital's offer, which might prevent the deferred-acceptance procedure from producing a stable matching. (The problem is not just with proving existence via the procedure: A stable matching may not exist.)

Stability plays two roles in welfare analyses of matching markets. If a marriage or college-admissions market is organized in a way that yields stable outcomes, participants have no incentive to recontract. Stability thus ensures that one important necessary condition for a workable market design is satisfied. Further, any stable matching is efficient in the restricted sense that there is no matching that is feasible without adjusting salaries that all participants would prefer to the current matching: If there were such a matching, those who would be matched in it could benefit bilaterally by matching with each other, contradicting stability. The qualification "without adjusting salaries" is important, because a matching that is efficient in that restricted sense may not be Pareto-efficient in the actual market with flexible salaries. It can therefore be Pareto-inferior to a matching produced by the Flexible-Salary Match, which does yield matchings that are (at least approximately) Pareto-efficient in the actual market.

Misrepresentation of Rankings in the Match

These results all concern participants' expressed rankings, but participants might have incentives to "game" the system by misrepresenting their rankings. In the deferred-acceptance procedure in a marriage or college-admissions market with students proposing, it is a *dominant strategy*—best for each participant no matter what the other participants do—for students to express their true rankings. But there is no way to make it a dominant strategy for colleges—which differ in that they match with more than one participant on the other side—to express their true rankings: They normally have incentives to misrepresent their rankings even in the deferred-acceptance procedure when they propose (Roth (1982, 1985)). Further, although in some circumstances equilibrium misrepresentation in a marriage market may make the procedure converge to a matching that is stable for the true as well as the expressed rankings (Roth (1984c)), this is not generally true in a college-admissions market (Roth (1985)).

Although the possibility of misrepresentation is an important theoretical qualification to the above welfare conclusions, it may have limited relevance for the Match. Roth and Peranson (1999) conducted computational experiments, running the pre- and post-1998 algorithms with submitted rankings from previous Matches. They found that when the number of positions for which a student can interview is realistically small, proposers' advantage was small both before and after the 1998 change, as were the opportunities to gain by misrepresenting rankings.

Bulow and Levin's Critique of the Current Match

I close this section by reviewing Bulow and Levin's (2003) critique of the current Match. They begin by arguing that setting salaries in advance, given the unpredictability of matching outcomes, in effect restricts hospitals to setting salaries for positions rather than individual students.⁵ As they 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'."

Setting salaries for positions rather than students would make little difference if students were a homogeneous commodity like wheat, so neither they nor hospitals cared which student filled a given position. A single salary for each position would then suffice to support Pareto-efficient allocation of students to positions. It may seem that even when students and positions are heterogeneous, hospitals will set salaries for positions that are competitive "on average" for the students they hope to attract, so that setting salaries in advance adds noise but otherwise does not matter much. Bulow and Levin show, however, that setting salaries for positions can make a large, systematic difference. They focus on the leading case in which all participants on a given side of the market agree on the rankings of potential partners, so that hospitals compete mainly against other hospitals like themselves. They show that higher-ranked hospitals then face less intense competition, which compresses salaries and drives them below those that would emerge in a competitive market with match-specific salaries. The result is like 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. The salary compression may also prevent the market from supporting a Pareto-efficient allocation of students to positions.

⁵In theory, match-specific salaries are now possible by prior agreement conditioned on the Match's results, but students are unlikely to negotiate aggressively on salaries before they have an offer and hospitals are unlikely to bid aggressively before they have concrete evidence of the strength of a student's market.

The Flexible-Salary Match and the Generalized Deferred-Acceptance Procedure

The Flexible-Salary Match begins like the current Match, with hospitals defining their positions; but a hospital now has the option of defining each position with a small number, say three, of alternative salaries of its own choosing.⁶ Students and hospitals then submit rankings of potential partners, but each now has the option of ranking a given partner differently at different salaries.⁷ A hospital could choose to specify one salary per position, and participants could choose to report rankings independent of salaries, as now. Finally, a generalized deferred-acceptance procedure is used to determine salaries, along with the matching.

The CKKR Algorithm

The generalized deferred-acceptance procedure was introduced by Crawford and Knoer (1981), Kelso and Crawford (1982), and Roth (1984a) ("CKKR") as a model of decentralized adjustment processes in matching markets with flexible salaries. For simplicity, I describe it in the marriage-market case in which each hospital has only one position; the generalization to the college-admissions case with substitutable preferences is straightforward. I also simplify by requiring participants to express strict rankings of partners as in the current Match.⁸ The CKKR algorithm can be run with students proposing as in the current Match, or hospitals proposing as in the pre-1998 Match. To minimize the difference from the current Match, I assume here that students propose; but letting hospitals propose may have some advantages, as explained below.

In the CKKR algorithm, each student begins by comparing positions, each at its highest alternative salary, and proposes to the position s/he ranks highest given those salaries. Each hospital then tentatively accepts its highest-ranked proposal, taking salaries into account, and rejects any others. Each rejected student then proposes to her/his next-highest-ranked position, excluding any position-salary combination(s) for which s/he has previously been rejected. (If a student has already been rejected by a position at all of its salaries, s/he is not permitted to propose to it again.) Each hospital then tentatively accepts its highest-ranked proposal and rejects any others, and the process continues as before until no more rejections are issued.

⁶More salaries are easily accommodated as long as they are discrete, and as CKKR noted, there is no obstacle to determining non-salary aspects of positions along with the matching, other than complexity. Hatfield and Milgrom (2004) discuss such models in more detail, and suggest an alternative approach to salary flexibility.

⁷It would be simpler to have hospitals rank students without regard to salary, and then specify which students they would be willing to hire at the highest salary, the next-highest salary, and so on, with analogous rankings for students. But this does not code enough information about student-salary tradeoffs to assure an efficient matching, which in general depends on hospitals' relative willingnesses to pay for different students, and vice versa. Neither would it suffice to have participants submit separate rankings for each salary level.

⁸See http://www.nrmp.org/res_match/tables/2005_Applicant.pdf, p. 57. If a participant's reported ranking does not respond to the salaries for a given position, the algorithm implicitly ranks them in the "normal" direction.

CKKR showed that with substitutable preferences, Gale and Shapley's (1962) main results for college-admissions markets extend to markets with flexible salaries: The CKKR algorithm converges in a finite number of steps to an *outcome*—a matching and match-specific salaries—that is *stable with flexible salaries* for participants' rankings as expressed in their proposal and acceptance decisions, in that no student and hospital can find a salary at which they would prefer to be matched with each other over their current partners. The proof follows Gale and Shapley's, with substitutability ensuring that no proposer becomes dissatisfied with an offer that has not been rejected (Kelso and Crawford (1982), Roth (1984a)).⁹ The Flexible-Salary Match thus limits participants' incentives to recontract just as the current Match does.

Although the Flexible-Salary Match is centralized, and constructs the matching behind the scenes from participants' submitted rankings, it mimics a decentralized salary adjustment process, which resembles an English (ascending-price) auction run upside-down, as students bid the match-specific salaries of the most desirable positions down until the market clears. This suggests that whether the market is decentralized or organized by a centralized procedure is less important than the details of the procedure.¹⁰ Further, just as stability is analogous to competitive equilibrium in college-admissions markets, stable outcomes in matching markets with flexible salaries are traditional competitive equilibria, with match-specific salaries as prices (Shapley and Shubik (1972), Kelso and Crawford (1982)). Stable outcomes with flexible salaries are also Pareto-efficient: If all participants would prefer some outcome to the current one, those who would be matched in that outcome could benefit bilaterally by matching with each other and adjusting salaries as needed, contradicting stability.¹¹ As illustrated in the next section, this may allow the Flexible-Salary Match to yield an outcome that is Pareto-superior to the one generated by the current Match, which is efficient only given predetermined salaries.

Finally, the Flexible-Salary Match with students proposing converges to an outcome that is stable with flexible salaries, and that all students rank at least as highly as any other such outcome; and its analog with hospitals proposing reaches a generally different outcome that is stable with flexible salaries, with an analogous tendency to favor hospitals (Roth (1985)).

⁹Failures of substitutability can cause genuine problems for the Flexible-Salary Match, but those problems are essentially the same as those encountered in the previous section for the current Match.

¹⁰Niederle and Roth (2003a) have shown that with predetermined salaries, there is little difference between Match and non-Match medical fellowship subspecialties, within hospitals or in average salaries across hospitals.

¹¹Strictly speaking, the Flexible-Salary Match yields an outcome that is Pareto-efficient relative to the required, given set of discrete salaries. This outcome may only approximate full Pareto-efficiency, which might require more flexible salaries. The approximation is better, the finer the divisibility of salaries; once it is fine enough, the matching must be the one associated with Pareto-efficient outcomes. Demange, Gale, and Sotomayor (1986) derive bounds on the inefficiency of the CKKR algorithm's outcomes with separable preferences.

Misrepresentation of Rankings in the Flexible-Salary Match

In the Flexible-Salary Match, as in the Match, participants might have incentives to misrepresent their rankings. The results here are similar to those without flexible salaries. With substitutable preferences, in the Flexible-Salary Match with students proposing it is again a dominant strategy for students to express their true rankings. But there is no way to make it a dominant strategy for hospitals to express their true rankings: They generally have incentives to misrepresent their rankings even when they propose (Hatfield and Milgrom (2004, Section IV.2)). However, although to my knowledge no one has done computational experiments like Roth and Peranson's (1999) for the Match (previous section) for markets with flexible salaries, it seems likely that for realistic specifications they would yield similar conclusions about the limited opportunities to gain by misrepresenting rankings.¹² Finally, as with the Match, equilibrium misrepresentation need not lead to an outcome that is stable for the true rankings. Overall, the possibility of misrepresentation does not seem to shift the balance significantly either for or against replacing the current Match with the Flexible Salary Match.

Reporting and Computation Costs

I now consider the increased cost to participants of reporting rankings involving salaries, and the related concern that even if participants had that option they would not take advantage of it. I also discuss the increased cost of computing the Flexible-Salary Match's outcomes.

Computation is probably not a serious obstacle, as running the algorithms should take computer time approximately quadratic in the number of salaries per position, a slow rate of increase. With three salaries per position, the effect is like tripling the number of students and positions, which requires roughly nine times more computer time than the current Match.

Reporting costs and inertia are also unlikely to be a serious problem. The Flexible-Salary Match would multiply the length of participants' ranking lists by three. But this may overstate the increase, because participants would not be required to report detailed salary rankings for partners they do not expect to be relevant (low-ranked hospitals for star students, and so on). To the extent that such expectations are correct, the failure to report detailed rankings would not alter the outcome. Overall, it seems likely that most participants, particularly once they have seen the Flexible-Salary Match in operation, would be willing to bear the additional reporting costs to reap the benefits of improved allocation. Other participants would still have the option of reporting rankings independent of salaries, just as in the current Match.

¹²Because no one has used a decentralized procedure with flexible salaries, there would be no data comparable to the rankings from previous Matches that Roth and Peranson used; but one could make realistic assumptions.

Peranson's Critique of the Flexible-Salary Match

I close this section by discussing a critique of the Flexible-Salary Match raised by Elliott Peranson (private communication). Because matching markets, with or without flexible salaries, normally have a range of stable outcomes, and the Flexible-Salary Match with students proposing produces a stable outcome that students prefer to any other stable outcome, it may yield higher salaries than would be negotiated between matched pairs in a decentralized competitive market, even if it produces the same matching of students to positions.¹³ Thus, if the goal is to mimic decentralized competitive market outcomes, the Flexible-Salary Match may be biased in favor of higher salaries. However, Roth and Peranson's (1999) computational experiments for the Match suggest that for realistic specifications, proposers' advantage may be small in the Flexible-Salary Match. Even if further research qualifies this view, the critique could be neutralized by a Flexible-Salary Match with hospitals proposing, as in the pre-1998 Match. Salaries would be lower than if students proposed, but the efficiency gains would be the same, plausibly large enough to allow an increase in students' as well as hospitals' welfares.

An Example Comparing the Match and the Flexible-Salary Match

To illustrate how the Flexible-Salary Match works and how its outcomes relate to those generated by the current Match, 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 this example I assume that the current Match would lead both hospitals to define their positions with salary M(edium), a moderate salary, not necessarily the same for A and B. I assume the Flexible-Salary Match would lead Hospital A still to define its position with salary M, yielding a single position called (A, M) ("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).¹⁴

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-

¹³This possibility is illustrated in the next section's example, where the Flexible-Salary Match matches hospital B and student 1 at salary H when students propose, but at salary L when hospitals propose. Bargaining between matched pairs in a decentralized competitive market might yield the same match, but at intermediate salary M.

¹⁴Implicit in hospital B's ability to fill its position at any of three salaries is a "soft" budget constraint. Hard budget constraints, or other interactions in hospitals' rankings, may cause violations of substitutability, and might require modifications of the Flexible-Salary Match. I leave these issues for future work.

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 given salary, and so reports ranking $2 > 1$ or equivalently, $(2, M) > (1, M)$, in the current Match; and in the Flexible-Salary Match, which requires more information about trade-offs, reports the more detailed ranking $(2, L) > (1, L) > (2, M) > (2, H) > (1, M) > (1, H)$.¹⁵

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.¹⁶ 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 matching by a different route: Both hospitals initially propose to student 2. Student 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 matching 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. This is a natural consequence of salary flexibility, which gives participants new ways to upset a matching. In this case such an offer leads directly to an outcome that is stable with flexible salaries: 2 to (B, H), 1 to (A, M). In this outcome there is no way to make student 2 better off. There is also no way to make hospital B better off

¹⁵Hospitals' rankings in this example are consistent with separable preferences, assigning a monetary value to each individual student (rather than to groups of students whose evaluations interact) and ranking students by those values net of salaries. Students' rankings are also consistent with assigning a value to each hospital and ranking hospitals by those values plus salaries. Thus a student (hospital) never ranks a position lower (higher) when its salary increases, and the rankings are well-behaved in any reasonable sense. Hospital B's rankings use salary flexibility to avoid a version of the "Mendoza" problem in the current Match, by expressing a willingness to pay a higher premium for student 2 than for student 1. To economize on reporting costs, student 1 could be allowed to report the ranking $(A, M) > (B, H) > (B, M) > (B, L)$ as $A > B$. In a more realistic setting, a hospital's rankings might interact across positions, as when it seeks to hire in only one of two related subspecialties.

¹⁶More precisely, hospital B cannot offer 2 salary H without risking ending up with student 1 at salary H, which the example assumes is enough of a risk to deter B from offering salary H, following Bulow and Levin (2003).

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 outcome with flexible salaries than in the stable matching with salaries fixed at M of 2 to (A, M), 1 to (B, M).

The Flexible-Salary Match with students proposing reaches the stable outcome 2 to (B, H), 1 to (A, M) 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 outcome 2 to (B, H), 1 to (A, M).

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 outcome 2 to (A, M), 1 to (B, L). Both students do worse than in the Flexible-Salary Match with students proposing, and student 1 does worse than in the current Match (whether students or hospitals propose). Both hospitals do better than in the Flexible-Salary Match with students proposing, and B does better than in the current Match.

The example illustrates the reason why the Flexible-Salary Match yields (approximately) Pareto-efficient outcomes by showing how it (but not the current Match) allows hospital B to outbid hospital A for student 2 when 2 is sufficiently more productive at B to make this worthwhile. It is noteworthy that this "outbidding" occurs even though the Flexible-Salary Match is centralized, and even when it is students who make offers. The proof of approximate Pareto-efficiency, however, depends on a more detailed mathematical analysis (footnote 11; Shapley and Shubik (1972); CKKR; and Demange, Gale, and Sotomayor (1986)).

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 examples where with fixed salaries as in the current Match, there are multiple stable matchings, which bear different relationships to the stable outcomes with flexible salaries; and where the Flexible-Salary Match (with students or hospitals proposing) yields outcomes with different relationships to those generated by the current Match.

The example also does not bear directly on the issues Bulow and Levin (2003) discuss. Their analysis turns on hospitals' optimal choices of salaries in the face of the uncertainty created by the current Match. The example makes illustrative assumptions about those salaries,

which do not fully reflect the logic of Bulow and Levin's argument that the Match depresses salaries below those that would arise in a decentralized competitive market.

Finally, the example understates the difficulty of using the CKKR algorithm to find stable matchings, which in more realistic settings typically takes many more than one or two steps.

Conclusion: The Case for the Flexible-Salary Match

I have argued that the Flexible-Salary Match could be implemented with a moderate, voluntary increase in the complexity of reported rankings over the current Match. It also limits participants' incentives to recontract just as the current Match does, and is no more vulnerable to the problems caused by incentives to misrepresent rankings and failures of substitutability. It would allow the market to be cleared by flexible match-specific salaries, promoting Pareto-efficient allocation while avoiding the main problems of a decentralized market. By contrast, the current Match yields matchings that are efficient only in a restricted sense, taking salaries as given; and may therefore be Pareto-inferior to the outcomes of the Flexible-Salary Match.

Probably the most important concern not yet discussed is the likely increase in the costs to hospitals of hiring residents. Here there is a useful analogy with academic labor markets for new Ph.D.s, which 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 effects of different salaries for different people in the same job are similar.

Among elite universities, there was for many years a "gentleman's agreement" not to compete on salary. The result was a decentralized, sequential process that cleared the market without adjusting salaries, much as the centralized Match does now. The custom then began to erode, and there was a gradual transition to more vigorous salary competition, making the process more like a decentralized market, or the 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 in the market for residencies, and that it would yield comparable benefits in allocative efficiency and increased salaries for students. The costs to hospitals would probably be higher, but they would be partly compensated by more efficient matching of students to their positions. In the end, the question is whether society should try to alleviate the budgetary problems of hospitals by organizing the market for residencies in a way that yields inefficient and, arguably, unfair allocations; or search instead for some better way.

References

- Jeremy Bulow and Jonathan Levin, "Matching and Price Competition," manuscript, Stanford University, 2003.
- Vincent P. Crawford and Elsie Marie Knoer, "Job Matching with Heterogeneous Firms and Workers," *Econometrica* 49 (March 1981), 437-450.
- Gabrielle Demange, David Gale, and Marilda Sotomayor, "Multi-Item Auctions," *Journal of Political Economy* 94 (August 1986), 863-872.
- David Gale and Lloyd Shapley, "College Admissions and the Stability of Marriage," *American Mathematical Monthly* 69 (1962), 9-15.
- John Hatfield and Paul Milgrom, "Auctions, Matching and the Law of Aggregate Demand," manuscript, Stanford University, 2004, *American Economic Review*, in press.
- Alexander S. Kelso, Jr. and Vincent P. Crawford, "Job Matching, Coalition Formation, and Gross Substitutes," *Econometrica* 50 (November 1982), 1483-1504.
- 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. (a)
- Muriel Niederle and Alvin Roth, "Unraveling reduces mobility in a labor market: Gastroenterology with and without a centralized match," *Journal of Political Economy*, December, 111, pp. 1342-1352. (b)
- Sara Robinson, "Antitrust Lawsuit Over Medical Residency System Is Dismissed," *New York Times*, August 14, 2004, A16. (a)
- Sara Robinson, "Tweaking the Math to Make Happier Medical Marriages," *New York Times*, August 24, 2004, D2, linked at <http://weber.ucsd.edu/~vcrawfor/24matc.htm> and <http://weber.ucsd.edu/~vcrawfor/match.jpg>. (b)
- Alvin Roth, "The Economics of Matching: Stability and Incentives," *Mathematics of Operations Research* 7 (1982), pp. 617-628.
- 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, "Misrepresentation and Stability in the Marriage Problem," *Journal of Economic Theory* 34 (December 1984), 383-387. (c)

Alvin Roth, "The College Admissions Problem is not Equivalent to the Marriage Problem,"

Journal of Economic Theory 36 (August 1985), 277-288.

Alvin Roth, "The Origins, History, and Design of the Resident Match," *Journal of the*

American Medical Association 289 (February 19, 2003), 909-912.

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 and Marilda Sotomayor, *Two-Sided Matching: a study in game-theoretic modeling*

and analysis, New York : Cambridge University Press, 1990

Lloyd Shapley and Martin Shubik, "The Assignment Game I: The Core," *International*

Journal of Game Theory 1 (1972), 111-130.