Matching Through Decentralized Markets

Muriel Niederle *　　　Leeat Yariv†‡

Preliminary Draft, November 12, 2008

Abstract. We study a simple model of a decentralized market game in which firms make directed offers to workers. We identify three components of the market game that are key in determining whether stable matches can arise as equilibrium outcomes. The first is related to the structure of preferences of agents. The second pertains to the agents' information on preferences. The third is whether there are frictions in the market, which in this paper take the form of discounting. Our results show that complete information, or at least frictionless economies are needed for the existence of equilibrium strategies that yield the stable outcome. In the presence of uncertainty, as soon as frictions are introduced, much harsher assumptions have to be made to guarantee existence of an equilibrium that yields the stable match.

Keywords: Decentralized Matching, Stability, Market Design.

*Department of Economics, Stanford University and NBER, http://www.stanford.edu/~niederle

†Division of the Humanities and Social Sciences, Caltech, http://www.hss.caltech.edu/~lyariv

‡We are grateful to Attila Ambrus, Federico Echenique, Fuhito Kojima, and Al Roth for many useful conversations and suggestions. This work is supported by NSF and the Sloan Foundation.

1. Introduction

1.1 Overview

The theoretical literature on two sided matching markets has focused predominantly on the analysis of outcomes generated in centralized markets. There are many examples in which two sided matching markets are centralized (e.g., the medical residency match, school allocations, the U.S. market for reform rabbis, etc.). Nonetheless, many markets are not fully centralized (for instance, college admissions in the U.S., the market for law clerks, junior economists, and so on). Furthermore, almost all centralized markets are preceded by decentralized opportunities for participants to match.[1] Understanding the outcomes generated by decentralized markets is therefore important to the design of institutions, both fully decentralized ones, as well as ones followed by centralized procedures.[2] The current paper offers a first step in that direction.

We provide a simple model of a decentralized *market game* in which firms make directed offers to workers. In our setup, a market game is identified by four components: the preference distribution of agents (workers and firms), the information agents have about their own and others' realized preferences, whether agents have full information about all offers or not, and whether there are frictions in the economy. While stability is the common notion used to describe market outcomes,[3] both positively as well as prescriptively, we show that the market has to satisfy rather harsh conditions for it to be a unique *equilibrium* prediction.

In more detail, we focus on markets in which firms can employ up to one worker, who can work for at most one firm. We consider environments in which preferences of firms and workers assure that there is a unique stable matching. This allows us to sidestep coordination problems that make achieving a stable outcome more difficult. There are special classes of preferences that guarantee that the stable matching is unique. For example, when firms' and workers' preferences are aligned,[4]

[1] There are a few exceptions of markets in which decentralized bargaining prior to the centralized match are prohibited by design, such as some residency matches in the UK (Roth, 1991).

[2] Indeed, the consequences of a decentralized matching processes prior to a centralized match can be large, as documented by the collapse of the market for gastroentorology fellows (see Niederle and Roth, 2003).

[3] A *stable matching* is a pairing of workers and firms (where some workers and some firms may be left alone), in which no firm (worker) who is matched to a worker (firm), prefers to be alone, and no firm and worker pair prefer to jointly deviate by matching to one another.

[4] We say that workers' and firms' preferences are *aligned* if, whenever a firm i prefers worker j to worker j', the utility worker j gets from matching with firm i is greater than the utility worker j' gets from matching with firm i, and vice versa. This is the case when, for instance, a worker's and firm's utility from a match coincide, which is a case heavily studied in the empirical literature.

markets have a unique stable match (Proposition 1). We will concentrate our analysis on the case of aligned preferences, and the general case. When preferences are aligned, the market is such that there is always either an agent whose most preferred option is to remain unmatched, or a firm-worker pair that are each other's first choice.

In our decentralized market game, firms and workers interact over time. At the outset of the game, preferences of workers (over firms) and firms (over workers) are realized. In every period, each firm can make up to one offer to a worker of her choice if she does not already have an offer out in the market. Workers can accept, reject, or hold on to an offer, in which case it is available also next period. In the market game, firms and workers share a common discount factor, and receive their match utilities as soon as they are matched (by having an offer accepted), or leave the market. This allows us to study frictionless economies in which the discount factor is 1, and economies with frictions, in which the discount factor is strictly lower than 1.

With respect to information agents have about their own and each others' preferences, we consider two potential structures: (a) *complete information* in which all market participants are fully informed of the realized match utilities; (b) *private information* in which each agent is fully informed only of her own match utilities. While extant literature has mostly focused on complete information environments (see below), we find the case of private information particularly important from an empirical point of view. Indeed, the larger the market, the more likely there is to be some incomplete information regarding participants' preferences. In particular, in many of the matching markets studied in the literature (e.g., the medical match, the law clerks match, etc.), the volume of communication required to make everyone's preferences common knowledge would be tremendous.

Throughout our analysis, we concentrate on equilibria in weakly undominated strategies.

Under complete information, all agents can compute the stable matching and we show that the stable matching is an equilibrium outcome in the market game we analyze (Proposition 3).

Nonetheless, even when information is complete, in general there are unstable matchings that can be achieved as equilibrium outcomes. Note that this is at odds with findings in the case in which agents use a centralized market, in which all equilibria in weakly undominated strategies yield the stable outcome. Nonetheless, when preferences are aligned, stable matching is the unique equilibrium outcome (Propositions 3 and 4).

When agents are informed only of the realization of their own match utilities and preferences are aligned then, in a frictionless economy, stable matchings may still be implemented as an equilibrium outcome (Proposition 5). Underlying this result is the idea that firms and workers can replicate, in essence, the firm proposing Gale-Shapley deferred acceptance algorithm (Gale and Shapley, 1962) as part of an equilibrium profile: firms make offers to workers, and workers accept offers when they are made by their most preferred firm.

However, as soon as there are frictions in the market, agents may have incentives to deviate from these strategies to speed up the matching process, or affect market participants' learning regarding their expected stable matches. We characterize the class of strategies that constitute an equilibrium in centralized markets (utilizing the deferred acceptance algorithm) that implements the stable match (Proposition 6). We then convert this class of strategies to decentralized market strategies. These strategies impose, in a sense, minimal conditions for the generation of stable outcomes in the dynamic market game, however need not be incentive compatible. One of our main results restricts the class of economies for which they are. In fact, in Proposition 7 we show that when the market is sufficiently rich in terms of the possible market realizations, there is a Bayesian Nash equilibrium that implements the unique stable match.

For general preferences (that are not necessarily aligned), the stable match may be impossible to achieve in equilibrium even in the absence of frictions. This is due to the fact that in decentralized markets agents do not know when the market has "settled down", that is, when there are no more offers being made. The scope of market monitoring then becomes crucial in whether stability can be achieved in equilibrium. In fact, when we assume full market monitoring, i.e., that agents can observe all offers and responses to those offers, then once more there exist strategies that replicate the deferred acceptance mechanism and constitute an equilibrium yielding the stable outcome (Proposition 8). With discounting, even harsher restrictions than in the case of aligned preferences have to be imposed on the economy to guarantee the existence of such strategies (Proposition 9).

1.2 Related Literature

With the progress of empirical techniques and theoretical results on matching markets, recent papers have used some insights on stability to deduce market participants' characteristics. For instance, Sorensen (2007) finds that companies funded by more experienced venture capitalists are more likely to go public. He does so by assuming a matching market between investors and companies which is characterized by aligned preferences and therefore generates a unique stable match (much like in our Proposition 1). Sorensen can the identify and estimate influence and sorting in the market.[5]

There are several recent theoretical advances that are related to the current paper. Haeringer and Wooders (2007) Study a model similar to our complete information case, only with exploding offers. While not requiring a unique stable matching, they do place several restrictions on firms' strategies. Namely, firms cannot make repeat offers to workers who have rejected them, firms must have an offer out or else leave the market. Most importantly, a firm making an offer must make an offer to her most preferred worker whom she is able to make an offer to. Under these assumptions, Haeringer and Wooders show that the unique equilibrium outcome (though not necessarily profile) is the worker preferred stable match.

Blum, Roth, and Rothblum (1997) study a similar model to ours toward the end of their paper, and allow for some uncertainty, though, again, restrict firms to follow "preference strategies," which prescribe the order of offers using some preference ordering. The paper characterizes the Nash equilibria in this model and illustrates that they coincide with the centralized market outcomes they propose (see Pais, 2006, for related work).

Diamantoudi, Miyagawa, and Xue (2007) study the role of commitment in dynamic settings such as the one underlying our analysis when information is complete. They point to the difference in outcomes arising from different commitment abilities of workers and firms.

In the context of implementation, Alcalde and Romero-Medina (2000) study a game where there are only two stages. First, firms make offers. Then, workers reply and the game concludes. They demonstrate that this game implements the stable matchings (see also Alcade, 1996, and Alcalde, Pérez-Castrillo, and Romero-Medina, 1998).

The search literature has considered setups that bear some similarity to ours. Particularly, Burdett and Coles (1997) and Eeckhout (1999) consider a setup in which, at each point in time, workers and firms randomly encounter each other. Such an encounter entails observing the resulting match utilities, and deciding jointly whether to pursue the match and leave the market or to separate and wait for future periods. Each side of the market therefore solves an option value problem. If there is some natural ranking of types, in which higher types (of either side of the market) are

[5] In a different context, Hitch, Hortacsu, and Ariely (2006) use data form an online dating service to estimate market participants' preferences. Part of their analysis assumes the Gale-Shapley algorithm is followed.

preferable to all, then mis-aligned preferences imply different thresholds for each side of the market and matching sets can be characterized (see also Shimer and Smith, 2000).

1.3 Structure of the Paper

We start with a formal description of the economy we analyze in Section 2. We then describe the implementation of stable matches through a centralized clearinghouse in Section **??**. Section 4 illustrates the capacity of a decentralized market composed of fully informed agents to generate stable matches as unique equilibrium outcomes in weakly undominated strategies. We then shift to studying the impacts of incomplete information. Section 5 investigates the case of private information and illustrates the wedge between centralized and decentralized markets, as well as the effect of frictions. In Section 6 we generalize our analysis to unrestricted preferences. We conclude with Section 7. The proofs of all our propositions are relegated to the Appendix.

2. The Model

2.1 The Economy

A *market* is a triplet $\mathcal{M} = (\mathcal{F}, \mathcal{W}, U)$, where $\mathcal{F} = \{1, ..., F\}$ and $\mathcal{W} = \{1, ..., W\}$ are disjoint finite sets of firms and workers, respectively, and $U = \left\{\left\{u_{ij}^f\right\}, \left\{u_{ij}^w\right\}\right\}$ are agents' match utilities.[6] Each firm $i \in \mathcal{F}$ has match utility u_{ij}^f from matching to worker $j \in \mathcal{W} \cup \varnothing$, where matching to $\varnothing$ is interpreted as no match. Similarly, for each worker $j \in \mathcal{W}$, u_{ij}^w is the match utility from matching to firm $i \in \mathcal{F} \cup \varnothing$. For presentational simplicity we assume that firms and workers have strict preferences. That is, for any firm i, $u_{ij}^f \neq u_{ij'}^f$ for any $j, j' \in \mathcal{W} \cup \varnothing$ and for any worker j, $u_{ij}^w \neq u_{i'j}^w$ for any $i, i' \in \mathcal{F} \cup \varnothing$. We also consider strictly positive match utilities: for each firm i and for all $j \in \mathcal{W} \cup \varnothing$, $u_{ij}^f > 0$, and similarly for each worker j and for all $i \in \mathcal{F} \cup \varnothing$, $u_{ij}^w > 0$. A market $(\tilde{\mathcal{F}}, \tilde{\mathcal{W}}, \tilde{U})$ is a submarket of $(\mathcal{F}, \mathcal{W}, U)$ if $\tilde{\mathcal{F}}, \subseteq \mathcal{F}$, $\tilde{\mathcal{W}}, \subseteq \mathcal{W}$, and $\forall i, j \in \left(\tilde{\mathcal{F}} \cup \varnothing\right) \times \left(\tilde{\mathcal{W}} \cup \varnothing\right) \setminus \{\varnothing, \varnothing\}$, $\tilde{u}_{ij}^w = u_{ij}^w$ and $\tilde{u}_{ij}^f = u_{ij}^f$

For fixed sets $\mathcal{F}$ and $\mathcal{W}$ of firms and workers, an *economy* is a finite collection of markets $\{(\mathcal{F}, \mathcal{W}, U)\}_{U \in \mathcal{U}}$ together with a distribution G over possible utility levels $U \in \mathcal{U}$.

A *matching* is a function $\mu : \mathcal{F} \cup \mathcal{W} \to \mathcal{F} \cup \mathcal{W} \cup \varnothing$ such that for all $i \in \mathcal{F}$, $\mu(i) \in \mathcal{W} \cup \varnothing$ and for all $j \in \mathcal{W}$, $\mu(j) \in \mathcal{F} \cup \varnothing$. Furthermore, if $(i, j) \in \mathcal{F} \times \mathcal{W}$ then $\mu(i) = j$ if and only if $\mu(j) = i$.

[6] Cardinal utilities are required to trade off matchings at different points in time and examine the impacts of discounting.

If $\mu(k) \neq \varnothing$ for $k \in \mathcal{F} \cup \mathcal{W}$, we say that k is matched under μ. For any firm i, we call worker j *unacceptable* if $u^f_{i\varnothing} > u^f_{ij}$. Similarly, for any worker j, we call firm i *unacceptable* if $u^w_{\varnothing j} > u^w_{ij}$. A matching μ is blocked by an individual $k \in \mathcal{F} \cup \mathcal{W}$ if $\mu(k)$ is unacceptable (and we then say that μ is not individually rational). A *blocking pair* for a matching μ is a pair $(i, j) \in \mathcal{F} \times \mathcal{W}$ such that $u^f_{ij} > u^f_{i\mu(i)}$ and $u^w_{ij} > u^w_{\mu(j)j}$. A matching is *stable* if it is not blocked by any pair or individual.

Gale and Shapley (1962) showed that any market has a stable matching, and provided an algorithm that identifies one. In the *firm proposing deferred acceptance algorithm*, in step 1, each firm makes an offer to its most preferred worker. Workers collect offers, hold the offer from their most preferred acceptable firm and reject all other offers. In a general step k, firms whose offer got rejected in the last step make an offer to the most preferred acceptable worker who has not rejected them yet. Workers once more collect offers, including, possibly, an offer held from a previous step, keep their most preferred offer from an acceptable firm and reject all other offers. The algorithm ends when there are no more offers that are rejected, that is, any firm either has their offer held by a worker, or has been rejected by all its acceptable workers. Once the algorithm ends, held offers turn into matches.

The resulting match is the firm optimal stable matching, i.e., for any firm it is the stable matching that is not dominated by any other stable matching. It is in turn the least preferred stable matching for workers. Similarly, there always exists a worker optimal stable match, which is the least preferred by firms. In general, these two matchings can be different, and many other stable matchings can exist.

In this paper we do not want coordination to be the hurdle to the existence of strategies that yield a stable outcome as equilibrium in the market game (the details of which we soon describe). Therefore, we only consider markets $\mathcal{M} = (\mathcal{F}, \mathcal{W}, U)$ that have a unique stable matching denoted as μ_M. In those markets we have the following lemma (see, e.g., Roth and Sotomayor, 1990).

Lemma *If all the markets of the economy have a unique stable matching, then it is a Bayesian Nash equilibrium for all agents to submit their preferences truthfully to a mechanism that implements the firm proposing deferred acceptance algorithm.*

The theoretical literature has not yet identified general conditions on fundamentals (i.e., match utilities) that ensure a unique stable match. Some of our analysis will concentrate on a particular setting, that of *aligned preferences*, in which a unique stable matching is guaranteed and that has additional characteristics that can facilitate the existence of strategies that yield the stable outcome as equilibrium in decentralized markets.

2.2 ALIGNED PREFERENCES

Some of our analysis will pertain to a particular class of preferences that guarantee a unique stable matching.

Definition (Aligned Preferences) Firms and workers have *aligned preferences* if:

1. For any firm i, and workers j, j', if $u^f_{ij} > u^f_{ij'} > u^f_{i\varnothing}$, then $u^w_{ij} > u^w_{ij'}$ whenever $u^w_{ij} > u^w_{\varnothing j}$ or $u^w_{ij'} > u^w_{\varnothing j'}$; and
2. For any worker j, for firms i, i', if $u^w_{ij} > u^w_{i'j} > u^w_{\varnothing j}$, then $u^f_{ij} > u^f_{i'j}$ whenever $u^f_{ij} > u^w_{i\varnothing}$ or $u^f_{i'j} > u^f_{i'\varnothing}$.

Condition 1 means that if the match utility of a firm i with worker j is higher than with worker j', then the same has to be true of the match utilities of workers j and j' whenever either worker finds i acceptable (condition 2 is a mirror image).[7] For example, when firms and workers receive the same match utility, or share the match utility in fixed proportions, preferences are naturally aligned. While certainly restrictive, most applied papers implicitly assume that preferences are aligned (starting from Becker, 1973, and going to Sorensen, 2007, see references therein).

Preference alignment entails several important implications. It is easy to see that condition 1 implies that when firms make offers in the order of their preferences, the rejected offer of a worker cannot trigger a chain that results in an offer from a more desirable firm. That is, if a worker j rejects the offer from firm i after receiving the offer from firm k, then the resulting chain of offers can only result in offers to worker j that he prefers less than the offer from firm k. In particular, in a firm proposing deferred acceptance algorithm workers have no incentive to reject offers from acceptable firms in the hopes of receiving superior ones as a consequence of a chain of reactions in the market. We call this the **no cycle property.**[8]

[7] Note that if preferences are aligned and there are no unacceptable agents for any market participant then for any firm i, and workers j, j', $u^f_{ij} > u^f_{ij'} \Leftrightarrow u^w_{ij} > u^w_{ij'}$, and for any worker j, for any firms i, i': $u^w_{ij} > u^w_{i'j} \Leftrightarrow u^f_{ij} > u^f_{i'j}$.

[8] In general, this property does not hold. Workers may reject dominated offers as they may trigger chains that result in even more desirable offers.

The second implication is that for any submarket $(\tilde{\mathcal{F}}, \tilde{\mathcal{W}}, \tilde{U})$ either (1) there is an agent $a \in \tilde{\mathcal{F}} \cup \tilde{\mathcal{W}}$ such that firm a's preferred outcome is to exit the market, or (2) there is a firm i and a worker j that form a $top - top$ match, that is, worker j is firm $i's$ most preferred acceptable worker among $\tilde{\mathcal{W}}$ and firm i is worker $j's$ most preferred acceptable firm among $\tilde{\mathcal{F}}$.

Furthermore, aligned preferences ensure that each market admits a unique stable outcome (see also Sorensen, 2007, for a slightly restricted version of this proposition).

Proposition 1 (Alignment – Uniqueness) *When preferences are aligned, there is a unique stable matching.*

Intuitively, the unique stable matching can be identified through a recursive process. In the initial steps, firms and workers whose highest match utilities are to remain unmatched exit the market. This must constitute part of any stable, or simply individually rational, matching. Once there are no more firms and workers that rather exit the market, we find the firm-worker pair that constitutes a $top - top$ match. The corresponding firm and worker must be matched to each other for the match to be stable. Looking at the remaining firms and workers, we can continue recursively in this way. By construction, this procedure generates the unique stable matching.

When preferences are not aligned (but there is a unique stable matching μ_M), any non-stable matching μ' will either not be individually rational, or allow for a blocking pair $(i, j) \in \mathcal{F} \times \mathcal{W}$ such that $u^f_{i\mu_M(i)} \geq u^f_{i\mu'(i)}$ and $u^w_{\mu_M(j)j} \geq u^w_{\mu'(j)j}$. That is, both i and j prefer the outcome in the stable match μ_M to the outcome in μ' (see Gale and Sotomayor, 1985a and Demange, Gale and Sotomayor, 1987). However, in general, it need not be the case that $\mu_M(i) = j$, unless, as it turns out, when preferences are aligned. Formally,

Proposition 2 (Stable Blocking Pair) *When preferences are aligned, for any matching $\mu' \neq \mu_M$ the following holds: either μ' is not individually rational, or there exists a blocking pair (f, w) such that $\mu_M(f) = w$.*

The proof tracks that of Proposition 1. Indeed, going through the recursive process described above, at some stage a discrepancy must arise between μ_M and μ'. At that stage, either an agent who is unmatched under μ_M suddenly gets matched, in which case μ' is not individually rational, or a match that occur under μ_M does not get formed, in which case it can be shown that the corresponding worker and firm form a blocking pair.

2.2 A Decentralized Market

For a given economy $\{(\mathcal{F}, \mathcal{W}, U)\}_{U \in \mathcal{U}}$ together with a distribution G over utility realizations U, we analyze the following *market game*. The economy, together with the distribution G, is common knowledge to all agents. At the outset of the game, the market is realized according to the distribution G. Firms make offers over time, indexed by $t = 1, 2, ...$ and workers react to them. Specifically, each period has three stages. In the first stage, firms simultaneously decide whether and to whom to make an offer and whether to exit the market. In the second stage of any period, workers observe which firms exited, and observe only the offers they received themselves. Each worker j who has received an offer from firm i can accept, reject, or hold the offer. Once an offer is accepted, worker j is matched to firm i. Workers can also decide to exit the market. In the third stage, firms observe rejections and deferrals of their own offers. Finally, all participants are informed of the agents who exited the market and the matchings that occurred.

We will consider market games without frictions, and market games with frictions, which take the form of discounting. If a firm i is matched to worker j at time t, firm i receives $\delta^t u_{ij}^f$ and worker j receives $\delta^t u_{ij}^w$, where $\delta \in [0, 1]$ is the market discount factor. As long as agents are unmatched, they receive 0 in each period. One interpretation is that once a worker and a firm are matched, they receive their match utility, or, equivalently, they receive a constant, perpetual stream of payoffs, the present value of which is their match utility. Alternatively, one can assume that the market has a probability δ of collapsing each period, or that each firm has a probability δ of losing its position and receiving a payoff of 0 (and, analogously, each worker has a probability δ of leaving the market and receiving 0 as well).

A major source of obstacles to the decentralized market game yielding a stable outcome is the fact that time is valuable as described through discounting. In the conclusions we address other ways in which congestion and market frictions might arise (e.g., fixed costs for making offers).

In each period t, each firm i that has not yet hired a worker and has no offer held by a worker can make up to one offer to any worker that has not yet been matched.

We focus on two configurations of information in the economy. The simplest is that of *complete information* in which both firms' and workers' match utilities are common knowledge.[9] This is the

[9] The crucial assumption in the analysis of the complete information case is that each agent knows all other agents' preference *ordering*.

case that most of the literature on decentralized matching has tackled.[10] Note that in this case, both firms and workers can deduce the stable outcome μ_M.

The second information structure we analyze is that of *private information*. In that setup, each agent is informed only of her own match utilities. This setup still allows agents to be informed about correlations in preferences, as they are informed of the distribution over possible markets G. However, agents are not fully informed of the exact realization of match utilities and corresponding rank orderings. Note that the complete information case is a special case of private information in which there is only one possible market realization in the economy. In the case of private information, for an equilibrium of the market game to be a stable outcome, information has to be transmitted to allow agents to deduce their stable matching partner.

Each of these information structures defines a Bayesian game where type spaces correspond to the available private information of each agent. Our analysis concentrates on equilibria of decentralized market games in which all agents use weakly undominated strategies. For all information structures, this imposes several restrictions on equilibrium play:

1. A worker who accepts an offer, always accepts his best available offer. In particular, a worker cannot exit and simultaneously reject an offer that leads to a higher payoff.
2. When $\delta < 1$, a worker who knows he received an offer from his most preferred unmatched firm accepts it immediately and, similarly, a firm and a worker who know that their best possible remaining outcome is to remain unmatched, exit the market immediately.

Note that restriction 2 is due to the fact that all payoffs are strictly positive.

3. Centralized Matching

Before analyzing decentralized markets, we show that in the presence of a centralized clearinghouse (such as ones used by medical markets and many school districts), in our environment it is always possible to elicit preferences of agents in a way to induce the stable matching as an equilibrium outcome. To make analogies easier between centralized and decentralized markets, we will assume that agents report match utilities that are then translated to ordinal preferences. That is, each agent

[10] See Alcalde, Pérez-Castrillo, and Romero-Medina (1998), Haeringer and Wooders (2007), and Diamantoudi, Miyagawa and Xue (2007). Our analysis differs from that appearing in most of this work in that we allow both firms and workers to be strategic agents.

submits a vector of positive match utilities.[11] Technically, this is equivalent to having agents simply report ordinal preferences directly.

For each type of agent $a \in \{f, w\}$, and each agent l, let $P(u^a_l)$ be the strict ordinal preferences associated with l's reported match utilities, in which ties are broken depending on the index of the relevant match partners.[12] Specifically, consider firm i, then $u^f_{ij} > u^f_{ik}$ implies $jP(u^f_{i\cdot})k$, for $j, k \neq \varnothing, u^f_{ij} = u^f_{ik}$ and $j < k$ implies $jP(u^f_{i\cdot})k$ and for $j \neq \varnothing$, and $u^f_{ij} = u^f_{i\varnothing}$ implies $jP(u^f_{i\cdot})i$.

We define a *deferred acceptance mechanism* as a mechanism in which each agent l of type a reports his or her match utilities u^a_l. The mechanism then computes the corresponding ordinal preferences P and associates them with the stable match induced by the firm proposing deferred acceptance algorithm on P. The payoffs of firms and workers correspond to their match utilities given by their matching partner. The associated *deferred acceptance mechanism game* is the game in which agents simultaneously decide which match utilities to report to the mechanism after receiving their private information.

It follows directly from incentive compatibility attributes of the deferred acceptance algorithm that the deferred acceptance mechanism game allows for a Bayesian Nash equilibrium in weakly undominated strategies in which the resulting matching corresponds to the unique stable matching in each market of the economy (see Roth and Sotomayor, 1990). That is,

Lemma (Centralized Matching)

1. *For any economy with a unique market all the Nash equilibria in the deferred acceptance mechanism game in weakly undominated strategies yield the unique stable matching* μ_M.

2. *For any economy there exists a Bayesian Nash equilibrium in weakly undominated strategies in the deferred acceptance mechanism game such that the matching determined by the deferred acceptance mechanism yields the unique stable matching* μ_M *in each market.*

It is important to note that even though implementing the stable match is always possible through the centralized clearinghouse, the stable match is not necessarily the unique equilibrium outcome in the presence of uncertainty, as the following example illustrates.

[11] The restriction to positive numbers is made only for presentation simplicity. The sufficient restriction throughout our analysis is that the set of available reports contains as many elements as the maximal number of different match utilities each agent can experience in the economy.

[12] Note that eventhough agents never experience indifference in their realized match utilities, they may still report indifference, which is why the tie breaking specification is necessary.

Example 1 (Equilibria Yielding an Unstable Match) Consider a simple market consisting of firms $\{F1, F2, F3\}$ and workers $\{W1, W2, W3\}$, and assume that all agents always prefer to be matched rather than unmatched. Suppose the economy entails two markets, M_1 and M_2, occurring with equal probability. Their corresponding match utilities are as follows, where in each rubric (i, j) the first number corresponds to firm Fi's utility from matching to worker Wj,and the second number corresponds to worker Wj's utility from that match (and we highlight the stable matches in bold):

$U_1 =$

3, 2	1, 4	2, 2
$\underline{2, 4}$	**3, 2**	1, 4
3, 1	1, 1	**2, 1**

, $U_2 =$

2, 2	1, 4	**3, 2**
1, 4	**3, 2**	$\underline{2, 4}$
2, 1	1, 1	3, 1

.

Since it is a dominant strategy for firms to submit preferences truthfully, that is what they do in any equilibrium of the deferred acceptance mechanism game. Note that $U_1^w = U_2^w$, hence workers cannot distinguish between the two markets. We claim that the following strategies of workers also form an equilibrium. $W1$ reports $\hat{u}^w_{F1W1} = 1, \hat{u}^w_{F2W1} = 4, \hat{u}^w_{F3W1} = 3, \hat{u}^w_{\varnothing W1} = 2$, that is $W1$ reports only $F2$ and $F3$ as acceptable and ranks $F2$ above $F3$.Similarly, worker $W2$ reports $\hat{u}^w_{F1W2} = 4, \hat{u}^w_{F2W2} = 1, \hat{u}^w_{F3W2} = 3, \hat{u}^w_{\varnothing W2} = 2$, and worker $W3$ reports $\hat{u}^w_{F1W3} = 1, \hat{u}^w_{F2W3} = 4, \hat{u}^w_{F3W3} = 3, \hat{u}^w_{\varnothing W3} = 2$. Note that $W1$ and $W2$ exclude their stable match partner of M_1, and workers $W2$ and $W3$ exclude their stable match partner of M_2. What are the matchings μ_1 and μ_2 corresponding to the two markets given the submitted preferences? In market M_1 the outcome is $\mu_1(W1) = F2, \mu_1(W2) = F1$, $\mu_1(W3) = F3$ and in M_2 it is $\mu_2(W1) = F3$, $\mu_2(W2) = F1$, and $\mu_2(W3) = F2$ (the deviations from the stable match are underlined in the match utility matrices). The unique blocking pair of μ_1 is $(F1, W3)$. However, if $W3$ lists $F1$ as the second preferred firm, $W3$ would be matched to $F1$ in both M_1 and M_2.[13] Since $\frac{1}{2}u^w_{33} + \frac{1}{2}u^w_{23} = \frac{5}{2} > u^w_{13} = 2$, then such a deviation is not profitable for $W3$. Similarly, the unique blocking pair of μ_2 is $(F1, W1)$. However, if $W1$ lists $F1$ as his second choice (all other deviations clearly being unprofitable), then $W1$ would be matched to $F1$ not only in market M_2, but also in market M_1. Since $\frac{1}{2}u^w_{21} + \frac{1}{2}u^w_{31} = \frac{5}{2} > u^w_{11} = 2$, this deviation is not profitable either. Interestingly, in this equilibrium workers are better off than in the

[13]Note that in order for $(F1, W3)$ to block the outcome at M_1, worker $W3$ has to rank $F1$ higher than $F3$, and it cannot help $W3$ to rank $F1$ even higher than $F2$.

equilibrium achieving the stable match.

4. Complete Information

We start by analyzing economies in which all participants are informed of the realized markets. That is, there is complete information regarding all match utilities. When information is complete, all agents can compute the stable matching. Our first result illustrates that this implies that the stable match is an equilibrium outcome.

Proposition 3 (Complete Information) *For any economy in the market game:*

1. *There exists a Nash equilibrium in weakly undominated strategies that generates the stable matching.*
2. *When preferences are aligned, any Nash equilibrium in weakly undominated strategies yields the stable matching.*

To glean some intuition for existence of a Nash equilibrium that yields the stable outcome, consider strategies where each firm makes an offer to its stable matching partner, and exits the market if the firm is unmatched under the unique stable match μ. Each worker accepts his best available offer from an acceptable firm in period 1, and if he receives no (acceptable) offers, exits the market. Note that this construction hinges on the fact that all agents are completely informed of the realized market, and hence each firm can compute the stable matching.

Nonetheless, there may still be Nash equilibria in weakly undominated strategies whose outcome is not the stable match. We first provide a specific example and later show that the construction can be generalized.

Example 2 (Multiplicity with Complete Information) Consider a simple market consisting of firms $\{F1, F2, F3\}$ and workers $\{W1, W2, W3\}$ in which all agents prefer to be matched rather than unmatched. Suppose match utilities are the following (using notation as before, where the stable match is indicated in bold):

$$U = \begin{array}{|c|c|c|}\hline \mathbf{2,3} & \underline{3,2} & 1,3 \\ \hline \underline{3,1} & \mathbf{2,3} & 1,1 \\ \hline 3,2 & 2,1 & \mathbf{1,2} \\ \hline \end{array}.$$

The unique stable matching μ_M is such that Wi is matched with Fi. Note that firm $F1$ and $F2$ strictly prefer matching $\tilde{\mu}$ in which they exchange their stable matching worker, and which leaves $F3$'s match unchanged (the swap is indicated with underlined entries in U). The matching $\tilde{\mu}$ is blocked by firm $F3$ and worker $W1$. Note that within the submarket M' consisting of $\{F1, F2\}$ and $\{W1, W2\}$ there are two stable matchings, where the matching μ' corresponds to the worker optimal stable match, while the matching $\tilde{\mu}$ implements the firm optimal stable match on M'.The weakly undominated strategies below yield $\tilde{\mu}$:

Firm $i \in \{F1, F2\}$ makes no offer in period 1. If $F3$ is matched to $W3$ in period 1, then from period 2 on, firm i makes an offer to $\tilde{\mu}(i)$ if that worker is unmatched, and otherwise to the sole remaining worker. If $F3$ is not matched to $W3$ in period 1, then firm i makes an offer to its stable match partner given the remaining set of firms and workers. Note that this implies that if no firm is matched after period 1, firm i makes an offer to $\mu_M(i)$ as μ_M is the unique stable matching of the whole market.

Firm 3 makes an offer to $\mu(F3) = W3$ in period 1, and in general makes an offer to its stable match partner, given the remaining set of firms and workers.

Workers accept an offer from their most preferred unmatched firm, and furthermore:

Worker $j \in \{W1, W2\}$ accepts an offer from $\tilde{\mu}(j)$ in period 2 if $F3$ is matched to $W3$ in period 1, and $\tilde{\mu}(j)$ is its only offer. Otherwise, j accepts an offer from its stable match partner given the remaining set of firms and workers starting in period 2.

Worker 3 accepts an offer from its stable match partner given the remaining set of firms and workers as soon as it receives such an offer.

It is straightforward to see that these strategies constitute an equilibrium implementing the matching $\tilde{\mu}$. Note that the total payoffs in $\tilde{\mu}$ are lower than in μ_M, however, when we change the highest payoffs of firms to 4 instead of 3, then the total payoffs in $\tilde{\mu}$ would be higher than in μ_M. This once more shows that there is no straightforward connection between stability, which is an ordinal concept, and total welfare. Note that the preferences of firms and workers are not aligned and, more generally, there is no *top-top* match among the whole set of agents.

For any realized market $\mathcal{M} = (\mathcal{F}, \mathcal{W}, U)$ with a unique stable matching μ_M, if $\mathcal{F}' \subseteq \mathcal{F}$ and $\mathcal{W}' \subseteq \mathcal{W}$ are such that $\mu_M(F') = W'$, then we call $\mathcal{M}' = (\mathcal{F}', \mathcal{W}', U')$, where U' is the restriction of U to $\mathcal{F}'$ and $\mathcal{W}'$, a μ_M*−induced submarket*. The example illustrates that even if $\mathcal{M}$ has a unique stable matching some μ_M-induced submarket $\mathcal{M}'$ may have several stable matchings. In fact, the following proposition provides the generalization of the example above:

Proposition 4 (Complete Information – Multiplicity) *Let $\mathcal{M} = (\mathcal{F}, \mathcal{W}, U)$ be a market game with a unique stable matching μ_M. Suppose there exists a μ_M-induced submarket $\mathcal{M}' = (\mathcal{F}', \mathcal{W}', U')$, with matching μ' that is not the firm-optimal stable matching. Then, for sufficiently high δ, there are Nash equilibria in weakly undominated strategies that implement an unstable matching.*

As in the example, the strategies that are used to implement an unstable matching require part of the market to match in period 1 (according to μ_M), and the remaining μ_M-induced submarket, that entails multiple stable matches, to match in period 2. The strategies rely on some firms *conditional* offers in period 2. Their offers are such that they implement the alternative stable match in the submarket only if the appropriate firms match in period 1. This highlights one stark contrast between centralized and decentralized markets. Indeed, if agents were to participate in a centralized clearinghouse that uses a firm proposing deferred acceptance algorithm they would not be able to implement such conditional strategies.

For markets with aligned preferences, when μ_M is the unique stable match, any μ_M-induced submarket has one stable match coinciding with that induced by μ_M. In particular, the condition of Proposition 4 is never satisfied, and from the first part of Proposition 3, all equilibria in weakly undominated strategies yield the unique stable match μ, illustrating the second part of Proposition 3.[14] That is, a non-cooperative market game results necessarily in a matching that is equivalent to the unique stable matching (the unique core outcome in this market). It is important to note that there may still be multiple equilibria generating it. Indeed, one way to establish the stable match is through an instantaneous match in period 1 as described above. An alternative way involves emulating the deferred acceptance algorithm. Since different equilibria may take different times to produce a full match, equilibrium *payoffs* are not necessarily unique even when preferences are aligned.

[14] When firms and workers can use weakly dominated strategies, there are many more equilibria. For instance, it is an equilibrium for all agents to exit the market in period 1, resulting in no matches. Weak dominance rules this out, as it does not allow a worker to exit the market when he has an acceptable offer in hand.

Note that, with or without aligned preferences, restricting the market game to end after one period (and using weakly undominated strategies) would assure a unique equilibrium outcome coinciding with the unique stable match for any realized market. This results from the fact that any other matching μ' has either a blocking individual or a blocking pair.[15] Alternatively, restricting firms to make their offers according to their preference list, or, more generally, restricting them to strategies that emulate their actions in the firm proposing deferred acceptance algorithm (as some of the literature has assumed, e.g., Haeringer and Wooders, 2007) would also assure the stable match is created as a unique equilibrium outcome.

5. Economies with Uncertainty – Aligned Preferences

We now turn to the analysis of economies entailing uncertainty regarding the realized market and, consequently, regarding the relevant stable match. For a decentralized market to reach a stable outcome, sufficient information has to be transmitted through the market to ensure that (i) workers only accept offers from firms that are their stable match, which means sufficient information has to be transmitted for workers to identify those firms, and (ii) firms make offers to those workers. Furthermore, the decentralized market has to allow for this information to be transmitted in an incentive compatible way.

There are basically three channels through which information flows in the market game. First, information is publicly transmitted when agents exit the market or form a match. Second, information is privately transmitted when workers receive offers from firms and workers respond to those offers (unless offers are accepted, in which case that information becomes public). The third component of information is time – all participants are aware of the period they are in.

We start by considering economies with aligned preferences. In such economies, there exists a *top – top* match for any subset of firms and workers (or an agent who prefers to be unmatched). In particular, if firms and workers follow strategies that resemble the deferred acceptance algorithm (namely, firms make offers to workers in order of their preferences and workers hold on only to their best available offer and accept an offer from their most preferred available firm), then in every period some agents would be matched or exit the market, and learning with time takes a rather simple form.

Intuitively, when there is sufficient uncertainty on the part of the firms, skipping workers, or

[15] In detail, suppose $\mu' \neq \mu_M$ were implemented. If μ' were not individuall rational, agents would exit the market rather than accept an unacceptable partner. Otherwise, the firm f of a blocking pair (f, w) would make an offer to its worker w rather than to $\mu'(f)$ given that w would accept f over $\mu'(w)$.

making offers out of order, can end up in a match to an inferior worker. On the worker's side, when preferences are aligned, the no cycle property assures that rejecting an offer cannot launch a cycle in the market the yields a superior offer. In that sense, the main hurdle for establishing stability with aligned preferences is that normal deferred acceptance strategies may not be incentive compatible since firms and workers may have incentives to deviate in order to *speed up* their matches. In particular, time discounting will play a crucial role in our analysis.

5.1 Learning in a Decentralized Market

Before investigating strategies of firms and workers, we describe the information agents have at each period t which they can and have to use to update their beliefs about potential stable match partners.

Let $M_t \subseteq (\mathcal{F} \cup \varnothing) \times (\mathcal{W} \cup \varnothing)$ denote the matches formed at time t (including firms and workers who leave the market by themselves), let $M_t^F \subseteq \mathcal{F} \times \varnothing$ be the set of firms that leave the market in the beginning of period t, and let the set of agents who exited the market up to, but excluding, period t be

$$\mathcal{E}^t \equiv \{j \mid \exists k \text{ s.t. } (j,k) \in M_\tau \text{ for some } \tau < t\} \cup \{i \mid \exists l \text{ s.t. } (l,i) \in M_\tau \text{ for some } \tau < t\}.$$

At the beginning of period t, each active firm i observes a history that consists of the (timed) offers the firm made, the responses of workers to those offers, denoted by r for rejection and h for holding (where we use the notational convention that an offer to no worker is denoted as an offer to $\varnothing$ that is immediately rejected), and the (timed) set of agents that have left the market:

$$h_{t,i}^f \in ((\mathcal{W} \cup \varnothing) \times \{r,h\})^{t-1} \times \prod_{\tau=0}^{t-1} M_\tau,$$

In addition, at each period t, suppose workers $j_1, ...j_{k(t-1)}$ rejected offers from firm i in periods $1, ..., t-1$. Denote by $\tilde{W}_i^t = \left\{j | j \notin \left\{j_1, ...j_{k(t-1)}\right\}\right\}$ the set of workers that have not rejected firm i yet.

Each unmatched worker acts in the interim stage of each period, and at period t observes a history that consists of all (timed) offers he received and a (timed) sequence of offers he held, the (timed) set of agents that have left the market, and the set of firms that left the market at time t:

$$h_{t,j}^{w} \in \left(2^{\mathcal{F}}\right)^{t} \times \left(2^{\mathcal{F}}\right)^{t} \times \prod_{\tau=0}^{t-1} M_{\tau} \times M_{t}^{F}$$

In addition, at each period t, suppose firms $i_1, ...i_{k(t)}$ made offers to worker j in periods $1, ..., t$. Denote by $\tilde{F}_j^t = \left\{i | u_{ij}^w \geqslant \max\left\{u_{lj}^w | l = i_1, ...i_{k(t)}\right\}\right\}$ the set of firms that have not made an offer to worker j yet and that he weakly prefers to any firm that has made him an offer thus far.

For a given prior distribution G over possible utility levels $U \in \mathcal{U}$, for any private information $\mathcal{I}$ regarding the realized market, let $G(\mathcal{I})$ denote the posterior distribution over utility realizations. For each type of agent $a \in \{f, w\}$ and each agent l, let $\mathcal{S}_l^a(\mathcal{I}) = \{\mu(U)(l) \mid U \in \text{supp } G(\mathcal{I})\}$. That is, $\mathcal{S}_i^f(\mathcal{I})$ and $\mathcal{S}_j^w(\mathcal{I})$ denote the set of all potential stable match partners that could conceivably constitute part of a stable match under the updated distribution over market match utilities, for firm i and worker j respectively.

5.2 Frictionless Economies

Suppose there is no discounting, i.e., $\delta = 1$. Then, one way in which information may be transmitted in the market is if agents simply follow **deferred acceptance strategies**. That is, firms make offers to workers according to their match utilities, and only exit the market when all acceptable workers had rejected them; Workers hold on to their best available offer, and accept an offer only once the offer is from the firm that yields the highest match utility given the set of firms that are still unmatched. These prescriptions are followed by all after any detectable deviations as well.

Proposition 5 (Aligned Preferences – No Discounting) *Suppose preferences are aligned, and $\delta = 1$. Deferred acceptance strategies constitute a Bayesian Nash equilibrium in weakly undominated strategies and yield the stable matching.*

Intuitively, when all agents use deferred acceptance strategies, workers ultimately hold offers from their stable match partners. Since preferences are aligned, every period has either an agent exiting, or a top-top match that is formed. In particular, the process stops in finite time. When $\delta = 1$, the timing of matches is of less importance and, from alignment, deviations (say, the rejection of the best available by a worker) cannot generate a better match.

5.3 ECONOMIES WITH FRICTIONS

In the absence of frictions, when $\delta = 1$, the stable match can be achieved in equilibrium when all market participants emulate the firm proposing deferred acceptance algorithm. This strategy profile, however, turns out not to be incentive compatible in a decentralized market in which there is discounting, i.e., when $\delta < 1$. Consider a complete information economy with two workers and two firms, who all prefer to be matched rather than unmatched, and for which $u_{ij}^f = u_{ij}^w$ for all i, j. Match utilities are given as follows (rubric (i,j) corresponds to utility from the match of firm i with worker j).

$$U_1 = \begin{array}{|c|c|} \hline \mathbf{4} & 1 \\ \hline 3 & \mathbf{2} \\ \hline \end{array}.$$

Firm 2 knows that worker 2 is the unique stable match partner and that, furthermore, worker 2 would accept an offer from firm 2 immediately, as firm 2 is worker $2's$ first choice. Hence, it cannot be an equilibrium for firm 2 to first make an offer to worker 1, and lose a period. Firms may therefore be tempted not to make all offers in order of their preferences, but rather concentrate on offers to potential stable match partners.

In order to identify circumstances in which the stable match can be implemented in equilibrium, we start by analyzing some minimal conditions strategies have to satisfy in a centralized mechanism in order to generate the stable match (this is a generalization of Proposition 1, which illustrated the existence of such a profile). We then characterize economies in which emulating strategies from that class is incentive compatible for all participants in the decentralized market game.

Definition (Reduced Deferred Acceptance Strategies) Let $\mathcal{E}$ be an economy, and assume agents participate in a centralized match. Agents use **reduced deferred acceptance strategies**, if

1. Every firm i submits utilities v corresponding to preferences $P(v)$ that satisfy the following. Let $P(v,i) = \{l : lP(v)i\} \cup \{i\}$, then:

   (a) $\mathcal{S}_i^f\left(u_{i\cdot}^f\right) \subseteq P(v,i)$.

   (b) For each for each $k, l \in \mathcal{S}_i^f\left(u_{i\cdot}^f\right)$, $kP(v)l \Leftrightarrow u_{ik}^f > u_{il}^f$.

   (c) For each $k \in P(v,i)\backslash\mathcal{S}_i^f\left(u_{i\cdot}^f\right)$ and each $l \in \mathcal{S}_i^f\left(u_{i\cdot}^f\right)$, $kP(v)l \Rightarrow u_{ik}^f > u_{il}^f$.

2. Every worker j submits utilities v corresponding to preferences $P(v)$ that satisfy the following. Let $P(v,j) = \{l : lP(v)j\} \cup \{j\}$, then:

   (a) $\mathcal{S}_j^w \left(u_{\cdot j}^w\right) \subseteq P(v,j)$.

   (b) For each for each $k, l \in \mathcal{S}_j^w \left(u_{\cdot j}^w\right)$, $kP(v)l \Leftrightarrow u_{jk}^w > u_{jl}^w$.

   (c) For each $k \in P(v,j) \backslash \mathcal{S}_j^w \left(u_{\cdot j}^w\right)$ and each $l \in \mathcal{S}_j^w \left(u_{\cdot j}^w\right)$, $kP(v)l \Rightarrow u_{jk}^w > u_{jl}^w$.

Reduced deferred acceptance strategies require that each agent submit preferences that satisfy three conditions. First, any potential stable match (provided the ex-ante information) are declared acceptable. Second, potential stable matches are ranked truthfully.[16] Third, rankings of agents who are not potential stable matches above potential stable matches must be truthful.

Certainly, if there is any hope to achieve the stable match for any market realization, agents must declare potential stable matches acceptable. In order to glean some intuition on the second and third conditions, consider, e.g., the firms. Since the centralized mechanism achieves the firm preferred stable match for the submitted preferences, changing the ranking of agents that are preferred to the stable match would not change the resulting match in the centralized market, nor ranking some of those agents below the stable match. However, any other deviation from the actual ranking may generate an unstable match. Assuring that agents ranked above *any* potential stable match are, in fact, preferred to that stable match boils down to the second and third conditions reduced deferred acceptance strategies impose.

We now show the sense in which reduced deferred acceptance strategies place minimal requirements for establishing stability in centralized markets. Note that the conditions identifying reduced deferred acceptance strategies pertained only to the *support* of the stable matches, and did not depend on the precise likelihood a potential stable match would realize. Indeed, we consider economies with a finite number of potential markets, and so strategy profiles that generate the stable match should be robust to how these markets are distributed.

Consequently, we say an agent l uses a **reduced deferred acceptance rule** if, for any economy $\mathcal{E}$ containing agent l, the agent uses a reduced deferred acceptance strategy that depends only on the set of market participants, their realized match utilities, and the set of potential stable matches.

[16] Note that if $f_i \in \mathcal{S}_i^f \left(u_{i\cdot}^f\right)$, then f_i has to be the least preferred element in $\mathcal{S}_i^f \left(u_{i\cdot}^f\right)$ of firm i, since a stable matching is always individually rational. Similarly for workers.

In other words, the agent always uses a deferred acceptance strategy that does not depend on the particular markets underlying the economy, only on the stable matches that can be induced in these markets.

Proposition 6 (Centralized Aligned Economies with Discounting)

1. *If all agents use a reduced deferred acceptance rule then for any economy $\mathcal{E}$ the outcome in the associated deferred acceptance game is the stable match.*

2. *Suppose there is an agent $a \in \mathcal{F} \cup \mathcal{W}$ who does not use a reduced deferred acceptance rule. Then there exists an economy $\mathcal{E}$ such that in the associated deferred acceptance game there is a market realization in which the outcome is not stable.*

3. *Given an economy $\mathcal{E}$, all agents using deferred acceptance strategies constitutes a Bayesian Nash equilibrium in the centralized deferred acceptance game.*

Proposition 6 illustrates the effectiveness of reduced deferred acceptance rules in generating the stable match outcome. Part 2 of the Proposition highlights the necessity of the conditions imposed by reduced deferred acceptance strategies for implementing stable outcomes in any economy.

When moving from a centralized mechanism to a decentralized market, we can translate reduced deferred acceptance strategies to decentralized market strategies. For firms the translation is straightforward. Whenever firm i would submit a utility profile v, firm i does the following in the decentralized market. In each period in which firm i is not matched and does not have an offer held by a worker, firm i makes an offer to its most preferred worker who has not rejected firm i yet and who is still unmatched according to v, where ties are broken according to the same rules determining $P(v)$ in the centralized setting. When firm i gets rejected by the last acceptable worker, that is the last worker who is still unmatched and has a higher utility than firm i itself according to v, firm i exits the market.

For workers, if in a centralized market worker j submits a utility profile v, then under the associated decentralized reduced deferred acceptance strategy, (1) worker j holds the best offer that is acceptable given v and rejects all other offers, and (2) worker j accepts an offer whenever it is an offer from the highest potential stable match partner who is still unmatched.

As discussed in Section 5.1, firms and workers may learn over time and refine their set of potential stable match partners. **Decentralized reduced deferred acceptance strategies** are ones that can be derived as above when considering the translation period by period, thereby allowing agents to submit reduced deferred acceptance strategies accounting for the information they have accumulated. Put another way, decentralized reduced deferred acceptance strategies are ones in which firms and workers formulate a ranking at each period that dictates their behavior. Namely, given the strategy profile of all other agents, at each period each agent can calculate the set of potential stable matches. Each firm then is required to never make an offer to anyone who is less preferred than her most preferred potential stable match. Each worker who receives offers that are weakly preferred to his most preferred potential stable match must hold one of those offers. Furthermore, upon receiving offers only from potential stable matches the worker must hold on to the best of those (and immediately accepts an offer from his most preferred potential stable match). Given these de-facto rankings, the agents all emulate the deferred acceptance algorithm.[17]

While reduced deferred acceptance strategies impose minimal restrictions on strategies submitted to a centralized mechanism that achieve stability, their decentralized counterpart is generally not incentive compatible in markets with frictions. In such markets, dynamic considerations alter incentives through two channels. First, agents care about *when* matches are created. Second, agents update the set of their potential matches using the history of play.[18]

Proposition 6 above suggested that strategy profiles that yield stable matches in general must be of the form of decentralized reduced deferred acceptance strategies. Our goal now is to characterize the class of economies for which these strategies are incentive compatible.

We start with motives that have to do with speeding up the matching process. When all agents in the market follow reduced deferred acceptance strategies, firms may at times be able to speed the process by altering the ranking of agents. Intuitively, suppose that all other players use strategies that implement the stable outcome. There are economies in which a firm's offer to a worker j who is not its first choice worker will be accepted only if that worker is actually its stable match partner. Then the firm may have incentive to make that offer first, in order to speed up the timing of its match, as such "out of order" offers entail no little risk of "wrong acceptance" if all other agents use strategies that implement the stable outcome. However, in such cases, other firms then in turn have

[17] As time progresses, learning can only lead to the elimination of previously perceived potential stable matches. In particular, the conditions imposed by reduced deferred acceptance strategies are weakened in the interim stages relative to the ex-ante stage dealt with in the centralized market.

[18] The learning that takes place during play may allow agents to eliminate certain agents from their ex-ante set of potential stable matches. In that respect, the conditions of reduced deferred acceptance strategies are weakened in interim stages.

a strict incentive to manipulate their offers, resulting in situations in which there are no strategies that always implement the stable outcome with certainty.[19] The following provides a simple example of such an economy.

Example 3 (Timing of Matches) Consider an economy with two firms $\{F1, F2\}$ and two workers $\{W1, W2\}$,who all prefer to be matched rather then unmatched, in which $u^f_{ij} = u^w_{ij}$ for each of 6 potential markets, described by the following match utilities (notation as before):

$$U_1 = \begin{array}{|c|c|}\hline \mathbf{3} & 6 \\ \hline 4 & \mathbf{7} \\ \hline \end{array},\ U_2 = \begin{array}{|c|c|}\hline 3 & \mathbf{6} \\ \hline \mathbf{4} & 5 \\ \hline \end{array},\ U_3 = \begin{array}{|c|c|}\hline \mathbf{3} & 2 \\ \hline 4 & \mathbf{8} \\ \hline \end{array},\ U_4 = \begin{array}{|c|c|}\hline \mathbf{3} & 2 \\ \hline 1 & \mathbf{7} \\ \hline \end{array},\ U_5 = \begin{array}{|c|c|}\hline \mathbf{9} & 6 \\ \hline 8 & \mathbf{5} \\ \hline \end{array},\ U_6 = \begin{array}{|c|c|}\hline 7 & \mathbf{3} \\ \hline \mathbf{8} & 5 \\ \hline \end{array}.$$

U_3 and U_4 guarantee that $F1$ makes an offer to $W1$ in any equilibrium when $W1'$s match utilities are $(3, 4)$.[20] Similarly, U_5 and U_6 guarantee that $W2$ with match utilities $(6, 5)$ will in equilibrium sometimes receive no offer in period 1, but only in period 2.[21]

We are now ready to show that there are no equilibria in weakly undominated strategies that always implement the stable outcome. From now on we focus on U_1 and U_2.[22]

Note that $W1$ and $W2$ always accept an offer from $F2$ immediately in U_1. Hence for any $\delta < 1$, $F2$ must make an offer to $W2$ when U_1 is realized. Suppose given $F1$'s match utilities $(3, 6)$, the probability of U_1 is p and that of U_2 is $1 - p$.

[19] Note that workers cannot speed things up by holding on to offers that are not potential stable match partners in the decentralized market. However, when preferences are aligned, holding such offers is not strictly harmful to the worker. Indeed, the no cycle property assures that rejection of an offer cannot trigger a chain in the market leading to a more preferred offer. It follows that this incentives issue is more relevant for firms than for workers in the case of aligned preferences.

[20] Note that in U_4: $W1$ accepts an offer from $F1$ immediately, therefore, $F1$ has an incentive to make an offer to $W1$ in period 1 whenever U_4 is realized. However, $F1$ cannot distinguish between U_3 and U_4, so what are possible consequences of an offer to $W1$ in U_3? Given an offer from $F1$, $W1$ canot exit simultaneously, he can at best reject the offer from $F1$. Note furthermore, that in U_3 it must be the case that $F2$ makes an offer to $W2$ in period 1 who accepts that offer. Therefore, in period 2, if $W1$ decided to reject the offer of $F1$, then $F1$ can remake that offer, in which case $W1$ has to acept it, as we impose participants to use weakly undominated strategies. Hence, $F1$ can only gain from making an offer to $W1$ when $F1's$ match utilities are $(3, 2)$ compared to not making an offer, or making an offer to $W2$.

[21] In U_5, $F1$ will make an offer to $W1$ who will accept immediately. In U_6, to guarantee a stable outcome, $F2$ with utilities $(8, 5)$ cannot make an offer to $W2$ in period 1, and hence has to make an offer to $W1$ in period 1. This implies that in U_5, $W2$ will in equilibrium not receive any offers in period 1. Therefore, $W2$ cannot exit the market when he receives no offer in period 1, as in period 2, along the equilibrium path, $F2$ makes an offer to $W2$ in period 2 when U_5 is realized.

[22] That is, $F1$ observes $(3, 6)$ and $W1$ observes $(3, 4)$, the only two agents that cannot distinguish these two markets. We keep in mind though that in equilibrium, $W1$ will observe an offer from $F1$ (because of U_3 that $W1$ cannot distinguish from U_1 and U_2), and that $W2$ with match utilities $(6, 5)$ will not always receive an offer in period 1 (because of U_5 that $W2$ cannot distinguish from U_2).

Suppose $F1$ **makes an offer to** $W2$ **(in** U_1 **and** U_2**)** in period 1. Then when U_2 prevails, $F2$, that is aware U_2 is realized, makes an offer to $W1$ in period 1, who will accept that offer. Furthermore, $F1$ makes an offer to $W1$ with match utilities of $(3,4)$ only in U_3, when $F1$ is the stable match, hence $W1$ accepts an offer from $F1$ whenever $W1's$ match utilities are $(3,4)$ (and that is the only offer he observes). These strategies generate a payoff for $F1$ of $6(1-p)+3p\delta$.

Consider the following deviation: $F1$ makes an offer to $W1$ in period 1. $W1$ accepts that offer (see above) in U_1. In U_2 the offer is rejected, and $F1$ matches to $W2$ in period 2 (as $W2$ does not leave the market in period 1 when observing match utilities $(6,5)$, see above), resulting in payoffs $6(1-p)\delta+3p$. This deviation is profitable, when $p > 2/3$ (independent of δ). The idea is that $F1$ can assure that when approaching $W1$ in period 1, its offer gets accepted only when $W1$ is the stable match. The effect of such a deviation is therefore to *speed up* the creation of its match when U_1 is realized. The cost is the delay of a match with $W2$ in U_2. However, when U_1 is sufficiently more likely ex-ante (given $F1$'s private information), the benefits outweigh the costs.

Suppose $F1$ **makes an offer to** $W1$ **with probability** $q \in (0,1]$ **(in** U_1 **and** U_2**)** in period 1. First, note that this implies that $W1$ has to accept the offer from $F1$ with positive probability, $m \in (0,1]$.[23] In order for the market to always yield a stable outcome, it has to be the case that $F2$ makes an offer to $W1$ with probability 1 in period 1 when U_2 prevails, which implies that once more $W1$ has to accept the offer from $F1$ whenever he receives that offer. Now, $F1$ has the same trade-off as before, and hence, with the same argument, $F1$ strictly prefers making an offer to $W1$, that is, $q = 1$. Can we induce $F2$ to make an offer to $W1$ with certainty? When U_2 prevails, an offer to $W1$ yields 4. An offer to $W2$ yields 5δ which is bigger than 4 for $\delta > 4/5$. Hence, for large δ, $F1$ making an offer to $W1$ with positive probability cannot be part of an equilibrium.

Suppose $F1$ simply delays making an offer and makes an offer to its most preferred available worker in period 2. This clearly cannot be part of an equilibrium since $F1$ can profitably deviate to making an offer to $W2$ in period 1, which will be accepted with probability $1-p$.

[23] Suppose $W1$ accepts $F1$ with probability 0 in period 1. Then $F1$'s payoff from making an offer to $W1$ in period 1 is $3p\delta + 6(1-p)\delta$, an offer to $W2$ yields however $3p\delta + 6(1-p)$. Hence, there does not exist an equilibrium such that $W1$ accpets an offer from $F1$ with probability 0, and $F1$ makes an offer to $W1$ with positive probability.

Assumption 1 makes sure that when all market participants follow the decentralized reduced deferred acceptance strategies, for any firm making an offer to a worker who is ranked below the favorite potential match partner among workers that have not rejected that firm yet runs the risk of that offer being held or accepted. This is because such offers have a chance of being eventually accepted, as they are better than any stable match outcome.

Assumption 1 *For any market realization, the corresponding match utilities U satisfy that, if all agents use decentralized reduced deferred acceptance strategies, then for any firm i, at each period t, for all available workers j that are worst than firm i's most preferred potential stable match, $\exists \tilde{U} \in$ supp $G\left(u^f_{i\cdot}, h^f_{t,i}\right)$ such that $\tilde{u}^w_{ij} > \tilde{u}^w_{\mu(\tilde{U})(j)j}$.*

If firms and workers use decentralized reduced deferred acceptance strategies, then Assumption 1 implies that firms do not have an incentive to make offers to workers who are not at least as valuable as their most preferred stable match.

However, firms may not only alter when they make an offer to each worker, they can also decide to completely scratch some workers from their preference list. Specifically, if firms do not remake offers that were rejected, and workers always hold the best available acceptable offer (unless they accept an offer), then the set of potential stable matches after some history is given by:

$$\mathcal{S}^f_i\left(u^f_{i\cdot}, h^f_{t,i}\right) \subseteq \left\{j | j \in \mathcal{S}^f_i\left(u^f_{i\cdot}\right) \cap \tilde{W}^t_i\right\} \backslash \mathcal{E}^t,\ \mathcal{S}^w_j\left(u^w_{\cdot j}, h^w_{t,j}\right) \subseteq \left\{i | i \in \mathcal{S}^w_j\left(u^w_{\cdot j}\right) \cap \tilde{F}^t_i\right\} \setminus \left(\mathcal{E}^t \cup M^F_t\right).$$

Note that Example 3 points to a case in which the inclusion is strict. Are there other instances in which the inclusion can be strict? Indeed, and we show that in such environments there may not exist strategies that always yield the stable outcome.

Example 4 (Manipulability of Offer Timing) Consider an economy with three firms $\{F1, F2, F3\}$ and three workers $\{W1, W2, W3\}$,who all prefer to be matched rather then unmatched, in which $u^f_{ij} = u^w_{ij}$ for each of 4 potential markets, described by the following match utilities:

$U_1 =$

5	4	3
7	**10**	2
6	8	**9**

, $U_2 =$

5	9	**4**
7	**10**	3
6	8	2

, $U_3 =$

5	4	3
9	7	**10**
6	**8**	2

, $U_4 =$

5	4	3
2	6	**9**
1	**8**	7

We look for weakly undominated strategies that in equilibrium always implement the stable match.

U_3 guarantees that in such an equilibrium $F3$ with match utilities $(6, 8, 2)$ always makes an offer to $W2$ in period 1.[24] Furthermore, U_4 guarantees that in any such equilibrium $F1$ with match utilities $(5, 4, 3)$ makes an offer to $W1$ with probability 1.[25] From now on we concentrate on U_1 and U_2.

Can there be an equilibrium in which F1 makes an offer to W3 in U_2 in period 1 with certainty? In this case, $W1$ receives an offer from $F1$ only when $F1$ is the stable match partner (as in U_2, $F3$ makes an offer to $W2$ in period 1), so in equilibrium $W1$ will accept an offer from $F1$ in period 1, if it is the only offer he receives. This of course provides strict incentives for $F1$ to make an offer to $W1$ even in U_2, resulting in an outcome that is not stable.

Can there be an equilibrium in which F1 makes an offer to W1 in $\mathbf{U_2}$ with probability $\mathbf{q} \in (\mathbf{0}, \mathbf{1}]$? When $F1$ makes an offer to $W3$, $F1$ receives a payoff of 4, as $W3$ accepts immediately. In order for the market to always yield a stable outcome, it has to be the case that $W1$ never accepts an offer from $F1$ in period 1 when he receives only that offer. Therefore, the expected payoff if $F1$ to make an offer to $W3$ is $4\delta < 4$, hence such an equilibrium does not exist.

The example suggests that in order to achieve stability, it must be the case that workers cannot infer what is the underlying state from the order (and timing) of offers. This imposes restrictions on the correlation between preferences, specifically on how agents can update their beliefs about the set of stable match partner. Specifically, the above examples illustrate the potential for manipulation of offers when some information is transmitted by the mere timing of offer acceptance or rejections. These transfers of information need to be restricted to allow for equilibria that yield the stable matching, as captured by Assumption 2.

Before presenting Assumption 2, we note one delicate point regarding workers learning in the economy. Suppose there are only 2 firms (left) in the market, and any worker who may receive an

[24] Such an offer will be immediately accepted in U_3. Furthermore, if $W2$ makes an offer to $W1$ with some probability $p > 0$ in period 1, then in U_3, $W1$ is aware that the stable match partner is $F1$ yelding match utility of 5, so $W1$ will accept that offer, yielding an unstable outcome.

[25] In U_4, $W1$ accepts an offer from $F1$ immediately. Note that in U_3 $F2$ matches with $W3$ and $F3$ with $W2$ in period 1 (see above). Hence, in U_3 in period 2, at most $W1$ can match to $F1$, so $F1$ does not lose anything from making an offer to $W1$ in period 1. The situation in the first market is analogous.

offer prefers to be matched rather than unmatched. Then, when firms use a decentralized reduced deferred acceptance strategy, a worker who receives an offer from a potential stable match partner can infer that he will not receive another offer. He can therefore accept that offer, even if it is not from his first choice firm. This suggests that in order to address what happens in interim stages of the market game we need to define perceived sub-economies when a particular profile of strategies is used.

Definition (Perceived Firm Sub-economy) Worker j, observing $u^w_{\cdot j}$, has a *perceived firm sub-economy* $\mathcal{F}\left(u^w_{\cdot j}\right)$ if $\mathcal{F}\left(u^w_{\cdot j}\right) \subseteq \mathcal{F}$ and there exists some $\mathcal{W}\left(u^w_{\cdot j}\right) \subseteq \mathcal{W}$ such that $j \in \mathcal{W}\left(u^w_{\cdot j}\right)$, and for each U consistent with $u^w_{\cdot j}$ all firms in $\mathcal{F}\left(u^w_{\cdot j}\right)$ prefer any worker in $\mathcal{W}\left(u^w_{\cdot j}\right)$ over any worker in $\mathcal{W} \setminus \mathcal{W}\left(u^w_{\cdot j}\right)$, and all workers in $\mathcal{W}\left(u^w_{\cdot j}\right)$ prefer any firm in $\mathcal{F}\left(u^w_{\cdot j}\right)$ over any other firm (including the option to exit the market), that is over all matching outcomes $\left(\mathcal{F} \setminus \mathcal{F}\left(u^w_{\cdot j}\right)\right) \cup \{\varnothing\}$.

Note that worker j may have several firm sub-economies $\mathcal{F}_l\left(u^w_{\cdot j}\right)$, for $l = 1,..,m_j$. However, for any two firm sub-economies $\mathcal{F}_l\left(u^w_{\cdot j}\right)$ and $\mathcal{F}_k\left(u^w_{\cdot j}\right)$, if $\left|\mathcal{F}_l\left(u^w_{\cdot j}\right)\right| \leq \left|\mathcal{F}_k\left(u^w_{\cdot j}\right)\right|$, either $\mathcal{F}_l\left(u^w_{\cdot j}\right) = \mathcal{F}_k\left(u^w_{\cdot j}\right)$ or $\mathcal{F}_l\left(u^w_{\cdot j}\right) \subset \mathcal{F}_k\left(u^w_{\cdot j}\right)$.[26] Therefore the set of firm sub-economies form a nested collection of subsets of $\mathcal{F}$.

The following assumption simply poses that when all market participants follow decentralized reduced deferred acceptance strategies, the ordering of offers and matches does not convey information in and of itself to either workers or firms. Formally,

Assumption 2 When all market participants follow decentralized reduced deferred acceptance strategies, for any realized match utilities U in the support of G, for any period t,

1. For any active firm $i, \mathcal{S}^f_i\left(u^f_{i\cdot}, h^f_{t,i}\right) = \left\{j | j \in \mathcal{S}^f_i\left(u^f_{i\cdot}\right) \cap \tilde{W}^t_i\right\} \setminus \mathcal{E}^t$.
2. For any active worker j, for any firm sub-economy $\mathcal{F}_l\left(u^w_{\cdot j}\right)$, let $\mathcal{F}_{1,l}\left(u^w_{\cdot j}, h^w_{1,j}\right) = \mathcal{F}_l\left(u^w_{\cdot j}\right) \setminus M^F_1$ and $\mathcal{F}_{t,l}\left(u^w_{\cdot j}, h^w_{t,j}\right) = \mathcal{F}_{t-1,l}\left(u^w_{\cdot j}, h^w_{t-1,j}\right) \setminus \left(\mathcal{E}^t \cup M^F_t\right)$ $\forall t > 1$. For any t, let $\mathcal{F}^t\left(u^w_{\cdot j}, h^w_{t,j}\right)$ be the firm sub-economy $\mathcal{F}_{t,l}\left(u^w_{\cdot j}, h^w_{t,j}\right)$ such that

$$\left|\mathcal{F}_{t,l}\left(u^w_{\cdot j}, h^w_{t,j}\right)\right| = \min\left\{\left|\mathcal{F}_{t,l}\left(u^w_{\cdot j}, h^w_{t,j}\right)\right| \text{ such that } \left|\mathcal{F}_{t,k}\left(u^w_{\cdot j}, h^w_{t,j}\right)\right| > 1 \text{ and } k = 1, ..., m_l\right\}.$$

[26] Suppose $\left|\mathcal{F}_l\left(u^w_{\cdot j}\right)\right| \leq \left|\mathcal{F}_k\left(u^w_{\cdot j}\right)\right|$ but $\exists i \in \mathcal{F}_l\left(u^w_{\cdot j}\right)$ s.t. $i \notin \mathcal{F}_k\left(u^w_{\cdot j}\right)$. Then, by assumption there has to $\exists h \in \mathcal{F}_k\left(u^w_{\cdot j}\right) \setminus \mathcal{F}_l\left(u^w_{\cdot j}\right)$, and by definition of $\mathcal{F}_l\left(u^w_{\cdot j}\right)$, $u^w_{ij} > u^w_{hj}$ and by definition of $\mathcal{F}_k\left(u^w_{\cdot j}\right)$, $u^w_{hj} > u^w_{ij}$, contradiction.

If $\left|\mathcal{F}^t\left(u^w_{\cdot j}, h^w_{t,j}\right)\right| > 2$, then

$$\mathcal{S}^w_j\left(u^w_{\cdot j}, h^w_{t,j}\right) = \left\{i | i \in \mathcal{S}^w_j\left(u^w_{\cdot j}\right) \cap \tilde{F}^t_i\right\} \setminus \left(\mathcal{E}^t \cup M^F_t\right).$$

If $\left|\mathcal{F}^t\left(u^w_{\cdot j}, h^w_{t,j}\right)\right| = 2$ and the worker receives at least one offer from a firm i in $\mathcal{F}^t\left(u^w_{\cdot j}, h^w_{t,j}\right)$, then worker j believes that his most preferred firm of these is his stable match partner, that is

$$\mathcal{S}^w_j\left(u^w_{\cdot j}, h^w_{t,j}\right) = \{i\}.$$

Assumption 2 assures that firms cannot *cross out* their favorite available worker from the set that has not rejected them yet as a potential match. Similarly, workers cannot *cross out* any of the remaining available firms they (at least weakly) prefer to the ones that had made them offers already as potential stable matches. An exception occurs when only two firms are left. In that case, when a worker receives only one offer, he believes the offering firm is his stable match.[27]

It is assumptions 1 and 2 that will prove crucial in assuring incentive compatibility of decentralized reduced deferred acceptance strategies. We therefore introduce the following definition.

Rich Economy An economy is **rich** whenever it satisfies Assumptions 1 and 2.

It is important to note that richness refers to the *support* of potential match utilities. It does not rule out probabilistic updating on the likelihood of different agents being one's stable match in the realized market.

While richness is certainly restrictive (we shall return to some important examples later on), it is in fact a restriction that is satisfied in several important and heavily studied examples:

Example 5 (Rich Economies)

1. **Complete Information Economies.** Under complete information, agents know their stable match at the outset (i.e., the set of stable matches for each agent is a singleton). Note that modified deferred acceptance strategies achieve the stable match quickly, namely in one period.

[27] Any decentralized reduced deferred acceptance strategy profile leads to the same learning pattern regarding stable matches, so the requirements are not of Assumption 2 are not affected by which particular such profile is used.

2. **Full Support Economies.** For a fixed set of firms $\mathcal{F}$ and workers $\mathcal{W}$, let $\mathcal{U}(\Pi)$ denote the set of all aligned match utilities in which each match utility is taken from a set of potential payoffs Π. Consider the economy in which full support is put on elements of $\mathcal{U}(\Pi)$. When Π contains enough elements, say, $2WF + W + F$ (so that some elements of $\mathcal{U}(\Pi)$ are such that all match utilities in the market are different from one another), the economy is rich.

We are now ready to state our existence result. Before doing that, we clarify the class of strategies that will constitute an equilibrium generating the stable match.

We say that workers use a **maximal decentralized reduced deferred acceptance strategy** whenever the associated ranking is truthful, i.e., ranks all of the firms in order of their match utilities, and when workers accept an offer from a firm who is their most preferred potential stable match. In particular, workers do not rank agents who are acceptable (according to the realized match utilities) as unacceptable.

So far, we have considered only pure strategies. Note that a pure strategy is essentially a prescription of the action to be taken for any combination of private information and observed histories. Mixed strategies are then identified by a probability distribution over pure strategies. In particular, a **mixed decentralized reduced deferred acceptance strategy** is one which puts positive probability only on pure decentralized reduced deferred acceptance strategies.

Decentralized reduced deferred acceptance strategies are weakly undominated and so we get the following proposition.

Proposition 7 (Aligned Preferences – Existence) *Suppose preferences are aligned and the economy is rich. For sufficiently high δ,there exists a Bayesian Nash equilibrium in weakly undominated strategies of the decentralized market game in which workers use maximal decentralized reduced deferred acceptance strategies and firms use mixed decentralized reduced deferred acceptance strategies.*

In the proof of Proposition 7, we first show that whenever workers use maximal decentralized reduced deferred acceptance strategies and firms use mixed decentralized reduced deferred acceptance strategies, the maximal decentralized reduced deferred acceptance strategy is a best response for each worker, and all best responses for firms are decentralized reduced deferred acceptance strategies (though not necessarily coinciding with those prescribed by the profile considered). We then look for a (possibly mixed) equilibrium in the decentralized market game in which agents' strategies are restricted to the class given by the Proposition.

Proposition 6 above guarantees that when all agents use any decentralized reduced deferred acceptance strategies, the market outcomes are stable. We therefore get the following Corollary.

Corollary (Aligned Preferences – Stable Implementation) *Suppose preferences are aligned and the economy is rich. Then for sufficiently high δ there exists a Bayesian Nash equilibrium in weakly undominated strategies that implements the unique stable matching.*

We note that while our existence proof is not constructive in nature, for particular classes of economies characterizing an equilibrium profile that implements the stable match can be done. For instance, suppose the markets composing the economy are sufficiently varied so that at each stage, any agent (say, any worker) who has not left the market is a potential stable match for any agent from the other side (say, any firm). In that case the class of decentralized reduced deferred acceptance strategies contains only profiles in which agents echo the full ranking of acceptable agents. In particular, in that case a pure equilibrium implementing the stable outcome exists. This class of economies will turn out important for the analysis of general preferences that follows.

With complete information, Proposition 3 assured that equilibrium outcomes all coincided with the unique stable match. In the presence of uncertainty, even when the conditions of Proposition 7 hold and there exists an equilibrium yielding the stable match for any market realization, there may exist equilibria in weakly undominated strategies that generate unstable matches for some markets in the economy, as the following example illustrates.

Example 6 (Multiplicity with Alignment and Uncertainty) Consider the following economy comprised of two firms $\{F1, F2\}$ and two workers $\{W1, W2\}$. Furthermore, assume that for all firms i and workers $j, u_{ij}^f = u_{ij}^w$. Similar to previous examples, each realization of utilities is then captured by a 3×3 matrix $U = (u_{ij})$, where u_{ij} denotes the utility of a match of firm i to worker j whenever $i, j \in \{1, 2\}$, u_{i3} denotes the utility of firm i from remaining unmatched, and similarly u_{3j} denotes the utility of worker j from remaining unmatched. There are three markets that occur with positive probabilities with corresponding utilities as follows.

$$U_1 = \begin{array}{|c|c|c|}\hline \mathbf{4} & 3 & 0.2 \\ \hline 1 & 2 & \mathbf{5} \\ \hline 0.3 & \mathbf{0.1} & - \\ \hline \end{array}, \quad U_2 = \begin{array}{|c|c|c|}\hline 4 & \mathbf{3} & 0.2 \\ \hline \mathbf{5} & 2 & 1 \\ \hline 0.3 & 0.1 & - \\ \hline \end{array}, \quad U_3 = \begin{array}{|c|c|c|}\hline 4 & \mathbf{1} & 0.2 \\ \hline \mathbf{5} & 2 & 1 \\ \hline 0.3 & 0.1 & - \\ \hline \end{array}.$$

Note that U_3 assures that the economy is rich.

The unique stable matching when U_1 is realized assigns worker 1 to firm 1 while worker 2 as well as firm 2 remain unmatched, generating a sum of utilities for workers and firms of 13.1. The unique stable matching when U_2 is realized assigns worker 1 to firm 2 and worker 2 to firm 1, with sum of utilities of 16 for firms and workers. Note that firm 1 and worker 2 do not know which market is realized. This can result in the following strategies and beliefs that form a Bayesian Nash equilibrium in weakly undominated strategies, that does not yield the stable matching:

- Firm 1 makes an offer to worker 2 whenever worker 2 is unmatched and makes an offer to worker 1 otherwise (when possible).
- Firm 2 leaves the market immediately when U_1 is realized. When U_2 is realized, Firm 2 makes an offer to worker 1 whenever possible and, otherwise, makes an offer to worker 2 (when possible).
- Worker 1 accepts his best available offer in period 1. Otherwise, worker 1 leaves the market in period 1.
- Worker 2 accepts his best available offer in period 1 and otherwise leaves the market immediately.

Worker 2 believes that U_1 is realized whenever he does not receive an offer from firm 1 in period 1. All other beliefs are determined according to Bayes rule or arbitrarily.

This equilibrium generates a matching between firm 1 and worker 2 for either of the utility realizations.

Intuitively, both firm 1 and worker 2 are not informed of the realized utilities. Whenever firm 2 plays a weakly undominated strategy it reveals the realized market and never makes an offer to worker 2. For firm 1 to wait for period 2, he needs to be certain that worker 2 will wait as well, and vice versa. However, in equilibrium, beliefs can be constructed so that this coordinated wait does not occur. Note that this equilibrium results in lower total utilities for workers and firms (relative to the stable match) when U_1 is realized, namely only 11.3 instead of 13.1. However, this comparative statics could go either way. For instance, if the payoff to worker 1 from being unmatched were 2.3 (instead of 0.3), the above profile would remain an equilibrium generating total payoffs of 13.3 conditional on the first market being realized.

6. General Economies with Uncertainty

General preferences present several crucial hurdles to the generation of stable equilibrium matchings. When preferences are aligned agents using reduced deferred acceptance strategies lead to an agent exiting or a top-top match being formed in every period. Consequently, information is revealed in each period and workers know when they can start accepting offers. When preferences are not aligned, this need not be the case. In fact, as we will see, decentralized reduced deferred acceptance strategies may not transmit enough information for the market to unravel, even when the (general) economy is rich or when there are no frictions.

6.1 General Frictionless Economies

We start with a simple example illustrating a difficulty encountered when preferences are unrestricted. Namely, the problem that a worker may not know when the market is "over", that is, when he should accept an offer. Recall that in the firm proposing deferred acceptance algorithm the market is over when no offer is rejected, as then, every firm either has its offer held by a worker, or has decided not to make an offer anymore. Can this information be transmitted in a decentralized market?

Example 7 (No Information Transmission) Consider an economy with three firms $\{F1, F2, F3\}$ and three workers $\{W1, W2, W3\}$, where all agents prefer to be matched rather than unmatched. Suppose firms and workers use decentralized deferred acceptance strategies (that implement the stable match as an equilibrium outcome when preferences are aligned). We now inspect when a worker, who does not receive any further public or private information, can be certain that the offer he received is his stable matching partner.

- Suppose a worker, say worker $W1$, who receives two offers in period 1 and receives no further information (that is no matches occur, nor does he receive any further offers). When worker $W1$ rejects one of his two offers, there are two options for period 2 (given that by assumption $W1$ does not receive another offer). First, it may be that each of the other workers now holds one offer, in which case no offer is rejected, and the market would be over when the deferred acceptance algorithm is emulated. The second possibility is that one of the other workers now holds two offers. This other worker, say $W2$, will now in turn reject one of the offers. This offer can either go to $W3$ in period 3, or in case it is not the offer from the firm that $W1$ rejected in period 1, can go to $W1$. To summarize, if $W1$ does not receive any information until period 3, he can conclude that he received his stable match, no more offers are made, and he can accept the offer he held from period 1.
- Suppose worker $W1$ receives only one offer in period 1, and receives no further information, public or private. If each worker received exactly one offer, then the market would be over when the deferred acceptance algorithm is emulated. Suppose one worker receives two offers. Then one of these firms will be rejected and can make an offer to another worker in period 2. Hence if $W1$ does not receive another offer in period 2, the offer is made to the third worker, and hence each worker holds exactly one offer. $W1$ will not receive another offer and can accept his offer. So, if $W1$ receives no new information, he can accept his offer in period 2.

Assume that many market realizations are possible, including market M_1 that delivers the preferences below (we use the convention that for firm i, $j \succ j'$ whenever $u^f_{ij} > u^f_{ij'}$ and similarly for worker j, $i \succ i'$ whenever $u^w_{ij} > u^w_{i'j}$), where the unique stable matching is indicated in bold.

$$M_1: \begin{array}{ll} \mathbf{F1}: & W1 \succ \mathbf{W2} \succ W3 \\ \mathbf{F2}: & W1 \succ \mathbf{W3} \succ W2 \\ \mathbf{F3}: & W3 \succ \mathbf{W1} \succ W2 \end{array}, \quad \begin{array}{ll} \mathbf{W1}: & \mathbf{F3} \succ F1 \succ F2 \\ \mathbf{W2}: & \mathbf{F1} \succ F2 \succ F3 \\ \mathbf{W3}: & \mathbf{F2} \succ F3 \succ F1 \end{array}.$$

Furthermore, assume that in M_1, $F1$ has a very high match utility for $W1$ compared to the other workers.

If many match utilities are possible, and occur with comparable probability, then $F1$ has a positive chance to match to a worker that is not $W1$. Suppose all other firms make offers according to decentralized deferred acceptance strategies. Can $F1$ improve the chance of receiving worker $W1$ by delaying her offer, if all other firms follow decentralized deferred acceptance strategies?

In market M_1 an offer by $F1$ to $W1$ triggers a chain of offers ($W1$ rejects $F2$, who then makes an offer to $W3$, who rejects $F3$ who then makes an offer to $W1$) that displaces $F1$ two periods later (note that when preferences are aligned, such a chain cannot occur).

Suppose $F1$ delays making an offer to $W1$, and only makes that offer in period 2. Then, $W3$ who received an offer only from $F3$ in period 1, believing that all firms use a decentralized deferred acceptance strategy, will accept that offer, in light of no new information in period 2. Therefore, if $F1$ makes an offer to $W1$ only in period 2, the offer of $F1$ cannot trigger a chain that eventually replaces $F1$.

It follows that $F1$ increases the chance of hiring $W1$ by delaying the offer for one period.

Intuitively, the problem illustrated in the example arises from the fact that workers cannot monitor whether firms make offers. Furthermore, simply observing whether a firm made an offer or not is not sufficient. Indeed, consider an extension of Example 7 in which the economy contains a fourth worker, who always prefers to exit the market rather than match to any firm, and this is common knowledge. Then, in a firm proposing deferred acceptance algorithm, the same set of outcomes should occur, in the same amount of time, as if this worker were not present. However, firm 1 in the example could make an offer to this fourth worker instead of not making any offer at all in period 1. Thus, in order to fully circumvent the issue presented in the example, market participants need to monitor also the target of firms' offers.

When considering general preferences, it will therefore be useful to assume the following.

Full Market Monitoring *Firms and workers observe all offers. That is, at the end of the first stage of each period t, all agents are informed regarding which firm made an offer to which worker. At the end of the third stage of any period t, not only are all agents informed of matches, but also of the workers' reaction to any specific offer they had in hand that period.*

Let a **delayed decentralized deferred acceptance strategy** be one where firms make offers according to the deferred acceptance strategy, that is, to the highest acceptable worker who has not rejected them yet. Furthermore, workers hold on to their most preferred acceptable offer, and reject all other offers. Finally, a worker accepts an offer only when there is a period in which no more offers are made, because either all the firms exited the market, or all the firms that have not exited, have their offer held in the previous period by a worker. If agents use these strategies, we have the following analogue of Proposition 5 for general preferences.

Proposition 8 (General Preferences – No Discounting) *Suppose $\delta = 1$. Delayed decentralized deferred acceptance strategies constitute a Bayesian Nash equilibrium in weakly undominated strategies and yield the stable matching in the decentralized market game with full market monitoring.*

The proof is similar in spirit to that of Proposition 5. Note that full market monitoring assures that firms cannot manipulate outcomes by delaying offers or making offers out of order undetected, and the structure of strategies allows the market to emulate the firm proposing deferred acceptance algorithm that generates the stable match for any market realization.

6.2 General Economies with Frictions

When preferences are aligned, we showed that when firms and workers use decentralized reduced deferred acceptance strategies, these provide sufficient information to determine the stable match μ_M. First, we show that this need not be the case anymore, when preferences are not aligned. The example below illustrates this and will suggest the requirements imposed on general economies that guarantee the existence of an equilibrium implementing the stable match.

Example 8 (Insufficient Information – Reduced Deferred Acceptance) Consider the economy with three firms $\{F1, F2, F3\}$, three workers $\{W1, W2, W3\}$, where all always prefer to be matched rather than unmatched, and match utilities of firms and workers coincide. Suppose there are four equiprobable markets $\{M_k\}_{k=1}^4$ corresponding to match utilities $\{U_k\}_{k=1}^4$ as follows (note that the preferences corresponding to M_1 are those discussed in Example 1).

$$U_1 = \begin{array}{|c|c|c|}\hline \mathbf{2,3} & 3,2 & 1,3 \\ \hline 3,1 & \mathbf{2,3} & 1,1 \\ \hline 3,2 & 2,1 & \mathbf{1,2} \\ \hline \end{array}, \quad U_2 = \begin{array}{|c|c|c|}\hline 2,1 & \mathbf{3,3} & 1,3 \\ \hline \mathbf{3,3} & 2,2 & 1,1 \\ \hline 3,2 & 2,1 & \mathbf{1,2} \\ \hline \end{array},$$

$$U_3 = \begin{array}{|c|c|c|}\hline 2,3 & \mathbf{3,3} & 1,3 \\ \hline \mathbf{3,1} & 2,1 & 1,1 \\ \hline 2,2 & 1,2 & \mathbf{3,2} \\ \hline \end{array}, \quad U_4 = \begin{array}{|c|c|c|}\hline 2,2 & \mathbf{3,2} & 1,3 \\ \hline \mathbf{3,3} & 2,3 & 1,1 \\ \hline 2,1 & 1,1 & \mathbf{3,2} \\ \hline \end{array}.$$

Note that in M_1, if the minimal reduced deferred acceptance strategies were followed (namely, ones that specify only potential stable matches as acceptable), then firm $F1$ would effectively rank $W2$ above $W1$ and rank $W3$ as unacceptable, firm $F2$ would rank $W1$ above $W2$ and rank $W3$ as unacceptable, and firm $F3$ would rank $W3$ as the only acceptable worker. Similarly, worker $W1$ would rank $F1$ above $F2$,and firm $F3$ as unacceptable, worker $W2$ would rank $F2$ above $F1$ and $F3$ as unacceptable, and worker $W3$ would rank only $F3$ as acceptable. The outcome in M_1 would not be the stable match μ_{M_1} but rather $F1$ and $F2$ would swap workers in a way that would actually increase their utility (as is the case in the equilibrium discussed in Example 1). Markets M_2, M_3, and M_4 assure that when discount factors are high enough, agents best respond with the above reduced deferred acceptance strategies. Indeed, market M_2 guarantees that $F1$ and $F2$ will rank $W1$ and $W2$ as they do, while market M_3 assures the optimality of the above strategy for worker $W1$ and M_4 for worker $W2$. Note that if we shift some probability from these markets to an additional market M_5, in which $F3$ has the same preferences as in M_1, and in which furthermore $F3$ is the first choice of $W3$, then there would be no Bayesian Nash equilibrium in which $F3$ would make an offer to $W1$ or $W2$, and hence no strategies would yield the stable matching in this economy.

The example highlights the strength of Proposition 2, that assured that when preferences are aligned, for any unstable match $\mu \neq \mu_M$, either there is an individual who blocks μ, or there is a pair (f, w) that blocks μ, such that $\mu_M(f) = w$. In particular, the agents composing this blocking pair (f, w) rank each other as acceptable in the decentralized reduced deferred acceptance strategies. When preferences are general, this need not be the case, and decentralized reduced deferred acceptance strategies may reduce restrictions sufficiently to allow for other matchings as in the example above.

Clearly, all the hurdles to stability appearing in decentralized markets when preferences are aligned are still hurdles in the general case. Therefore, to show existence of a Bayesian Nash equilibrium that yields a stable outcome we will assume the following.

Definition (Super-rich Economy) *An economy is super-rich if it is a rich economy in which for every firm i, if $j \in W \cup \varnothing$ and $j \in S_i^f(u_{i\cdot}^f)$ then for all $j' \in W \cup \varnothing$ such that $u_{ij'}^f > u_{ij}^f$, $j' \in S_i^f(u_{i\cdot}^f)$.*

When the economy is super-rich, the decentralized reduced deferred acceptance strategy for each firm is defined uniquely, essentially forcing them to make offers at each period to their most preferred acceptable worker who has not rejected them yet. In particular, as in the aligned case, if an unstable match is created agents composing a blocking pair effectively rank each other as acceptable. Together with full market monitoring this assures that manipulations as illustrated in Example 8 are not possible. Using the proof of Proposition 7 we then get the following.

Proposition 9 (General Economies – Existence) *Suppose the economy is super-rich. Then, for sufficiently high δ, in the decentralized market game with full market monitoring there exists a Bayesian Nash equilibrium in weakly undominated strategies that implements the unique stable matching for each utility realization.*

In fact, one equilibrium justifying the Proposition is that entailing all workers using the maximal decentralized reduced deferred acceptance strategy and all firms using the (unique) decentralized reduced deferred acceptance strategy.

7. CONCLUSIONS

The paper analyzed decentralized market games in which firms and workers interact dynamically. We showed that there are general environments in which any equilibrium outcome of the market game coincides with the unique stable match. Indeed, this is the case whenever (i) the market is one in which firms and workers have complete information, and (ii) firms and workers' preferences are aligned. When these two conditions are satisfied, the assumption that a market outcome is stable is valid. In particular, any econometric exercise that deduces restrictions on underlying preferences by observing a specific market outcome and assumes that it is stable is justified.

While stability is a (very particular) coalitional concept, (Bayesian) Nash equilibrium is not. It is thus not surprising that the two notions generally do not coincide. Indeed, there are many environments in which equilibria other than the stable equilibrium exist, even when there is complete information and a unique stable match. In these cases, inferring underlying preferences of agents from observed market outcomes and the restrictions imposed by stability is generally not valid.

In that respect, our results illustrate the wedge between centralized markets, that always entail an equilibrium implementing the stable outcome, and decentralized markets, that generally do not.

The analysis also highlights the importance of frictions. Indeed, even a small amount of discounting makes achieving stability in equilibrium so much harder in the presence of uncertainty. While throughout the paper we have operationalized frictions through time discounting, or the probability of market breakdown, we note that frictions can take a variety of forms that would lead to similar conceptual insights. For instance, frictions could manifest through a cost of making offers.

To conclude, the paper shows that when studying markets, it is generally crucial to understand market characteristics that go beyond the identification of market participants and their preferences. Indeed, it is important to describe markets in detail, in terms of the information available to participants and the plausibility of frictions, in order to be able to predict which outcomes they may achieve. This, in turn, implies that channels by which information can be transmitted among market participants can be a critical element of market design.

8. Appendix - Proofs

Proof of Proposition 1. A stable matching can be constructed recursively as follows. First, denote by

$$\mathcal{F}^1 = \{i \in \mathcal{F} : \forall j \in \mathcal{W} : u^f_{i,\varnothing} > u^f_{i,j}\}$$

and

$$\mathcal{W}^1 = \{j \in \mathcal{W} : \forall i \in \mathcal{F} : u^w_{\varnothing,j} > u^w_{i,j}\}.$$

In a stable match μ, firms $\mathcal{F}^1$ and workers $\mathcal{W}^1$ must remain unmatched, and exit the market.

Iterate this process, where step i is:

$$\mathcal{F}^i = \{i \in \mathcal{F} \backslash \left\{\mathcal{F}^1 \cup .. \cup \mathcal{F}^{i-1}\right\} : \forall j \in \mathcal{W} \backslash \left\{\mathcal{W}^1 \cup .. \cup \mathcal{W}^{i-1}\right\} : u^f_{i,\varnothing} > u^f_{i,j}\}$$

and

$$\mathcal{W}^i = \{j \in \mathcal{W} \backslash \left\{\mathcal{W}^1 \cup .. \cup \mathcal{W}^{i-1}\right\} : \forall i \in \mathcal{F} \backslash \left\{\mathcal{F}^1 \cup .. \cup \mathcal{F}^{i-1}\right\} : u^w_{\varnothing,j} > u^w_{i,j}\}.$$

The process ends, when there is a k, such that $\mathcal{F}^k = \mathcal{W}^k = \varnothing$. Let $\mathcal{F}^\varnothing = \mathcal{F} \backslash \left\{\mathcal{F}^1 \cup .. \cup \mathcal{F}^{k-1}\right\}$ and let $\mathcal{W}^\varnothing = \mathcal{W} \backslash \left\{\mathcal{W}^1 \cup .. \cup \mathcal{W}^{k-1}\right\}$. Then let

$$M^f = \arg \max_{(i,j) \in \mathcal{F} \backslash \mathcal{F}^\varnothing \times \mathcal{W} \backslash \mathcal{W}^\varnothing} u^f_{i,j},$$

which, by assumption is a singleton. Because preferences are aligned, any stable matching μ must assign $\mu(i) = j$.

Denote

$$\mathcal{F}^* = \mathcal{F} \setminus \left\{\{i\} \cup \mathcal{F}^\varnothing\right\}, \mathcal{W}^* = \mathcal{W} \setminus \left\{\{j\} \cup \mathcal{W}^\varnothing\right\}.$$

This whole procedure can be replicated for $\mathcal{F}^*$ and $\mathcal{W}^*$ and continued until all firms and workers are assigned. The generated μ is a matching and any other stable matching must coincide with μ's assignments. Since stable matchings exist (Gale and Shapley, 1962), μ is the unique stable matching. ■

Proof of Proposition 2. Let μ_M be the stable match, and μ' be a matching that is not stable. Reconstruct μ_M as in the proof of Proposition 1 above. At some iterative state a discrepancy must occur between μ_M and μ'. This can manifest itself through an individual who is unmatched in μ_M

but matched in μ', in which case μ' is not individually rational. Alternatively, there is a pair (f, w) that is matched under μ_M and not under μ'. In that case, (f, w) blocks μ' since at the iterative stage at which f and w get matched, they form a top-top match. As this is the first discrepancy between μ_M and μ' in the iterative process, the match partner $\mu'(f)$ of f and $\mu'(w)$ of w are part of the remaining set of firms and workers and hence inferior to the match partner in μ_M. ■

Proof of Proposition 3.

1. Let μ be the unique stable matching. Let

$$\mathcal{F}^{\varnothing} = \{i \in \mathcal{F} : \mu(i) = \varnothing\}.$$

   Consider the following profile of strategies (which will be appropriate for any length of offers). In the first period, all firms in $\mathcal{F}^{\varnothing}$ leave the market immediately and any firm $i \in \mathcal{F} \setminus \mathcal{F}^{\varnothing}$ makes an offer to $\mu(i)$. From the second period and on, each firm makes an offer to her most preferred acceptable unmatched worker, and leaves the market if all her acceptable workers are matched. Each worker accepts his most preferred acceptable offer at hand in period 1 and leaves the market upon receiving no offers from acceptable firms. This profile clearly constitutes an equilibrium implementing μ.

2. Suppose now that preferences are aligned. Consider any equilibrium of the market game (with any type of offers). We use similar notation to that used in the proof of Proposition 1. First, denote $\mathcal{F}_{(0)} \equiv \mathcal{F}$ and $\mathcal{W}_{(0)} \equiv \mathcal{W}$. For any $k = 0, 1, ...$, denote

$$\mathcal{F}^{\varnothing}_{(k)} = \{i \in \mathcal{F}_{(k)} : \forall j \in \mathcal{W}_{(k)} : u^f_{i,\varnothing} > u^f_{i,j}\}$$

   and

$$\mathcal{W}^{\varnothing}_{(k)} = \{j \in \mathcal{W}_{(k)} : \forall i \in \mathcal{F}_{(k)} : u^w_{\varnothing,j} > u^w_{i,j}\}.$$

   Then, let

$$M^f_{(k)} = \arg \max_{(i,j) \in \mathcal{F}_{(k)} \setminus \mathcal{F}^{\varnothing}_{(k)} \times \mathcal{W}_{(k)} \setminus \mathcal{W}^{\varnothing}_{(k)}} u^f_{i,j}$$

   (and recall that this is a singleton). If $M^f_{(k)} = \{(i_k, j_k)\}$, define

$$\mathcal{F}_{(k+1)} = \mathcal{F}_{(k)} \setminus \left\{ \{i_k\} \cup \mathcal{F}^{\varnothing}_{(k)} \right\}, \qquad \mathcal{W}_{(k+1)} = \mathcal{W}_{(k)} \setminus \left\{ \{j_k\} \cup \mathcal{W}^{\varnothing}_{(k)} \right\}.$$

Notice that by construction $\mu(i_k) = j_k$ for all k, and $\mu(i) = \varnothing$ for all $i \in \mathcal{F}^{\varnothing}_{(k)}$ for some k.

We show that in any period k, firm i_k cannot make an offer to a worker j for which $u^f_{ij} < u^f_{i\mu(i)} \equiv u^f_{ij_k}$. Indeed, suppose there is a minimal $k \geqslant 0$ at which i_k has an offer out to worker j for which $u^f_{ij} < u^f_{ij_k}$. We now show that firm i has a profitable deviation – namely, making an offer to worker j_k at an appropriate period. Indeed, from the minimality of k, for any $k' < k$, firm $i_{k'}$ has an offer to worker j such that $u^f_{ij} \geqslant u^f_{ij_{k'}}$. From construction, in period $k' = 0$, firm i_0 and worker j_0 get matched and all firms in $\mathcal{F}^{\varnothing}_{(0)}$ and workers in $\mathcal{W}^{\varnothing}_{(k)}$ leave the market. If $k = 1$, an offer to j_1 in periods 0 and 1 assures worker j_1 accepts firm i_1 in period 1 at the latest, and is therefore a profitable deviation. Assume then that $k > 1$. By definition of k, it must be the case that firm i_l is matched to worker j_l by the end of period l for all $l = 1, ..., k-1$ and that any worker $j_{l'}, l' \geqslant k$ is not matched up to period k. Furthermore, in equilibrium, firms in $\cup^k_{t=0}\mathcal{F}^{\varnothing}_{(t)}$ and workers in $\cup^k_{t=0}\mathcal{W}^{\varnothing}_{(t)}$ must be out of the market by the end of period k. Thus, if firm i_k does not have an offer out at the beginning of period k, its best response is to make an offer to worker j_k (which will be accepted immediately, since firm i_k is worker j_k's most preferred effectively available firm). Suppose then that firm i_k makes an offer to worker $j_{k'}, k' > k$ at period $l < k$, which is held up to the end of period k. Suppose worker j_k accepts his equilibrium offer at period t. Consider the following deviation for firm i_k. If $t \leqslant k$, i_k makes an offer to worker j_k from period t and on. If $t > k$, i_k makes an offer to no one until period k, at which point she makes an offer to worker j_k. By construction, this deviation will surely generate a match between firm i_k and worker j_k and is therefore profitable, in contradiction to our original profile constituting an equilibrium. In particular, each firm i_l is matched to worker j_l for all l. ■

Proof of Proposition 4. Let $\mathcal{M}' = (\mathcal{F}', \mathcal{W}', U')$ be a μ_M-induced submarket with matching μ' induced by μ_M and the firm optimal stable matching $\tilde{\mu} \neq \mu'$ with the maximal number of firms. Consider the following strategy profile:

Period 1: Firms $\mathcal{F} \setminus \mathcal{F}'$ make offers to $\mathcal{W} \setminus \mathcal{W}'$ according to μ and these offers get accepted.

Period 2: Firms $\mathcal{F}'$ make offers to $\mathcal{W}'$ according to $\tilde{\mu}$ and those offers get accepted.

Upon Deviation: Let $\hat{\mu}$ be the firm preferred stable match in the sub-market containing the remaining firms and workers. Each firm i without an offer out makes an offer to the worker prescribed to her by $\hat{\mu}$, or leaves the market if $\hat{\mu}(i) = \varnothing$. Each worker j accepts an offer from the most preferred firm whom he values at least as much as the firm prescribed by $\hat{\mu}$, leaves the market if $\hat{\mu}(j) = \varnothing$, and if having only offers from firms inferior to that prescribed by $\hat{\mu}$, holds on to all of them.

The strategies comprising the above profile are weakly undominated. Furthermore, they constitute an equilibrium. Indeed, in period 1, a firm in $\mathcal{F} \setminus \mathcal{F}'$ that does not make an offer, would lead to an observable deviation and a μ-induced sub-market with more than $|\mathcal{F}'|$ firms. Therefore, from maximality of $\mathcal{F}'$, the firm-preferred match will be given by μ, and she can at most get the worker she could have gotten in period 1 at a one period delay, which is costly due to discounting. Similarly, if a firm i in $\mathcal{F} \setminus \mathcal{F}'$ makes an offer to a worker who leads to a greater match utility than that prescribed to her by μ, then since μ is the unique stable match, firm i cannot be preferred over the firm (or market exit) prescribed by μ to the worker. Thus, that firm will not be matched in period 1, and as before, will at most get the original worker assigned to her at a delay. As for workers, if any worker in $\mathcal{W} \setminus \mathcal{W}'$ holds or rejects an offer from a firm assigned to him by μ, then he generates an observable deviation to a μ-induced market with more than $|\mathcal{F}'|$ firms. As for firms, the most the worker can gain is a delayed match with the firm he could have matched with in period 1. For sufficiently high discount factors, workers are indeed willing to reject offers as prescribed.

In period 2, a firm who does not make any offer or exit the market as prescribed to her simply delays the realization of her match utility, which cannot be profitable. If a firm makes an offer to a worker whom she prefers over that prescribed to her by $\hat{\mu}$, from stability the worker will not accept that offer and, again, the firm delays the realization of her match utility. Similarly, a worker who does not accept an offer or exits the market according to $\hat{\mu}$ simply delays the realization of his match utilities, which cannot be profitable.

Thus, the above profile constitutes an equilibrium in weakly undominated strategies that generates an outcome that does not coincide with μ over the submarket $\mathcal{M}'$. ■

Proof of Proposition 5. First we show that deferred acceptance strategies yield the stable outcome. Note that workers will eventually hold offers that correspond to their stable match partner by construction. Furthermore, since preferences are aligned, in every period, there is either a firm or a worker that exits, or there is a top-top match that is formed. In particular, the process stops in finite time.

We now show that deferred acceptance strategies constitute a Bayesian Nash equilibrium. Namely, that there are no profitable deviations for any agent.

Suppose a worker j rejects an offer from firm i instead of holding it. From the no cycle property, such a rejection cannot launch a cycle generating a superior offer for the worker. Therefore, if the offer is from a potential stable match partner, then the worker may have rejected his best offer, and be strictly worse off. Such a deviation is therefore profitable if it makes the worker strictly better off in a market realization in which firm i is not his stable match. In that market realization, it cannot be that firm i is strictly better than j's stable match partner, as then j should never receive an offer from firm i (else, firm i and worker j would form a blocking pair to the stable match). It follows that rejecting an offer from a firm that is the current superior offer is never a strictly profitable deviation.

The other potential deviation by worker j is the acceptance of an offer that is not from the highest potential stable match partner. In that case, if the realized market is one in which the stable match leads to the preferred stable match, the worker strictly loses. Such an acceptance cannot lead to a match with firm i that is preferable to j's stable match in some realization. Indeed, if that were the case worker j and firm i would form a blocking pair in that market. Therefore, such a deviation leads the worker to either match with a worse firm than he otherwise would have, or speed up his ultimate match, which does not affect his payoff.

Consider now the firms. Suppose firm i deviates and makes an offer to worker j who is not the the most preferred worker among workers who have not rejected that firm. Suppose there is a market in which firm i strictly benefits from making an offer to that worker. Because there is no discounting, this can only imply that firm i ends up matching with a worker different, and strictly preferable, to her stable match partner. Suppose the resulting matching in the market (assuming all other agents follow the deferred acceptance strategies) is μ'. The matching μ' has the property that the set of firms F', who prefer this match to the stable match μ, is not empty, as it contains at least firm i. By the Blocking Lemma there exists a blocking pair (i^*, j^*) with i^* not in F'such that j^* is matched in μ' to a firm in F'. However, since i^* and j^* follow deferred acceptance strategies, i^* must have made an offer to j^*, who must have rejected that offer, which would only happen had he received an offer he strictly prefers, in contradiction to μ' being formed. ■

Proof of Proposition 6.

1. Assume all agents use a reduced deferred acceptance rule and suppose $\mathcal{E}$ is an economy with a market realization in which the outcome μ' is different than the stable match μ. By Proposition 2, either μ' is not individually rational, in which case it could not be the result of all agents in $\mathcal{E}$ using reduced deferred acceptance strategies, or there exists a pair (i^*, j^*) that blocks μ' and that furthermore is matched to each other under μ. That is, $\mu(i^*) = j^*$. This implies, $u^f_{i^*j^*} > u^f_{i^*\mu'(i^*)}$ and $j^* \in \mathcal{S}^f_{i^*}(u^f_{i^*\cdot})$. Hence if firm i^* uses a reduced deferred acceptance strategy, i^* must rank j^* above $\mu'(j^*)$. Similarly, for worker j^*, since $i^* \in \mathcal{S}^w_{j^*}\left(u^w_{\cdot j^*}\right)$, j^* cannot be matched with someone other than i^* unless, within the deferred acceptance algorithm, he receives a better offer. The fact that $u^w_{i^*j^*} > u^w_{\mu'(j^*)j^*}$ implies that j^* does not reject i^*'s offer, in contradiction.

2. Suppose $a \in \mathcal{F} \cup \mathcal{W}$ is an agent who does not use a deferred acceptance rule. That is, agent a either ranks some agent (including potentially a themselves) as preferred to a potential stable match when match utilities prescribe the reverse. Whenever there is only one worker or only one firm, the claim follows trivially. Assume then that $|\mathcal{F}|, |\mathcal{W}| \geqslant 2$.

Certainly, if a ranks a potential stable match as unacceptable for some set of potential stable matches, then whenever the market in which the corresponding match is stable is realized, the centralized outcome is not stable.

Suppose that $a \in \mathcal{W}$ ranks a potential stable match i below a firm i' who is not a potential stable match when observing $u^w_{\cdot a}$ and the set of potential stable matches is S. Assume $u^w_{ia} > u^w_{i'a}$. Let $j \in \mathcal{W}$ be another worker (other than a).

Consider an economy in which there are three markets characterized by match utilities $U, \tilde{U}$, and $\hat{U}$ in which all agents other than i, i', a, and j have no acceptable agents other than themselves. It therefore suffices to focus on the match utilities corresponding to agents $\{i, i', a, j\}$.

We construct U and $\tilde{U}$ so that they satisfy the following:[28]

[28] These conditions are consistent with alignment. Indeed, assuming without restriction that $u^w_{ia}, u^w_{ia} > 1$, the reader can think of the following manifestation of $U, \tilde{U}$ in which we summarize preferences through the following two matrixes, where the first number in each rubric corresponds to the firm's preference and the second number to the appropriate worker:

U:

	a	j
i	u^w_{ia}, u^w_{ia}	$u^w_{ia} - 1, u^w_{ia} - 1$
$i\prime$	$u^w_{i'a}, u^w_{i'a}$	$u^w_{i'a} - 1, u^w_{i'a} - 1$

$\tilde{U}$:

	a	j
i	$\tilde{u}^f_{i\varnothing} - 1, \tilde{u}^w_{\varnothing a} - 1$	$\tilde{u}^f_{i\varnothing} - 2, \tilde{u}^w_{\varnothing j} - 1$
$i\prime$	$u^w_{i'a}, \tilde{u}^w_{\varnothing a} + 1$	$u^w_{i'a} - 1, \tilde{u}^w_{\varnothing j} + 1$

.

a. Firm i' cannot distinguish between the two markets, while all other agents can.

b. Firm i' prefers worker a to worker j in both markets.

c. Under U, i and a, and i' and j, are part of the stable match, while under $\tilde{U}$, i' and a are part of the stable match.

d. Under $\tilde{U}$, i' is both a's and j's only acceptable firm.

$\hat{U}$ is such that $\hat{u}^w_{\cdot j} = \tilde{u}^w_{\cdot j}$, so that worker j cannot distinguish $\tilde{U}$ from $\hat{U}$, j is the only acceptable worker to i' and all other agents prefer to stay unmatched than be matched to any market participant.

Each of the remaining markets in the economy is one in which a's match utilities are given by $u^w_{\cdot a}$ and the stable match is an element $i'' \in S \setminus \{i\}$ (for instance, by making a the only acceptable worker for i'' and having all other agents prefer to be by themselves over being matched with any market participant).

If the stable match is achieved under $\hat{U}$, worker j must rank firm i' as acceptable when observing $\hat{u}^w_{\cdot j} = \tilde{u}^w_{\cdot j}$. Therefore, if the stable match is achieved under $\tilde{U}$, it must be the case that i' ranks a higher than j (and acceptable) when observing $u^f_{i'\cdot}$. But then, under U, it cannot be the case that the stable match is established. Indeed, the centralized mechanism generates a stable match for the submitted preference rankings, and i' and a would form a blocking pair.

Suppose that $a \in \mathcal{W}$ ranks a potential stable match i below a potential stable match i' when observing $u^w_{\cdot a}$ and the set of potential stable matches is S. Assume $u^w_{ia} > u^w_{i'a}$. As above, let $j \in \mathcal{W}$ be another worker (other than a).

Consider an economy in which there are three markets characterized by match utilities $U, \tilde{U}$, and $\hat{U}$ in which all agents other than i, i', a, and j have no acceptable agents other than themselves. It therefore suffices to focus on the match utilities corresponding to agents $\{i, i', a, j\}$.

We construct U and $\tilde{U}$ so that they satisfy the following:[29]

a. Worker a and firm i' cannot distinguish between the two markets, while all other agents can.

b. Firm i' prefers worker a to worker j (in both markets).

[29] Again, these conditions are consistent with alignment. Using the same notation as in the previous footnote (and assuming without restriction that $u^w_{ia}, u^w_{i'a} > 1$), we can look at:

U:

	a	j
i	u^w_{ia}, u^w_{ia}	$u^w_{ia} - 1, u^w_{ia} - 1$
i'	$u^w_{i'a}, u^w_{i'a}$	$u^w_{i'a} - 1, u^w_{i'a} - 1$

$\tilde{U}$:

	a	j
i	$\tilde{u}^f_{i\varnothing} - 1, u^w_{ia}$	$\tilde{u}^f_{i\varnothing} - 2, \tilde{u}^w_{\varnothing j} - 1$
i'	$u^w_{i'a}, u^w_{i'a}$	$u^w_{i'a} - 1, \tilde{u}^w_{\varnothing j} + 1$

.

c. Under U, i and a, and i' and j, are part of the stable match, while under $\tilde{U}$, i' and a are part of the stable match.

d. Under $\tilde{U}$, i' is j's only acceptable firm.

$\hat{U}$ is such that $\hat{u}^w_{\cdot j} = \tilde{u}^w_{\cdot j}$, so that worker j cannot distinguish $\tilde{U}$ from $\hat{U}$, j is the only acceptable worker to i' and all other agents prefer to stay unmatched than be matched to any market participant.

As before, the other remaining markets in the economy are ones in which a's match utilities are given by $u^w_{\cdot a}$ and the stable match is an element $i'' \in S \setminus \{i, i'\}$ (for instance, by making a the only acceptable worker for i'' and having all other agents prefer to be by themselves over being matched with any market participant).

If the stable match is achieved under $\hat{U}$, worker j must rank firm i' as acceptable when observing $\hat{u}^w_{\cdot j} = \tilde{u}^w_{\cdot j}$. Therefore, if the stable match is achieved under $\tilde{U}$, it must be the case that i' ranks a higher than j (and acceptable) when observing $u^f_{i'\cdot}$. But then, under U, it cannot be the case that the stable match is established, just as before.

Analogous constructions follow when agent a is a firm that does not follow a reduced deferred acceptance rule.

3. Let v be a profile of reduced deferred acceptance strategies. Suppose they do not form a Bayes Nash equilibrium. That is, there is agent $a \in \mathcal{F} \cup \mathcal{W}$ who strictly benefits submitting some other strategy, say σ_a.

First, suppose $a \in \mathcal{F}$. Since truthful revelation is a weakly dominant strategy in the firm-proposing deferred acceptance algorithm, which the centralized market emulates, it follows directly that σ_a either leads to the same outcome or to a less preferred one.

Suppose $a \in \mathcal{W}$ and σ_a is a strictly profitable deviation. If σ_a is also a reduced deferred acceptance strategy, then by part 1. above the outcome is unchanged, and the deviation cannot be strictly profitable.

Suppose that the deviation to σ_a yields a match μ' such that $u^w_{\mu'(a)a} > u^w_{\mu(a)a}$. Then, by Proposition 2, one possibility is that μ' is not individually rational. However, since worker a is the only agent not using a reduced deferred acceptance strategy, he could be the only one whose outcome violates individual rationality. Since worker a strictly prefers μ' to μ, this cannot be the case (and μ, by definition, is individually rational). The only other possibility by Proposition 2 is that there exists a pair (i^*, j^*) that blocks μ' such that $\mu(i^*) = j^*$. First, it is clear that $j^* \neq a$ since a strictly prefers

μ' to μ. Since both i^* and j^* submit reduced deferred acceptance strategies, i^* must rank j^* above $\mu'(i^*)$. Hence, it must be that j^* rejects i^* through the centralized mechanism, contradicting the fact that (i^*, j^*) are a blocking pair to μ'. ■

Proof of Proposition 7. We first note that for sufficiently high δ, whenever all other workers use maximal reduced deferred acceptance strategies and firms use mixed decentralized reduced deferred acceptance strategies, using a maximal reduced deferred acceptance strategy is a best response for a worker.

Indeed, at each period t, a worker with no acceptable offers can deviate by exiting the market when there are still active conceivable potential matches that are preferred to exiting. This cannot be part of a best response for sufficiently high discount factors. Similarly, accepting an offer from an acceptable firm who is not the worker's most preferred stable match cannot be part of a best response when discount factors are high enough.

Alternatively, a worker having offers at hand can deviate by rejecting a set of firms that does not coincide with the firms to be rejected according to the maximal reduced deferred acceptance strategy, namely all firms but his most preferred. From alignment, rejection of firms cannot generate the arrival of an offer from a preferred firm, and assures that rejected firms will not make future (repeat) offers. In particular, such deviations cannot speed up matches, nor alter positively the ultimate match.

We now show that richness assures that whenever workers use a maximal decentralized reduced deferred acceptance strategy and firms use mixed decentralizes reduced deferred acceptance strategies, a firm's best responses are within the class of reduced deferred acceptance strategies.

Consider first a firm i that in period t has no outstanding offers, and whose updated strategies suggest worker j as the most preferred stable match. There are two kinds of deviations from a decentralized reduced deferred acceptance prescription: (1) Make no offer; or (2) make an offer to some other worker k who generates worse match utility than j. The benefits of such deviations can be either through speeding up the time at which the firm's offer is accepted, or through generating a preferred ultimate match.

Regarding (1), if firm i does not make an offer at period t, there are three potential implications. First, if making an offer according to any decentralized reduced deferred acceptance strategy would not have affected market participants' history following period t,[30] then the only effect of this deviation could be the prolonging of its match creation. If not making an offer affects certain participants' histories, then due to Assumption 2, this cannot affect the firm's final match. Again, such a deviation can only prolong the timing of its match. Finally, suppose that the firm's most preferred potential stable match has a sub-economy at period t which consists of only 2 firms, one of which is firm i. In that case, if the worker receives an offer from the other firm in his perceived sub-economy, he will accept that offer immediately, even if he prefers firm i, in which case firm i is strictly worse off.

Regarding (2), suppose firm i makes an offer to a worker k who is lower ranked than her most preferred potential stable match, worker j. By Assumption 1, there is a positive probability that her offer is accepted in a state of the world in which she would have otherwise gotten worker j. Furthermore, such a deviation can never lead to a better ultimate match from the incentive compatibility inherent in the firm-proposing deferred acceptance algorithm. Indeed, note that such a deviation would be tantamount to submitting an untruthful preference list when the firm deferred acceptance algorithm is used. However, revealing preferences truthfully is a dominant strategy for firms.

Furthermore, from Assumption 2, such an offer will not make other participants change their effective rank orderings. Thus, it may only delay the time at which i's offer gets accepted.

Consider now the restricted centralized market game in which workers' strategy set is restricted to maximal decentralized reduced deferred acceptance strategies, and firms' strategy set is restricted to decentralized reduced deferred acceptance strategies (mixed or pure). Since there is a finite number of firms' decentralized reduced deferred acceptance strategies, an equilibrium (possibly mixed) exists in this restricted game. From the above, for sufficiently high δ, the corresponding strategy profile is also an equilibrium in our original decentralized market game, as required. ■

Proof of Proposition 8. First we show that delayed deferred acceptance strategies yield the stable outcome. We show that workers will eventually hold offers that correspond to their stable match partner. This follows by assumption. Second, in every period, there is either a firm that makes an offer or exits the market, hence the process stops in finite time.

Now we need to show that no agent has an incentive to deviate:

Suppose a worker j rejects an offer from firm i instead of holding it. If the offer is from a potential

[30] For instance, in the case in which any such strategy suggests an offer to j, who gets matched in that period with probability 1 (to firm i or to some other firm).

stable match partner, then the worker may have rejected his best offer, and be strictly worse off in some realization, so, this is only profitable, if he is strictly better off in some other market realization. It cannot be that in that market realization firm i is strictly better than j's stable match partner, as then j should never receive an offer from firm i (since else firm i and worker j would form a blocking pair to the stable match). Hence, firm i is strictly worse than j's stable match partner. In this case, eliminating firm i from the list of acceptable firms does not change the stable match, hence does not affect the final outcome.

Suppose a worker accepts an offer, before he has received an offer from his highest potential stable match partner. It is easy to see that there is no market realization in which he can strictly benefit from such behavior.

Suppose firm i deviates and makes an offer to worker j who is not the the most preferred worker among workers who have not rejected that firm. Suppose there is a market in which firm i strictly benefits from making an offer to that worker. Because there is no discounting, this can only imply that firm i receives a different match, hence receives a match that is strictly better than his stable match partner. Let this new matching be μ', which has the property that the set of firms F', who prefer this match to the stable match μ is not empty, since it contains at least firm i. By the blocking lemma there exists a blocking pair (f, w) with f not in F'such that w is matched in μ' to a firm in F'. However, since f and w follow deferred acceptance strategies, f made an offer to w, who must have rejected that offer, which only happens if he has an offer he strictly prefers. contradiction. ■

REFERENCES

[1] Alcalde, José, David Pérez-Castrillo, and Antonio Romero-Medina. 1998. "Hiring procedures to implement of stable allocations," *Journal of Economic Theory*, **82**, 469–480.

[2] Avery, Christopher, Christine Jolls, Richard A. Posner, and Alvin E. Roth. 2001. "The Market for Federal Judicial Law Clerks," *University of Chicago Law Review,* **68(3)**, 793-902.

[3] Becker, Gary S. 1973. "A theory of marriage: part I," *Journal of Political Economy*, **81(4)**, 813-846.

[4] Blum, Yosef, Alvin E. Roth, and Uri G. Rothblum. 1997. "Vacancy chains and equilibration in senior-level labor markets," *Journal of Economic Theory*, **76**, 362–411.

[5] Bulow, Jeremy and Jonathan Levin. 2006. "Matching and Price Competition," *American Economic Review,* **96(3)**, 652 - 668.

[6] Burdett, Ken and Melvyn G. Coles. 1997. "Marriage and Class," *The Quarterly Journal of Economics*, **112(1)**, 141-168.

[7] Crawford, Vincent P., and Elsie Marie Knoer. 1981. "Job Matching with Heterogenous Firms and Workers," *Econometrica,* **49(3)**, 437 - 450.

[8] Demange, Gabrielle, David Gale and Marilda Sotomayor. 1987. "A further note on the stable matching problem," *Discrete Applied Mathematics*, **16**, 217-222.

[9] Diamantoudi, Effrosyni, Eiichi Miyagawa, and Licun Xue. 2007. "Decentralized Matching: The Role of Commitment," mimeo, http://pareto.uab.es/matching/program_files/dynamic(07)0416.pdf.

[10] Eeckhout, Jan. 1999. "Bilateral Search and Vertical Heterogeneity," *International Economic Review*, **40(4)**, 869-887.

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

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

[13] Gale, David, and Marilda Sotomayor, 1985b, "Ms Machiavelli and the stable matching problem," American Mathematical Monthly, 92, 61-8.

[14] Haeringer, Guillaume and Myrna Wooders. 2007. "Decentralized Job Matching", mimeo, http://idea.uab.es/~r1g/docs/matching.pdf.

[15] Hitsch, Guenter J., Ali Hortacsu, and Dan Ariely. 2006. "What Makes You Click? Mate Preferences and Matching Outcomes in Online Dating," mimeo, http://home.uchicago.edu/~ghitsch/Hitsch-Research/Guenter%20Hitsch_files/Online-Dating.pdf.

[16] Kelso, Alexander S., and Vincent P. Crawford. 1982. "Job Matching, Coalition Formation, and Gross Substitutes," *Econometrica*, **50(6)**, 1483 - 1504.

[17] Konishi, Hideo and Margarita Sapozhnikov. 2006. "Decentralized Matching Markets with Endogeneous Salaries," mimeo, http://www2.bc.edu/~konishih/dcmatch5.pdf.

[18] Monderer, Dov and Lloyd S. Shapley. 1994. "Potential Games," *Games and Economic Behavior*, **14**, 124-143.

[19] Pais, Joana V. 2006. "Incentives in Decentralized Random Matching Markets," mimeo, http://ideas.repec.org/p/ise/isegwp/wp122006.html.

[20] Niederle, Muriel. 2007. "Competitive Wages in a Match with Ordered Contracts," *American Economic Review*, vol. 97, No.5, 1957 - 1969.

[21] Niederle, Muriel and Alvin E. Roth. 2003. "Unraveling reduces mobility in a labor market: Gastroenterology with and without a centralized match," *Journal of Political Economy*, **111(6)**, 1342 – 1352.

[22] Niederle, Muriel, and Alvin E. Roth. 2007. "Market Culture: How Norms Governing Exploding Offers Affect Market Performance," mimeo, http://www.stanford.edu/~niederle/MakingMarketsThick.May2007.pdf.

[23] Pycia, M. 2007. "Many-to-One Matching with Complementarities and Peer Effects," mimeo, http://www.econ.psu.edu/~pycia/pycia-matching-with-complementarities.pdf.

[24] Roth, Alvin E., 1982, "The economics of matching: stability and incentives," *Mathematics of Operations Research*, 7, 617-628.

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

[26] Roth, Alvin E., 1989, "Two-sided matching with incomplete information about others' preferences," *Games and Economic Behavior*, 1, 191-209.

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

[28] Roth, Alvin E. and Xiaolin Xing. 1994. "Jumping the Gun: Imperfections and Institutions Related to the Timing of Market Transactions," *American Economic Review*, **84**, 992-1044.

[29] Roth, Alvin E. and Xiaolin Xing. 1997. "Turnaround Time and Bottlenecks in Market Clearing: Decentralized Matching in the Market for Clinical Psychologists," *Journal of Political Economy*, **105**, 284-329.

[30] Segal, Ilya. 2007. "The Communication Requirements of Social Choice Rules and Supporting Budget Sets," *Journal of Economic Theory*, 136, 341-378.

[31] Sorensen, Morten. 2007. "How Smart is Smart Money? A Two-Sided Matching Model of Venture Capital," *Journal of Finance*, forthcoming.