

Competitive Wages in a Match with Ordered Contracts

By MURIEL NIEDERLE*

In response to a recently dismissed antitrust lawsuit against the National Residency Matching Program (NRMP), Jeremy Bulow and Jonathan Levin (2006) show that when firms set impersonal salaries simultaneously, before matching with workers, then such a match leads to lower aggregate wages compared to any competitive outcome. Vincent P. Crawford (forthcoming) shows that this concern can be addressed by incorporating flexible salaries in the centralized match, that is, the possibility for each position to have more than one potential salary, with the final salary to be determined together with the worker-firm pairing. This builds on earlier work that shows that a match in which each position can have a large number of contracts, which I will call a multiple contract match, allows for competitive outcomes (Crawford and Else Marie Knoer 1981; Alexander S. Kelso and Crawford 1982; Alvin E. Roth 1984b; and John Hatfield and Paul Milgrom 2005).

Here, I observe that the NRMP has a feature that I will call ordered contracts, which destroys the low-wage equilibrium of Bulow and Levin and in fact allows for competitive outcomes. In a match with ordered contracts, firms can specify several possible contracts, just as in a match with flexible contracts. However, firms can also specify the order in which they prefer to fill these contracts. Specifically, at any point in the match, only one type of contract is available, which is in contrast to a match with multiple contracts. This additional control of firms changes which stable outcomes are reached through appropriately modified deferred acceptance algorithms. Most importantly, the set of contracts reached through either a firm or worker proposing modified deferred acceptance algorithm is the same,

which in general is not the case in a multiple contract match.

The results of this paper will cast the lesson we learned from Bulow and Levin in a new light. In particular, the relative lack of personal wages in the match appears to be an outcome of the market, not a constraint of the match. This also becomes apparent when we look at salaries in medical specialties that do, and do not, use a match.

Since 1951, the market for medical residents has been organized through a centralized matching procedure that assigns medical students to residency programs using a variant of a deferred acceptance algorithm.¹ In 1998, Roth and Elliot Peranson introduced a new algorithm, which switched from a hospital proposing to a student proposing deferred acceptance algorithm. The new algorithm also incorporates several special features, such as accommodating couples who want two jobs. A second special feature is to allow for ordered contracts, or “reverting positions.” Programs that try to fill a position under a certain contract can, in case they do not find a suitable candidate, change (or revert) that contract to a position with a different contract (see Roth and Peranson 1999). This allows programs to effectively have more than one contract for any position, while being able to control the order in which they want to try to fill those positions.

In the 1990s, about 7 percent of the three to four thousand programs that participate each year have contracts that could revert to other contracts if they remain unfilled (accounting for almost 6 percent of the total quota of positions). Roth and Peranson (1999) note that such reversions typically occur when, for example, a director of a second-year postgraduate program arranges with the director of a first-year prerequisite program for his residents to spend their first year in that prerequisite program. If the

* Department of Economics, Stanford University, Stanford, CA 94305, and NBER (e-mail: niederle@stanford.edu). I thank Georgy Artemov, Jonathan Levin, Eric Maskin, Al Roth, and Michael Schwarz for helpful discussions, the National Science Foundation and the Sloan Foundation for generous support, and the Institute for Advanced Study at Princeton and the Economics Department at Harvard University for their hospitality.

¹ The match was introduced to ensure a uniform appointment date, control the unraveling of hiring decisions, and reduce congestion problems of other, decentralized plans that tried to promote late and uniform hiring (see Roth 2003).

second-year program fails to fill all its positions, then the vacancies can “revert” to the first-year program to be filled by other applicants. More recently, in the reinstatement of the fellowship match for gastroenterologists, this feature is especially advertised to, for example, allow programs to try to fill a slot first with a research fellow, and in case no suitable research fellow can be attracted, the program can decide to fill this position with a more clinically oriented fellow instead (see Niederle, Deborah D. Proctor, and Roth 2006). While ordered contracts have been used in practice in the NRMP, theoretical properties have not been analyzed before.² In Section IV, I’ll fill this theoretical gap and explore the additional options that ordered contracts give to market designers.

In 2002, an antitrust lawsuit was filed, charging that the main effect of the match is to suppress wages of medical residents. Bulow and Levin develop a stylized model to analyze the effects on wages of using a centralized match such as the NRMP. In their simple model, firms simultaneously announce a wage at which they are willing to hire any worker. Workers care only about salaries and firm preferences over firms after the announcement. Then an assortative matching occurs, which matches more productive workers with firms that offer a higher wage. The main result of the paper is that this model yields wage compression, sub-competitive average wages, and higher profits for firms compared to any competitive outcome. The intuition for the result is that, compared to a competitive market, firms cannot change their salaries depending on the worker they end up hiring.³

² Roth and Peranson (1999) show how applicants who want two kinds of contracts, namely a first-year contract and a second-year contract, as well as couples who want two jobs, present complementarities that may make the set of stable matchings empty. They fail to note that ordered contracts, or reversions of positions in themselves, are not a source of complementarities, rather substitutes, and hence do not pose any problems for the existence of stable matchings and are amenable to theoretical analysis.

³ The fact that firms cannot change their salary when they try to hire different workers becomes important when firms are asymmetric. For example, some firms are more productive than others. In this case, a competitive outcome calls for gaps in the wages paid to different workers, which are not reproduced by equilibrium strategies when firms simultaneously set one wage for their position independent of the worker they end up hiring.

As such, the paper can be seen as providing support for the contention that a centralized match, such as the NRMP, may indeed be used to reduce wages. However, the NRMP allows for ordered contracts, which allow for wage competition that can restore competitive outcomes. In a stylized model of the NRMP with ordered contracts, each firm i , instead of advertising only one position at one contract (wage) p_i , can create a second contract p_i^* for the same position, and decide which workers are eligible for each contract. Firm i first tries to fill the position at contract p_i^* . If it fails to fill the position, then the contract is changed (or reverted) to a new contract p_i and firm i tries to fill the position at this new contract with a new set of eligible workers. Ordered contracts have two main effects on wages.

First, when other firms play the mixed strategies in the wage-setting equilibrium of Bulow and Levin, every firm has a strict incentive to use ordered contracts, and can strictly increase its expected payoff. Second, if all firms use ordered contracts, there exists an equilibrium in which wages are competitive. The actual NRMP algorithm is therefore able to achieve competitive outcomes in the model of Bulow and Levin with the use of ordered contracts.

Empirically, Niederle, Proctor, and Roth (2006) show that in the labor market for gastroenterology fellows, after their market operated without a match for nearly a decade, most programs offered the same wage to all their fellows (i.e., impersonal wages). That is, wage data show less differentiation than a competitive model in which gastroenterology fellows care only about wages would predict.⁴ In that case, we would not necessarily expect wages to be very different in a market with and without a match, even if the match would work as described by Bulow and Levin. This implies that we may not observe every program submitting two or more contracts, in a model in which, on the one hand, workers do not care most about wages but care, rather, for the firm where they work, and, on the other hand, multiple contracts serve only to discipline potential wage manipulations by competing firms. Hence, this paper does not

⁴ In Section V, I will discuss the effects of eliminating a match, such as the NRMP, on the competitiveness of a market.

predict that a host of programs submit two or more contracts in the NRMP. Since the NRMP does not in fact force wages to be impersonal, when we observe impersonal wages (e.g., all new hires in a given program receive the same wage), this is due to other factors. The results of Bulow and Levin may still apply, i.e., wages may be more compressed than if each worker were paid his marginal product. However, the observed uniformity of non-match wages, and the lack of price competition through ordered contracts even though they are available, suggests that the match is not the cause.

In the next two sections, I adopt the assumption of Bulow and Levin: workers care only about salaries, and not for which firm they work. I then discuss that this may not be an adequate assumption in order to understand the wages in the market of medical residents and fellows. The last section discusses theoretical properties of a match with ordered contracts. I conclude by discussing uses of ordered contracts in the NRMP and the positive effect of a match on the competitiveness of the market of medical residents.

I. A Simple Example

The following simple example illustrates the idea behind the low wage equilibrium of Bulow and Levin, and how ordered contracts undermine this equilibrium and can restore competitive outcomes.

Consider a market with three firms and three workers, where each firm wants to match with one worker, and every worker can work for at most one firm. Firm n 's profit from hiring worker m at wage p is $m \cdot n - p$, and worker m 's profit is p , that is, workers care only about the wage they receive, not for which firm they work. Wages in a competitive equilibrium have to be such that the matching is efficient, that is, worker i works for firm i .

Worker 1's competitive wage is $p_1 \in [0, 1]$, such that both firm 1 and worker 1 receive a nonnegative surplus. The wage p_2 of worker 2 has to be high enough that firm 1 does not prefer to employ worker 2 at p_2 , and has to be low enough that firm 2 prefers employing worker 2 at that wage to employing worker 1 at the wage p_1 , that is, $p_2 \in [p_1 + 1, p_1 + 2]$. Similarly, $p_3 \in [p_2 + 2, p_2 + 3]$. The lowest competitive wages are therefore 0, 1, and 3.

Now consider a centralized match where firms simultaneously announce wages, followed by an assortative matching by worker productivity and wages. Clearly none of the three firms will post the lowest competitive wage, since, for example, firm 3 has an incentive to lower the wage closer to p_2 . Bulow and Levin (2006) show that, in equilibrium, firms use mixed strategies for wages. Firm 1 offers 0 with a probability $5/6$, and any wage between $(0, 1/3]$ with a density of $1/2$. Firm 2 offers any wage between $(0, 1/3]$ with a density of 1 and a wage between $(1/3, 7/3]$ with a density of $1/3$. Firm 3, finally, makes offers between $[1/3, 7/3]$ with a density of $1/2$.

These strategies result in wage compression and lower average wages for workers, namely 0.02, 0.73, and 1.56, though worker 1 receives a higher wage than in the lowest-wage competitive equilibrium. Strategies also result in higher profits for firms, namely 1, 3.67 and 6.67 compared to 1, 3, and 6 in the competitive equilibrium.

Now I introduce the change that firm i announces two wages, p_i and p_i^* , and determines which of the workers are eligible for p_i^* , the contract it prefers when trying to fill the position. Workers observe all four contracts, that is, they observe p_j for $j \neq i$, p_i , and p_i^* and rank all four contracts according to the announced wage (and in case of ties, prefer more productive firms).

In the simple example, a matching algorithm that yields a stable outcome is as follows. The centralized procedure uses p_i^* , p_j for $1 \leq i \neq j \leq 3$ to create an assortative match just as before (though p_i is not used yet, as firm i has only one position and can hire at most one worker). This interim match is the final match if firm i fills its position at p_i^* . If firm i is unmatched, then p_i^* gets replaced by p_i and the new assortative match is the final allocation.

Suppose firm 1 is the firm that can make a normal wage offer p_1 , and a wage offer connected to a "star-position" p_1^* . Then, the following strategy makes firm 1 strictly better off than simply announcing one wage p_1 according to the Bulow and Levin equilibrium. Let $p_1^* = 1/6$, the midpoint of the highest interval on which firm 1 is randomizing over wages, and competing for worker 1 and worker 2, and let only worker 2 (and 3) be eligible for this contract p_1^* , while p_1 is determined as before. Then the expected surplus of firm 1 is 1.08, strictly higher than 1. In the Bulow and Levin model, firm 1 is indifferent

among all strategies, and hence its payoff is, for example, determined by offering wage $p_1 = 1/3$, and hiring worker 2 (with probability $1/3$) or worker 1 (with probability $2/3$). Here firm 1 can use its star-position to have a positive chance to hire worker 2 at a lower wage. The expected wage of worker 2 is 0.86.

Similarly, if either only firm 2 or only firm 3 can have a star-position contract, they can announce $p^* = 4/3$, and only worker 3 is eligible for this job, and otherwise announce p_2 or p_3 , respectively, as in the Bulow-Levin equilibrium.

When all firms can have a star-position contract, there exists an equilibrium in which workers receive their competitive salaries, the lowest of which are 0, 1, and 3, respectively. Each firm i can discipline the wage offer of firm $i + 1$ to worker $i + 1$ with the star-position contract. Specifically, let firm 1 have $p_1^* = 1$, for which only workers 2 and 3 are eligible, and in case the position does not get filled, it reverts to a contract $p_1 = 0$, for which all workers are eligible. Let firm 2 have $p_2^* = 3$, for which only worker 3 is eligible, and in case the position does not get filled, $p_2 = 1$ for which workers 2 and 3 are eligible. Firm 3 has no real use for a star contract, so let firm 3 use only its normal position and let $p_3 = 3$, for which only worker 3 is eligible. When we assume that, in case of ties, workers prefer more productive firms, then the star-position contracts ensure that firms cannot lower their salaries for the standard position. Similarly no worker can gain from declining to rank some contracts, the strategies form a Nash equilibrium.

Next, I show that the results of the example generalize. For these results, I assume that stable matchings with ordered contracts exist, and can be reached through suitably modified deferred acceptance algorithms. Later on, I will discuss these previously unexplored theoretical properties of ordered contracts. This algorithm is a closer version of the actual algorithm used by the NRMP and enhances the toolbox of market designers.

II. Matching with Ordered Contracts in the Bulow-Levin Environment

The general model of Bulow and Levin has N firms and N workers, where each firm wants

to hire one worker, and every worker can work for at most one firm. Firm n 's profit from hiring worker m at wage p is $\Delta_n \cdot m - p$, where $\Delta_N \geq \dots \geq \Delta_1 \geq 0$, and worker m 's utility is p ; that is, workers care only about the wage they receive, not for which firm they work. All preferences and productivities are common knowledge. In their model of a match, firms first make simultaneous salary offers, followed by an assortative match, which matches high-productivity workers to high-salary firms, and, in case of salary ties, to more productive firms. Firms are required to be willing to employ any worker, which is inconsequential when $\Delta_n = n$, but not in general.⁵

In a pricing equilibrium, firms choose an offer to maximize their expected profits given the other firms' choices and the matching process. An equilibrium involves mixed strategies. For firm n , a mixed strategy is a distribution where $G_n(p)$ denotes the probability that firm n offers a salary less than or equal to p . Let g_n denote the density of firm n 's offer distribution. In equilibrium, firms compete over ranges of salaries with each other, where more productive firms offer higher salaries and compete for more productive workers. To describe the equilibrium, Bulow and Levin first identify the set of firms that compete on each salary level, the density with which they offer wages, and—once all but the lowest productivity firm is dealt with—specify that firm 1 may offer a wage of zero with a positive probability. Firm n 's offer density at wage p , if it competes with firms $l \leq n \leq m$, is

$$(1) \quad g_n(p) = q(l, m) = \frac{1}{m - l} \sum_{k=l}^m \frac{1}{\Delta_k} - \frac{1}{\Delta_n}.$$

For any given highest firm m that offers p , let $l(m)$ be the lowest firm that offers p , for which $q(m) \equiv q(l(m), m) > 0$. Let p_{N+1} be the highest salary offered, and p_n denote the lowest salary offered by firm n . Then firms $l(N) \leq n \leq N$ compete on $(p_N, p_{N+1}]$ with offer densities $q_n(N)$, where p_N is such that firm N exhausts its offer probability. Then, below offer p_N , firms $l(N - 1), \dots, N$ compete on $(p_{N-1}, p_N]$ such that firm $N - 1$ exhausts all its remaining offer probability.

⁵ However, in an example, Bulow and Levin show that the fact that firms cannot announce which (subset of) worker is eligible for their unique wage offer cannot account for the whole wage reduction, though it has some effect.

The process continues until the behavior of firms 2, ..., N is specified. If $\Delta_1 = \Delta_2$, then firm 1's behavior is also specified. Otherwise, firm 1 offers zero with its remaining offer probability, namely $G_1(0) = 1 - \sum_n q_1(n) \cdot (p_{n+1} - p_n)$.

THEOREM 1 (Bulow and Levin 2006, 659): *There is a unique price equilibrium. Let $q_n(\cdot)$ and p_1, \dots, p_{N+1} and $G_1(0)$ be defined as above. Then, for each firm n , and each nonempty interval $[p_m, p_{m+1}]$, $g_n(p) = q_n(m)$ for all $p \in (p_m, p_{m+1}]$.*

First, I show that the equilibrium of Bulow and Levin does not survive the introduction of ordered contracts (proofs of theorems are in the Appendix).⁶

THEOREM 2: *Suppose all firms have only one position, and offer wages p_i according to the Bulow and Levin equilibrium. If some firm i can offer two wages, p_i^* —for which it can restrict which workers are acceptable—and p_i —for which any worker is eligible—then the firm makes strict positive gains from using that possibility.*

Furthermore, if every firm has an ordered star-position contract, then competitive wages are an equilibrium outcome. Let c_i be the lowest competitive wage of worker i ; then $c_1 = 0$, $c_{i+1} = c_i + \Delta_i$.

THEOREM 3 (Competitive Equilibrium Wages): *The following strategies form a Nash equilibrium. Every firm announces $p_i = c_i$ for $1 \leq i \leq N$ and $p_N^* = c_N$ with only worker w_N being eligible, and $p_j^* = c_{j+1}$ for $j < N$; and the workers being eligible for p_j^* are workers w_{j+1} and higher. The workers report their preferences truthfully, that is, they rank all contracts such that they prefer higher wages to lower wages,*

⁶ Because ordered contracts have not been discussed and theoretically analyzed before, the last section of the paper provides the algorithm for a match with ordered contracts and explores theoretical properties, most notably the fact that the algorithm yields a stable matching, that is, an outcome in which every firm employs at most one worker, and each worker works at most for one firm, such that no firm-worker pair (including the option of unemployment and an empty position) can be found that would rather be matched together than with their current partner.

and for a given wage they prefer more productive to less productive firms.

Therefore a match with ordered contracts, which provides a description of the actual possibilities offered by the NRMP algorithm, allows for competitive outcomes, and does not necessarily result in lower wages.

III. Wages of Medical Fellows

In the 1990s, about 7 percent of the three to four thousand programs that participate each year had positions with contracts that could revert to other contracts if they remain unfilled (accounting for almost 6 percent of the total quota of positions).⁷ In the reinstatement of the fellowship match for gastroenterology fellows, this feature is advertised in a way to allow programs to try to fill a slot first with a research fellow, and in case no suitable research fellow can be attracted, the program can decide, instead, to fill this position with a more clinically oriented fellow (see Niederle, Proctor, and Roth 2006).

Do we observe price competition through ordered contracts in the NRMP? There are two main empirical questions to solve before it is clear how contracts in the NRMP should look to result in competitive wages.

First, note that firm 1, in order to discipline the wage offer of firm 2, needs to offer the star-contract only with probability $1/2$.⁸ Similarly, firm 2 has to offer the star-contract only with probability $2/3$, to effectively discipline firm 3's wage offer.

Second, several firms and workers may be of identical or very similar quality. In that case, a competitive wage equilibrium with ordered contracts can be achieved even if only a subset

⁷ This feature is also used by the Internship Matching Program sponsored and supervised by the Association of Psychology Postdoctoral and Internship Centers (APPIC).

⁸ Firm 2 prefers to announce the contract $p_2 = 1$ versus $p_2 = 0$, if firm 1 offers $p_1^* = 1$ with probability $q \geq 1/2$. Furthermore, firm 2 cannot use his star-contract to exploit the fact that firm 1 offers $p_1^* = 1$ with probability q , and whenever firm 1 does not offer a star-contract, try to recruit worker 2 with $p_2^* = 0$ (where only worker 2 is eligible), and in case the position is not filled, that is, whenever firm 1 does offer the star-contract, revert the contract to $p_2 = 1$. The reason this is not robust is that worker 2 can simply announce that the low-wage star-contract is unacceptable, and so always ensure a high salary of 1.

of firms offers a second contract. Consider the initial example, only now assume that there are four firms and four workers, of “quality” 1, 2, 2, 3 each, where the profits of firm m from hiring worker n at wage p are $mn - p$. Then, a competitive wage equilibrium can be sustained even if only firm 1 and one firm 2 offer a second contract. This is true even as we increase the number of firms of quality 1, 2, or 3. That is, in general the number of firms that need to offer two contracts is highly dependent on the quality distributions of hospitals, residents, and fellows.

More generally, it is not clear that wages are the major determinant of hospital choice of residents and fellows, rather than, say, the actual hospital at which the resident or fellow receives his education. That is, it is not clear how competitive wages would actually look. This is obviously hard to explore directly with wage data obtained through the match, when the question is whether the match does or does not affect wage distributions. However, not all fellowship markets operate with a match, allowing for direct evidence on the importance of impersonal wages. Furthermore, we can compare wages in markets with and without a match.

Niederle, Proctor, and Roth (2006) studied the market for gastroenterology fellows when a centralized match was not used. A survey of gastroenterology program directors reveals that only a few program directors (6 percent of respondents) did not offer the same wage to all their incoming fellows, and they all responded that wages were not adjusted to outside offers. Similar results hold for other terms of fellowship contracts, such as hours on call. It does not appear that in a market without a match wages are highly personalized.

The survey results are supported by data on the internal medicine fellowship market. One can compare wages of fellows whose specialty participates in a match with wages of fellows whose specialty hires in a decentralized way. In internal medicine, of all subspecialties that require three years of prior residency, in the years between 2002 and 2004 four specialties used the Medical Specialties Matching Program (MSMP), while ten did not. Niederle and Roth (2003a) compare wages of all programs that report positive wages, excluding those from Puerto Rico, using the data from the Graduate Medical Education Library 2002–2003. A simple regression of the

wage on a match dummy (which is one when the specialty uses the match) reveals no significant effect of the match. Similarly, a comparison of wages within hospitals for specialties that use a match and that do not yields a small *positive*, significant (but not economically significant) effect of a match on wages. Similar results were found for the next year, using the Graduate Medical Education Library 2003–2004 data (Niederle and Roth 2004).

It is not clear that a match compresses or lowers wages, because, for one thing, the ordered contracts match used by the NRMP allows for wage competition. Furthermore, the market of internal medicine subspecialties fellows strongly suggests that wages are not different for specialties that use a match compared to those that do not. While the NRMP does not, in fact, force wages to be impersonal, we still observe a lot of impersonal wages, both in centralized and decentralized markets. In this respect, residents or first-year fellows may not be unique. For example, in many economics departments, first-year salaries of junior faculty hired in the same year are often the same. As such, the lessons of Bulow and Levin may still apply: wages are more compressed than if each worker were paid their marginal productivity; however, the match does not seem to be the cause of impersonal wages.

IV. Matching with Ordered Contracts

In this final section, I analyze matching with ordered contracts. I show the existence of stable matchings and how deferred acceptance algorithms have to be modified to account for ordered contracts. I point out similarities and differences between a match with ordered contracts and a match with multiple contracts, as introduced by Crawford and Knoer (1981) and Kelso and Crawford (1982). (For an overview, see Roth and Marilda Sotomayor (1990).)

In a matching with ordered contracts, every firm can have several contracts, for each of which the firm specifies a strict preference ordering over the set of workers eligible for this contract (formal definitions are in the Appendix).⁹ Furthermore, each firm has a strict ordering over

⁹ For example, a firm may decide that a specific worker may be eligible for the first contract but not for the second contract.

which contract should be filled first. Firms have preferences over worker-contract pairs and not filling their position at all, while workers have preferences over firm-contract pairs or remaining unemployed. In a matching, each worker can work for at most one firm at a specific contract, and every firm can employ at most one worker and so can have at most one of its contracts accepted. A matching is stable if there exists no firm, worker, or contract triplet such that the firm would rather fill its position with that worker at that contract, and the worker would rather accept that contract from that firm, than stay with their current match.

First, I show that a stable matching always exists, by providing a modified deferred acceptance (MDA) algorithm whose outcome is always a stable match. There are two versions of MDA, worker and firm proposing. The *Firm Proposing MDA* is basically a succession of standard deferred acceptances (DA) (see David Gale and Lloyd S. Shapley 1962). First, in *MDA step 1*, all firms have only their first contract available. Then comes the DA part. In *DA step 1*, firms make offers to their most preferred worker. Workers collect all their offers, keep their most preferred acceptable offer, and reject any other offers. More generally, in *DA step k*, firms whose offer was rejected in *step k - 1* make an offer to their next most desirable worker. Workers collect all their offers, keep their most preferred acceptable offer, and reject any other offers. The DA subalgorithm ends when either no firm has its offer rejected, or all rejected firms have no more workers they want to make an offer to, at the current contract. In the general *MDA step k*, any firm i that has its position at the j -th contract p_i^j unfilled changes the contract to p_i^{j+1} in case it has another contract to revert to. Then, the algorithm continues with a DA subalgorithm, where all previous offers are cancelled and have to be remade.¹⁰ The algorithm ends when all firms

that have no offer held by an applicant have no more contracts to change or revert to. Workers who hold an offer from a firm at a contract are matched to that firm at that contract; remaining firms and workers are unmatched.

A worker proposing MDA is actually a simplified version of the current algorithm used by the NRMP, which has, however, not been fully described (or analyzed) in this respect before.¹¹

In a *Worker Proposing MDA* in *MDA step 1*, all firms have only their first contract available. Then comes the DA step, just as before. At the general *MDA step k* any firm i that has its position at j -th contract p_i^j unfilled reverts the contract to p_i^{j+1} in case it has another contract to change to. Then the algorithm continues with the DA steps, where all previous offers are annulled and have to be remade. The algorithm ends when all the firms that have no offer have no more position to revert to. At this point, any worker whose offer is held by a firm at a specific contract is matched to that firm at that contract.

Unlike in the firm proposing algorithm, in the worker proposing MDA, interim offers have to be annulled, because some worker, who has an offer held by a firm, may prefer one of the new contracts that are introduced when some other firm changed its contract.

THEOREM 4 (Stability): *Whenever firms have a strict ordering over a finite number of contracts, which are ordered contracts, and for each contract a strict preference ordering over the workers, and workers have a strict ordering over firm-contract pairs, then both the firm and worker proposing MDA yield a stable outcome.*

THEOREM 5 (Firm-Optimal Stable Match): *The firm proposing MDA yields the firm-optimal stable match.*

I now describe properties that are different from other models of matchings with multiple contracts, in which firms cannot enforce the order in which contracts should be filled (see Crawford and Knoer 1981; Kelso and Crawford 1982; Roth 1984b; and Hatfield and Milgrom

¹⁰ Note that when some firm f reverts its position to a new contract, the deferred acceptance algorithm part can simply continue at whichever offers are held right now, instead of canceling all offers and restarting the whole process anew. The reason is that no firm that has not changed a contract can gain by remaking an offer that was rejected, as it was rejected because the worker either finds the offer unacceptable, or has a better offer in hand, and will have so once more as the algorithm unfolds, now that there are even more desirable contracts than before.

¹¹ The algorithm designed by Roth and Peranson, and used by the NRMP, works like a worker proposing MDA (private communication from Alvin Roth) (see Roth and Peranson 1999 for other aspects of the algorithm).

2005). In models of multiple contracts, a firm has a finite set of contracts and preferences over contract-worker pairs, and every worker w has preferences over contracts.¹² One difference of standard models of multiple contracts from a model with ordered contracts is that, in the standard model, all contracts are potentially present simultaneously and available immediately. Therefore, in a deferred acceptance algorithm where workers make offers, they can immediately offer to work for the highest wage or the most desirable contract. In contrast, an ordered contracts match has a sequential dimension built in, in which firms, even when they are not the ones making offers, decide on the order in which to try to fill contracts. The effect of this seemingly small difference can best be seen in a simple example.

Suppose there are two firms f_1 and f_2 with two contracts for one position each, one at low wage and one at high wage, and two workers w_1 and w_2 . Every firm prefers to fill the position at the low wage, and prefers, for a given wage, worker w_1 over w_2 . Every worker prefers high over low wages, but for a given contract prefers f_1 over f_2 .

In standard models with multiple contracts, firms have both wage offers available immediately. Then, in a worker proposing deferred acceptance algorithm, both workers make an offer to firm 1 at a high wage. Firm 1 keeps the offer of worker 1 and rejects the offer of w_2 , who makes an offer to work for firm 2 at the high wage, which is a stable outcome. The difference in a match with ordered contracts is that firms first submit the order of contracts to be used when trying to fill the position. In this case, the first contract is the low-wage contract and the second contract is the high-wage contract, where for each contract, firms prefer w_1 to w_2 . In a worker proposing MDA, similar to before, both workers apply to firm 1, although this time at the low-wage contract, the only contract available. And similarly to before, the final matching has worker 1 work for firm 1, and worker 2 for firm 2, although now both workers work at the low-wage contract. Note that this is the outcome

¹² Note that a firm that has multiple contracts does not necessarily have more possible rankings over contract-worker pairs compared to a firm that has ordered contracts, as long as the firm with ordered contracts has enough of those.

of the firm proposing MDA (and also the firm proposing algorithm with multiple contracts).

In a standard matching model where firms can have simultaneous multiple contracts, the firm and worker proposing deferred acceptance algorithm lead in general to different contracts being filled, even though the same set of firms will be matched (Hatfield and Milgrom 2005). In the case of a match with ordered contracts, not only the same set of firms, but also the same set of contracts, are matched.

THEOREM 6: *The two (possibly different) stable outcomes reached through the worker proposing MDA and the firm proposing MDA have the same set of workers and the same set of firms matched at the same contracts.*

The theorem has a few immediate implications. On a first note, a firm proposing MDA seems computationally easier than the worker proposing MDA, since there, after every revision of contracts, all the existing offers have to be cancelled and the whole offering process has to start anew. Theorem 6 implies, however, that a simple way to determine the outcome of a worker proposing MDA is the following: First, with the use of a firm proposing MDA, determine the set of contracts that will be used in a worker proposing MDA. Then, with those contracts, simply run a regular worker proposing deferred acceptance algorithm.

To apply the theory of ordered contracts to actual markets, it is important to understand the incentives firms and workers face when submitting their preference list to a centralized system. The next corollary shows that firms cannot gain from manipulating the timing at which they reach various contracts in either the worker proposing or firm proposing MDA.

COROLLARY 1: *In both firm and worker proposing MDAs, firms cannot gain by adding contracts that will never be filled on top of their list or scratching contracts that are completely undesirable. Furthermore, in both MDAs the order in which firms revert contracts is irrelevant to the final outcome.*

Theorem 6 also immediately implies, however, that workers may have an incentive to

manipulate their preferences, which is in contrast to standard models of multiple contracts.¹³

COROLLARY 2: *The worker and firm proposing MDA both result in matchings that use the contracts filled in the firm optimal stable match. That is, the worker optimal stable match is not reached by the worker proposing MDA, unless the worker proposing optimal match has the same contracts filled as the firm optimal MDA. This implies that workers may have an incentive to misrepresent their preferences even in the worker proposing MDA.*

V. Conclusions

In the present paper, I show how the institutional details of the NRMP support a different conclusion than the simple model of Bulow and Levin, in which a centralized match is modeled as impersonal wages. Attention to details is especially important when research is used for specific policy implications, such as whether a match, and the NRMP specifically, homogenize and reduce wages. Within the Bulow and Levin framework, Kojima (2007) shows that the wage reduction may not be robust to the case in which firms employ several workers, such as in the NRMP. Similarly, in order to understand the effect of a match on wages, attention has to be paid to a realistic model of an alternative decentralized market.¹⁴ Will a decentralized market be a highly competitive market?

The history of the market for medical residents (Roth 1984a) itself casts doubt that a market without a clearinghouse should be thought of as a competitive market. Niederle and Roth (2003b, 2004, 2005) and Niederle, Proctor, and Roth (2006) show that the labor market for gastroenterology fellows, after they stopped using a centralized match, once more unraveled, with

thin and dispersed markets, and reduced mobility.¹⁵ It seems that the market for gastroenterology fellows is one in which programs scramble to deal with congestion issues, programs make exploding offers, and start making offers earlier and earlier (Niederle, Proctor, and Roth 2006).¹⁶ Motivated by these facts, Siva Anantham and Jennifer Stack (2006) and Hideo Konishi and Margarita Sapozhnikov (2006) show in their models of decentralized market how wages may be lower than competitive wages as well.

Because decentralized markets impose frictions because offers take time, and wages are nonetheless impersonal, centralizing the offer process may very well increase the competitiveness of a market, since, in contrast to a market with exploding offers, in a match, all residents can compare different offers, and are not forced to decide on offers without knowing whether other offers may come through.

In a market in which wages are not personalized, what is the role played by ordered contracts, and how do they affect the competitiveness of a market? Ordered contracts, as opposed to multiple contracts, allow a program to replicate how they may try to fill positions in a decentralized market. For example, in the market for gastroenterology fellows, they try to find a research fellow and hire a clinical fellow only in the event no suitable candidate can be attracted.¹⁷ In an ordered contracts match, a program can do that without compromising the quality of the applicant pool they can consider for their second contract (Corollary 1). It is not clear, however, that in a decentralized market a program may not be hurt if it tries to fill a position first under a certain contract, and then only later under a second contract. If some applicants

¹³ Peter Coles (2005) discusses the extent of manipulation of strategies in simple matching environments. On the other hand, Fuhito Kojima and Parag A. Pathak (2006) show that as markets are large but rank order lists are short, there is not much room for manipulation. The extent of the possibilities of manipulation in ordered contracts, as workers may or may not be eligible for more than one contract for each firm, is still an open question.

¹⁴ Georgy Artemov (2006) points out that noise in standard matching models with multiple contracts may have large effects on wages.

¹⁵ The failure of the gastroenterology fellowship match is one of the rare instances in which a match that produces stable outcomes has been abandoned. C. Nicholas McKinney, Niederle, and Roth (2005) argue that this failure was due to an unusual event, and shed light not only on why this market failed, but also on why such failures are so rare.

¹⁶ Christopher Avery, Andrew Fairbanks, and Richard Zeckhauser (2003) show that in the market for college students, colleges use the option of early decision to secure students and limit their availability to competitors. Niederle and Roth (2007) show how markets may unravel as exploding offers and binding agreements are introduced.

¹⁷ Many economics departments that may have tried to hire a senior faculty in a field instead opt for a junior person once they are not so successful.

have already been hired (for example, by programs that have not tried to fill the position first under another contract), then a program may actually lose some potential candidates simply by having tried to fill a position first under a different contract. And this is exactly how ordered contracts are used in the NRMP, and how they have been advertised for use in the new gastroenterology fellowship match.

Some of the success of economic theory comes from the ability to make vastly simplified

models that capture some essential, general properties of markets. On the other hand, some of the recent successes in the market design literature reflect an attention to the details of particular markets.¹⁸ And sometimes case studies on details that are important in specific markets can lead to new general insights and theory.

¹⁸ For an overview on the importance of the attention to details in market design, and its successes, see Roth (2002).

APPENDIX

PROOF OF THEOREM 2:

Every firm i has a highest wage interval that it offers with a constant density, let it be $[p_i^L, p_i^H]$, on which it competes for several workers, the highest being w_H and the lowest w_L . The highest (and lowest) worker are easily determined by determining the highest (or respectively, lowest) firm that is offering a wage on this highest wage interval of firm i . Suppose that all other firms use the mixed strategies from before; then the following strategy makes firm i in expectation strictly better off. Let $p_i^S = (p_i^H - p_i^L)/2$ and the only worker eligible for that wage be w_H . It is easy to see that firm i is strictly better off with this strategy than foregoing the possibility to use a p_i^S job at all. The reason is that firm i is indifferent among all wages it offers in the Bulow and Levin equilibrium, so its profit is determined by, for example, offering wage $p_i = p_i^H$ and hiring any of the workers w_H, \dots, w_L with (different) positive probability. So, trying to hire worker w_H at a lower wage first, with the use of the star-contract, which is successful with positive probability, strictly increases expected payoffs.

PROOF OF THEOREM 3:

First I show that a firm i cannot gain by deviating. Without firm i , resulting wages would be c_j for worker w_j , and workers w_j with $j > i$ work for firm j , while workers w_j with $j \leq i$ work for firm $j - 1$. If firm i submits the strategies suggested by the theorem, firm i hires worker w_i at the lowest competitive wage for w_i (displacing firm $i - 1$). Firm i cannot hire any workers w_j with $j > i$, unless firm i is willing to pay ϵ more than competitive wages, and may hire workers w_j with $j \leq i$ at competitive wages. So, given the definition of lowest competitive wages, firm i cannot make higher profits than hiring worker w_i at c_i .

Now I show that workers cannot gain by deviating either. Given the strategies of firms, and workers $j > i$, worker i is the highest ranked worker of the standard contract of firm i at c_i , the star-contract of firm $i - 1$ at c_i , and contracts at wages lower than c_i . Any higher-wage contract is not achievable for worker w_i . Hence, worker w_i , by reporting truthfully, receives the highest wage he can receive given the strategies of other firms and workers.

In a model with ordered contracts, every firm i can have up to K contracts p_i^1, \dots, p_i^K . Let P_i be the set of contracts and K_i the number of contracts of firm i . For each contract p_i^k , firm i specifies a strict preference ordering over the set of workers eligible for this contract $W_i^k \subseteq W$. Furthermore, firm i has a strict ordering over which contract should be filled first. Let the first contract be p_i^1 , and only if firm i cannot fill the position at p_i^1 will firm i try to recruit workers at p_i^2 , and so on.

Firm f has preferences over $\{f\} \cup P_i \times W$, where, by definition, for any k, j such that $k + j < K_i$, $\forall w \in W_i^k, \forall w' \in W_i^{k+j}, (p_f^k, w) \succ_f (p_f^{k+j}, w')$.

Let P_F be the total set of contracts, where $p_f \in P_f \iff f \in F$ and p is a contract that f offers, that is, $P_F = \cup_{f \in F} P_f$. A worker w has preferences over $\{w\} \cup P_F$.

A matching is a function $\mu : P_F \cup W \rightarrow P_F \cup W$ such that $\forall w, p_f(i) |\mu(w)| = |\mu(p_f)| = 1$, (ii) $\mu(w) \in P_F \cup \{w\}$ and $\mu(p_f) \in W \cup \{p_f\}$ and (iii) $\mu(w) = p_f \iff \mu(p_f) = w$ and (iv) $\forall f: |\{p_f : \mu(p_f) \in W\}| \leq 1$.

For any matching μ , in slight abuse of notation, let $\mu(f)$ be the contract worker pair in case f has one of its contracts filled, and otherwise let $\mu(f)$ be f .

A matching is *stable* if (i) $\forall w, p_f, f$: If $\mu(w) = p_f$ then $\mu(w) \succ_w w$, and $\mu(f) \succ_f f$; and (ii) $\nexists f, p_f, w$ such that $p_f \succ_w \mu(w)$ and $(p_f, w) \succ_f \mu(f)$.

PROOF OF THEOREM 4 (Stability):

Stability in the case of firm proposing MDA is trivial, as any firm makes an offer to any worker it prefers more, and gets rejected by that worker (which implies the worker has a better offer in hand). To show stability in the worker proposing MDA outcome, note that for a given set of contracts used in the final MDA step, the outcome is stable to deviations that use only these contracts, because the DA yields stable outcomes (see Gale and Shapley 1962). Therefore, I need only show that no worker prefers a position p_i^j that got reverted to p_i^{j+1} . Suppose that at some step in the MDA, at the end of the DA part, a position p_i^j is unfilled. Let the interim matching be μ , where μ is the worker optimal stable match given the contracts available. Then, at μ , for any worker w eligible for p_i^j , $\mu(w) \succ_w p_i^j$. Now p_i^j gets reverted into p_i^{j+1} . Technically, this is equivalent to adding a new firm to an existing market. By Gale and Sotomayor (1985), adding a firm implies that the new worker optimal stable match μ' satisfies for any worker w : $\mu'(w) \succ_w \mu(w)$. That is, every worker eligible for p_i^j still has $\mu'(w) \succ_w p_i^j$. This is true for any reversion, that is, no worker would accept a contract that was reverted into another contract.

PROOF OF THEOREM 5 (Firm-Optimal Stable Match):

For a firm f , define a worker contract pair (p_f, w) to be achievable if there exists a stable matching at which firm f is matched to worker w at contract p_f . I show, by induction, that the stable outcome produced by the firm proposing MDA matches every firm to their most preferred achievable worker contract pair, and is therefore the (unique) firm optimal stable matching. Assume that up to a given step in the procedure, no firm has yet been rejected at a contract by a worker who is achievable. At this step, suppose that worker w rejects firm f at contract p_f . If worker w rejects firm f at contract p_f as unacceptable (i.e., $w \succ_w p_f$), then this worker is unachievable at this contract and I'm done. If worker w rejects firm f at contract p_f in favor of a firm g at contract p_g , I show that w is not achievable for firm f at contract p_f .

Firm g prefers w at p_g to any other worker contract pair except for those workers who have already rejected firm g at contract p_g and at any contracts in place before the contract got reverted into p_g , and hence (by the inductive assumption) are unachievable to firm g . Consider a hypothetical matching μ that matches firm f to worker w at contract p_f and everyone else to an achievable worker contract pair. Then firm g prefers w at contract p_g to the achievable worker-contract pair at μ . So, the matching μ is unstable, since it is blocked by (g, p_g, w) who prefer each other to their match at μ . Therefore, there is no stable matching that matches f to w at p_f , and so worker w is not achievable to firm f at contract p_f , which completes the proof.

PROOF OF THEOREM 6:

The worker proposing and firm proposing MDA follow the same steps, the only difference being that the interim matching is reached by either a firm proposing or a worker proposing DA. However, D. G. McVitie and L. B. Wilson (1970) and Roth (1984a) showed that for a given set of workers and firms, all stable matchings have the same workers and positions matched and share the set of unmatched workers and positions. This implies that at any interim match at the end of a DA step, in both MDA algorithms, the same positions are unfilled and get reverted into the same set of new contracts. I have already made the argument that the DA part of firms can as easily be thought of as one in which all former offers are annulled and remade.

PROOF OF COROLLARY 1:

By adding undesirable contracts at the top of the preference list (or scratching them), a firm does not influence the set of stable matchings. Hence, a firm does not influence the set of contracts that are the outcome of both the firm and worker proposing MDA.

In the worker proposing MDA, no firm can benefit from delaying its reversion of positions, as the more steps of the MDA pass, the more desirable the competing positions become. Since the DA step restarts whenever there is a change in a contract, delaying to revert a position, that is, having a round in which a position by a firm is unfilled has no effect. In the firm proposing algorithm, the statement is equivalent to the statement that, in a regular DA algorithm, some firms may start making offers only after others firms have already made offers. This does not affect the outcome of a DA algorithm.

REFERENCES

- Anantham, Siva, and Jennifer Stack.** 2006. "Unraveling and Wage Formation in Entry-Level Markets." Unpublished.
- Artemov, Georgy.** 2006. "Matching and Price Competition: Would Personalized Prices Help?" http://www.econ.brown.edu/students/Georgy_Artemov/artemov_personalized_mistakes_IJGT_f.pdf.
- Avery, Christopher, Andrew Fairbanks, and Richard Zeckhauser.** 2003. *The Early Admissions Game: Joining the Elite*. Cambridge, MA: Harvard University Press.
- ▶ **Bulow, Jeremy, and Jonathan Levin.** 2006. "Matching and Price Competition." *American Economic Review*, 96(3): 652–68.
- Coles, Peter.** 2005. "Strategic Behavior in Matching Markets." Unpublished.
- Crawford, Vincent P.** Forthcoming. "The Flexible-Salary Match: A Proposal to Increase the Salary Flexibility of the National Resident Matching Program." *Journal of Economic Behavior and Organization*.
- ▶ **Crawford, Vincent P., and Elsie Marie Knoer.** 1981. "Job Matching with Heterogeneous Firms and Workers." *Econometrica*, 49(2): 437–50.
- ▶ **Gale, David, and Lloyd S. Shapley.** 1962. "College Admissions and the Stability of Marriage." *American Mathematical Monthly*, 69(1): 9–15.
- ▶ **Gale, David, and Marilda Sotomayor.** 1985. "Some Remarks on the Stable Matching Problem." *Discrete Applied Mathematics*, 11: 223–32.
- ▶ **Hatfield, John William, and Paul R. Milgrom.** 2005. "Matching with Contracts." *American Economic Review*, 95(4): 913–35.
- ▶ **Kelso, Alexander S., Jr., and Vincent P. Crawford.** 1982. "Job Matching, Coalition Formation, and Gross Substitutes." *Econometrica*, 50(6): 1483–1504.
- ▶ **Kojima, Fuhito.** 2007. "Matching and Price Competition: Comment." *American Economic Review*, 97(3): 1027–31.
- Kojima, Fuhito, and Parag A. Pathak.** 2006. "Incentives and Stability in Large Two-Sided Matching Markets." <http://www.people.fas.harvard.edu/~kojima/largematchingmarket.pdf>.
- Konishi, Hideo, and Margarita Sapozhnikov.** 2006. "Decentralized Matching Markets with Endogeneous Salaries." <http://www2.bc.edu/~konishih/dcmatch4.pdf>.
- ▶ **McKinney, C. Nicholas, Muriel Niederle, and Alvin E. Roth.** 2005. "The Collapse of a Medical Labor Clearinghouse (and Why Such Failures Are Rare)." *American Economic Review*, 95(3): 878–89.
- ▶ **McVitie, D. G., and L. B. Wilson.** 1970. "Stable Marriage Assignments for Unequal Sets." *BIT Numerical Mathematics*, 10: 295–309.
- Niederle, Muriel, Deborah D. Proctor, and Alvin E. Roth.** 2006. "What Will Be Needed for the New GI Fellowship Match to Succeed?" *Gastroenterology*, 130: 218–24.
- ▶ **Niederle, Muriel, and Alvin E. Roth.** 2003a. "Relationship between Wages and Presence of a Match in Medical Fellowships." *Journal of the American Medical Association*, 290(9): 1153–54.
- ▶ **Niederle, Muriel, and Alvin E. Roth.** 2003b. "Unraveling Reduces Mobility in a Labor Market: Gastroenterology with and without a Centralized Match." *Journal of Political Economy*, 111(6): 1342–52.
- Niederle, Muriel, and Alvin E. Roth.** 2004. "The Gastroenterology Fellowship Match: How It Failed, and Why It Could Succeed Once Again." *Gastroenterology*, 127: 658–66.
- ▶ **Niederle, Muriel, and Alvin E. Roth.** 2005. "The Gastroenterology Fellowship Market: Should There Be a Match?" *American Economic Review*, 95(2): 372–75.
- Niederle, Muriel, and Alvin E. Roth.** 2007. "Making Markets Thick: Designing Rules for Offers and Acceptances." <http://www.stanford.edu/~niederle/MakingMarketsThick.May2007.pdf>.
- ▶ **Roth, Alvin E.** 1984a. "The Evolution of the Labor Market for Medical Interns and Residents: A

Case Study in Game Theory.” *Journal of Political Economy*, 92: 991–1016.

- ▶ **Roth, Alvin E.** 1984b. “Stability and Polarization of Interests in Job Matching.” *Econometrica*, 52(1): 47–57.
- ▶ **Roth, Alvin E.** 2002. “The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics.” *Econometrica*, 70(4): 1341–78.
- ▶ **Roth, Alvin E.** 2003. “The Origins, History, and Design of the Resident Match.” *Journal of the*

American Medical Association, 289(7): 909–12.

- ▶ **Roth, Alvin E., and Elliott Peranson.** 1999. “The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design.” *American Economic Review*, 89(4): 748–80.

Roth, Alvin E., and Marilda A. Oliveira Sotomayor. 1990. *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis (Econometric Society Monographs No. 18)*. Cambridge, MA: Cambridge University Press.