

Better-reply Dynamics in Deferred Acceptance Games*

Guillaume Haeringer

Hanna Hałaburda

Universitat Autònoma de Barcelona

Harvard Business School

Friday 15th April, 2011

— PRELIMIARY AND INCOMPLETE —

Abstract

In this paper we address the question of learning in two-sided matching mechanism that utilizes the Deferred Acceptance algorithm. We consider a repeated matching game where at each period agents observe their match and have the opportunity to revise their strategy (i.e., the preference list they will submit to the mechanism). We focus in this paper on better-reply dynamics. To this end, we first provide a characterization of better-replies and a comprehensive description of the dominance relation between strategies. Better-replies are shown to have a simple structure and can be decomposed into four types of changes. We then present a simple better-reply dynamics with myopic and boundedly rational agents and identify conditions that ensure that limit outcomes are outcome equivalent to the outcome obtained when agents play their dominant strategies. Better-reply dynamics may not converge, but if they do converge then the limit strategy profiles constitute a subset of the Nash equilibria of the stage game.

JEL codes C72, D41.

Keywords: Better-reply dynamics, Deferred Acceptance, two-sided matching.

*Guillaume Haeringer acknowledges support of Barcelona GSE Research Network and of the Government of Catalonia, Spanish Ministry of Science and Innovation through grant “Consolidated Group-C” ECO2008-04756 and the *Generalitat de Catalunya* (SGR2009-419).

1 Introduction

Mechanism design theory usually assumes that agents know the mechanism they face and have enough cognitive resources to respond optimally to the incentives and constraints that are imposed by the designer. In practice this may not be the case. Agents may not have access to a clear and detailed description of the mechanism, and/or may have difficulties in finding their best choices. Whether agents are able to learn their optimal strategies under a mechanism is thus of paramount importance when implementing mechanisms in real-life settings. It is therefore surprising that the issue of learning in mechanism design has not been a focus of the recent theoretical literature.¹ This paper addresses this research question by considering a mechanism that is increasingly used in real-life settings, namely the mechanism built upon Gale and Shapley’s (1963) Deferred Acceptance (DA) algorithm for two-sided matching markets.

Real life use of the DA algorithm (or its variants) now abounds. For example, the entry market for medical interns in the US, school admission in several US cities, academic hiring in France or college admission in Hungary all use this algorithm. In some cases the DA algorithm was chosen because of its theoretical properties —e.g., school choice in Boston— and in other cases the choice was “accidental” —e.g., medical interns in the US until 1997.² Although those markets use (almost) the same algorithm, they often differ in the way the mechanism is presented to market participants. For instance, parents in Boston are given a precise description of the algorithm, while job candidates in the French academic job market have little knowledge of the algorithm used to match candidates and departments.³

The DA algorithm is a multistep process that works as follows. There is a side of the market that makes match proposals, and the other side either rejects or accepts the proposals they receive. We keep with the tradition in the matching literature, calling the proposing side *men* and the accepting side *women*. In the DA algorithm each man makes an offer to one woman at a time (i.e., one per step). At each step of the algorithm a

¹Most of the work on learning and mechanism design have been undertaken by the computer science (e.g., Balcan *et al.* (2005)) or by the experimental literatures (e.g., Healy (2006)).

²See Roth and Peranson (1999).

³In fact, the French administration does not directly utilize the DA algorithm but an algorithm that is outcome equivalent — see Iehlé and Haeringer (2010).

woman can hold at most one offer. Thus, she has to choose between the offer received in the given round and the offer she holds from the previous step. One of the attractive features of this algorithm is that it naturally describes a direct mechanism, in which each agent submits to a central clearinghouse a preference list over potential partners to match with. The final matching is then computed using the DA algorithm with those preference lists as the input. The preference lists of men indicate in which order the algorithm should make offers to women, and the preference lists of women indicate which offers are to be held or rejected.

In line with well known results, the dominant strategy for the proposing side is to list the true preferences. However, with limited information about the applied mechanism, or because of bounded rationality, fully informed participants may not correctly respond to the incentives they face, even when being recommended to be truthful (if such recommendations are made). This is because in spite of the relative simplicity of the DA mechanism, the existence of a dominant strategy for the proposing side is not straightforward. Moreover, understanding what outcome a given strategy profile leads to is somewhat difficult.⁴ This is a reason for concern because participants with little or wrong understanding of the mechanism may try to strategize and thus affect market outcomes. It is then crucial to know (i) whether agents can learn how to play and converge to some sort of stable behavior, and (ii) the type of outcome we obtain in the long run.

In a two-sided matching market an “equilibrium” is often better described by the notion of *stability*, that is, a matching between the agents from both sides of the market where no agents who are not matched together would both prefer to be matched to one another.⁵ In the presence of frictions (e.g., the existence of a deadline to make transactions or the absence of monetary transactions), stability becomes one of the most important property sought in matching markets. Roth (1991) showed that matching markets that do not produce stable matchings tend to perform worse or simply collapse whereas markets that produce stable matchings do not experience failures.⁶ In the context of matching markets with strong frictions, the DA algorithm is now considered to be a

⁴Note, however, that from a computational point of view the DA algorithm is very simple (and polynomial).

⁵Stability also requires that matching is individually rational.

⁶McKinney, Niederle and Roth (2005) show that centralized matching markets that utilize Gale and Shapley’s Deferred Acceptance are not always immune to market failure.

serious contender, as it always produces a stable matching.⁷ The mechanism built upon the DA algorithm has strong appealing properties. First, as we just intuited, the DA algorithm produces a stable matching with respect to the submitted preferences. It is also the most preferred stable matching by all men among all stable matchings. Second, and not less interesting from the mechanism design perspective, it is a (weakly) dominant strategy for men to submit their true preferences. One drawback, however, is that women may have an incentive to manipulate the mechanism by reporting false preferences — Dubins and Freedman (1981), Roth (1982).

To streamline our analysis of learning in matching games we shall consider the simplest matching environment and use the insights of the game theoretic literature on learning. More precisely, we shall consider a one-to-one matching market with strict preferences between men and women, and assume that agents repeatedly play this matching mechanism.⁸ Throughout this paper, our focus will be the proposing side of the matching mechanism. That is, we shall assume that only the individuals on the proposing side face strategic choices. Agents on the other side of the market are simply assumed to always submit their true preferences. This case fits well the so-called school choice problem, where the proposing side are students’ parents and the other side are schools.⁹ Of course, in real-life settings agents often do not participate repeatedly in a matching game. We argue however that a repeated setting can be understood as a proxy for social learning. For instance, parents participating in a school choice program usually seek information from their acquaintances who participated in the past and who have similar preferences. While we assume that the same population with the same (true) preferences repeatedly plays a matching game, the constraints we impose on the updating behavior can dispense with this assumption. That is, we shall consider that agents in our model only know their true preferences, the preferences they submitted in the previous period, and the

⁷This is true for “classic” matching markets. Many-to-one markets where there are complementarities between agents constitute an example where the existence of stable matchings is not guaranteed — see Roth and Sotomayor (1990).

⁸All our results easily carry out to more complex environments such as school choice —see Abdulkadiroğlu and Sönmez (2003).

⁹Schools in a school choice problem are not agents per se, but rather perceived as “goods to be consumed,” and schools’ “preferences” (needed to run the algorithm) are in fact exogenous rankings of students imposed by the authorities built, for instance, upon students’ grades, social characteristics (distance to school, presence of siblings in the school, etc.)

identity of the partners they were matched with. In other words, an agent updating in our model may well think that other agents' preferences are changed or that some agents were replaced by new ones with different preferences.

While our approach in this paper is theoretical, it is our contention that our learning model should take into account the specificities of a matching game. First, an important aspect of a matching mechanism is that strategies consist of orderings. Consequently, strategy sets are large. While the presence of large strategy sets is not a real issue in most games, it makes our task difficult here because it is difficult (if not impossible) to reduce strategic choice to that of a parameter. The size of the strategy sets also implies that approaches based on statistics about the performance of past actions, e.g., reinforcement learning (Roth and Erev (1995)), are not well suited to our context. Second, and equally important aspect of a matching mechanism is the frequency of play. Most centralized markets operate once a year. It is therefore natural to consider Cournot type dynamics, i.e., situations where at each period many (if not all) agents update their strategies at the same time. Hence, we will need to address whether dynamics with simultaneous updating converge. In this paper we focus on better-reply dynamics. As we shall see, better-replies in a matching game are not very demanding, and have a strong intuitive interpretation. In particular, an attractive feature of our approach is that better-replies do not necessitate knowing the strategy profile of the other players.

Models of learning in games are often presented as equilibrium selection devices, and this question does not lose its importance here. Indeed, although the game we study admits, for each man, a weakly dominant strategy, Haeringer and Klijn (2009) showed that in fact it admits many Nash equilibria. In particular, any stable matching can be sustained as an equilibrium outcome.¹⁰ Thus, we need to verify whether simple dynamics lead to a Nash equilibrium. If so, the second and perhaps more important question is whether in the limit individuals play their dominant strategy (or at least whether the limit profile is outcome equivalent to the dominant strategy profile).

We first characterize the better-replies in a matching game. Better-replies can be described by a combination of four types of changes: *reshuffle below*, *reshuffle above*, *move-up* and *move-down*. Reshuffling below (resp. above) consist of changing the relative

¹⁰In fact, the set of stable matchings is usually a strict subset of the set of equilibrium outcomes —see Haeringer and Klijn (2009).

order of the agents that are declared less preferred (resp. more preferred) than the current matched partner. For instance, if an individual submitted the preference list a, b, c, d, e (in this order) and is matched to, say, c , then submitting the list b, a, c, d, e is a reshuffling above change (and the list a, b, c, e, d is a reshuffling below). Moving-down consists of declaring less preferred than the current match an individual that was declared more preferred than the current match but truly less preferred. For instance, if an individual submitted the list a, b, c, d, e and is matched to c and truly prefers c to, say, b , then the list a, c, b, d, e is a move-down. It is important to note that we do not specify *where* below c the individual b is moved. A move-up is a converse: if there is an individual that is truly preferred to the current match but declared less preferred, then moving him above is a move-up. We show that, holding the strategy of other agents fixed, only a move-up can improve the outcome of an agent. That is, moving-down and reshuffling (below or above) have no impact on the outcome (for the individual changing). When individuals do not know the strategy profile of the other individuals, this changes uniquely characterizes the better-replies.

We then consider a dynamic model of repeated matching where at each period individuals use a better-reply to update their strategies. Our main results are the following. When we consider only move-ups, convergence of the better-reply dynamics is guaranteed. However, the limit outcome can be large, but it always coincides with a stable matching. When move-downs are also allowed, then the dynamic process can cycle. However, if it does not cycle then the limit outcome is unique and corresponds to the man-optimal matching.

In Section 2 we present the model. In Section 3 we characterize the better-replies and in Section 4 we consider the better-reply dynamics.

2 Framework

The market consists of two finite disjoint sets: the men, $M = \{1, \dots, m\}$, and the women, $W = \{1, \dots, n\}$. Each man $m \in M$ is endowed with a strict preference relation P_m over the set $W \cup \{m\}$.¹¹ Similarly, each woman $w \in W$ has a strict preference relation P_w over $M \cup \{w\}$. We denote by $P = (P_v)_{v \in M \cup W}$ a profile of preferences, and use the usual

¹¹When an agent remains unmatched, we denote it as being matched to himself or herself.

notation P_{-v} to denote the profile $P \setminus P_v$. We denote by \mathcal{P}_v the set of all preference relation of individual $v \in M \cup W$, and by \mathcal{P} the set of all preference profiles. For a preference profile $P_v \in \mathcal{P}_v$, let R_v denote the weak preference relation associated with P_v . Similarly, for a preference profile $P \in \mathcal{P}$, R denotes the weak preference profile associated to P . Given a set $G \subseteq M \cup W$, P_G and R_G denote the profile of strict and weak preferences restricted to individuals in G .

Assumption 1 (Market thickness) *There are at least as many women as men.*

We need the market thickness assumption only for expositional convenience. The results provided in the paper would not change qualitatively in absence of this assumption.

A matching is a one-to-one mapping $\mu : M \cup W \rightarrow M \cup W$ such that

- For each man $m \in M$, $\mu(m) \in W \cup \{m\}$.
- For each woman $w \in W$, $\mu(w) \in M \cup \{w\}$.
- For each agent $v \in M \cup W$, $\mu(\mu(v)) = v$.

Given a set $G \subset M$ and a matching μ we denote by $\mu(G)$ the set of individuals to whom the members of G are matched, i.e., $\mu(G) = \{v \in G \cup W : \mu(m) = v \text{ for some } m \in G\}$.

A matching μ is *individually rational* if for each $v \in M \cup W$, $\mu(v)R_v v$. A matching μ is *blocked* by a pair (m, w) if $wP_m\mu(m)$ and $mP_w\mu(w)$. A matching μ is *stable* if it is individually rational and it is not blocked by any pair $(m, w) \in M \times W$. Given a preference profile P we denote by $S(P)$ the set of stable matchings.

It is well known that for any (strict) preference profile P the set of stable matchings is nonempty (Gale and Shapley, 1962). A stable matching can be obtained using Gale and Shapley's *deferred acceptance* (DA) algorithm. Their algorithm, with men making proposals to women, works as follows:

Step 1: Each man m proposes to his most preferred woman among the ones that are acceptable for him. If there is no such woman then the man is matched to himself.

Each woman declines all but her most preferred man among the men who proposed to this woman and are acceptable to it (if any).

Step k, k ≥ 2: Each man who has been declined in the previous step proposes to his most preferred woman among the women that have not yet declined him and are acceptable for him. If there is no such woman then the man is matched to himself. Each woman declines all but her most preferred man among the men who proposed to this woman and are acceptable to it and the man it did not decline in the previous step (if any).

The algorithm stops when every man is either matched to a woman or to himself. Given a preference profile P , we denote by $\varphi(P)$ the man-optimal stable matching, i.e., matching obtained by the DA algorithm we just described.

In the sequel, P^* denotes the true preference profile, which is fixed throughout the paper, and μ^* the man-optimal stable matching under the true preferences.

It is easy to see that the set of men and women, the true preference profile P^* , the set of all preference profiles \mathcal{P} and the DA algorithm defines a strategic form game where the set of players is $M \cup W$, the set of strategies of player v is \mathcal{P}_v , outcomes are given by φ and are evaluated by players using their true preferences.

A preference relation P_m *weakly dominates* the preference relation P'_m if, for any profile P_{-m} , man m is always at least as well off with P_m as with P'_m ,

$$\varphi(P_m, P_{-m})(m) R_m^* \varphi(P'_m, P_{-m})(m).$$

Theorem 1 (Dubins and Freedman (1981), Roth (1982)) *For each man $m \in M$, the preference P_m^* weakly dominates any other preference $P_m \in \mathcal{P}$.*

3 Comparing strategies

In order to build our learning model we need first to characterize better-replies in a matching game, or — said differently — we need to understand how to compare strategies. This step is even more necessary since most of the matching literature has ignored other strategies than a weakly dominant one — the truthful preferences.

Since being truthful is a weakly dominant strategy one may intuit that the “closer” we are to the true preferences the better and hence that for a man m a better-reply with respect to some strategy profile P_m is any preference P'_m that is “closer” to P_m^* than is

P_m . As we shall see in Section 3.1 a notion of distance between preferences is in fact not related to the concept of better-replies but rather that of dominance.

3.1 Dominance relations

To determine whether P_m or P'_m is closer to P_m^* , we need a metric to compare preferences. A common way and natural way to do so is by considering the so-called *Kemeny distance* (Kemeny, 1959), which consists of comparing the number of pairs of alternatives that are ranked differently between P_m and P_m^* to determine the distance between P_m and P_m^* . This in turn allows us to see whether the distance between P_m and P_m^* is smaller or larger than the distance between P'_m and P_m^* . In our context, it is more useful to consider a related concept, which we call *Kemeny set*.

Definition 1 Given a preference relation P_v^* , the Kemeny set of the preference relation P_v with respect to P_v^* is the set of all pairs (v, v') that are not ordered identically in P_v and P_v^* .

$$K(P_v, P_v^*) = \{(v, v') : vP_v v' \text{ and } v'P_v^* v\}$$

It is readily verified that the Kemeny distance between the preferences P_v and P_v^* is simply the cardinality of the Kemeny set.

We are now ready to state the main result of this section: for each man $m \in M$, if the Kemeny set of a preference ordering P_m is a subset of the Kemeny set of another preference ordering P'_m , then for any profile of preferences of the other men and of the women, the strategy P_m weakly dominates the strategy P'_m .

Theorem 2 *A preference ordering P_m dominates another preference ordering P'_m if, and only if, $K(P_m, P_m^*) \subset K(P'_m, P_m^*)$.*

The proof of Theorem 2 is in the Appendix. An immediate consequence of Theorem 2 is that we can order preference orderings when comparing their Kemeny sets. Note however that this ordering is not complete, since it is possible for that two preference orderings P_m and P'_m that neither $K(P_m, P_m^*) \subset K(P'_m, P_m^*)$ nor $K(P'_m, P_m^*) \subset K(P_m, P_m^*)$ hold.

This characterization can be for instance useful in data analysis to measure how far agents are from their true preferences.

3.2 Better replies

Recall that we denote m 's true preferences by P_m^* . The reported preferences, P_m may but do not need to be truthful.

Definition 2 P'_m is a *better reply* than P_m to P_{-m} if and only if

$$\varphi(P'_m, P_{-m}) R_m^* \varphi(P_m, P_{-m}),$$

according to m 's real preferences P_m^* .

Consider then a man m who submitted a preference list P_m and is matched to some individual v . Suppose that the man m would like to submit now a different preference list, P'_m . For any P_m and P'_m that list the same set of women, P'_m is achievable by making one or more changes of the four types:

- (i) Changing ranking of a woman originally below v , to a different position still below v (i.e., *reshuffling below*).
- (ii) Changing ranking of a woman originally above v , to a different position still above v (i.e., *reshuffling above*).
- (iii) Changing ranking of a woman originally below v to a position above v (i.e., *move up*).
- (iv) Changing ranking of a woman originally above v to a position below v (i.e., *move down*).

We show below what effect each type of change in ranking has on the matching, φ . And in particular that some of the change types have no bearing on which woman the man m is matched to under φ .

Notice that in the DA algorithm as soon as a man is matched to his final partner he is never asked again to propose to a lower ranked individual. That is, the relative ranking of the individuals that are ranked below his match in his submitted list is irrelevant. So, if all men re-arrange the relative ranking of the individuals that are ranked below their match (but keeping all these individuals still ranked below their match), then the final matching is the same.

Observation 1 Let P and P' such that for each man m , $L(P_m, \varphi(P)(m)) = L(P'_m, \varphi(P)(m))$, and otherwise P and P' are the same. Then $\varphi(P) = \varphi(P')$.

In the observation, L denotes the lower-contour set, i.e., $L(P_m, \varphi(P)(m))$ denotes the set of all women ranked lower than $\varphi(P)(m)$ in ranking P_m . Similarly we later use U to denote an upper-contour set. That is, $U(P_m, \varphi(P)(m))$ denotes the set of all women ranked above $\varphi(P)(m)$ in ranking P_m .

The next proposition states that the a property similar to Observation 1 also holds if all (or some) men change the relative ranking of the women ranked above their current match.

Proposition 1 Let P and P' such that for each man m , $U(P_m, \varphi(P)(m)) = U(P'_m, \varphi(P)(m))$ and $L(P_m, \varphi(P)(m)) = L(P'_m, \varphi(P)(m))$. Then $\varphi(P) = \varphi(P')$.

Proof Let $P^1 = (P'_{m_1}, P_{-m_1})$ and $\mu^1 = \varphi(P^1)$. Since φ is strategy-proof $\mu^1 R_{m_1}^1 \mu$. Suppose that $\mu^1(m) \neq \mu(m)$. So, $\mu^1 P_{m_1}^1 \mu$ and thus $\mu^1 P_{m_1} \mu$, which contradicts the strategy-proofness of φ . Hence, $\mu^1(m) = \mu(m)$.

Also, notice that $\mu \in S(P^1)$. So, $\mu^1 R_M^1 \mu$, where R_M^1 means that μ^1 is weakly preferred to μ by all men $m \in M$. Since only man m_1 changed his (declared) preferences between P and P^1 , and since $\mu^1(m) = \mu(m)$, $\mu^1 R_M \mu$. Suppose that there exists a man \hat{m} such that $\mu^1 P_{\hat{m}} \mu$. By the Blocking Lemma¹² there exists a pair (m, w) such that (m, w) block μ^1 under P . If $m \neq m_1$, then (m, w) also block μ^1 under P^1 , a contradiction. So, $m = m_1$, i.e., $w P_{m_1} \mu^1(m_1)$. Since $\mu^1(m) = \mu(m)$ it follows that $w P_{m_1}^1 \mu^1(m_1)$. So, (m_1, w) block μ^1 under P^1 , a contradiction. Hence, $\mu R_M \mu^1$, and thus $\mu^1 = \mu$.

Notice that in this argument, we do not employ any element of the lower-contour set. Therefore, $\mu^1 = \mu$ for any order of preferences in the lower-contour set.

It suffices now to repeat the same reasoning with the profiles $P^2 = (P'_{m_1}, P'_{m_2}, P_{-m_1, m_2})$, $P^3 = (P'_{m_1}, P'_{m_2}, P'_{m_3}, P_{-m_1, m_2, m_3})$, etc. until we attain the profile P' to obtain the desired result. ■

Proposition 1 implies that reshuffling below and reshuffling above do not affect the matching outcome. And therefore such changes do not constitute a better reply. Now, we

¹²Blocking Lemma is a standard result in the literature, and we restate it in the Appendix (see Lemma 6).

look into the impact of the other two possible changes: *move up* and *move down*. Recall that a woman is moved down when she was declared more preferred than v in P_m , and in P'_m she is declared less preferred than v . Observe that — provided the other men do not change their submitted preference lists — any change of this type alone is inconsequential. That is, as Lemma 1 states, if m only moved down some women, then it is easy to see that he remains matched to the same individual.

Lemma 1 *Let P and let $\mu = \varphi(P)$. Let P'_m be such that for some woman w ranked above $\mu(m)$ in P_m , $wP_m\mu(m)$, this woman is moved down below $\mu(m)$ in P'_m , $\mu(m)P'_mw$. And otherwise P'_m is the same as P_m . Preferences for all other agents stay the same $P'_{-m} = P_{-m}$. Let $\mu' = \varphi(P')$. Then $\mu'(m) = \mu(m)$.*

We omit the formal proof of Lemma 1 because the argument is straightforward: Since DA is strategy-proof the man cannot end up being matched to a woman ranked higher in the submitted preference list. Since the previous matching is still stable once the man has changed his preferences the result follows.

Therefore, if other men do not change their submitted preferences, only the change involving a *move up* may lead to a better outcome for man m . Recall that a woman is moved up when she was declared less preferred than v in P_m , and in P'_m she is declared more preferred than v . However, moving up may constitute a better or a worse reply. This is because, as Lemma 2 states, if the man moves a woman w above his current match v , he is either matched with w or with the same v . If he moves up a truly less preferred woman, i.e., vP_m^*w , he risks being worse off.

Lemma 2 *Let P and let $\mu = \varphi(P)$. Let P'_m such that for some woman w such that $\mu(m)P_mw$ this woman is moved up above $\mu(m)$, $wP'_m\mu(m)$. And otherwise P'_m is the same as P_m . Preferences for all other agents stay the same $P'_{-m} = P_{-m}$. Let $\mu' = \varphi(P')$. Then either $\mu'(m) = \mu(m)$ or $\mu'(m) = w$.*

Therefore, the only possible better reply is for man m to move up a truly preferred woman. This is the only type of change that can constitute a better reply. Other types of changes — reshuffling below and above, as well as moves down — do not change the match. Moving up a truly less preferred woman may either make the man worse off or leave his match unchanged.

It is important to notice that a better reply is not necessarily equivalent to choosing dominating strategy. To see this, suppose that a man changes his submitted list from P_m to P'_m by moving up a woman, say, w , above his current match, say, v . This implies that the pair (v, w) was in the Kemeny set $K(P_m, P_m^*)$ and this pair is not in the set $K(P'_m, P_m^*)$. However, since we do not specify *where* is w in the preference list P'_m , it may well be the case that we have w declared preferred to some woman w' under P'_m , while it is the contrary under P_m and P_m^* . In other words, if P'_m is a better reply to P_m is may not be the case that P'_m dominates P_m .

We say that P_m is a best reply to P_{-m} if there does not exists a better reply to P_{-m} than P_m .

Definition 3 P'_m is a *best reply* against P_{-m} if and only if

$$\nexists P''_m \text{ such that } \varphi(P''_m, P_{-m}) P_m^* \varphi(P'_m, P_{-m}).$$

In the reminder of the paper, we assume that the man does not know the reported preferences of other agents, P_{-m} , while changing his rankings. However, it is useful to compare our results to a benchmark case where the man m knows the strategies of other players.

If the strategies of other players, P_{-m} , are known, a sufficient condition for a strategy to be best reply is that all women less preferred than $\varphi(P_m^*, P_{-m})(m)$ are listed *below* in m 's preference ranking. It is not a necessary condition. A necessary condition would require that only those less preferred women have to be listed below $\varphi(P_m^*, P_{-m})(m)$ that could be matched to m if m listed them above $\varphi(P_m^*, P_{-m})(m)$. Notice that a best reply is not unique. In fact, there is a whole set of best replies. They all yield the same match for m : $\varphi(P_m^*, P_{-m})(m)$.

Without knowing the strategies of other player, a man m does not know $\varphi(P_m^*, P_{-m})(m)$. In the following section we investigate the case when the preferences of other agents are not known to m .

3.3 Other Agents' Strategies Unknown

Our approach of learning in the matching game assumes — we believe somewhat realistically — that men do not know other agents' submitted preferences. Moreover, we also

assume that they are not fully aware of the mechanism providing the match. They gain the information by trial-and-error. They observe only their own match and the preference list they submitted. In other words, we shall assume that for a man who submitted the preference list P_m and is matched to $\varphi(P_m, P_{-m})(m)$, the other men's and women's submitted preferences could be *any* preference \widehat{P}_m such that $\varphi(P_m, \widehat{P}_m)(m) = \varphi(P_m, P_{-m})(m)$.

In this section, we investigate characteristics of better and best reply when the strategies of other agents are unknown. In such a case, each man needs to depend on his expectations when deciding on his strategy.

Assumption 2 *For any man $m \in M$, m does not know P_{-m} .*

We assume such beliefs of m where any preference profile of other agents, P_{-m} , can occur with some positive probability.

Suppose P_{-m} is fixed, even if unknown to m . However, man m observes $\varphi(P_m, P_{-m})(m) = v$, i.e., his own match given the preference he submitted, P_m , and P_{-m} . We show that strategy P_m is a best reply to P_{-m} if all women truly preferred to v by m are listed above v , i.e., $\forall w (wP_m^*v \implies wP_mv)$. Moreover, if P_m is a best reply to P_{-m} , then $v = \varphi(P_m^*, P_{-m})$. Thus, listing true preferences is a special case of best reply.

To show this property of best reply, we first need to establish certain properties of better replies in expected terms.

Suppose now that the man only makes changes of one type. It follows directly from the results in the previous section that if the change is either type of reshuffling or a move down, then his match remains unchanged with certainty (so also in expected terms). When, however, the man moves up a woman or several women, it affects his match in expected terms. If he only moves up truly less preferred women, it constitutes a strictly worse reply in expected terms. If he only moves up truly more preferred women, it constitutes a strictly better reply in expected terms. If the man moves several women out of which some are truly less and some are truly more preferred, the result is ambiguous.

However, if a man contemplates moving up a subset of women G such that some of them are truly more and some of them are truly less preferred, he is strictly better off in expected terms, to move up only the subset $G' \subset G$ where all women are truly more preferred.

This is because, by Lemma 2, when m moves up a (truly) less preferred woman he

may be matched to her. Since he does not know the preferences of other agents, and believes that any preference profile of other agents, P_{-m} , can occur with some positive probability, such a move-up makes him strictly worse off in expected terms. Conversely, moving up a truly more preferred woman makes m better off in expected terms. Therefore, it constitutes a better reply (in expected terms).

Proposition 2 *Suppose that other agents' preferences P_{-m} are fixed, but not known. Suppose that P'_m is a better reply (in expected terms) than P_m to P_{-m} . Then it must be that P'_m was created from P_m by moving up a set of truly more preferred women. I.e., there exists a set of women G such that for all $w \in G$, $\mu(m)P_m w$, $wP_m^* \mu(m)$ and $wP'_m \mu(m)$.*

Note that the relative order of women in S may change between P_m and P'_m and may not be related to their relative order in P_m^* .

Suppose first that we allow for only one woman's (relative) ranking in m 's preferences to change. Then moving up a more preferred woman is a necessary and sufficient condition for P'_m to be a better reply than P_m in expected terms.

However, when we allow for ranking of multiple women to change between P_m and P'_m , it is no longer the case. Moving up some (truly) more preferred women is still a necessary condition for a better reply. But it is no longer sufficient. When both (truly) less preferred and more preferred women are moved up, the result is ambiguous. This because any of those women could be matched to m under such P'_m , which could result in a better or in a worse match.

Suppose that we restrict that only truly more preferred women are moved up, and no less-preferred women. Then it is a sufficient condition for a better reply, but not a necessary one.

Nonetheless, from the necessary condition for a better reply stated in Proposition 2, we can derive a characterization of a best reply.

Corollary 1 *Suppose that other agents' preferences P_{-m} are fixed, but not known. P_m is a best reply to P_{-m} if and only if all women that are truly more preferred to $\mu(m)$ are also listed above $\mu(m)$ in P_m . I.e., $\forall w(wP_m^* \mu(m) \implies wP_m \mu(m))$.*

Corollary 1 follows directly from Proposition 2: The only way to create a better reply to P_{-m} than P_m is to move up truly preferred women. But if there are no truly preferred

women listed below $\mu(m)$, there does not exist a better reply to P_{-m} than P_m . Thus, P_m is a best reply.

Moreover, let $\mu^*(m) = \varphi(P_m^*, P_{-m})$. If P_m is a best-reply against P_{-m} then $\varphi(P_m, P_{-m}) = \varphi(P_m^*, P_{-m}) = \mu^*(m)$. That is, under any best-reply the outcome is the same as if m reported his true preferences.

3.4 Simultaneous Rankings Updating

In this section suppose that multiple men change their rankings. Clearly, if multiple men only reshuffle below, the matching does not change. It is also the case when multiple men reshuffle above (by Blocking Lemma).

Notice that if m only “moves down” and reshuffling then $\mu'(m) = \mu(m)$.

Lemma 3 *Let P and let $\mu = \varphi(P)$. Let $\widetilde{M} \subset M$ and P' be a profile such that for each man $m \in \widetilde{M}$, $U(P'_m, \mu(m)) \subseteq U(P_m, \mu(m))$, and $P'_m = P_m$ for each man $m \notin \widetilde{M}$. Let $\mu' = \varphi(P')$. Then $\mu' R_{\widetilde{M}} \mu$ and there is at least one man $m \in \widetilde{M}$ such that $\mu'(m) = \mu(m)$.*

Proof Notice first that $\mu \in S(P')$, for any blocking pair under P' is also a blocking pair under P and if $\mu \notin IR(P')$ then obviously $\mu \notin IR(P)$. So, $\mu' R_M \mu$ and thus $\mu' R_{\widetilde{M}} \mu$. If $\mu'(m) \neq \mu(m)$ for each $m \in \widetilde{M}$, then $\mu' P_{\widetilde{M}} \mu$. But this contradicts the fact that φ is group-strategyproof. ■

Lemma 4 *If the active man moves up a woman, passive men may be worse off, better off or unchanged (according to submitted preferences).*

Proof We prove the statement by constructing examples.

In the examples below all women are acceptable to all men, and all men are acceptable to all women. Moreover, all women have the same preferences (stated and true):

$$P_{w_i} : m_1, m_2, m_3 \quad \text{for } i = 1, 2, 3.$$

Example: Passive men are better off

Suppose that all men submit the same preferences

$$P_{m_i} : w_1, w_2, w_3 \quad \text{for } i = 1, 2, 3.$$

In the stable matching, according to P_{w_i} and P_{m_i} , $i = 1, 2, 3$, m_i is matched with w_i . Now, suppose that m_1 (and only him) changes his submitted preferences to $P'_{m_1} : w_3, w_2, w_1$, i.e., moves up w_3 . The stable matching according to the new preferences matches m_1 with w_3 , m_2 with w_1 , and m_3 with w_2 . Both passive men are better off according to their submitted preferences P_{m_2} and P_{m_3} .

Example: Passive men are worse off or unchanged Suppose that men submit following preferences

$$P_{m_1} : w_2, w_3, w_1$$

$$P_{m_2} : w_1, w_2, w_3$$

$$P_{m_3} : w_1, w_2, w_3.$$

The stable matching assigns m_1 and w_2 , m_2 and w_1 and w_3 and m_3 .

Now, suppose that m_1 (and only him) changes his submitted preferences to $P_{m_1} : w_1, w_2, w_3$, i.e., moves up w_1 . The stable matching under the new preferences assigns m_i with w_i for $i = 1, 2, 3$. Man m_2 is strictly worse off, while m_3 is matched to the same woman. ■

4 Repeated matching

From now on we fix men and women's preferences P . Hereafter, the profile P^* will be referred to the *true* preference profile.

The game we shall consider consists of a repeated matching game between men. Women are assumed to play truthfully. At each period $t \geq 1, \dots$, men have to submit a preference relation P_M^t over potential mates. For each period t the outcome given P_M^t is $\varphi(P_M^t, P_W)$. To avoid cumbersome notation we shall sometimes denote the matching $\varphi(P_M^t, P_W)$ simply by μ^t .

For now we shall consider better-reply dynamics only, but considering two cases. In one case, men can only better-reply by doing “move-ups”, and in the other case men can better-reply by doing “move-ups” and “move-downs”. The case when men do only “move-downs” is easily discarded. Indeed, consider any profile P^0 where all men list a different woman as their top choice. So, the matching μ^0 is such that each man is matched to his top choice and thus have no opportunity to make a “move-down” at $t = 1$.

We also discard “reshuffling” in the updating process. It will be clear from our analysis that adding reshuffling to the updating process will considerably weaken our convergence results (and thus add little to our analysis of better-reply dynamics).

Traditional learning models of better-reply dynamics consider situations in which at each period only one player updates his strategy. One of the main reasons for this assumption is that players are assumed to be boundedly rational and thus may not necessarily be able to compute a better-reply (let alone a best-reply). Better-reply dynamics then naturally emerge from a “trial-and-error” dynamic. At each period, one player has the opportunity to try a new strategy. If this strategy yields a higher payoff, i.e., it is a better-reply, then the player adopts this new strategy, and goes back to his previous strategy otherwise. Clearly, in matching game, such as school choice, the assumption that only one player updates his strategy at each period is difficult to sustain. One we then need to consider dynamics where more than one player can update his strategy at each period. For the moment, we shall focus however on the case where only one man can update his strategy, and discuss later in the section the case when several men can update their strategy at the same time.

Our first result is about the convergence when only move-ups are allowed.

Proposition 3 *Let P^0 be any preference profile and suppose that at each period only one man updates his strategy and only “move-ups” are allowed. Then there exists $t^* < \infty$ such that for all $t > t^*$, $\mu^t = \mu^{t^*}$ and $\mu^{t^*} \in S(P_M^*, P_W)$.*

Proposition 3 states two things. First, a move-up-only dynamic necessarily converge, and the limit outcome is a stable matching. This shows then that a simple dynamic process as this one is a first refinement of the the Nash equilibrium. Indeed, Haeringer and Klijn (2009) showed that the set of Nash equilibrium outcomes includes but nor may necessarily coincide with the set of stable matchings.

Proof We first show that the dynamics converge. To this end, suppose we have a better-reply cycle involving two or more men. let m be one of these men, and let $P_m^1, P_m^2, \dots, P_m^T$ (with $P_m^T = P_m^1$) be the preference orderings submitted by m that belong to the cycle.

So, for instance, P_m^2 is a better-reply of m with respect to P_m^1 against the profile P_{-m} extracted from the cycle. That is, the cycle consists of a sequence of profiles, P^1, P^2, \dots and there is some t such that $P^t = (P_m^1, P_{-m})$ and $P^{t+1} = (P_m^2, P_{-m})$.

Let v_k be the individual (a woman or m himself) that has been moved up between P_m^1 and P_m^2 , and let $v_h = \mu^1(m)$, i.e., m 's match. So,

$$v_k P_m^* v_h.$$

At some point (before T), say, t_1 , we should obtain again $v_h P_m^{t_1} v_k$. That is, v_h has been moved up. It is necessarily with respect to $\mu^{t_1}(m) = v_1$, so,

$$v_k P_m^* v_h P_m^* v_1.$$

Notice that at t_1 we have $v_1 P_m^{t_1} v_h$ and $v_h P_m^{t_1+1} v_1$. So, there is a step t_2 in the cycle where v_1 is moved up (above v_h). If v_1 is moved up it has to be above $\mu^{t_2}(m) = v_2$. So we have

$$v_k P_m^* v_h P_m^* v_1 P_m^* v_2.$$

It suffices to continue the argument (with individuals v_3, v_4 , etc.) to deduce that the only possibility to obtain after some periods the preference ordering P_m^1 is to have an infinite number of woman, which is impossible. So, the dynamic process necessarily converge to some profile P^{t^*} .

We now show that the limit outcome is necessarily stable with respect to the true preference profile (P_M^*, P_W) . Let $\mu = \mu^{t^*}$. Suppose first that there is a man m such that $m P_m^* \mu(m)$. Notice that we necessarily have $\mu(m) P_m^{t^*} m$, so man m can still update his strategy at t^* , a contradiction. Suppose now that there is a pair (m, w) that blocks μ , i.e., $m P_w \mu(w)$ and $w P_m^* \mu(m)$. So, we have $w P_m^{t^*} \mu(m)$. Hence, it must be that $\mu(w) P_w m$, a contradiction. So, $\mu \in S(P_M^*, P_W)$. ■

Notice that the convergence result of Proposition 3 is still valid when we consider only move-downs. However, as we commented before, the limit matchings may not necessarily be stable matchings. If we consider move ups and move downs, the dynamics properties differ substantially. First, as the following example shows, even if we consider that at each period only one man updates his submitted preferences, the dynamic may cycle.

Example 1 Consider three men and three women, $m_i, w_i, i = 1, 2, 3$. Let the women's preferences be

$$P_{w_1} : m_1, m_2, m_3$$

$$P_{w_2} : m_2, m_3, m_1$$

$$P_{w_3} : m_3, m_1, m_2$$

Consider the following preference lists for the men:

$$\begin{array}{ll}
U_{m_1} : w_3, w_2, w_1 & D_{m_1} : w_2, w_1, w_3 \\
U_{m_2} : w_1, w_3, w_2 & D_{m_2} : w_3, w_2, w_1 \\
U_{m_3} : w_2, w_1, w_3 & D_{m_3} : w_1, w_3, w_2
\end{array}$$

The following sequence of strategy profiles constitutes a cycle where at each step the man changing his strategy is better-replying to the current strategy profile. $(D_{m_1}, D_{m_2}, U_{m_3})$, $(U_{m_1}, D_{m_2}, U_{m_3})$, $(U_{m_1}, D_{m_2}, D_{m_3})$, $(U_{m_1}, U_{m_2}, D_{m_3})$, $(D_{m_1}, U_{m_2}, D_{m_3})$, $(D_{m_1}, U_{m_2}, U_{m_3})$, $(D_{m_1}, D_{m_2}, U_{m_3})$, \dots

Although better-reply dynamics may cycle, we can easily show, however, that whenever they do converge the limit matching is always the same.

Proposition 4 *Let P^0 be any preference profile and suppose that at each period men update their strategy by better-replying (i.e., move ups and downs are allowed). If there exists $t^* < \infty$ such that for all $t > t^*$, $\mu^t = \mu^{t^*}$ then $\mu^{t^*} = \mu_M^*$.*

Proof Suppose the dynamic converge to P_M and let μ be the limit matching—i.e., $\mu = \varphi(P, P_W)$. Let

$$\begin{aligned}
M^+ &= \{m \in M : \mu P_m^* \mu_M\}, \\
M^\circ &= \{m \in M : \mu(m) = \mu_M(m)\}, \\
M^- &= \{m \in M : \mu_M P_m^* \mu\}.
\end{aligned}$$

Suppose $M^+ \neq \emptyset$. By the blocking Lemma, there exists $m \notin M^+$ and $w \in \mu(M^+)$ such that $w P_m^* \mu(m)$ and $m P_w \mu(w)$. So, $w P_m \mu(m)$ —for otherwise m could update again his revealed preferences. It follows that m made an offer to w and this offer was rejected. Hence, $\mu(w) P_w m$, which contradicts $m P_w \mu(w)$. So, $M^+ = \emptyset$.

Suppose $M^- \neq \emptyset$. So, $\mu_M R_M \mu$, and $\mu_M(M^-) = \mu(M^-)$ —because $M^+ = \emptyset$. By the Blocking Lemma, there exists $m \in M^-$ and $w \in \mu(M)$ such that $m P_w \mu_M(w)$ and $w P_m \mu_M(m) = \mu(m)$. Since μ is the limit matching, $w P_m^* \mu_M(m)$. So, $\mu_M \notin S(P^*)$, a contradiction. Hence, $M^- \neq \emptyset$, and thus $M^\circ = M$, i.e., $\mu = \mu_M$. \blacksquare

We turn now to the case where several men update their submitted preferences at the same time. When only one man can update his submitted preferences at a time, a better-reply makes necessarily the man who updates his strategy weakly better off. In this case no man will want to revert to his previous strategy. If simultaneous updating occurs, if all the updating men end up worse off then they may want to revert to their previous strategy, which in turn may create a cycle. Here again we observe differences between the case when only move-ups are allowed and when both move-ups and move-downs are permitted.

The next result shows that move-ups only dynamics retain their convergence properties when simultaneous updating occurs, i.e., there is at least one man (among those who have updated) who will not wish to revert to the strategy used before updating.

Lemma 5 *Let $P'_M \neq P_M$, and let $\widetilde{M} = \{m \in M : P'_m \neq P_m\}$. Suppose that for each man $m \in \widetilde{M}$, P'_m is a move-up of P_m . Then there exists at least one man $m \in \widetilde{M}$ such that $\mu' R_m \mu$.*

Proof Let P_M and P'_M satisfy the conditions of the Lemma and suppose that for all men $m \in \widetilde{M}$, $\mu P_m^* \mu'$. For each man m such that $P_m \neq P'_m$, let C_m be the set of individuals that have been moved-up between P_m and P'_m .

Since DA is group-strategy proof for the men there exists a least one man $\widetilde{m} \in \widetilde{M}$ such that $\mu' R'_{\widetilde{m}} \mu$. Since for each man $m \in \widetilde{M}$, $\mu P_m^* \mu'$, $\mu'(m) \notin C_m$. So, $\mu' R'_{\widetilde{m}} \mu$ implies $\mu' P_{\widetilde{m}} \mu$. Hence, by the blocking lemma we have $\mu' \notin S(P)$.

Since $\mu' \notin S(P)$, either $\mu' \notin IR(P)$ or there is a pair (m, w) such that $w P_m \mu'(m)$ and $m P_w \mu'(w)$. Suppose first that $\mu' \notin IR(P)$, i.e., there exists an individual v such that $v P_v \mu'(v)$. Clearly, $v \in \widetilde{M}$, for otherwise we would have $\mu' \notin IR(P')$ —because $P'_v = P_v$ for $v \notin \widetilde{M}$. Since $\mu'(m) \notin C_m$, $m P_m \mu'$ implies $m P'_m \mu'$. Hence, $\mu' \notin IR(P')$, contradiction. Hence, there exists a blocking pair (m, w) . Again, $m \in \widetilde{M}$, for otherwise (m, w) would also block μ' under P' . Since $\mu'(m) \notin C_m$, $w P_m \mu'$ implies $w P'_m \mu'$, i.e., (m, w) also block μ' under P' , a contradiction. Hence, $\mu' \in S(P)$. This is a contradiction, the desired result. ■

As the following example shows, the previous result no longer holds under simultaneous updating when both move-ups and move-downs are allowed.

Example 2 Consider a market with four men and women, whose true preferences are depicted below.

$P_{m_1}^*$	$P_{m_2}^*$	$P_{m_3}^*$	$P_{m_4}^*$	P_{w_1}	P_{w_2}	P_{w_3}	P_{w_4}
w_3	w_1	w_2	w_3	m_3	m_1	m_2	m_4
w_4	w_2	w_3	w_1	m_2	m_4	m_3	m_3
w_1	w_3	w_1	w_2	m_1	m_3	m_1	m_1
w_2	w_4	w_4	w_4	m_4	m_2	m_4	m_2

Let P be the following preference profile. Men 1 and 4 submit their true preferences, man 2 submits the preference $P_{m_2} : w_4, w_3, w_1, w_2$ and man 3 submits the preference $P_{m_3} : w_3, w_2, w_1, w_4$. With the profile P men 2 and 3 are matched to women 3 and 2, respectively.

Observe that m_2 truly prefers w_1 and w_2 to his match, w_3 . So a better reply for m_2 is $P'_{m_2} : w_1, w_4, w_3, w_2$. As for m_3 , w_3 is less preferred than w_2 , so a better reply for m_3 is $P'_{m_3} : w_2, w_1, w_3, w_4$. Let $P' = (P_{m_1}^*, P'_{m_2}, P'_{m_3}, P_{m_4}^*)$. Under the profile P' men 2 and 3 are matched to women 4 and 1, respectively. That is, both men are strictly worse off under P' than under P .

Appendix

Lemma 6 (Blocking Lemma, Gale and Sotomayor, 1985) *Let μ be any individually rational matching and let W° be the set of all workers who prefer μ to μ_W . If W° is nonempty, there is a pair (f, w) satisfying $w \in W \setminus W^\circ$ and $f \in \mu(W^\circ)$ that blocks μ .*

Consider two preference orderings P_v and P'_v , and suppose that we wish to “reach” the preference P'_v starting from the preference P_v . Clearly, it may well be the case that we need to rearrange the position of many alternatives to connect P_v and P'_v . In these manipulations, some alternatives may be re-ranked higher while other alternatives may be re-ranked lower.

The next Lemma states that if P_v and P'_v are two preference orderings such that P_v 's Kemeny set (with respect to some preference ordering P_v^*) is a subset of P'_v 's Kemeny set then we can find a sequence of preference orderings such that starting from P'_v we can reach P_v in a finite number of steps where each step consists of moving up in the ordering just *one* alternative.

Lemma 7 *Let P_v, P'_v and P_v^* be such that $K(P_v, P_v^*) \subset K(P'_v, P_v^*)$. Then there exists a finite chain of preference orderings P_v^1, \dots, P_v^k such that*

$$(i) \ P_v^1 = P'_v, \ P_v^k = P_v;$$

$$(ii) \ K(P_v^h, P_v^*) \supset K(P_v^{h+1}, P_v^*), \ h = 1, \dots, k-1;$$

(iii) *For each $h, h = 1, \dots, k-1$, there is v such that for each $v', v'' \neq v, v' P_v^h v''$ if, and only if, $v' P_v^{h+1} v''$ and there is v' such that $v' P_v^h v$ and $v P_v^{h+1} v'$.*

Proof Let P_v, P'_v and P_v^* such that $K(P_v, P_v^*) \subset K(P'_v, P_v^*)$. Compare P_v and P'_v starting from the most preferred alternatives and go down in the preferences until there is an alternative \bar{v} in P_v that has not the same rank in P_v and P'_v . Let \hat{v} be the alternative that ranks just above \bar{v} in P_v . That is, for the most preferred alternatives up to \hat{v} the preference orderings P_v and P'_v coincide. Construct the profile P_v^2 from P'_v in the following manner. For each $v', v'' \neq \bar{v}, v' P_v^2 v''$ if, and only if, $v' P'_v v''$, and let $\hat{v} P_v^2 \bar{v}$ such that there is no v' for which $\hat{v} P_v^2 v' P_v^2 \bar{v}$, i.e., alternative \bar{v} is moved up just below \hat{v} .

We claim that $K(P_v, P_v^*) \subseteq K(P_v^2, P_v^*) \subseteq K(P'_v, P_v^*)$. To see this, observe that for each pair (v', v'') such that $\bar{v} \neq v', v'', (v', v'') \in K(P_v^2, P_v^*)$ if, and only if, $(v', v'') \in$

$K(P'_v, P_v^*)$. Also, for each pair (v', v'') such that $\bar{v} \in \{v', v''\}$, $(v', v'') \in K(P_v^2, P_v^*)$ if, and only if, $(v', v'') \in K(P_v, P_v^*)$. Suppose now that there exists a pair (v', v'') where $\bar{v} \neq v', v''$ such that $(v', v'') \notin K(P_v^2, P_v^*)$ yet $(v', v'') \in K(P_v, P_v^*)$. So, $(v', v'') \notin K(P'_v, P_v^*)$, which contradicts $K(P_v, P_v^*) \subset K(P'_v, P_v^*)$. Hence, $K(P_v, P_v^*) \subseteq K(P_v^2, P_v^*)$. Suppose now that there exists a pair (v', v'') where $\bar{v} \in \{v', v''\}$ such that $(v', v'') \notin K(P'_v, P_v^*)$ yet $(v', v'') \in K(P_v^2, P_v^*)$. This latter implies that $(v', v'') \in K(P_v, P_v^*)$, which again contradicts $K(P_v, P_v^*) \subset K(P'_v, P_v^*)$. Hence, $K(P_v^2, P_v^*) \subseteq K(P'_v, P_v^*)$. If $P_v^2 = P_v$ then we are done. Otherwise, construct P_v^3 from P_v^2 in the same way P_v^2 was constructed from P'_v , and keep doing so until we reach some P_v^k such that $P_v^k = P_v$. Since there is a finite number of alternatives this procedure eventually stops. \blacksquare

Theorem 3 *For any man $m \in M$, let P_m^* , P_m and P'_m be preference relations over $W \cup \{m\}$. Then for any P_{-m} , $\varphi(P_m, P_{-m})R_m^* \varphi(P'_m, P_{-m})$ if, and only if, $K(P_m, P_m^*) \subset K(P'_m, P_m^*)$.*

Proof Let P_m and P'_m be such that for any P_{-m} , $\varphi(P_m, P_{-m})R_m^* \varphi(P'_m, P_{-m})$. So, $P_m \neq P'_m$, and thus there exists at least one pair v, v' such that vP_mv' and $v'P'_mv$. Notice that all such pairs (v, v') necessarily belong to either $K(P_m, P_m^*)$ or $K(P'_m, P_m^*)$.¹³ Consider any such pair (v, v') . So it suffices to show that $(v, v') \in K(P'_m, P_m^*)$. To begin with, we claim that there exists a profile P_{-m} such that $\varphi(P_m, P_{-m})(m) = v$ and $\varphi(P'_m, P_{-m})(m) = v'$. An example of such a profile is the following. For each \hat{w} such that either $\hat{w}P_mv$ or $\hat{w}P'_mv'$ pick a man \hat{m} such that \hat{w} is \hat{m} 's most preferred partner in $P_{\hat{m}}$, and \hat{m} is \hat{w} 's most preferred partner in $P_{\hat{w}}$. Let m be v and v' 's most favourite partner according to P_v and $P_{v'}$, respectively. It is easy to see that for any profile P_{-m} following these specifications we have $\varphi(P_m, P_{-m})(m) = v$ and $\varphi(P'_m, P_{-m})(m) = v'$. Since P_m dominates P'_m , it follows that vP'_mv' , and thus $(v, v') \notin K(P_m, P_m^*)$ and $(v, v') \in K(P'_m, P_m^*)$, the desired result.

Let P_m and P'_m be such that $K(P_m, P_m^*) \subset K(P'_m, P_m^*)$. Using Lemma 7 it suffices to assume that P and P' differ only by one alternative, say v . That is, for each $v', v'' \neq v$, $v'Pv''$ if, and only if $v'P'v''$, and v ranks higher in P than in P' . Consider any profile P_{-m} . Observe that if alternative v is ranked below alternative $\varphi(P'_m, P_{-m})(m)$ in

¹³Indeed, if $(v, v') \in K(P_m, P_m^*) \cap K(P'_m, P_m^*)$ or $(v, v') \notin K(P_m, P_m^*) \cup K(P'_m, P_m^*)$, then vP_mv' if, and only if, vP'_mv' .

both P_m and P'_m then $\varphi(P'_m, P_{-m})$ and $\varphi(P_m, P_{-m})$ coincide. So, in this case we obviously have $\varphi(P_m, P_{-m})R_m^*\varphi(P'_m, P_{-m})$. Similarly, if v is ranked above $\varphi(P'_m, P_{-m})(m)$ in P'_m , we also have $\varphi(P'_m, P_{-m}) = \varphi(P_m, P_{-m})$ and thus $\varphi(P_m, P_{-m})R_m^*\varphi(P'_m, P_{-m})$. In other words, the only possibility for m to change his match between P_m and P'_m is when $\varphi(P'_m, P_{-m})(m)$ is ranked above v in P'_m but ranked below v in P_m . We claim that in this case $\varphi(P_m, P_{-m})(m) \in \{\varphi(P'_m, P_{-m})(m), v\}$. To see this, suppose that $\varphi(P_m, P_{-m})(m) \notin \{\varphi(P'_m, P_{-m})(m), v\}$. Notice first that by strategy-proofness, $\varphi(P_m, P_{-m})R_m\varphi(P'_m, P_{-m})$, and thus $\varphi(P_m, P_{-m})P_m\varphi(P'_m, P_{-m})$ —because $\varphi(P_m, P_{-m})(m) \neq \varphi(P'_m, P_{-m})(m)$. Since for each $v', v'' \neq v$, $v'P_mv''$ if, and only if, $v'P'_mv''$, $\varphi(P_m, P_{-m})(m) \neq v$ implies

$$\varphi(P_m, P_{-m})P'_m\varphi(P'_m, P_{-m}).$$

This contradicts the strategy-proofness of φ , so $\varphi(P_m, P_{-m})(m) \in \{\varphi(P'_m, P_{-m})(m), v\}$. If $\varphi(P_m, P_{-m})(m) = \varphi(P'_m, P_{-m})(m)$, then we obviously have $\varphi(P_m, P_{-m})R_m^*\varphi(P'_m, P_{-m})$. Suppose then that $\varphi(P_m, P_{-m})(m) = v$. Recall that $vP_m\varphi(P'_m, P_{-m})$ and $\varphi(P'_m, P_{-m})P'_mv$, and $K(P_m, P_m^*) \subset K(P'_m, P_m^*)$. So, $vP_m^*\varphi(P'_m, P_{-m})v$. It follows that $\varphi(P_m, P_{-m})(m)P_m^*\varphi(P'_m, P_{-m})$, the desired result. ■

References

- [1] Abdulkadiroğlu, A. and T. Sönmez (2003) “School Choice: A Mechanism Design Approach,” *American Economic Review*, **93**, 729–747.
- [2] Balcan, M.-F., Blum, A.; Hartline, J.D. and Y. Mansour (2005) “Mechanism design via machine learning.” *Proceedings of the 2005 46th Annual IEEE Symposium on Foundations of Computer Science (FOCS’ 05)*, 605–614.
- [3] Dubins, L.E. and D.A. Freedman (1981) “Machiavelli and the Gale-Shapley Algorithm,” *American Mathematical Monthly*, **88**, 485–494.
- [4] Kemeny J. G. (1959) “Mathematics without numbers,” *Daedalus*, **4**, 577–591.
- [5] Gale, D. and L.S. Shapley (1962) “College Admissions and the Stability of Marriage,” *American Mathematical Monthly*, **69**, 9–15.
- [6] Iehlé, V. and G. Haeringer (2009) “Enjeux stratégiques du concours de recrutement des enseignants-chercheurs,” *Revue Économique*, **61**, 697–721.
- [7] Haeringer, G. and F. Klijn (2009) “Constrained School Choice.” *Journal of Economic Theory*, **144**, 1921–1947.
- [8] Healy, J. H (2006) “Learning dynamics for mechanism design: An experimental comparison of public goods mechanisms.” *Journal of Economic Theory*, **129**, 114–149.
- [9] McKinney, C. N., M. Niederle and A. E. Roth (2005) “The collapse of a medical labor clearinghouse (and why such failures are rare),” *American Economic Review*, **95**, 878–889.
- [10] Pathak, P. A. and T. Sönmez (2008) “Comparing Mechanisms by their Vulnerability to Manipulation” *mimeo*.
- [11] Roth, A.E. (1982b) “The Economics of Matching: Stability and Efficiency,” *Mathematics of Operations Research*, **92**, 617–628.

- [12] Roth, A.E. (1991) “A Natural Experiment in the Organization of Entry Level Labor Markets: Regional Markets for New Physicians and Surgeons in the UK,” *American Economic Review*, **81**, 415–440.
- [13] Roth, A. E. (2008) “What have we learned from Market Design?,” *Economic Journal* **118**, 285–310.
- [14] Roth, A. E. and I. Erev (1995) “Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term,” *Games and Economic Behavior*, **8**, 164–212.
- [15] Roth, A. E. and E. Peranson (1999) “The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design,” *American Economic Review* **89**, 748–780.
- [16] Roth, A. E. and M. A. O. Sotomayor (1990) *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Econometric Society Monograph Series. New York: Cambridge University Press.