NBER WORKING PAPER SERIES

COMPETITIVE WAGES IN A MATCH WITH ORDERED CONTRACTS

Muriel Niederle

Working Paper 12334
http://www.nber.org/papers/w12334

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2006

webpage: www.stanford.edu/~niederle, email: niederle@stanford.edu. I thank Georgy Artemov, Jonathan Levin, Eric Maskin, Al Roth and Michael Schwarz for helpful discussions, the NSF for generous support and the Institute for Advanced Study at Princeton for its hospitality.

Competitive Wages in a Match with Ordered Contracts
Muriel Niederle
NBER Working Paper No. 12334
June 2006
JEL No. D4, J3, J4

ABSTRACT

A recent antitrust lawsuit against the National Residency Matching Program renewed interest in understanding the effects of a centralized match on wages of medical residents. Bulow and Levin (forthcoming) propose a simple model of the NRMP, in which firms set impersonal salaries simultaneously, before matching with workers, and show that a match leads to lower aggregate wages compared to any competitive outcome.

This paper models a feature present in the NRMP, ordered contracts, that allows firms to set several contracts while determining the order in which they try to fill these contracts. I show that the low wage equilibrium of Bulow and Levin is not robust to this feature of the NRMP, and competitive wages are once more an equilibrium outcome. Furthermore, a match with ordered contracts has different properties than former models of centralized matches with multiple contracts.

Muriel Niederle
School of Social Science
Institute for Advanced Study
Einstein Drive
Princeton, NJ 08540
and NBER
niederle@stanford.edu

1 Introduction

In response to a recent antitrust lawsuit against the National Residency Matching Program, Bulow and Levin (forthcoming) 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. 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 Knoer 1981; Kelso and Crawford, 1982, Roth 1984b, and see also, Hatfield and Milgrom 2005).

Here, I observe that the NRMP has a feature that I will call ordered contracts, that 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 determine 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 new and additional control firms have over contracts results in new properties concerning which stable outcomes are reached through different adequately modified deferred acceptance algorithms. Most importantly, the set of contracts reached through either a firm or worker proposing suitably modified deferred acceptance algorithm is the same, which is in general not the case in a multiple contract match.

These results will cast the lesson we learned from Bulow and Levin in a new light.

Since 1951 the market for medical residents has been organized through a centralized matching procedure which assigns medical students to residency programs using a variant of a deferred acceptance algorithm. 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). In 1998 Roth and Peranson introduce a new algorithm, that 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 1990's about 7 percent of the three to four thousand programs that participate in 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 that his residents will spend their first year in that prerequisite program. If the 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 reinstitution 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, Proctor and Roth 2006).

In 2002 an antitrust law suit 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 of using a centralized match such as the NRMP on wages. In their model, the surplus for firm n ($1 \leq n \leq N$) from hiring worker m is $\Delta_n \cdot m$, where $\Delta_N \geq \Delta_{N-1} .. \geq \Delta_1$. Workers care only about their salary, p. In their simple model of the NRMP firms simultaneously announce a wage at which they are willing to hire any worker. Workers form preferences over firms after the announcement, where each worker prefers firms with higher wages. Then an assortative matching occurs, that 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.[1]

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, in fact the NRMP allows for ordered contracts, and in this paper I show that a match with ordered contracts allows for more wage competition that can restore competitive outcomes.

For a stylized model of the NRMP with ordered contracts, I introduce a small change in the Bulow and Levin framework, which will have a big effect on wages. Each firm i, instead of advertising only one position at one contract (wage) can create a second contract for the same position, and decide which subset of workers is eligible. Specifically, firm i, instead of announcing only one wage p_i can announce two (or more) wages p_i and p_i^S, and determine which workers are eligible for each contract. Firm i first tries to fill the position at contract p_i^S. If it

[1] The problem that firms cannot change their salary when they try to hire different workers becomes important when firms are asymmetric, that is when, for example, some firms are clearly 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 indendent of the worker they end up hiring. Furthermore, while the fact that firms cannot announce which (subset of) worker is eligible for their unique wage offer has some effect in reducing average wages, it cannot account for the whole problem.

fails to fill this 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 for that contract.

While Roth and Peranson (1999) mention that they include ordered contracts in their redesign of the NRMP, they did not present a formal model. Issues such as the appropriate definition of stability, existence of stable matches, incentives of applicants and firms to submit various contracts and rank order lists, and the difference from matches with multiple contracts as introduced by Crawford and Knoer (1981) have not been analyzed.[2] While a model with ordered contracts shares the definition of stability with a multiple contract match, I propose a modification to the deferred acceptance algorithm to accommodate ordered contracts. These modifications imply that the contracts reached are the same, whether a firm or applicant proposing algorithm is used. The reason is that a firm can control the order in which positions are made available as opposed to simply having a multitude of simultaneous potential contracts. After analyzing a model of ordered contracts, I show the two main effects on wages which are:

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.

A separate issue concerns what the wages of medical residents would be in a market without a match. 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, 2004a and 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. A survey of program directors in this market reveals that even after the market operated without a match for nearly a decade, most programs offer the same wage to all their fellows (i.e. impersonal wages). Furthermore, a comparison of wages of internal medicine subspecialty fellows in specialties that do and do not use a match reveals that wages are not different for specialties that use the match (Niederle and Roth 2003a and 2004a.) It seems that Gastroenterology programs prefer to use a dimension of contracts different from the terms of employment, namely exploding offers, or the amount of time an offer is available to close contracts, and to avoid losing fellows to their

[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, and hence do not pose any problems for the existence of stable matchings.

competitors (Niederle, Proctor and Roth 2006).[3]

The NRMP, a matching program that allows for ordered contracts, therefore allows for wage competition. Furthermore, eliminating a centralized match, such as the NRMP, may result in an overall reduction of competition, as wages seem not to be affected, but hospitals use a decentralized unraveled market to exercise local market power in time to limit the positions candidates can contemplate (see also Niederle and Roth 2004b.) That is, it appears that the NRMP does not in fact force wages to be impersonal. Therefore, 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. that wages may be more compressed then if each worker were paid his marginal product. However, the match is not the cause.

2 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 3 firms, and 3 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 firm 1 and worker 1 both receive non-negative 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 compared to employing worker 1 at p_1, that is $1 - p_1 \geq 2 - p_2$. Furthermore, p_2 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 $4 - p_2 \geq 2 - p_1$, hence $p_2 \in [p_1 + 1, p_1 + 2]$. 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. In the example, 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. This of course in turn will make firm 2 want to compete for worker 3 and hence increase its wage. Bulow and Levin (forthcoming) 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

[3] Avery, Fairbanks and Zeckhauser (2003) show that similarly in the market for college students, colleges use the option of early decision, to secure students and limit their availability to competitors.

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 clearly result in wage compression, and in 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, instead of advertising only one position at one wage can create a second contract for that same position, for which only a subset of workers is eligible. Specifically, firm i, instead of announcing only one wage p_i announces two wages p_i and p_i^S, and determines which of the workers are eligible for p_i^S, 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^S 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^S, 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^S. If firm i is unmatched, then p_i^S 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^S. Then the following strategy makes firm 1 strictly better off than simply only announcing one wage p_1 according to the Bulow and Levin equilibrium. Let $p_1^S = 1/6$, 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^S. Then the expected surplus of firm 1 is 1.08, strictly higher than 1. The intuition is that compared to Bulow and Levin, where firm 1 is indifferent between all strategies, and hence its payoff is for example determined by offering wage $p_1 = 1/3$, and hiring worker 2 and worker 1, firm 1 with the use of its star-position has 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 $p^S = 4/3$, 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, then 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 $p_1^S = 1$, 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 $p_2^S = 3$, 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 of a star contract, so, let firm 3 only use 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, and that the strategies form a Nash Equilibrium. Given the contracts of all other firms, it is easy to see that no firm can gain by deviating, as each firm is disciplined in its wage offer by the slightly less productive firm, and no worker can gain from declining to rank some contracts. 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.[4] Similarly, firm 2 has to offer the star-contract only with probability 2/3, to effectively discipline firm 3's wage offer.

Before I show that the results of the example generalize, I analyze matching with ordered contracts. I show existence of stable matchings, how deferred acceptance algorithms have to be modified to account for ordered contracts, and how a match with ordered contracts differs from a match with multiple contracts as introduced by Crawford and Knoer (1981) and Kelso and Crawford (1982).

3 Matching with Ordered Contracts

In a first section I summarize known findings on simple markets in which each firm offers only one fixed wage for their position (see also Roth and Sotomayor 1990). Then I formally define markets with ordered contracts, show existence of stable matchings and how some standard results fail to be true in this more general framework.

3.1 Stable Matching in a simple fixed-wage framework

The market consists of a set F of firms and a set W of workers, where every firm has specified one fixed contract for the position it advertises. Every worker w has strict preferences $\succ_w$ over all firms and herself (if she prefers at some point to rather be unmatched), that is, $\succ_w$ implies a strict ordering on the set $F \cup \{w\}$. Similarly, every firm has strict preferences over $W \cup \{f\}$. Each worker can work for (be matched to) at most one firm, and every firm can employ (be matched to) at most one worker. Formally, a *matching* μ is a function $\mu : F \cup W \to F \cup W$ such that $\forall w, f$: (i) $|\mu(w)| = |\mu(f)| = 1$, (ii) $\mu(w) \in F \cup \{w\}$ and $\mu(f) \in W \cup \{f\}$ and (iii) $\mu(w) = f \Leftrightarrow \mu(f) = w$. A matching is *stable*, if every firm and worker is matched to an

[4] Firm 2 prefers to announce the contract $p_2 = 1$ versus $p_2 = 0$, if firm 1 offers $p_1^S = 1$ with probability q such that $(2 \cdot 2 - 1) \geq (1 - q)(2 \cdot 2 - 0) + q(2 \cdot 1 - 0)$, that is $q \geq 1/2$. Furthermore, firm 2 cannot use his star-contract to exploit the fact that firm 1 offers $p_1^S = 1$ with probability q, and whenever firm 1 does not offer a starcontract, try to recruit worker 2 with $p_2^S = 0$ (where only worker 2 is eligible) and in case the position is not filled, that is, whenever firm 1 does offer the starcontract, 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.

acceptable partner (that is, no firm or worker is rather unmatched than remaining with the current match) and if there does not exist a firm-worker pair that are currently not matched to each other, but prefer each other to their current match. Formally, a matching μ is *stable* if (*i*) $\forall w, f$: If $\mu(w) = f$ then $\mu(w) \succ_w w$ and $\mu(f) \succ_f f$, and (*ii*) $\nexists f, w$ such that $f \succ_w \mu(w)$ and $w \succ_f \mu(f)$.

In this simple model, the set of stable matchings is always nonempty, and there exists a stable match that all firms prefer over any other stable match (and which all workers like less than any other stable matching). This firm optimal stable matching can be obtained by the *Firm Proposing Deferred Acceptance Algorithm DA* (Gale and Shapley 1962).

> 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.
>
> 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 algorithm ends when either no firm has its offer rejected, or all rejected firms have already been rejected by all the workers to which they are willing to make an offer.

The key to stability of this outcome is that no worker ever regrets having rejected a firm's offer, since she does so only when she has an offer she prefers, and she will be matched to that preferred firm, unless she receives an offer she prefers even more. This implies that no firm whose offer is rejected in the algorithm has any hope that the offer would be accepted at any later stage. Furthermore, the stable outcome reached is the firm optimal stable outcome, that is the stable outcome that any firm weakly prefers over any other stable outcome.

Furthermore, it is a dominant strategy for firms to submit their true preferences, that is they cannot gain by misrepresentation (Dubins and Freedman, 1981, and Roth 1982). This is not true for workers, who can gain by truncating their preferences. In theory the amount of optimal truncation can be substantial (see Coles 2005), though in practice they may not be that large (Roth and Peranson 1999). Similarly, there exists a worker optimal stable matching that is reached by a worker proposing deferred acceptance algorithm.

3.2 Stable Matching with Ordered Contracts

In a model with ordered contracts every firm i can have up to K contracts $p_i^1, ..., p_i^K$, 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 $W_i^k \subseteq W$. For example, firm i may decide that a specific worker w may be eligible for p_i^1, but not for p_i^2.

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 $p_f \in P_F \leftrightharpoons f \in F$ and p is a contract that f offers. A worker w has preferences over $\{w\} \cup P_F$.

The definition of a matching is similar to before, only that we replace firms with contracts, and additionally ensure that each firm has only one of its contracts filled. Formally, 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) $\mu(w) = p_f \Leftrightarrow \mu(p_f) = w$ and (iv) $\forall f : |\{p_f : \mu(p_f) \in W\}| \leq 1$.

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

A matching is stable if there exists no firm, worker, 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, than stay with their current match. Formally, 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 $(p_f, w) \succ_f \mu(f)$.

To show that a stable matching always exists, I provide a modified deferred acceptance algorithm whose outcome is always a stable match.

I start with the simpler firm proposing modified deferred acceptance algorithm, before I show the worker proposing modified deferred acceptance algorithm. Then I show that they both yield stable outcomes.

The *Firm Proposing Modified Deferred Acceptance Algorithm: MDA*

MDA Step 1: All firms have only their first contract available.

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.

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 sub-algorithm 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.

MDA Step k: Any firm that has its position at 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 sub-algorithm, where all previous offers are cancelled and have to be remade.

The algorithm ends, when all firms that have no offer held by an applicant have no more contract to change or revert to. Workers that hold an offer from a firm at a contract are matched to that firm at that contract, remaining firms and workers are unmatched.

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.

The *Worker Proposing modified Deferred Acceptance algorithm: MDA*

MDA Step 1: All firms have only their first contract available.

DA: Worker proposing DA sub-algorithm which ends when either no worker has her offer rejected, or all rejected workers have no more firms they want to make an offer to, at the current contract.

MDA Step k: Any firm that has its position at 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 a firm changed its contract.

Theorem 1 ***Stability:*** *Whenever firms have a strict ordering over a finite number of contracts, that is 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 2 ***Firm-optimal stable match:*** *The firm proposing MDA yields the firm optimal stable match.*

But, as we will see below, the worker proposing *MDA* need not yield the worker optimal stable match.

That is in a two-sided matching market with ordered contracts, just as in simple matching markets, we can define the set of stable outcomes. Note that for stability the special dynamic feature of ordered contracts is not taken into account, and so, theorems that use stability in a market with multiple contracts still apply. Finally, a version of a Gale Shapley deferred acceptance algorithm yields stable outcomes, and in case firms propose, indeed the stable outcome that all firms prefer over any other stable outcome.

3.3 Some incentive properties of matchings achieved through MDA

Before analyzing incentive properties the modified deferred acceptance algorithm imposes on firms and workers when submitting their preferences, I investigate in more detail how the outcome of the firm and worker proposing *MDA* differ. In simple matching environments, McVitie and Wilson (1970) and Roth (1984a) showed that all stable matchings have the same workers and positions matched and share the set of unmatched workers and positions. We can find a version of this theorem in case firms have ordered contracts about the set of contracts matched when either the firm or worker proposing *MDA* is used.

Theorem 3 *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 reversion of contracts, all the existing offers have to be cancelled and the whole offering process has to start anew. However, Theorem 3 implies 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 do not have any incentive to add (or scratch) undesirable positions (i.e. positions that will not be filled), on the top of their preference list. That is 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 MDA's, 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 MDA's the order in which firms revert contracts is irrelevant to the final outcome.*

However, Theorem 3 also immediately implies that workers may have an incentive to manipulate their preferences, as, whichever *MDA* is used, firm or worker proposing, the set of contracts filled is the same.

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 than the firm optimal MDA. This implies that workers may have an incentive to misrepresent their preferences even in the worker proposing MDA.*

To illustrate the statements of corollary 2, an example is most useful.

Example: Suppose there are two firms f_1 and f_2 with two contracts for one position each, one at low wage L and one at high wage H, 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. There are four stable matchings:

$$\mu_W = \begin{pmatrix} h_1 & h_2 \\ w_1 & w_2 \end{pmatrix}, \mu_F = \begin{pmatrix} l_1 & l_2 \\ w_1 & w_2 \end{pmatrix}, \nu = \begin{pmatrix} h_1 & l_2 \\ w_1 & w_2 \end{pmatrix}, \rho = \begin{pmatrix} l_1 & h_2 \\ w_2 & w_1 \end{pmatrix}$$

The worker optimal stable match μ_W cannot be reached through the worker proposing modified deferred acceptance, in fact both firm and worker proposing *MDA*'s lead to μ_F. To see that, note that in a match with ordered contracts, firms submit first the order of contracts to be used when trying to fill the position, in this case, the first contract is the low wage contract L, and the second contract is H. Then, for each contract the firms submit the rank order over workers, in this case, for both L and H that would be w_1, w_2. In a worker proposing *MDA*, in the first round, the available contracts are, for each firm, the low wage contract L. Given these contracts, the first deferred acceptance subalgorithm yields a matching of worker w_i with firm f_i, , for $i = 1, 2$ and hence matching μ_F.

Furthermore, each worker has an incentive to misrepresent her preferences. If any of the two workers submits preferences such that only high wage contracts are acceptable, otherwise the worker prefers to be unmatched, then the worker can guarantee herself a high wage.

The example also highlights the differences of a match with ordered contracts, to a multiple contracts match in which salaries are determined using a standard version of a deferred acceptance algorithm. In this alternative way to introduce wages (Crawford and Knoer 1981, and Kelso and Crawford 1982) or also general contracts (Roth 1984b and Hatfield and Milgrom 2005) a firm i has a finite set of contracts $q_i \in Q_i$, and preferences over $Q_i \times W$, and every worker w has preferences over $\cup_i Q_i \cup \{w\}$. 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 (at most $|Q_i| \cdot |W|$ for each firm i). One difference of standard models of multiple contracts to a model with ordered contracts is that in the standard model, all contracts are potentially present simultaneously and available immediately. Therefore, if workers were to make offers, they could 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 once more best be seen in the continuation of the former example:

Example continued: Suppose firms now have both wage offers available immediately, that is each firm i submits a preference list $(w_1, l), (w_2, l), (w_1, h), (w_2, h)$, that is each firm most prefers a worker at a low wage, and for a given wage prefers w_1 to w_2. Analogously, each worker w submits preferences $(f_1, h), (f_2, h), (f_1, l), (f_2, l)$. Given these submitted rank order lists, a worker proposing deferred acceptance algorithm will yield μ_W, as first, both workers make an offer to firm 1 at a high wage, firm 1 keeps the offer of worker 1 and rejects w_2's offer, who makes an offer to work for firm 2 at the high wage, which is a stable outcome. The difference to the match with ordered contracts is that in this multiple contract match the high wage contracts are available immediately, that is workers can apply to them right away. In the case of ordered contracts, every firm has only one contract available at any round of a deferred acceptance algorithm, and only if a firm failed to match, does the position change to a new contract.

That is, 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 is the same set of firms matched, but also the same set of contracts, between a firm proposing and a worker proposing modified deferred acceptance algorithm.

The robustness of this feature to a worker and firm proposing algorithm is helpful in actual applications, as the NRMP switched to a student proposing algorithm in 1998.[5] Introducing the possibility to have several contracts ranked simultaneously would greatly affect the outcome of the match. However, the contracts filled in a match with ordered contracts are not affected by the choice of which side of the market makes offers.

[5]The algorithm designed by Roth and Peranson, and used by the NRMP, works like a worker proposing *MDA* (private communication from Alvin Roth.)

4 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 $\Delta_N \geq .. \geq \Delta_1 \geq 0$, 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 rank all workers, that is, required to be willing to employ any worker, which is inconsequential when $\Delta_n = n$, but not in general.

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 or equal than 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

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

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 $l(N-1), .., N$ compete on $(p_N, p_{N-1}]$ such that firm $N-1$ exhausts all its remaining offer probability. 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 4 ***Bulow and Levin (forthcoming):*** *There is a unique price equilibrium. Letting $q_n(.)$ and $p_1, .., p_{N+1}$ and $G_1(0)$ be defines as above, then for each firm n, and each non-empty 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.

Theorem 5 *Suppose all firms have only one position, and offer wages p_i according to Bulow and Levin. If some firm i can offer two wages, p_i^S - 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 6 ***Competitive Equilibrium Wages**: The following strategies form a Nash equilibrium: Every firm announces $p_i = c_i$ for $1 \leq i \leq N$ and $p_N^S = c_N$ with only worker w_N being eligible, and $p_j^S = c_{j+1}$ for $j < N$ and the workers being eligible for p_j^S 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, and for a given wage more productive firms.*

That is 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.

5 Wages of Medical Fellows with and without a Match

The National Residency Matching Program allows for ordered contracts. In the 1990's about 7 percent of the three to four thousand programs that participate in each year have positions with contracts that could revert to other contracts if they remain unfilled (accounting for almost 6 percent of the total quota of positions).[6] In the reinstitution of the fellowship match for Gastroenterology fellows, this feature is also 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, Proctor and Roth 2006).

Ordered contracts therefore allow a program to replicate how they may try to fill positions in a decentralized market, namely to try to find a research fellow, and only go for a clinical fellow in case no suitable candidate can be attracted.[7] In an ordered contracts match a program can do

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

[7] Many economics departments that may have tried to hire a senior faculty in a field, opt for a junior person once they were not so successful.

that without compromising the quality of the applicant pool they can consider for their second contract (Corollary 1). It is however not clear 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 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.

Indeed, Niederle and Roth (2004), and Niederle, Proctor and Roth (2006) showed that the market for gastroenterology fellows, after they stopped using the centralized match, once more unraveled.[8] The market fails to be one big market in which all positions are offered simultaneously. A survey of gastroenterology program directors reveals that programs use terse deadlines on offers that limit the possibilities of candidates to simultaneously consider other offers. And while a few program directors (9% of respondents) did not offer the same wage to all their incoming fellows, 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.

The survey results are supported by data on the internal medicine fellowship market. Niederle and Roth (2003b) show that the market is more localized without a centralized match, that is, more gastroenterology fellows remain at the same hospital, the same city and the same state where they did their internal medicine residency, than with a match. Furthermore, one can compare wages of fellows whose specialties participate in a match, with wages of fellows whose specialty matches 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 MSMP (Medical Specialties Matching Program) 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, comparing wages within hospitals for specialties that use a match and that do not, finds only 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).

That is, it is not clear that a match compresses or lowers wages, because on the one hand, the ordered contracts match used by the NRMP allows for wage competition. Furthermore, the gastroenterology market, as well as the history of the residency market strongly suggest that a market without a centralized match may not be a competitive market. While the NRMP

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

does not in fact force wages to be impersonal, we still observe a lot of impersonal wages. In that, residents or first year fellow may however not be unique. For example, in many economic 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, that wages are more compressed than if each worker were paid their marginal productivity, however the match does not seem to be the major culprit.

6 Conclusions

The NRMP is a match with ordered contracts, which allows for programs to effectively have more than one contract per position. This paper analyzes the properties of a match with ordered contracts, and as such fills a gap in the existing matching literature by analyzing the actual algorithm used by the NRMP, and providing a new tool to be used for future designs of centralized matches. I show that a match with ordered contracts has novel implications, most notably that the set of contracts filled is the same, whether a firm or worker proposing modified deferred acceptance algorithm is used.

In light of the controversy about the effects of a centralized match on wages, this paper shows that a match with ordered contracts allows for competitive equilibrium wages. Furthermore, ordered contracts allow programs to try to hire different kinds of candidates, without taking any penalty on the set of applicants available once they unsuccessfully tried to fill a position under their first contract. Compared to a decentralized market a centralized match may increase competition, as it does not allow programs to lock in candidates early, before other programs have finished reverting their contracts to different contracts.

7 Appendix

Proof of Theorem 1: Stability: Stability in the case of firm proposing *MDA* is trivial, as any firm made an offer to any worker it preferred more, and got 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 only use these contracts, because the *DA* yields stable outcomes (see Gale and Shapley 1962, and section 3.1). Therefore I only need to show that no worker prefers a position p_i^j that got reverted into 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: $\mu'(w) \succeq_w \mu(w)$. 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 2: 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 give 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 that 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 3: 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, McVitie and 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 Deferred Acceptance algorithm, some firms may start making offers only after some others firms already made offers. This does not affect the outcome of a Deferred Acceptance algorithm. ■

Proof of Theorem 5: 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 one w_L. The highest (and lowest) worker is 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, then foregoing the possibility to use a p_i^S job at all. The reason is that firm i is indifferent between 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 $w_H, ..., w_i$ 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 6: 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 eligible, and 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.■

References

[1] Avery, Christopher, Andrew Fairbanks and Richard Zeckhauser, *The Early Admissions Game: Joining the Elite*, Harvard University Press, Cambridge, MA 2003.

[2] Bulow, Jeremy and Jonathan Levin, "Matching and Price Competition", *American Economic Review,* forthcoming.

[3] Coles, Peter, "Strategic Behavior in Matching Markets", working paper 2005.

[4] Crawford, Vincent P., "The Flexible-Salary Match: A Proposal to Increase the Salary Flexibility of the National Resident Matching Program," forthcoming, *Journal of Economic Behavior and Organization.*

[5] Crawford, Vincent P., and Elsie Marie Knoer, "Job Matching with Heterogenous Firms and Workers," *Econometrica* 49, March 1981, 437 - 450.

[6] Dubins, L.E., and Freedman, "Machiavelli and the Gale-Shapley algorithm," *American Mathematical Monthly*, 88, 1981, 485-494.

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

[8] Gale, David, and Marilda Sotomayor, "Some remarks on the stable matching problem," *Discrete Applied Mathematics*, 11, 1985, 223-232.

[9] Hatfield, John, and Paul Milgrom, "Matching with Contracts," *American Economic Review,* Vol. 95, No. 4., September 2005, 913-935.

[10] Kelso, Alexander S., and Vincent P. Crawford, "Job Matching, Coalition Formation, and Gross Substitutes," *Econometrica*, 50 (November 1982, 1483 - 1504.

[11] McKinney, C. Nicholas, Muriel Niederle, and Alvin E. Roth, "The collapse of a medical labor clearinghouse (and why such failures are rare)," *American Economic Review*, vol. 95, No.3, June 2005, 878-889.

[12] McVitie, D.G., and L.B. Wilson, "Stable marriage assignments for unequal sets," *BIT,* 10, 1970, 295-309.

[13] Niederle, Muriel, Deborah D. Proctor and Alvin E. Roth, "What will be needed for the new GI fellowship match to succeed?", *Gastroenterology*, 130, January 2006, 218-224.

[14] Niederle, Muriel, and Alvin E. Roth, "Relationship Between Wages and Presence of a Match in Medical Fellowships," *JAMA, Journal of the American Medical Association*, vol. 290, No. 9, September 3, 2003a, 1153-1154.

[15] Niederle, Muriel, and Alvin E. Roth, "Unraveling reduces mobility in a labor market: Gastroenterology with and without a centralized match," *Journal of Political Economy*, vol.111, No. 6, December 2003b, 1342 – 1352.

[16] Niederle, Muriel, and Alvin E. Roth, "The Gastroenterology Fellowship Match: How it failed, and why it could succeed once again", *Gastroenterology*, 127, August 2004, 658-666.

[17] Niederle, Muriel, and Alvin E. Roth, "Market Culture: How Norms Governing Exploding Offers Affect Market Performance", January 2004b, and NBER working paper 10256.

[18] Niederle, Muriel, and Alvin E. Roth, "The Gastroenterology Fellowship Market: Should there be a Match?" *American Economic Review Papers & Proceedings*, vol. 95, No.2, May 2005.

[19] Roth, Alvin E., "The economics of Matching: Stability and Incentives", *Mathematics of Operations Research*, 7, 1982, 617 - 628.

[20] Roth, Alvin E., "The evolution of the labor market for medical interns and residents: A case study in game theory", *Journal of Political Economy*, vol. 92, 1984a, 991-1016.

[21] Roth, Alvin E., "Stability and Polarization of Interests in Job Matching," *Econometrica*, 1984b, 52(1), 47-57.

[22] Roth, Alvin E. "The Origins, History, and Design of the Resident Match," *JAMA, Journal of the American Medical Association*, vol. 289, No. 7, February 19, 2003, 909-912.

[23] Roth, Alvin E. and Marilda Sotomayor, *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*, Econometric Society Monograph Series, Cambridge University Press, 1990.

[24] Roth, Alvin E. and Elliot Peranson, "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design", *American Economic Review,* vol 89, no. 4, September 1999, 748 - 779.