

CROSSCUTTING AREAS

Stability in Large Matching Markets with Complementarities

Itai Ashlagi

Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Massachusetts 02142, iashlagi@mit.edu

Mark Braverman

Computer Science, Princeton University, Princeton, New Jersey 08544, mbraverm@cs.princeton.edu

Avinatan Hassidim

Computer Science, Bar-Ilan University, Ramat Gan, 52900, Israel, avinatanh@gmail.com

Labor markets can often be viewed as many-to-one matching markets. It is well known that if complementarities are present in such markets, a stable matching may not exist. We study large random matching markets with couples. We introduce a new matching algorithm and show that if the number of couples grows slower than the size of the market, a stable matching will be found with high probability. If however, the number of couples grows at a linear rate, with constant probability (not depending on the market size), no stable matching exists. Our results explain data from the market for psychology interns.

Subject classifications: stability; matching; couples; deferred acceptance; market design.

Area of review: Games, Information, and Networks.

History: Received July 2012; revisions received March 2013; accepted February 2014. Published online in *Articles in Advance* June 13, 2014.

1. Introduction

School assignments, college admissions, and labor markets can be viewed as two-sided matching markets. Stability has been shown to be an important property of two-sided matching markets, and as long as agents' preferences are substitutable, a stable matching always exists (see, e.g., Hatfield and Milgrom 2005). Agents' preferences, however, often have complementarities, and understanding when a stable matching exists, as well as finding one when it exists, are important market design issues.

In some labor markets, such as the National Resident Matching Program (NRMP), couples participate and naturally introduce complementarities by searching for a pair of jobs. In this paper we characterize the (non)existence of a stable matching in large random markets with couples. Furthermore, the results can be extended, as shown, to more general matching markets. The existing results are supported by data from the U.S. market for psychology interns.

The NRMP and the market for psychology interns are two examples for markets with couples that use clearinghouses to perform the matching. Approximately 16,000 doctors participated in the NRMP in 2010 as singles¹ and approximately 1,600 doctors participated as part of a couple (800 couples). In the same year, approximately 3,000 psychologists participated in the psychology clearinghouse as singles, and 38 participated as couples (19 couples). Until 1983 doctors that were part of a couple had to participate as singles, since clearinghouses for these markets used the deferred acceptance algorithm to find a matching. These two markets

as well as many other markets currently use the algorithm designed by Roth and Peranson (1999) (henceforth called RP) that allows couples to express their joint preferences. This algorithm has had great success in practice: every year since it has been used, the NRMP has found a stable matching with respect to the reported preferences. For a comprehensive background and history of these markets see Kojima et al. (2013), Roth (2009).

We consider many-to-one matching markets, in which one side of the market consists of hospitals that have responsive preferences² and the other side consists of doctors. Gale and Shapley (1962) introduced the well-known deferred acceptance algorithm, and showed that if every doctor is *single* (and in particular demands a single hospital), the algorithm will always find a stable matching. Naturally, when couples are present in the market, they may view pairs of hospitals as complements, and a stable matching may not exist (Roth 1984, Klaus and Klijn 2005).³

The Gale-Shapley algorithm simulates an application process by having doctors applying to hospitals according to their preference lists. Each doctor, when applying, “influences” other doctors through “rejection chains” (in order to accept a doctor, a hospital might reject some other doctor who will apply to a different hospital who might reject some other doctor to accept the applying doctor and so forth). When only singles are present, rejection chains never cycle because of the singles' substitutable preferences, and therefore the Gale-Shapley ends in a stable matching. The main difficulty that rises with the presence of couples is

that rejection chains may cycle. To understand when such cycles evolve in large markets, we introduce a notion of an “influence tree” for a couple, which roughly captures the set of possible doctors the couple may influence by entering the market. This allows us to investigate the way complementarities may interfere with each other (causing cycles) thus precluding a stable matching from being found.

We introduce a new matching algorithm called *sorted deferred acceptance* (SoDA). We show that when the number of couples is growing in a near-linear rate, the probability that a stable matching exists, and SoDA will find one, converges to one as the market grows.⁴ This result is tight: we show that if the number of couples grows in a linear rate, with constant probability *no* stable matching exists.

Our positive result is supported by real data from the market for psychology interns, in which we demonstrate not only that SoDA performs well, but also that influence trees can be used to predict whether an instance will have a stable outcome or not. According to the data, it is often the case that couples are influenced by other couples, but cycles of such influences do not exist. We also simulate the behavior of SoDA in various large random markets based on real data.⁵

When the number of couples grows in a near-linear rate, we also show that truth telling is an approximated Bayes-Nash equilibrium in the game induced by the SoDA algorithm for any large enough n , and further that the ex ante probability that a single doctor or a couple will obtain its best stable matching tends to one as n tends to infinity. In particular, we show that most agents do not have multiple stable outcomes. Finally, we show how the results generalize to many-to-one matching markets similar to the one in Roth and Sotomayor (1990) with workers and firms (limiting the number of firms that have nonsubstitutable preferences).⁶

SoDA, which runs in almost linear time, is the first matching algorithm for matching markets with couples that is proven to find a stable outcome in very general settings (it can also be adapted to more general worker-firm settings). This helps explain the fact that algorithms, RP in particular, have been successful in finding stable matchings in real life.⁷

1.1. Related Work

The closest work is by Kojima et al. (2013) who initiated the study of existence of stable matchings in large matching markets with couples. They showed that if the number of couples is asymptotically smaller than \sqrt{n} , where n is the number of singles, a stable matching exists with high probability (they also came with a simpler algorithm that also extends deferred acceptance for finding a stable matching in such markets). Although this might resemble our positive result, there is a striking difference. Using our language, they essentially show that no couple influences any other couple with high probability. In other words, imagine that Alice, who is part of a couple, applies to Mt. Sinai hospital (under their algorithm), which has one position, and Mt. Sinai temporarily assigns Alice (till they get a better application if at all). Under their assumptions it follows that with high probability Mt. Sinai will never get an application from an

applicant they like better than Alice if that applicant is either a member of a couple or a doctor that is part of a rejection chain beginning from a member of another couple (who got temporarily assigned in some other hospital). Eliminating such influences is too strong and is indeed violated in data. In particular, cyclical influences should be avoided rather than any one way influences.

In their model doctors have short preference lists, which essentially implies an excess number of positions. Instead, we assume an excess number of available positions⁸ and further do not limit a priori the length of preference lists.⁹

One can view the results in this paper as a characterization for the existence of a stable matching in large random matching markets. The positive results, under which the fraction of couples grows at a sublinear rate, add to a long list of works that adopt large random markets to show that various desired economic properties hold with high probability in matching and auction settings; Immorlica and Mahdian (2005) and Kojima and Pathak (2009) studied random growing matching markets to show that the lack of strategy-proofness in small markets vanishes in large random markets (see also Lee 2011). Che and Kojima (2010) and Kojima and Manea (2010) showed that in one-sided matching markets, under the mechanism induced by probabilistic serial dictatorship the desire to manipulate vanishes. Roberts and Postlewaite (1976) and Jackson and Manelli (1997) showed similar results in large exchange economies. Rustichini et al. (1994) and Cripps and Swinkels (2006) show that the lack of efficiency is small in large double auctions (see also Jackson and Swinkels 2005 and Satterthwaite and Williams 1989). All these works assumed, however, that agents' preferences have no complementarities. When the fraction of couples is assumed to grow in a linear rate, our result implies that the issues caused by complementarities do not vanish with the size and randomness of the market.

Finally, this paper adds to a short but growing literature for positive results in settings in which agents' preferences may have complementarities. See, e.g., Milgrom (2004), Gul and Stacchetti (1999), Sun and Yang (2006), and Lahaie and Parkes (2009) for auction settings, and Hatfield and Kominers (2014) and Pycia (2012) for matching settings. We believe the notion of influence trees introduced in this paper is of independent interest for analyzing other settings with complementarities.

2. Matching Markets with Couples

2.1. The Model

In a matching market there is a set of hospitals H , a set of single doctors S , and a set of couples of doctors C . Each single doctor $s \in S$ has a strict preference relation \succ_s over $H \cup \{\phi\}$, where ϕ denotes an outside option for doctors. If $h \succ_s \phi$, we say that hospital h is *acceptable* for s . Each couple $c \in C$ denoted by $c = (f, m)$ has a strict preference relation \succ_c over $(H \cup \{\phi\}) \times (H \cup \{\phi\})$, i.e., over pairs of hospitals including the outside option. If $(h, h') \in H \cup \{\phi\} \times H \cup \{\phi\}$,

$(h, h') \succ_s (\phi, \phi)$ and $(h, h') \neq (\phi, \phi)$ we say that the pair (h, h') is *acceptable* for c . For every couple c we denote by f_c and m_c the first and second members of c . Denote by D the set of all doctors. That is, $D = S \cup \{m_c \mid c \in C\} \cup \{f_c \mid c \in C\}$. Each hospital $h \in H$ has a fixed capacity $k_h > 0$ and a strict preference relation \succ_h over the set D . For any set $D' \subseteq D$ hospital h 's choice given D' , $\text{Ch}_h(D')$, i.e., the most preferred doctors h can employ from the set D' , is induced by \succ_h and k_h as follows: $d \in \text{Ch}_h(D')$ if and only if there exists no set $D'' \subseteq D' \setminus \{d\}$ such that $|D''| = k_h$ and $d' \succ_h d$ for all $d' \in D''$. Note that we assume that all doctors are acceptable for each hospital.¹⁰

A matching μ is a function from $H \cup C \cup S$ such that $\mu(s) \in H \cup \{\phi\}$ for every $s \in S$, $\mu(c) \in (H \cup \{\phi\}) \times (H \cup \{\phi\})$ for every $c \in C$, $\mu(h) \in 2^D$ for every $h \in H$, and

- (i) $s \in \mu(h)$ if and only if $\mu(s) = h$;
- (ii) $\mu(c) = (h, h')$ if and only if $f_c \in \mu(h)$ and $m_c \in \mu(h')$;
- (iii) $\mu(c) = (h, \phi)$ if and only if $f_c \in \mu(h)$ and $m_c \notin \mu(h')$ for all $h' \in H$;
- (iv) $\mu(c) = (\phi, h)$ if and only if $m_c \in \mu(h)$ and $f_c \notin \mu(h')$ for all $h' \in H$.

Here, $\mu(s) = \phi$ means that s is unassigned under μ , and similarly $\mu(c) = (\phi, \phi)$ means that the couple c is unassigned.

We proceed to define stability. It will be convenient to define $\text{Ch}_\phi(D) = D$ for every set of doctors D (capturing that the outside option has sufficient capacity for all doctors). Blocking coalitions for a given matching can be formed in several ways:

- $(s, h) \in S \times (H \cup \{\phi\})$ is a *block* of μ if $h \succ_s \mu(s)$ and $s \in \text{Ch}_h(\mu(h) \cup \{s\})$;
- $(c, h, h') \in C \times (H \cup \{\phi\}) \times (H \cup \{\phi\})$ (where $h \neq h'$) is a *block* of μ if $(h, h') \succ_c \mu(c)$, $f_c \in \text{Ch}_h(\mu(h) \cup \{f_c\})$, and $m_c \in \text{Ch}_{h'}(\mu(h') \cup \{m_c\})$;
- $(c, h) \in C \times (H \cup \{\phi\})$ is a *block* of μ if $(h, h) \succ_c \mu(c)$ and $\{f_c, m_c\} \subseteq \text{Ch}_h(\mu(h) \cup c)$.¹¹

Finally a matching is *stable* if there is no block of μ .

Gale and Shapley (1962) showed that the (doctor proposing) the deferred acceptance algorithm described below always produces a stable matching in a matching market without couples. They further showed that the stable matching produced by this algorithm is the one that is weakly preferred by all single doctors. Roth (1982) showed that the mechanism induced by this algorithm makes it a dominant strategy for all single doctors to report their true preferences.

Doctor-Proposing Deferred Acceptance Algorithm (DA):

Input: A matching market (H, S, \succ_H, \succ_S) without couples.

Step 1. Each single doctor $s \in S$ applies to her most preferred hospital. Each hospital rejects its least preferred doctor in excess of its capacity among those who applied to it, keeping the rest of the doctors temporarily.

Step t. Each doctor who was rejected in Step $t - 1$ applies to her next highest choice if such exists. Each hospital

considers these doctors as well as the doctors who are temporarily held from the previous step, and rejects the least-preferred doctors in excess of its capacity keeping the rest of the doctors temporarily.

The algorithm terminates at a step where no doctor is rejected.

In the next section we introduce a new algorithm for finding a matching in a market with couples. Roth (1984) showed that when there are couples, sometimes a stable matching does not exist. In §4 we show that as the size of the market grows, for (essentially any) sublinear growth of the number of couples the probability that the algorithm produces a stable matching approaches one.

2.2. A New Matching Algorithm

The matching algorithm presented here first finds the stable matching in the market without couples (using DA) and then attempts to insert the couples, while maintaining the deferred acceptance idea of letting rejected singles further propose according to their preference lists. An informal description of the *sorted deferred acceptance algorithm* is provided below, and the formal definition is deferred to Appendix A.

Sorted Deferred Acceptance Algorithm (SoDA):

Step 1. Find a stable matching in the submarket without couples using the DA algorithm.

Step 2. Fix an order π over the couples. In the order π , each couple c applies to pairs of hospitals according to its preference list \succ_c (beginning with the most preferred) and once it finds a pair of hospitals that accepts it, we assign the couple to the pair of hospitals and stabilize the current matching as follows:

Step 3 (Stabilize). Continue the DA algorithm, with the doctors that were rejected from the their positions in the pair of hospitals that the last couple c was assigned to (at most two doctors).

(i) If during stabilizing one of the members the couple c (the last to apply) gets rejected the algorithm fails.

(ii) If some other couple $c' \neq c$ gets rejected during stabilizing, the order π is changed so that c will move and apply one place ahead of c' , and Step 2 begins again with the altered permutation; if however the new order π' has been attempted previously (at Step 2) the algorithm fails.

Note that if the algorithm terminates without failure it produces a stable matching.

Kojima et al. (2013) used a similar algorithm but allowed couples to apply in only one order, i.e., if some couple is evicted their algorithm fails, even though there might be a different order of couples' applications that will not lead to such a failure. This is too conservative as naturally couples are expected to get evicted in the application process. In SoDA, if some couple has been evicted the algorithm allows couples to apply again using a different ordering. This difference is shown in real data in §6.1.

Observe that the SoDA algorithm fails to produce a matching in two cases: first, if a couple c that finds a pair of positions causes a “rejection chain” leading to a member of the couple c being rejected (Step 3(i)),¹² and second, it attempts to let couples apply in an order that has already been attempted before (Step 3(ii)).¹³ Again, if the algorithm does not fail, the matching produced is stable.

The following definition will be useful throughout the paper.

DEFINITION 1 (EVICTING). Let $d \in D$ be a doctor and suppose that d is (temporarily) assigned to some hospital h . Let $c \in C$. If during the execution of the SoDA algorithm some member of the couple c who is not assigned to h applies to h and causes d to be rejected by h , we say that d was evicted by c . Furthermore, if d was evicted by c , applies to some hospital h' and causes some other doctor d' who is assigned to h' to be rejected, we also say that d' is evicted by c , and so forth. Finally, if d was evicted by c and d belongs to a couple c' we say that c' was evicted by c .¹⁴

REMARK 1. According to this definition c can evict itself. Such a phenomenon may occur since one member of a given couple can evict the other member of the couple (in the algorithm this happens in Step 3(i)).

3. A Large Market Model

A *random market* is a tuple $\Gamma = (H, S, C, \succeq_H, Z, Q)$, where $Z = (z_h)_{h \in H \cup \{\phi\}}$ and $Q = (q_h)_{h \in H \cup \{\phi\}}$ are probability distributions over H .

The preference list of each single doctor $d \in S$ is independently drawn as follows: for each $k = 1, \dots, |H|$ given d 's preference list up to her k th most preferred hospital, draw a hospital h independently according to Z until h does not appear in d 's k most preferred hospitals and let it be d 's $(k + 1)$ th most preferred hospital. The preference list for each couple $c = (f, m)$ is drawn from the distribution $Q \times Q$.¹⁵

We will consider a *sequence of random markets* $\Gamma^1, \Gamma^2, \dots$, where $\Gamma^n = (H^n, S^n, C^n, \succeq_H^n, Z^n, Q^n)$. The sequence of markets will have growing size according to the following definition.

DEFINITION 2. A sequence of random markets $\Gamma^1, \Gamma^2, \dots$ is called *regular* if there exist $0 < \epsilon < 1$, $\lambda > 1$, $\bar{k} > 0$, and $\sigma \geq 1$ such that for all n

1. $|S^n| = n$ and $|C^n| = O(n^{1-\epsilon})$ (the number of couples grows almost linearly);
2. for each hospital $h \in H^n$, $k_h < \bar{k}$ (bounded capacity);
3. $\sum_{h \in H^n} k_h \geq \lambda n$ (excess number of positions);
4. $q_h/q_{h'} \in [1/\sigma, \sigma]$ and $z_h/z_{h'} \in [1/\sigma, \sigma]$ for every $h, h' \in H^n$ (uniformly bounded preferences).

Importantly our results are true even if ϵ is a “slow” decreasing function of n converging to zero. The exact rate is discussed in the last section. It will be useful to define $\gamma_{\max} = \max_{h \in H} \max(q_h, z_h)$.

In their model, Kojima et al. (2013) assumed that each doctor’s preference list is bounded by a constant, i.e., the preference list contains a constant number of acceptable hospitals, whereas in our setting preference lists can be of any length. A key step in their proof is to show that the number of unfilled positions grows *linearly* in n with high probability. Instead, we make a weaker assumption by starting with a linear excess number of positions.

4. Stability

In this section we show the following:

THEOREM 1. Let $\Gamma^1, \Gamma^2, \dots$ be a regular sequence of random markets. Then the probability that there exists a stable matching tends to 1 as n goes to infinity.

We defer the formal proof to Appendix A.2 and give in the next section an intuition and a brief outline of the proof. In the proof we show that for random doctors preferences the probability that the SoDA algorithm (outlined in §2.2) yields a stable matching converges to 1 as n goes to infinity. In Appendix A.2. we also show that the SoDA algorithm runs in near linear time. Furthermore, the rate at which the probability converges to one is approximately $1/n^\epsilon$.

REMARK 2. Kojima et al. (2013) proved a similar result allowing less than \sqrt{n} couples in the market. The difference in the growth rate in both results is not cosmetic; they do not allow couples to appear in any rejection chain that other couples cause. In other words, if a couple has been rejected at some point in the application process their algorithm fails, even though some other order of couples application may still result in a stable matching. By definition, SoDA and our analysis allow such rejection chains as long as they do not cycle (see the following section for further discussion).

REMARK 3. Although the growth rate is not linear, the correctness of the proof of Theorem 1 is only a lower bound on the performance of the algorithm, and it may perform even better in practice. Technically, the proof will show that the random market has a stable matching with probability at least $1 - (\log n)^{O(1/\epsilon)} / n^{\Omega(\epsilon)}$, which converges to 1 even if $\epsilon = \Omega(\log \log n / \sqrt{\log n})$, and not just when ϵ is constant.¹⁶ This means that the algorithm finds a stable outcome with probability approaching 1 even when the number of couples grows like $n/2^{\sqrt{\log n \cdot \log \log n}}$. Such growth is close to linear. Empirically it is indeed hard to distinguish between such subpolynomial factors and constant factors when there are $n = 16,000$ doctors.

It is well known that in a matching market without couples, under the doctors proposing deferred acceptance algorithm, each doctor obtains his best stable match, i.e., for no doctor there exists a stable matching in which he obtains a better hospital than under DA.¹⁷ When couples are present one can show that doctors and couples will not always obtain their best stable matches. However, at least ex ante this holds in

a large market, a corollary that follows from the proof of Theorem 1:

COROLLARY 1. *Let $\Gamma^1, \Gamma^2, \dots$ be a regular sequence of random markets. The ex ante probability that a doctor or a couple obtain their best stable match tends to 1 as n tends to infinity.*

A proof sketch is given in §A.3.

4.1. Intuition and Proof Outline for Theorem 1

The goal is to show that if the number of couples is $m = n^{1-\epsilon}$ (for any $0 < \epsilon < 1$) then as n approaches infinity the probability a stable matching exists approaches 1. To better understand our approach we begin with the intuition for why the result holds for any $\epsilon > \frac{1}{2}$ (essentially this will provide the intuition for the result by Kojima et al. 2013), and then for any $\epsilon < 1$.

1. *Number of couples is $n^{1/2-\delta}$.* Consider the following simplified version of the SoDA algorithm, which we call the *direct algorithm*: after finding the stable matching in the market without couples, the couples apply one by one and if some couple evicts another couple (directly or indirectly as defined in §2.2) the algorithm fails (i.e., it does not attempt to change the permutation over the couples). Observe that if the algorithm does not fail, it outputs a stable matching. Note that the direct algorithm is equivalent to the algorithm presented by Kojima et al. (2013), in which couples apply one by one after a stable matching was found in the submarket without couples, and only then all singles that got evicted can apply.

When the first couple applies it can evict itself. When the second couple applies it can evict itself or the first couple, and so on. We bound the probability that a couple evicts some other couple. By way of example, we analyze the probability that the second couple c evicts the first couple.

The second couple c creates a rejection chain, which can cause several doctors who were temporarily assigned to continue applying. To bound the length of this chain consider f_c . At some point she is temporarily assigned to a hospital h . If h had a vacant position she did not evict any doctor and therefore also no other couple and we are done. Since there are more positions than doctors, the probability that the hospital has a vacancy is $1 - 1/\lambda$ (for simplicity we assume here that each hospital has capacity one and the preference distributions are uniform). If the hospital has no vacancy, she evicts a doctor d_1 who enters some hospital h_1 . If h_1 has a vacancy, we are done. If h_1 is full, a doctor d_2 gets evicted, and looks for a new position. Say d_2 is assigned to h_2 . Again, h_2 can have a vacancy, or be full, and this goes onward. However, since at every step of the chain there is a constant probability for a vacancy, one can show that with probability $1 - 1/n^3$ the number of hospitals h, h_1, h_2, \dots in the chain is upper bounded by $3\lambda \log n / (\lambda - 1)$.

Now, we can estimate the probability that the second couple evicts the first or itself. The second couple evicts doctors from at most $6\lambda \log n / (\lambda - 1)$ hospitals. If this

list includes the hospitals that admitted the first couple, or occupied by one of the spouses in c , we could be in trouble. But since preferences are random, the chances that the second couple influences any of these hospitals are upper bounded by

$$4 \cdot \frac{6\lambda \log n}{(\lambda - 1)n} = \frac{24\lambda \log n}{(\lambda - 1)n}.$$

What about the third couple? Again, it influences at most $6\lambda \log n / (\lambda - 1)$ hospitals. But now there are six hospitals that must not be influenced: four hospitals (at most) for each previously assigned couple and two hospitals for the third couple. Generalizing this for the k th couple and summing the probabilities we get

$$\sum_{k=1}^m \frac{12\lambda k \log n}{(\lambda - 1)n} < \frac{12\lambda m^2 \log n}{(\lambda - 1)n} = O\left(\frac{\log n}{n^{2\delta}}\right),$$

which goes to zero as n goes to infinity.

Note that if $m = \sqrt{n}$ this argument would not hold. In fact one can show that the direct algorithm fails with high probability if the number of couples is a large multiple of \sqrt{n} (we shall see that this is also supported by data and simulation results).

The direct algorithm attempts to insert the couples according to a single permutation. A natural attempt to find a stable matching when more couples are in the market is to change the permutation each time a couple evicts another couple. Suppose for example that couple $c_1 = (m_1, f_1)$ applies first and is temporarily assigned to (h_1, h_2) . Then couple $c_2 = (m_2, f_2)$ applies to (h_1, h_3) and moreover, suppose h_1 prefers m_2 over m_1 . In this case, c_1 would be evicted, and the direct algorithm would fail. However, if c_2 applies before c_1 , the algorithm would temporarily assign (m_2, f_2) to (h_1, h_3) , and when c_1 applies, h_1 would not accept m_1 , so c_1 would not be temporarily assigned to (h_1, h_2) , and just keep applying according to its list of preferences. This simple example motivates changing the insertion order of couples when some couple gets evicted, with the evicting couple coming before the evicted one.

2. *Number of couples is $n^{1-\epsilon}$ (sketch of proof of Theorem 1).*

The SoDA algorithm attempts to find an ordering of the couples, such that if couples apply one by one according to this order, no couple gets evicted by another couple. Whether or not a couple c evicts another couple c' depends on the (current) matching and the preference profile. Identifying worst-case scenarios, such as where c could “possibly” evict c' if there exists a configuration in which this happens, are too weak to prove our result. Instead, we devise a notion of whether c is “likely” to evict c' , and use this notion to analyze the algorithm. To do so we define for each couple c an *influence tree*; roughly speaking the influence tree of c consists of the hospitals and doctors who are likely to be part of rejection chains because of the presence of c . In

other words, the influence tree of c is supposed to capture all the dynamics that are likely to be caused by c 's assignment.

We want to show that there are not “many” influence tree intersections, since an intersection implies that two couples might influence the same hospital, and more importantly might evict each other. A first key step in this direction is the following:

(i) *With high probability each influence tree is very small (compared to n).*

If influence trees had not intersected each other, one could have shown that any insertion order of the couples would yield a stable matching with high probability. Essentially Kojima et al. (2013) showed that if $\epsilon > 0.5$ then the probability that no two influence trees intersect approaches 1 as $n \rightarrow \infty$. This however is not the case for all $\epsilon < 1$.

Influence trees, their intersections, and hospital preferences induce a useful structure in the form of a directed graph, which we call the *couples graph*; Informally speaking, in the couples graph each couple is a node, and there is a directed edge from couple c to another couple c' if their influence trees intersect at some hospital h and c can possibly evict some doctor that caused h to be in the influence tree of c' (the doctor can be a member of the couple c'). We will show that the couples graph is sparse:

(ii) *With high probability all weakly connected components in the couples graph are small.*¹⁸ Furthermore:

(iii) *With high probability there are no directed cycles in the couples graph.*

Recall that an influence tree for one couple does not involve other couples. In the next step we verify that influence trees are indeed the right structure:

(iv) *With high probability if in the algorithm a couple c influences a hospital h under any ordering π over the couples, then that hospital will also belong to the influence tree of c .*

Finally, by (iii) and (iv), if one can find a topological sort π in the couples graph,¹⁹ then letting couples apply one by one according to π yields a stable matching.

In the language of couples graphs, Kojima et al. (2013) essentially showed that with less than \sqrt{n} couples, the couples graph is empty with high probability, i.e., possible rejection chains initiated by couples will never intersect each other. In particular, not only that this is not expected in practice, it misses the real difficulty couples bring to this setting—rejection cycles. In the following sections we discuss and define formally influence trees and the couples graph.

REMARK 4. In our model, each couple's preference list is drawn from a product distribution. The weakness of this model is the lack of correlation in each couple's preferences. Our results hold, however, for a quite general *city model* in which couples wish to be assigned to hospitals in the same city. Suppose there is a constant number of cities k , each with $\delta\lambda n$ hospitals (for some small δ) and each hospital belongs to a given city. Each single chooses a list of hospitals

at random, ignoring geographic constraints. Each couple $c = (m, f)$, repeatedly, first randomly selects a city w and then m and f independently draw a pair of hospitals that are located in w . Suppose that every couple chooses one city and then randomly chooses pairs of hospitals only from that city. In this case, every city will essentially be an independent market with its own couples. Although singles may create some correlation (since they can move between cities), the history of each single during the SoDA algorithm will not make a difference: since each city is of size $\delta\lambda n$, with high probability at any given stage of the SoDA algorithm the number of singles in the city will be at most $\delta n + \sqrt{\delta n} \log(\delta n)$, implying it will have many vacancies with high probability. If couples can choose at each iteration when drawing their preferences of a different city, our result will still hold since the probability for a stable matching, is higher than in a similar model with $kn^{1-\epsilon}$ couples in which each couple is restricted to choose one city.

REMARK 5. Another model that allows correlation in the preferences of each couple is the following. Assume singles and couples both have constant length preferences. Singles preferences are drawn as usual. Each couple c , however, first randomly draws a small (constant size) set of potential hospitals S_c , and then a joint preference list of constant length is created by choosing arbitrary tuples of pairs of potential hospitals. For example, if $S_c = \{h_1, h_2, h_3\}$, then the joint preference list $(h_1, h_2), (h_1, h_1), (h_3, h_1), (h_2, h_3), (h_3, h_3)$ is a valid preference list for c .

Note that this model allows pairs to list the same hospital for both members at the same entry. Under our original model (as is shown in the proof) couples do not apply to the same hospital at the same time with high probability. This is not the case here, and we need another condition for hospitals' preferences over sets of doctors. Assume each hospital h has lexicographical preferences, i.e., it prefers to always hire a better doctor even it means not to fill capacity or to hire a “bad” doctor. For example, suppose (d_1, d_4) and (d_4, d_3) are couples, and suppose h has capacity two and has the ranking d_1, d_2, d_3, d_4 over doctors. Then h will prefer to hire (d_1, d_4) over (d_2, d_3) . Such a condition is needed to prevent trivial cycles; if (d_2, d_3) are preferred over (d_1, d_4) then if d_1 applies to h when d_2 and d_3 are assigned to h , then they both will be rejected and then d_4 can join d_1 in h . But now d_2 and d_3 are preferred by h over d_2 and d_3 .

Under this model, Theorem 1 will hold. The proof will follow under very minor modifications since couples' influence trees are still small. Further types of correlations in preferences are left for future research.

4.2. Influence Trees

A main concept we introduce is that of an *influence tree*. Note that whenever a couple applies and is assigned to a pair of hospitals it causes a rejection chain of doctors that apply down their preference lists to find new assignments after they have been evicted from their previous assignments.

The influence tree of a couple c will consist of the set of pairs of doctors and hospitals that form new assignments in the rejection chain c causes.

Observe that the order of applications of couples may change the rejection chains a couple c causes. To avoid this difficulty, we (constructively) define the influence tree for a couple c on the submarket without all other couples, i.e., as if the couple c is the first to apply, but allowing $r > 0$ adversarial rejections. In particular the r rejections simulate the presence of other couples and the rejection chains they cause by an adversary that can reject doctors from settling in hospitals. Thus the adversary is allowed to reject r times (in addition to the natural rejections). In the analysis we will choose $r > 0$ such that with high probability the presence of other couples will only affect the influence path of the couple c in $\leq r$ places. Thus, the influence tree definition allows us to analyze the dynamic setting in which at each point a different number of couples already applied by upper bounding it with a “pessimistic” static setting.

Next we formally define influence trees (these will be defined for a fixed realization of the preferences). First, let $\Gamma = (H, S, C, \succ_H, \succ_S, \succ_C)$ be a matching market and let μ be a matching. Denote by $o_h(\mu)$ and by $f_h = k_h - o_h(\mu)$ the number of assigned doctors to hospital h and the number of available positions in h under μ , respectively. Let $d^j(\mu, h)$ be the j th least preferred doctor according to \succ_h that is assigned to h under μ .

DEFINITION 3 (INFLUENCE TREE). Let $\Gamma = (H, S, C, \succ_H, \succ_S, \succ_C)$ be a matching market with couples and let μ be the matching produced by the DA algorithm for the submarket without couples. Let $d \in D$ and let r be any integer. An *influence subtree* of doctor d with root h and with up to r rejections, denoted by $IT(d, r, h)$ is defined recursively as follows:

- (a) If $f_h(\mu) = 0$ and $d^{k_h}(\mu, h) \succ_h d$ then let h' be the next preferred hospital by d after h and let $IT(d, r, h) = IT(d, r, h')$. Otherwise
- (b) change μ such that d is assigned to h and
 - (b1) add (h, d) to $IT(d, r, h)$;
 - (b2) if $r > 0$ or $f_h(\mu) = -1$ then for each $j = 1, \dots, \min(o_h(\mu), r - f_h(\mu))$ let h_j be the most preferred hospital by $d^j(\mu, h)$ after h , and add to $IT(d, r, h)$ the influence subtree $IT(d^j(\mu, h), r - j - f_h(\mu), h_j)$.

For a couple $c = \{f, m\}$, let $(h_f^1, h_m^1), \dots, (h_f^r, h_m^r)$ be the top r pairs of hospitals according to \succ_c in which the couple c can be accepted. That is, either

- $h_f^i = h_m^i$ and $c \subseteq \text{Ch}_{h_f^i}(\mu(h_f^i) \cup c)$, or
- $h_f^i \neq h_m^i$ and $f \in \text{Ch}_{h_f^i}(\mu(h_f^i) \cup \{f\})$ and $m \in \text{Ch}_{h_m^i}(\mu(h_m^i) \cup \{m\})$.

The influence tree for the couple c is defined to be

$$IT(c, r) := \bigcup_{i=1}^{r+1} (IT(f, r+1-i, h_f^i) \cup IT(m, r+1-i, h_m^i)).$$

Note that we allow $f_h(\mu)$ to be -1 in the definition of an influence tree (this is possible since under this definition we first assign a doctor to a hospital and only then reject from that hospital). Also observe that each time a hospital h is inserted to the influence tree, a doctor d is associated with it. In this case we say that h was *inserted* to $IT(c, r)$ by d .²⁰ With a slight abuse of notation we will write $h \in IT(c, r)$ if h was inserted to $T(c, r)$ by some doctor d .

The next example illustrates the definition of an influence tree.


EXAMPLE 1. Consider a market with six hospitals each with capacity of two; five single doctors, d_1, d_2, \dots, d_5 ; and two couples $c_1 = (d_6, d_7)$ and $c_2 = (d_8, d_9)$; and let the preferences be as in Table 1. The deferred acceptance algorithm for the submarket without couples produces the matching given in the boxes as in Table 1. The influence tree for $c_1 = (d_6, d_7)$ with $r = 0$ is constructed as follows (see Figure 1(a)). The most preferred pair of hospitals by c_1 that it can be accepted to is (h_3, h_4) . Therefore (h_3, d_6) and (h_4, d_7) are added to $IT(c_1, 0)$. The assignment of c_1 to (h_3, h_4) causes d_4 to be evicted and assigned to h_5 , hence $(h_5, d_4) \in IT(c_1, 0)$. The tree $IT(c_1, 1)$ consists of $IT(c_1, 0)$ but also assignments that are caused because of a rejection chain resulting from a single rejection by an adversary anywhere down the chain in $IT(c_1, 0)$. If c_1 is rejected from (h_3, h_4) it can be accepted to (h_4, h_5) , thus $(h_4, d_4), (h_5, d_5) \in IT(c_1, 1)$. Furthermore, no other doctor is evicted by assigning c_1 to (h_4, h_5) . Although a single rejection can be used anywhere down the path in $IT(c_1, 0)$, d_4 cannot be evicted from h_5 since the capacity of h_5 is two and it is the only doctor assigned to that hospital.

The influence trees $IT(c_2, 0)$ and $IT(c_2, 1)$ are given in Figure 1(b). The construction is similar, but for $r = 1$ an adversarial rejection can be used to evict also d_5 , which in turn can be accepted to h_5 . Therefore $(h_5, d_5) \in IT(c_2, 1)$.

Table 1. Preference lists.

Doctors							Hospitals				
d_1	d_2	d_3	d_4	d_5	(d_6, d_7)	(d_8, d_9)	h_1	h_2	h_3	h_4	h_5
$\boxed{h_1}$	$\boxed{h_1}$	h_1	$\boxed{h_3}$	$\boxed{h_3}$	(h_1, h_2)	(h_1, h_1)	$\boxed{d_1}$	d_1	d_1	d_1	d_1
h_2	h_2	$\boxed{h_2}$	h_5	h_5	(h_2, h_1)	(h_2, h_2)	d_8	d_8	d_8	d_8	d_8
h_3	h_3	h_3	h_1	h_1	(h_3, h_4)	(h_3, h_4)	d_9	d_9	d_9	d_9	d_9
h_4	h_5	h_4	h_4	h_2	(h_4, h_5)	(h_4, h_3)	$\boxed{d_2}$	d_2	d_3	d_3	d_6
h_5	h_6	h_5	h_2	h_4	(h_5, h_5)	(h_4, h_2)	d_5	d_5	d_6	d_5	d_4
							d_3	$\boxed{d_3}$	d_2	d_4	d_2
							d_6	d_6	$\boxed{d_5}$	d_6	d_5
							d_4	d_4	d_7	d_2	d_7
							d_7	d_7	$\boxed{d_4}$	d_7	d_3

Figure 1. Influence trees with parameters $r = 0$ and $r = 1$.



4.3. The Couples Graph


Recall that the definition of an influence tree for a couple c does not involve any other couple and therefore the tree captures only what possibly could have happened had there been other couples. The SoDA algorithm inserts couples one by one after the DA algorithm has terminated, and if some couple c_1 evicts another couple c_2 the order of their insertions is altered so that c_1 is moved ahead of c_2 . Intuitively the intersection of two influence trees, of c_1 and of c_2 , together with the hospital preferences will provide a good guess for which couple to insert first. This motivates the definition of the *couples graph*; each couple in the graph is a node and a directed edge will exist from couple c_i to couple c_j if both couples influence some common hospital h by causing doctors d_i and d_j to apply to it, respectively, and h prefers d_i over d_j . Formally, we have the following:

DEFINITION 4. Let $\Gamma = (H, S, C, \succ_H, \succ_S, \succ_H)$ be a matching market and let $r \geq 0$. In the (directed) *couples graph* for depth $r \geq 0$, denoted by $G(C, r)$ the set of vertices is C and for every two couples $c_1, c_2 \in C$ there is a directed edge from c_1 to c_2 if and only if there exist $h \in H$ and $d_1, d_2 \in D$ ($d_1 \neq d_2$) such that $(h, d_1) \in IT(c_1, r)$, $(h, d_2) \in IT(c_2, r)$ and $d_1 \succ_h d_2$.

Before we continue we illustrate a couples graph.

EXAMPLE 2. Consider the same market as in Example 1 (see Table 1). Note that the influence trees with $r = 1$ intersect at h_3 twice since $(h_3, d_8) \in IT(c_2, 1)$, $(h_3, d_3) \in IT(c_2, 1)$ and $(h_3, d_6) \in IT(c_1, 1)$. Since $d_3 \succ_{h_3} d_6$ and $d_8 \succ_{h_3} d_6$ the couples' graph with $r = 1$ will have an edge from c_2 to c_1 as in Figure 2. One can verify that other intersections at h_4 and h_5 also result in the same directed edge. Indeed, letting c_1 apply before c_2 (after the DA stage) results in a stable matching.

Figure 2. Couples graph for $r = 1$.



As mentioned above, the idea of the proof of Theorem 1 is to show that the couples' graph does not contain cycles with high probability. Let $k = n^{1-\epsilon}$ be the number of couples. In one of the first steps we will show that the probability for an edge in the couples' graph, i.e., that some couple can possibly evict some other couple, is of order $1/n$. Since there are $n^{2-2\epsilon}$ possible edges, this means that for $\epsilon > \frac{1}{2}$ the graph will be empty with high probability (which provides the result by Kojima et al. 2013). However for $\epsilon < \frac{1}{2}$ the graph will contain many edges, and we will show it contains no cycles.

5. Incentive Compatibility

In this section we will show the following:

THEOREM 2. *Ex post truthfulness: the probability that any doctor can gain by misreporting her preferences is at most $O(n^{-\epsilon/2})$, even if the doctor knows the entire preference list.*

A similar result can be shown for hospitals using similar techniques as in the proof of Theorem 2. We avoid the exact details here.²¹ Together with Theorem 2 we obtain that reporting truthfully is a δ -Bayes-Nash equilibrium in the Bayesian game induced by the SoDA algorithm (assuming bounded utilities). We refer the reader for exact definitions of the Bayesian game to Kojima et al. (2013). The proof of Theorem 2 is deferred to Appendix A.5.

6. Empirical and Experimental Results

In this section we provide empirical and simulation results for matching markets with couples, beginning with the market for psychology internships, followed by more general simulations for large(r) markets.

6.1. The Clinical Psychology Market

Since 1999 the Association of Psychology Postdoctoral and Internship Centers (APPIC) used a centralized computerized clearing house every year to match applicants to internships, allowing also for couples to express their preferences. In this section we will use the data from the period 1999–2007 to illustrate the underlying couples graphs, and provide additional intuition for the SoDA algorithm. Size characteristics of the market in each year are summarized in Table 2 (a similar table appears in Kojima et al. 2013). Note that

Downloaded from informs.org by [18.154.1.73] on 14 August 2014, at 09:40. For personal use only, all rights reserved.

Table 2. Size of the psychology interns market each year between 1999–2007.

Year	1999	2000	2001	2002	2003	2004	2005	2006	2007
Single doctors	2,896	2,916	2,904	2,809	2,926	3,006	3,086	3,169	3,391
Couples	14	21	22	17	19	17	16	21	20
Positions	2,631	2,713	2,763	2,752	2,718	2,732	2,757	2,779	2,884
% of unfilled positions	8	11	12	13	11	11	11	11	10

Note. The last describes the percentage of unfilled positions each year in the stable matching that was found.

although the number of positions each year is smaller than the number of doctors, a substantial number of positions are unfilled in the stable matching that was found. The fact that there were positions that remained unfilled, as well as observed success probabilities (which we will see below) is consistent with the predictions of our large market model.


Recall that an edge in the couples graph from couple c_i to couple c_j captures the idea that couple c_i could potentially evict couple c_j if c_i applied after c_j . As we shall see, the couple graphs that we find predict quite accurately the probability that a stable outcome will be found in the direct matching algorithm (i.e., by allowing one permutation order of application by couples—with no subsequent reordering of couples).

To construct the couples graph, we ran 500 iterations of the direct algorithm, while recording all different evictions in which some couple evicts itself or another couple (note that these are also the cases in which the direct algorithm fails). The couples graph of different years are given in Figure 3. For example, in 2006 couple 20 was evicted by couples 1 and 11, and couple 13 was evicted by couples 19 and 12. Interestingly, in 2007 the graph we found consisted of one self edge of the graph, i.e., couple number 9 evicted itself. In the years not shown on the diagram no couple was evicted by another couple.

The proof for Theorem 1 suggests that if couple 20 applied before couples 1 and 11, and couple 13 applied before couples 19 and 12, then a stable matching would be found. Note that by randomly choosing only one permutation over the couples for the direct algorithm, this event happens with probability $\frac{1}{9}$. In Table 3 we observe that the fraction of failures out of the 500 iterations we observe is approximately captured by the structure of the couples graphs:

In the year 2007, a stable matching was not found by direct algorithm, and also not surprisingly by the SoDA algorithm.

Figure 3. Couples graphs constructed from realized evictions.



This is indeed consistent with the couples’ graph realized in the same year. By letting the couple that evicted itself (couple 9), continue applying according to its preference list after evicting itself, we were able to find a stable matching. Thus one can augment the SoDA algorithm with a component similar to the last part in the RP algorithm; if a “loop” involving couples is detected (for example a couple evicts itself, or a couple c_1 evicts couple c_2 , and vice versa), choose one of the couples involved in that loop at random, and proceed by having it apply according to its preference list (see, e.g., Roth and Peranson 1999).

The APPIC data is rather limited and the number of couples in it is relatively small. To better understand the existence of a stable matching with a larger fraction of couples, we added artificial couples into the data. We created these artificial couples by randomly merging singles who expressed “geographically similar” preferences. Two singles could potentially be merged in our simulation if most of the programs they rank belong to the same two geographic regions. The joint preference of a new artificial couple is generated as to respect the preferences of the individuals, as well as the geographic constraints, as follows.

Let (d_1, d_2) be an artificial couple. First each d_i chooses a function f_i that assigns scores to all hospitals on her list (including not being matched) such that higher ranked hospitals are given higher scores. The joint preference list includes only pairs $h_1, h_2 \in H \cup \{\phi\}$ that satisfy the two conditions below:

- (i) h_1 (h_2) is ranked by the first (second) member;
- (ii) h_1 and h_2 either both belong to one of the two main regions the members prefer, or one belongs to one of these regions and the other is ϕ .

Finally, each pair of hospitals h_1, h_2 that satisfies (i) and (ii) is given the sum of scores $f_1(h_1) + f_2(h_2)$ and the joint preference list is created by sorting the pairs according to their scores.

In each experiment we generated a different number of extra couples, and for each such number we ran 100 iterations,

Table 3. Fraction of failures under the direct algorithm, and the couples’ graph failure prediction.

Year	2000	2004	2006	2007
Fraction of failures	0.858	0.498	0.884	1
Graph prediction	0.875	0.5	0.889	1

Downloaded from informs.org by [18.154.1.73] on 14 August 2014, at 09:40. For personal use only, all rights reserved.

Table 4. Average percentage of failures to find a stable matching for different number of extra pairs.

Extra couples ↓	Year →	1999	2000	2001	2002	2003	2004	2005	2006
0	Direct alg. fail. prob.	0	0.879	0	0	0	0.507	0	0.894
	SoDA fail. prob.	0	0	0	0	0	0	0	0
20	Direct alg. fail. prob.	0.788	0.98	0.771	0.738	0.748	0.866	0.596	0.939
	SoDA fail. prob.	0	0	0	0	0	0	0	0
40	Direct alg. fail. prob.	0.966	0.997	0.985	0.981	0.975	0.981	0.948	0.995
	SoDA fail. prob.	0	0	0	0	0	0	0	0
60	Direct alg. fail. prob.	0.997	1	0.998	0.997	0.999	0.999	0.993	0.999
	SoDA fail. prob.	0	0.002	0.001	0.001	0.001	0	0	0.001
80	Direct alg. fail. prob.	1	1	1	1	1	1	0.999	1
	SoDA fail. prob.	0	0.003	0.002	0.001	0.001	0.001	0.001	0.001
160	Direct alg. fail. prob.	1	1	1	1	1	1	1	1
	SoDA fail. prob.	0.006	0.008	0.008	0.008	0.008	0.005	0.004	0.008

Note. For each number of extra pairs the first row provides the results under the direct algorithm and the second row provides the result under the SoDA algorithm.

i.e., in each iteration we drew a different set of couples. Note that if the SoDA algorithm finds a stable matching, one might expect that the direct algorithm would also find a stable matching by just running the algorithm again and again each time using a random permutation over the couples. In Table 4 we present the probability of failure of the SoDA matching algorithm as well as the probability of failure under the direct algorithm as a function of the number of extra couples. Whereas for SoDA, for each draw of extra pairs we just run the algorithm once, when running the direct algorithm we tried to find a stable matching with 100 random permutations over the generated couples. The results show that as the number of couples grows, the chance of randomly finding a permutation that will result in a stable matching quickly tends to 0. At the same time, in the majority of the cases, SoDA is successful in finding a stable matching.


6.2. Large Market Simulations

In this section we provide further simulations results for larger markets—markets of size similar to the NRMP market—using the SoDA algorithm. In particular we performed sensitivity analysis on various parameters of the problem. For each configuration we ran 600 trials. We assumed there are $n/2$ hospitals, where n is the number of singles and each hospital has capacity of 3.²²

In the first simulation we fixed the percentage of couples in the market and found the success rate of finding a stable matching. For comparison, in the NMRP match in 2010 the number of (United States) doctors was about 16,000 where as the number of couples was about 800.²³ As Figure 4 shows that the ratio of doctors that are members of couples plays a crucial role in the probability that a stable matching will be found. Note that although the number of singles grows (and the number of couples is linear) the probability for finding a stable matching appears to remain unchanged.

Next we fixed ϵ , i.e., the number of couples is $n^{1-\epsilon}$. Figure 5 shows that the probability for finding a stable

Figure 4. The success rate of SoDA in finding a stable outcome given the number of singles (x -axis), for different couples percentages (5% means that 10% of the doctors are members of couples).



matching with SoDA increases and is roughly concave in the number of singles. Observe that the rate of convergence is different for various ϵ 's.

Figure 5. The success rate in finding a stable outcome given the number of singles (x -axis), where the number of couples is $n^{1-\epsilon}$ for three different ϵ 's.






Figure 6. The histogram shows the percentage of singles and couples that got their k th favorite choice for each $k = 1, \dots, 8$.



Note. The left bar is for singles and the right bar is for couples.

Figure 7. The histogram shows the percentage of singles and couples that got their k th most preferred choice for each $k = 1, \dots, 8$ when hospitals have a ranking score.



Note. The left bar is for singles and the right bar is for couples.

In the next simulation (see Figure 6) we fixed the number of singles and the number of couples to be 16,000 and 800, respectively, as in the NMRP, and found the percentage of singles and couples that get their k th most preferred choice. We assumed that there is no common ranking, i.e., preference distributions of both doctors and hospitals are uniform.

In Figure 7 we provide the same histogram but adding ranking to hospitals; each hospital has been assigned a score uniformly at random from the interval $[0.2, 1]$. To decide the next preference of a doctor, she randomizes uniformly a hospital h and a number from $[0.2, 1]$, and if h 's score is below the number, the doctor resamples such a pair.

7. “Almost” Linear Is Necessary

In §4 we showed that the SoDA algorithm finds a stable matching with probability approaching 1 as n tends to infinity assuming the number of couples is growing at a rate of $n^{1-\epsilon}$ (for any $0 < \epsilon < 1$). In §6 we saw that when the number of couples is a constant fraction of the total capacity, the SoDA algorithm fails to produce a stable matching in a constant fraction of the cases. A natural question is whether a stable matching—one that SoDA fails to find—may still exist with probability approaching 1. In this section we will show that the answer to this question is negative, i.e., as the number of couples grows linearly, *no stable matching exists* with constant probability.

For simplicity we will consider only uniformly distributed preferences and a capacity of one for each hospital.

THEOREM 3. *There exists $\bar{\lambda} \geq 1$ such for all $\lambda > \bar{\lambda}$ in a matching market with n couples, n singles, λn hospitals each of capacity one, and uniformly distributed preferences, with probability $\delta = \delta(\lambda) > 0$ not depending on n , no stable matching exists.*

If the preference list of each single and each couple is bounded by a constant length $k > 0$, then the above statement holds for $\bar{\lambda} = 1$.²⁴

The proof of Theorem 3 is deferred to the appendix. Although the proof is not immediate, the idea (for both parts) is relatively simple and we present it here. Consider a submarket with one couple $c(f_c, m_c)$ and one single s and suppose the preferences of the couple c is such that its first and second preferred pairs of hospitals are (h_1, h_2) and (h_3, h_4) , respectively (where h_1, h_2, h_3 , and h_4 are pairwise distinct), and suppose that the following hold:

- (i) $h_2 \succ_s h_1 \succ_s h$ for any $h \notin \{h_1, h_2\}$;
- (ii) $s \succ_{h_1} m_c$;
- (iii) $f_c \succ_{h_2} s$.

Observe that if the entire market only consisted of the couple c and the single s , the market would have no stable matching. In the first step of the proof of Theorem 3 we show that such a submarket exists with a constant probability. To complete the proof we show that there is a constant probability that no other doctor other than s, m_c and f_c will ever apply to one of the hospitals $h_1 - h_4$. In other words, we embed a “small bad example” that does not admit a stable outcome and show that all doctors that do not participate in this example will never have to use the hospitals in that submarket.

The first step of the proof is similar for both parts of the theorem. However the second step of the proof—i.e., showing that there is a constant probability that no doctors except s, m_c , and f_c will ever be interested in the “bad” submarket—is much simpler for the second part. If the lists have constant lengths, the second step just follows since with constant probability the hospitals h_1, h_2, h_3 , and h_4 do not even appear in the preference lists of all other singles and couples.

The submarket that “possess” the instability is small. An interesting question is whether such instability exists in “substantial size of the market.”²⁵

An open issue that follows from Theorem 3 is the following. We have used a large excess number of hospitals to obtain our negative result when preference lists are long. We do not expect however that a small excess of hospitals will improve the chances of obtaining a stable matching.

8. Beyond Couples—Many-to-One Matching Markets with Complementarities

Complementarities in the preferences of couples are the cause of possible nonexistence of a stable matching. By adopting a large market approach, Theorem 1 provides the existence of a stable matching given that the number of couples is not “too big.”

In this section we extend this result to a more general matching model studied by Roth and Sotomayor (1990). Consider the many-to-one matching market in which one side consists of a set of workers W and the other side consists of a set of n firms, F , and each firm is interested in hiring up to $k > 0$ workers and each worker is interested in getting a single job. Roth and Sotomayor (1990) showed that if all firms’ preferences are substitutable a stable matching exists, and it need not exist even if one firm does not have substitutable preferences.²⁶ (Other works have produced similar results for other matching markets; see, e.g., Ostrovsky 2008 and Hatfield and Kojima 2010.) We will show that a similar result as Theorem 1 holds for this model.

Formally, a preference relation for firm $f \in F$, denoted by \succ_f , is *substitutable* if for all $w, y \in W$ and $W' \subseteq W$, if $y \in \text{Ch}_f(W' \cup \{w, y\})$, then $y \in \text{Ch}_f(W' \cup \{y\})$.

Note that in the matching market with couples, we in fact draw an ordered list, i.e., a responsive preference, for each single doctor. However, not all substitutable preferences are responsive. Therefore in order to analyze large random markets, first we need to define a probability measure P over the substitutable preferences. In fact, for our main result to hold, we do not require P to be a particular distribution. Rather, we only specify a property that P needs to satisfy. Informally, we require that P is a distribution over substitutable preferences such that if a firm currently has a set S of workers, which is its most preferred subset of a set T of workers, and worker $i \in S$ resigns, then the next most desirable set by the firm is $S \setminus \{i\} \cup \{j\}$, where the worker j is drawn uniformly from the set $T \setminus S$ of the remaining workers. Observe that responsive preferences satisfy this property, but also every distribution over substitutable preference that is symmetric with respect to permutations on the workers.

As for firms with nonsubstitutable preferences, we draw preferences in a similar fashion as for couples: there exists a probability distribution Q , and a firm who is interested in k workers draws its preference list from the distribution Q^k

(a similar aggregation of the k lists can be done as described in §4 for couples).

We will consider a *sequence of random markets* Ψ^1, Ψ^2, \dots of growing size, where each market is a tuple $\Psi^n = (W^n, F_S^n, F_C^n, \succeq_W^n, P^n, Q^n)$, where W^n is the set of workers, F_S^n is set of firms with substitutable preferences, and F_C^n is the set of firms with nonsubstitutable preferences, and preferences of firms are drawn from P^n and Q^n as described above.

DEFINITION 5. A sequence of random markets Ψ^1, Ψ^2, \dots is called *regular* if there exist $0 < \epsilon < 1$, $\lambda > 1$, $k > 0$ such that for all n

1. $|F_S^n| = n$ and $|F_C^n| = O(n^{1-\epsilon})$ (the number of firms with nonsubstitutable preferences grows almost linearly);
2. each firm has capacity at most $k > 0$;
3. $|W^n| \geq \lambda kn$ (excess number of workers).

We can now state our main theorem, which says that in a large random market with an excess number of workers, even if the number of firms that have nonsubstitutable preferences grows at a nearly linear rate, the probability that a stable matching exists approaches 1.

THEOREM 4. *Let Ψ^1, Ψ^2, \dots be a regular sequence of random markets. Then the probability that there exists a stable matching tends to 1 as n goes to infinity.*

The proof is similar to the proof of Theorem 1 and thus we omit it.²⁷ In particular the SoDA algorithm can be adapted to this setting, by letting firms propose workers (the first step of the algorithm finds the stable matching for the submarket containing only firms with substitutable preferences).

9. Conclusion

This paper characterizes the (non)existence of a stable matching with high probability in large random markets with bounded complementarities. The main difficulty couples introduce by their complementarities is that rejection chains may cycle. We present a new matching algorithm, SoDA, and show that if the number of couples grows at a sublinear rate, rejection cycles will not evolve with high probability and thus a stable matching will be found. However, if the number of couples grows in a linear rate, with constant probability no stable matching exists. This negative result is in contrast to many positive results about desired economic properties (such as efficiency and strategy-proofness) that do not hold in general but do hold in large markets. In particular, complementarities are a first-order difficulty.

For the existence result, we develop a novel concept, namely, influence trees, which allows us to analyze when couples interfere with each other in a way that rejection cycles may evolve. We believe this concept can be used in other markets with complementarities.

We show that when the growth rate of the number of couples is of a near-linear rate, the ex ante probability that each doctor and each couple will get their best stable

matching tends to one, and further that truth telling is an approximated Bayes-Nash equilibrium in such large enough markets.

We strongly believe that one can also use the Roth-Peranson (RP) algorithm to show the positive results using the same tools. However, it may still be the case that on some instances SoDA succeeds and RP does not, and vice versa. As the number of couples grows each year it is important to better understand the differences between the two algorithms.

Acknowledgments

The authors thank Shahar Dobzinski, Federico Echenique, Itay Fainmesser, John Hatfield, Jacob Leshno, Brendan Lucier, Scott Kominers, David Rahman, Alvin Roth, and Itai Sher for very helpful discussions and comments. The first author acknowledges the research support of the National Science Foundation [Grant SES-1254768]. The second author thanks the Alfred P. Sloan Fellowship, the National Science Foundation [Grant CCF-0832797], an NSF CAREER Award [CCF-1149888], NSF Award [CCF-1215990], a Turing Centenary Fellowship, and a Packard Fellowship in Science and Engineering. The third author thanks the United States-Israel Binational Science Foundation (BSF) [Grant 2012344] and the Israel Science Foundation [Grant 1241/12].

Appendix A

A.1. Formal Definition of the Sorted Deferred Acceptance Matching Algorithm

Sorted Deferred Acceptance Matching Algorithm:

Input: A matching market $(H, S, C, \succ_S, \succ_H, \succ_C)$ and a default permutation π over the set $\{1, 2, \dots, |C|\}$. Let $\Pi = \phi$.

Step 1. Find the stable matching μ produced by the DA algorithm in the matching market (H, S, \succ_S, \succ_H) without couples.

Step 2 [Iterate through the couples]. Let $i = 1$ and let $B = \phi$.

(a) Let $c = c_{\pi(i)}$ be the $\pi(i)$ th couple.

Let c apply to the most preferred pair of hospitals $(h, h') \in H \times H$ that has not rejected it yet. If such a pair of hospitals does not exist, modify μ such that $c = (f, m)$ is unassigned and go to Step 2(a) with $i + 1$. If such a pair (h, h') exists then

(a1) If $h = h'$ and $\{f, m\} \subseteq \text{Ch}_h(\mu(h) \cup c)$ then

let $R = \mu(h) \setminus \text{Ch}_h(\mu(h) \cup c)$ be the rejected doctors from h .

(a11) If there exist a couple $c' \neq c$ for which $\{f_{c'}, m_{c'}\} \cap R \neq \phi$ then let $j < i$ be such that $c_{\pi(j)} = c'$. Let π' be the permutation obtained by π as follows:

$\pi'(j) = \pi(i)$, $\pi'(i) = \pi(j)$ for all l such that $l < j$ or $l > i$ and $\pi'(l) > \pi(l - 1)$ for other $j + 1 \leq l \leq i$.

If $\pi' \in \Pi$ terminate the algorithm. Otherwise add π' to Π and go to Step 1 setting $\pi = \pi'$.

(a12) Modify μ by assigning c to h , remove R from $\mu(h)$. Add R to B and do Step 3 (Stabilize) with the couple c .

(a2) If $h \neq h'$, $f \in \text{Ch}_h(\mu(h) \cup f)$, and $m \in \text{Ch}_{h'}(\mu(h) \cup m)$ then

let $R_h = \mu(h) \setminus \text{Ch}_h(\mu(h) \cup \{f\})$ and $R_{h'} = \mu(h') \setminus \text{Ch}_{h'}(\mu(h') \cup \{m\})$.

(a21) If there exist a couple $c' \neq c$ for which $\{f_{c'}, m_{c'}\} \cap (R_h \cup R_{h'}) \neq \phi$ then let $j < i$ be such that $c_{\pi(j)} = c'$, change π as in Step 2(a11). If $\pi \in \Pi$ terminate the algorithm. Otherwise add π to Π and go to Step 1.

(a22) Modify μ by assigning f to h and m to h' , remove R_h from $\mu(h)$ and remove $R_{h'}$ from $\mu(h')$. Add $R_h \cup R_{h'}$ to B and go to Step 3 (Stabilize) with the couple c .

(a3) Otherwise, let h and h' reject the couple c and go to Step 2(a).

Step 3 [Stabilize]. Let $j = |B|$. As long as $j \geq 0$,

(a) if $j = 0$ increment i by one and go to Step 2;

(b) otherwise pick some $s \in B$ and

(b1) let h be the most preferred hospital s has yet to apply to. If such a hospital does not exist then modify the matching μ such that s is unassigned and go to Step 2(a). Otherwise

let $R = (\mu(h) \cup \{s\}) \setminus \text{Ch}_h(\mu(h) \cup \{s\})$.

(b21) If $\{f_c, m_c\} \cap R \neq \phi$ then the algorithm fails.

(b22) If there exist a couple $c' \neq c$ for which $\{f_{c'}, m_{c'}\} \cap R \neq \phi$ then let i and j be such that $c_{\pi(i)} = c$ (c is the last couple that applied) and $c_{\pi(j)} = c'$. Change π as in Step 2(a11). If $\pi \in \Pi$ terminate the algorithm. Otherwise add π to Π and go Step 1.

(b23) If $s \in R$ then go to Step 3(b1).

(b24) Modify μ by assigning s to h , remove R from $\mu(h)$. Add R to B and go to Step 3.

A.2. Proof of Theorem 1

Throughout the proof we will fix r to be $r = 4/\epsilon$ for some fixed $0 < \epsilon < 1$. One should interpret this r as a “small” number of possible rejections (relative to n). In a random market the influence trees are random variables. In the following Lemma we show that the influence trees are small:

LEMMA 1. (a) For every hospital h and couple c , $\Pr(h \in \text{IT}(c, r)) = O((\log n)^{r+1}/n)$.

(b) The probability that the size of every influence tree $\text{IT}(c, r)$ is $O((\log n)^{r+1})$ is at least $1 - n^{-3}$.

(c) The probability that for all couples c , each hospital h appears in $\text{IT}(c, r)$ at most once is at least $1 - n^{-\epsilon/2}$.

PROOF. We begin with the second part. Let c be a couple. For each of the two $d \in c$ and for each $h' \neq h$ we will give an upper bound of $O((\log n)^r/n)$ on $\Pr(h \in \text{IT}(d, r, h'))$. The claim will follow from the definition of $\text{IT}(c, r)$ and union bound.

An alternative way of viewing the recursive definition of $\text{IT}(d, r, h')$, is as follows: doctor d proceeds down his list beginning with h' until he finds the first hospital willing to accept him. If d is accepted into a hospital h_1 and h_1 was full to capacity, then some doctor d' is evicted and goes to a hospital h_2 , and we add $\text{IT}(d', r, h_2)$ to $\text{IT}(d, r, h')$. In this case, continuing the “rejection chain” did not require any arbitrary rejections. We call the hospitals added into $\text{IT}(d, r, h')$ with parameter r the main path of $\text{IT}(d, r, h')$. We then also allow the adversary to introduce up to r arbitrary rejections (for example, precluding d from being accepted into h_1). Thus the influence tree is composed of the main path, with lower-order influence trees (i.e., influence trees with a strictly smaller value of r) attached along it.

We first show by induction that with probability at least $1 - n^{-6}$ the length of the main path in $\text{IT}(d, r, h')$ is at most $b \log n$, where $b = 6 \cdot (c_{\max} \cdot \gamma_{\max}) / (\lambda - 1)$. At any step along the main path, for the main path to continue, the currently evicted doctor d needs to choose a full hospital h . Because of the way the doctors’ preferences are sampled, the probability of this happening is bounded by $1 - (\lambda - 1) / (c_{\max} \cdot \gamma_{\max})$. Since each subsequent step along the path is independent from the previous ones, the bound follows.

By union bound, we see that with probability at least $1 - n^{-4}$ all potential main paths contain at most $b \log n$ hospitals. Each main

path of length ℓ recursively gives rise to at most $r \cdot \ell$ lower-order influence trees (i.e., influence trees with smaller r) that are added to $IT(d, r, h')$. Thus we can prove by induction that for each r , the size $S(r)$ of the largest order- r influence tree is bounded by $(1 + br \log n)^{r+1} = O((\log n)^{r+1})$. For the base case, an influence tree with $r = 0$ only contains the main path, and thus $S(0) \leq b \log n$. For the step, we get

$$\begin{aligned} S(r) &\leq b \log n + (b \log n) \cdot r \cdot S(r - 1) \\ &\leq b \log n + (b \log n) \cdot r \cdot (1 + br \log n)^r \\ &< (1 + br \log n)^r + (b \log n) \cdot r \cdot (1 + br \log n)^r \\ &= (1 + br \log n)^{r+1}. \end{aligned}$$

Next, the first part of the lemma follows from the proof of the second part and the fact that the hospitals that are added to $IT(c, r)$ are hospitals on the doctors' preference lists and are chosen independently. Thus the probability of h to be added to $IT(c, r)$ at some point is bounded by $S(r) \cdot (c_{\max} \cdot \gamma_{\max}/n) = O((\log n)^{r+1}/n)$.

Finally, we show that $IT(c, r)$ does not “intersect itself” except with probability $< n^{\epsilon/2}$. Note that this means that the members of the couple may not apply into the same hospital or evict each other.²⁸ We have seen that the probability of a hospital h belonging to $IT(c, r)$ is bounded by $O(S(r)/n)$. Similarly, the probability of h to be added twice or more to $IT(c, r)$ is bounded by $O(S(r)^2/n^2)$. Taking a union bound over all possible hospitals h and all possible couples c , we see that the probability that any hospital appears in any $IT(c, r)$ twice or more is bounded by

$$O(S(r)^2/n^2) \cdot n \cdot n^{1-\epsilon} < n^{-\epsilon/2}. \quad \square$$

Throughout the remainder of the proof, we will assume that each hospital appears in each $IT(c, r)$ at most once, neglecting an event of probability $< n^{-\epsilon/2}$.

In fact, in Lemma 1, one can prove a stronger bound of $O(\log n/n)$ for the probability that a hospital belongs to an influence tree. Although we do not prove or use the stronger bound in the rest of the paper, it provides intuition for why the SoDA algorithm works well even in a rather small market (e.g., when $n = 256$ we have $(\log 256)^3 = 8^3 = 512$, which does not explain why the algorithm works).

Next we analyze how much influence trees intersect with each other. Let c_1 and c_2 be two different couples. We say that two influence trees $IT(c_1, r)$ and $IT(c_2, r)$ intersect at hospital h if there exist d' and d'' such that $d' \neq d''$, $(h, d') \in IT(c_1, r)$ and $(h, d'') \in IT(c_2, r)$.²⁹

LEMMA 2. *No two influence trees intersect more than once, except with probability $< n^{-\epsilon/2}$.*

PROOF. By Lemma 1, we can assume that for every couple c the size of $IT(c, r)$ is at most $O((\log n)^{r+1})$. For the remainder of the proof, we will denote this upper bound on the size of $IT(c, r)$ by $S(r) = O((\log n)^{r+1})$. Recall also that we have assumed that no $IT(c, r)$ intersects itself.

We prove that with high probability no two influence trees intersect exactly two times. A similar proof shows that for every $3 \leq k \leq S(r)$ no two influence trees intersect exactly k times. The proof will then follow by a union bound on k (since the size of each tree is $\leq S(r)$ with high probability they cannot intersect more than $S(r)$ times).

Let c_1, c_2 be two couples, and h_1, h_2 be two hospitals. We want to bound the probability of the event

$$\begin{aligned} &\Pr(h_1, h_2 \in IT(c_1, r) \cap IT(c_2, r)) \\ &= \Pr(h_1, h_2 \in IT(c_1, r)) \\ &\quad \cdot \Pr(h_1, h_2 \in IT(c_2, r) \mid h_1, h_2 \in IT(c_1, r)). \end{aligned} \tag{A1}$$

We first note that if h_1 is an ancestor of h_2 in, e.g., $IT(c_1, r)$, and $IT(c_1, r)$ intersects $IT(c_2, r)$ in both h_1 and h_2 , then the influence tree $IT(c_2, 2r + c_{\max})$ will self-intersect at h_2 . The hospital h_2 will be added to $IT(c_2, 2r + c_{\max})$ twice: once following the path in $IT(c_2, r)$, and a second time through h_1 and then following the path from h_1 to h_2 in $IT(c_1, r)$. Since $2r + c_{\max}$ is a constant, by Lemma 1 the probability that any $IT(c, 2r + c_{\max})$ will self-intersect is smaller than $n^{-\epsilon/2}$, and can be disregarded. Thus we can assume that h_1 and h_2 are not each other's ancestors in either $IT(c_1, r)$ or $IT(c_2, r)$.

We begin by calculating the probability of the first event in (A1). A similar proof to that of Lemma 1 gives that the probability for this event is

$$\Pr(h_1, h_2 \in IT(c_1, r)) = O\left(\frac{S(r)^2}{n^2}\right).$$

Rather than compute $\Pr(h_1, h_2 \in IT(c_2, r) \mid h_1, h_2 \in IT(c_1, r))$ directly, to avoid the conditioning, we consider inserting c_2 into a modified world, in which all hospitals in $IT(c_1, r)$ except for $\{h_1, h_2\}$ and all the doctors in these hospitals do not exist. We argue that in this case,

$$\Pr(h_1, h_2 \in IT(c_2, r)) = O\left(\frac{S(r)^2}{n^2}\right)$$

using similar reasoning.

The influence tree generated in the modified algorithm (where we took out some of the hospitals) may differ from the one in the “real” algorithm. Note however that if removing $IT(c_1, r)$ affects the generation of the tree $IT(c_2, r)$ before it reaches h_1, h_2 , then it is the case that $IT(c_2, r)$ intersects $IT(c_1, r)$ at another hospital (which comes before h_1, h_2). But this is a contradiction, since we assumed $IT(c_1, r)$, $IT(c_2, r)$ intersect exactly twice.

Multiplying the probabilities, we get that

$$\Pr(h_1, h_2 \in IT(c_1, r) \cap IT(c_2, r)) = O\left(\frac{S(r)^4}{n^4}\right).$$

Taking a union bound over $O(n)$ hospitals and $n^{1-\epsilon}$ couples, bounds the probability that two couples exist that intersect exactly twice is at most

$$O\left(\frac{S(r)^4}{n^{2\epsilon}}\right).$$

We do not present the proof for exactly k intersections, and only state that the probability for that event drops at a rate of

$$\frac{S(r)^{2k}}{n^{k-\epsilon}} < \frac{S(r)^4}{n^{2\epsilon}}.$$

Taking a union bound over all possible values of k , we get that the probability that any two couples intersect strictly more than once is at most

$$O\left(\frac{S(r) \cdot S(r)^4}{n^{2\epsilon}}\right) = \frac{\text{polylog}(n)}{n^{2\epsilon}}.$$

as required.³⁰ \square

Our goal will be to show that with high probability the graph $G(c, r)$ can be topologically sorted; such a sorting corresponds to a good insertion order of the couples in the SoDA algorithm. In Example 2 the order c_1, c_2 is a topological sort.

At this point we show that with high probability weakly connected components in the couples graph are small. Formally, in a couples graph $G = G(C, r)$ a *weakly connected component* is defined to be a connected component in the graph obtained from G by removing the direction of the edges.³¹

LEMMA 3. *With probability $> 1 - 1/n$ the largest weakly connected component of the couples' graph has size at most $3/\epsilon$.*

PROOF. We will first consider an arbitrary set of $3/\epsilon$ couples and show that the probability that they form a weakly connected component is very small. The statement of the lemma will follow through union bound. Let $I = (c_1, c_2, \dots, c_{\lfloor 3/\epsilon \rfloor})$ be a sequence of couples with no repetitions: $c_i \neq c_j$. Let A_I be the event that for every $1 < i \leq \lfloor 3/\epsilon \rfloor$ the influence tree of c_i intersects with one of the previous influence trees, that is

$$IT(c_i, r) \cap \left(\bigcup_{j < i} IT(c_j, r) \right) \neq \emptyset.$$

We first show that

$$\begin{aligned} \Pr(A_I) &\leq \frac{(S(r))^2 \cdot c_{\max} \cdot \gamma_{\max} \cdot (3/\epsilon)^{\lfloor 3/\epsilon \rfloor}}{n^{\lfloor 3/\epsilon \rfloor - 1}} \\ &\leq \frac{(S(r))^2 \cdot c_{\max} \cdot \gamma_{\max} \cdot (3/\epsilon)^{3/\epsilon}}{n^{3/\epsilon - 2}}, \end{aligned} \quad (A2)$$

where $S(r)$ is the bound on the size of the influence trees $IT(c_i, r)$ as in Lemma 1.

Let

$$IT_i = \bigcup_{j \leq i} IT(c_j, r)$$

be the union of the influence trees of the first i couples. The probability of A_I can be written as

$$\begin{aligned} \Pr(A_I) &= \Pr(IT(2, r) \cap IT_1 \neq \emptyset) \\ &\quad \cdot \Pr(IT(3, r) \cap IT_2 \neq \emptyset \mid IT(2, r) \cap IT_1 \neq \emptyset) \\ &\quad \dots \cdot \Pr(IT(\lfloor 3/\epsilon \rfloor, r) \cap IT_{\lfloor 3/\epsilon \rfloor - 1} \neq \emptyset \mid \forall j \leq \lfloor 3/\epsilon \rfloor - 1, \\ &\quad \quad IT(j, r) \cap IT_{j-1} \neq \emptyset). \end{aligned} \quad (A3)$$

All the interactions that cause the influence trees within IT_{j-1} to intersect happen within IT_{j-1} , and conditioned on the set IT_{j-1} of hospitals do not affect the probability of $IT(c_j, r)$ intersecting IT_{j-1} . Hence for every $j = 2, \dots, \lfloor 3/\epsilon \rfloor$,

$$\begin{aligned} \Pr(IT(c_j, r) \cap IT_{j-1} \neq \emptyset \mid \forall 2 \leq l \leq j-1, IT(l, r) \cap IT_{l-1} \neq \emptyset) \\ = \Pr(IT(c_j, r) \cap IT_{j-1} \neq \emptyset \mid IT_{j-1}). \end{aligned}$$

Furthermore, from Lemma 1 it follows that the probability that $|IT(c_j, r)| < S(r)$ is at least $1 - 1/n^3$ and therefore $|IT_j| < j \cdot S(r)$. Hence,

$$\begin{aligned} \Pr(IT(c_j, r) \cap IT_{j-1} \neq \emptyset \mid IT_{j-1}) \\ \leq \frac{(j-1) \cdot S(r)^2 \cdot \gamma_{\max}}{\lambda n / c_{\max}} + \frac{1}{n^3} < \frac{j \cdot S(r)^2 \cdot \gamma_{\max}}{\lambda n / c_{\max}}. \end{aligned}$$

Since there are $\lfloor 3/\epsilon \rfloor - 1$ terms in (A3), we derive inequality (A2).

To finish the proof, observe that if there is a connected component of size at least $3/\epsilon$ then there exists a sequence I such that A_I holds. Since there are $n^{1-\epsilon}$ couples there exists fewer than

$$(n^{1-\epsilon})^{3/\epsilon} = n^{3/\epsilon - 3}$$

such possible sequences I . Therefore using a union bound over all of them proves the lemma. \square

Recall that we ignore all realizations of preferences at which two influence trees intersect more than once (in particular there is at most a single edge between every two couples in the couples graph). From now on we also ignore realizations where the largest weakly connected component of the couples graph contains more than $3/\epsilon$ couples.

LEMMA 4. *With probability $1 - O(1/n^\epsilon)$ the couples graph has no directed cycles.*

PROOF. We first prove the following claim, that is basically a simple general statement about directed graphs:

CLAIM 1. *If the shortest directed cycle has length k , it involves k different hospitals.*

PROOF. Suppose the shortest directed cycle is of length k and consider such a cycle $c_1 \rightarrow c_2 \rightarrow \dots \rightarrow c_k \rightarrow c_1$. Suppose couples c_1 and c_2 intersect at h because of d_1 and d_2 , respectively, i.e., $(h, d_1) \in IT(c_1, r)$, $(h, d_2) \in IT(c_2, r)$ and $(h, d_2) \in IT(c_2, r)$. Assume for contradiction that for some $2 \leq i \leq k$, c_i and c_{i+1} (i is taken modulo k) intersect at hospital h because of some doctors d_i and d_{i+1} , i.e., $(h, d_i) \in IT(c_i, r)$, $(h, d_{i+1}) \in IT(c_{i+1}, r)$ and $d_i \succ_h d_{i+1}$. Consider the case in which $d_i \succ_h d_2$. In this case a cycle of length less than k exists that consists of $c_2 \rightarrow c_3 \rightarrow \dots \rightarrow c_i \rightarrow c_2$. If $d_2 \succeq_h d_i$, i.e., either $d_2 \succ_h d_i$ or $d_2 = d_i$, then $d_1 \succ_h d_2 \succeq_h d_i \succ_h d_{i+1}$ implying that $c_1 \rightarrow c_{i+1} \rightarrow \dots \rightarrow c_k \rightarrow c_1$ is a shorter cycle. \square

To prove the lemma it is sufficient to show that the probability that the shortest directed cycle has length k is $O(S(r)^{2k}/n^{\epsilon k})$ since by taking the sum of these probabilities over all values of k gives the result (note that the dominant term in this sum is when $k = 2$).

We proceed in a manner similar to that of the proof of Lemma 3. Let $I = (c_1, c_2, \dots, c_k)$ be a sequence of couples without repetitions $c_i \neq c_j$. Let $J = (h_1, h_2, \dots, h_k)$ be a sequence of k hospitals without repetitions $h_i \neq h_j$. Let $A_{I,J}$ be the event that for every $i = 1, \dots, k$, $IT(c_i, r)$ and $IT(c_{i+1}, r)$ intersect at hospital h_i . Applying Lemma 1, and using reasoning similar to the proof of Lemma 3 the probability of the event $A_{I,J}$ can be bounded by

$$\Pr(A_{I,J}) < \frac{(2S(r) \cdot \gamma_{\max})^{2k}}{(\lambda n / c_{\max})^{2k}}.$$

Since there are $\leq \lambda n$ positions and $n^{1-\epsilon}$ couples, there are $\lambda^k n^k n^{(1-\epsilon)k}$ such different events $A_{I,J}$. A union bound over all these events implies the lemma. \square

For the analysis we will consider the event that the couples' graph contains a cycle as a failure.³² If the couples graph does not have cycles, then it has a topological sort. Let π denote any topological sort of G . We claim that inserting the couples according to π will result in a stable matching with couples. Moreover, we will show that a failure of the SoDA algorithm corresponds to a backward edge in the couples graph.³³

The next lemma shows that the influence trees indeed capture "real influences."

LEMMA 5. *Suppose we insert the couples as in the SoDA algorithm according to some order π until a couple evicts another couple or until all couples have been inserted. If a couple c is inserted and influences hospital h , then $h \in \text{IT}(c, r)$.*

PROOF. Recall that we consider only small weakly connected components (Lemma 3 upper bounded the probability that such a component is large). Let c be the couple currently being inserted, and assume that the statement of the lemma was true for couples inserted before c . Let $\{c_1, \dots, c_k\}$ be c 's weakly connected component in the couples graph, where $k \leq 3/\epsilon$, ordered according to their insertion order in π . We prove by induction a stronger claim, namely, that if $c = c_i$ influenced a hospital h , then $h \in \text{IT}(c, i - 1)$.

Suppose that $c = c_i$ is currently being inserted and that its insertion affects a hospital h . Consider the path of evictions that was started by c and led to hospital h being affected. There are two types of evictions along this path: the first type would have occurred even without any other couples present. The second type occurs because a hospital h' on the path has already been affected by a previously inserted couple c_j . If this happens, then the influence tree of c intersects the influence tree of c_j and thus c_j belongs to the weakly connected component of c in the couples graph. Moreover, since influence trees intersect only once, evictions due to influences from previously inserted couples happen at most $i - 1$ times: at most once for each previously inserted couple in the weakly connected component of c . By the definition of $\text{IT}(c, i - 1)$ this implies $h \in \text{IT}(c, i - 1)$. \square

As an immediate corollary of Lemma 5 we obtain that a couple causing another couple to be evicted corresponds to an edge in the couples graph.

COROLLARY 2. *If in an insertion order π inserting the couple $c_{\pi(i)}$ causes the couple $c_{\pi(j)}$ to be evicted ($j < i$) then in the couples' graph there is an edge from $c_{\pi(i)}$ to $c_{\pi(j)}$.*

Since there exists a topological sort, with a high probability Theorem 1 follows from the following corollary:

COROLLARY 3. *Inserting the couples according to any topological sort π of the couples graph gives a stable outcome.*

A.3. Proof Sketch of Corollary 1

We assume the couples' graph has no cycles and that every connected component of the couples' graph is of size less than $3/\epsilon$ (we showed that these events occur with probability tending to 1 as n tends to infinity). We will prove the result assuming capacity is 1 in each hospital.

First we show that as n tends to infinity, every couple that is not pointed by another couple in the couples graph, and every single that is not evicted after couples begin to apply in SoDA, obtain their best match with probability converging to 1.

Note that for any topological order of the couples' graph, the same set of singles will be evicted after the couples start applying. Denote this set of singles by E . Let M be the stable matching with couples that SoDA outputs. Suppose that there is another stable matching M' , such that there is at least one single not belonging to E who is better off in M' .

Let H' be the set of hospitals in M' in which couples are assigned to. Let $\tilde{\Gamma}$ denote the market that has only singles and the set of hospitals $H \setminus H'$.³⁴ Let \tilde{M}' denote the matching in the market $\tilde{\Gamma}$ in which every single s is matched to the same hospital it is matched to under M' . Observe that \tilde{M}' is a stable matching

for $\tilde{\Gamma}$ since any blocking pair would also be a blocking pair in M' . Let \tilde{M} be the matching in the market $\tilde{\Gamma}$ obtained by running doctor-proposing deferred acceptance on the market $\tilde{\Gamma}$. Then by the doctor optimality of the doctor-proposing deferred acceptance algorithm, all singles are at least as well off in \tilde{M} as in \tilde{M}' . Let Γ_s be the market with all hospitals, but with couples removed. Then it is not hard to see that running doctor-proposing deferred acceptance on Γ_s will result in a matching M_s where each doctor is at least as happy as in \tilde{M} , and therefore is at least as happy as in M' . We conclude the proof for the singles by observing that if a doctor $s \notin E$, then the hospital she is matched to under M_s is identical to the matching she receives under M and is thus at least as good as the matching she receives under M' .

A similar proof works for couples. Let c be a couple who has no incoming edges in the couples' graph, and is better off under a matching M' . Denote by h'_1, h'_2 the pair of hospitals c is matched to in M' .

Note that under all topological orderings of the couples' graph c is assigned to the same hospitals. Consider a topological order in which c is the first couple to enter the system.

Let H' be the set of hospitals who are matched to members of couples in M' . Let $\hat{\Gamma}$ be the market with all the singles, c , and the hospitals $(H \setminus H') \cup \{h'_1, h'_2\}$. Let \hat{M}' be the matching in $\hat{\Gamma}$ that matches every doctor d with the hospital d is assigned to under M' . As in the previous case, \hat{M}' is stable.

If c is matched to h'_1, h'_2 in M' , then it must be the case that c applied to them in SoDA, and was rejected. Let d_1, d_2 be the doctors who are matched to h'_1 and h'_2 , respectively, in SoDA when c applies (recall that the singles finished applying before the couples). We assume without loss of generality that h'_1 prefers d_1 to c_1 .

Since \hat{M}' is stable and h'_1 prefers d_1 to c_1 , it must be the case that d_1 prefers the hospital he is matched to under \hat{M}' to h'_1 . A contradiction is reached from the following claim.

Claim. Consider the submarket without couples but with all hospitals. Suppose that in the best stable matching for doctors (the one that SoDA outputs for this submarket), a doctor d is matched to hospital h . In any stable matching in the market with couples, d will not be assigned to a hospital that he strictly prefers to h .

PROOF. Consider a stable matching M with couples. Let M' be the matching obtained from M by restricted only to singles. Let Γ' be the submarket that contains only singles and all hospitals that do not accommodate members of couples under M (for capacities more than one, one needs only to reduce capacities in hospitals that have couples' members). Note that M' is stable under Γ' . But in Γ' the number of hospitals is smaller, and thus even in the best stable matching Γ' , d cannot get a hospital that he prefers to h . \square

Finally, for every single the ex ante probability that he will be evicted when couples enter the system tends to zero, and the ex ante probability that a couple will have an incoming edge also tends to zero as n grows, thus completing the proof.

A.4. SoDA Runs in Near-Linear Time

From the proof of Theorem 1 we can analyze the running time of (a slight modification of) the SoDA algorithm. Note that with high probability we have that the couples graph has small connected components (of size $< 3/\epsilon$) and can be topologically sorted. According to Corollary 3 each failed iteration of the SoDA algorithm is due to a backward edge in the insertion order π . By recording the backward edge, and ensuring that all future attempts

are consistent with it, we can guarantee that at most $(3/\epsilon)^2 \cdot n^{1-\epsilon}$ permutations will be tried before either a topologically sorted order is arrived at, or a cycle in the couples' graph is found.³⁵ \square

A.5. Proof of Theorem 2

We will use the same assumptions on influence trees as in the proof of Theorem 1, i.e., the assumptions hold except with probability $O(n^{-\epsilon/2})$. Informally, we will show that if a doctor or a couple does not interact with any other couple's influence tree, then she does not have an incentive to deviate. To this end we show the following:

LEMMA 6. *Let $d \in S$ be any doctor. Suppose that the SoDA algorithm terminates and assigns d to a hospital h in the first (deferred acceptance) stage of the algorithm. Suppose that h does not belong to any of the couples' influence trees. Then d may not improve her allocation under SoDA by misrepresenting her preferences.*

Similarly, if $c \in C$ is a couple whose influence tree is disjoint from all other influence trees, then c may not improve their allocation under SoDA by misrepresenting their preferences.

PROOF. We start with the first statement. At the end of the execution of the first stage of the SoDA algorithm d ends up in h . By Lemma 5, if d was moved from h , in the second stage, then h must belong to the influence tree of one of the couples, contradicting the assumption. Hence at the end of the SoDA algorithm d is still assigned the hospital h .

Suppose that d misrepresents her preferences and obtains a hospital h' such that $h' \succ_d h$ in a valid execution of the SoDA algorithm. It is well known that the outcome of the (regular) deferred acceptance algorithm on singles does not depend on the insertion order. Hence we can execute the SoDA algorithm so that d is the last single doctor to be inserted. Just before d is inserted, for all doctors d' that are assigned to h' , $d' \succ_{h'} d$, otherwise d would have been assigned h' when stating her true preferences. From that point on, a valid execution of the SoDA algorithm does not lead to any couples being evicted, and hence the quality of the least preferred doctor in h' according to $\succ_{h'}$ may only improve. Hence d may not be assigned to h' in the second phase of the SoDA algorithm. Contradiction.

Next, let $c = (f, m)$ be a couple such that $IT(c, r)$ is disjoint from all other influence trees. Suppose that c is assigned the hospitals (h_1, h_2) in a valid execution of the SoDA algorithm with an ordering π on couples. Since $IT(c, r)$ is disjoint from other influence trees, by Lemma 5 we see that inserting the couples in the order π' obtained from π by putting c first, leads to another valid execution that results in the same allocation.

Suppose that c misrepresents their preferences and obtains the hospitals $(h'_1, h'_2) \succ_c (h_1, h_2)$ in a valid execution of the SoDA algorithm. Note that the couple c was the first to be inserted under π' and did not get accepted into (h'_1, h'_2) because one of the hospitals preferred all the doctors that were assigned to it in the DA stage of the algorithm to the corresponding couple member. Without loss of generality, assume that h'_1 preferred all of its assigned doctors to f . As in the single doctor case above, in the second phase of the SoDA algorithm the least preferred doctor according to $\succ_{h'_1}$ that is assigned to h'_1 may only improve. Thus f may never be assigned to h'_1 . Contradiction. \square

Using Lemma 6 we can now prove Theorem 2.

PROOF (OF THEOREM 2). Fix any doctor $d \in S$ and the hospital h it is assigned in the DA stage of the SoDA algorithm. By an argument

very similar to Lemma 1 we can show that the probability that any influence tree contains h (or any other hospital in the influence tree of d) is bounded by $O(S(r)^2/n^\epsilon) < n^{-\epsilon/2}$. By Lemma 6, if this is the case, d does not have an incentive to deviate.

Similarly, the probability of the influence trees of two couples intersecting is bounded by $O(S(r)^2/n)$, and thus for each couple c , the probability that $IT(c, r)$ is disjoint from all other influence trees—and thus c has no incentive to deviate—is at least $1 - O(S(r)^2/n^\epsilon) > 1 - O(n^{-\epsilon/2})$. \square

Appendix B

B.1. Proof of Theorem 3

Consider the following event E : there exist a couple $c = (f_c, m_c) \in C$, a single doctor $s \in S$, and four hospitals $h_1 \neq h_2 \neq h_3 \neq h_4$ so that the most preferred pair of hospitals by c is (h_1, h_2) , the second most preferred by c is (h_3, h_4) , and the following properties hold:

- (i) $h_2 \succ_s h_1 \succ_s h$ for any $h \notin \{h_1, h_2\}$;
- (ii) $s \succ_{h_1} m_c$;
- (iii) $f_c \succ_{h_2} s$.

Observe that if only the couple c and the single doctor s existed no stable matching would exist.

The proof will follow by first bounding (from below) the probability of the event E and then bounding (from above) the event that some other doctor except those in the event E ever obtains either h_1 or h_2 in any stable matching.

Fix a couple $c \in C$ and a single s and let (h_1, h_2) be the pair of hospitals most preferred by and (h_3, h_4) be the second most preferred hospitals by c . The probability that $h_1 \neq h_2 \neq h_3 \neq h_4$, and properties (ii) and (iii) hold is $\delta > 1/32 \cdot 1/2^2$. The probability that $h_1 \neq h_2 \neq h_3 \neq h_4$ and Properties (i)–(iii) hold is $\Omega(\delta(1/(\lambda n)^2)) = \Omega(1/n^2)$ (see footnote 16 regarding Ω). Therefore, since there are n couples the probability that for a given single s there exists a couple c such that $h_1 \neq h_2$ and properties (i)–(iii) hold is $\Omega(1/n)$. Therefore since there are n singles, the probability that there exists a single s such that the event E holds is some constant $\gamma > 0$.

Suppose the event E occurs with the couple c' and doctor s' and let $D' = D \setminus \{f_{c'}, m_{c'}, s'\}$. Consider the following application/rejection algorithm in which doctors are assigned to $l > 0$ positions (rather than 1):

***l*-Pessimistic DA:** At each step $t = 1, 2, \dots$, either a single doctor $s \in S$ or a couple $c \in C$ that has less than l temporary assignments are chosen at random and applies to the most preferred hospital or pair of hospitals on their list, respectively, that they have not applied so far. Each hospital h assigns a doctor d if and only if no other doctor is currently assigned to h and no other doctor applied at this step to h . If some doctor d applies to h and some doctor d' (could be that $d' = d$) is temporarily assigned to h , h rejects both d and d' and no doctor is ever allowed to apply to h .³⁶

We will first show that the probability that any doctor but $f_{c'}$, $m_{c'}$ and s' ever applies to h_1 or h_2 in the 3-Pessimistic DA process is bounded from above by a small constant. We will show a stronger lemma:

LEMMA 7. *With constant probability no more than αn hospitals are visited in the process 3-Pessimistic DA for some $\alpha < \lambda$. In particular the 3-Pessimistic process terminates with constant probability.*

PROOF. Let $L = \{0, 1, 2, 3\}$. For every $q \in L$ we say that a doctor is q -settled if it is temporarily assigned to exactly q positions and we

say that a hospital h is *visited* if some doctor applied to it during the 3-Pessimistic DA process.

For every $t = 0, 1, 2, \dots$, and every $q \in L$ denote by A_t^q the number of q -settled doctors at step t , by V_t the number of visited hospitals up to step t , where $A_0^0 = 3n$, and $A_0^1 = A_0^2 = A_0^3 = V_0 = 0$. Let $Y_t = V_t + 15A_t^0 + 10A_t^1 + 5A_t^2$ and consider the process $X_t = Y_t + t$ for every $t = 0, \dots, \min(J, K)$, where K is the first step in which $V_K = (\lambda n)/10$ and J is the first step in which $A_J^0 = A_J^1 = A_J^2 = 0$ (i.e., $A_J^3 = 3n$).

Claim. X_1, X_2, \dots , is a super martingale, that is for every $t > 0$, $E[X_{t+1} | X_1, \dots, X_t] \leq X_t$.

PROOF. Suppose a couple c is chosen at step t and has $q \in L \setminus \{3\}$ temporary assignments. If it applies to two unvisited hospitals then $A_{t+1}^{q+1} = A_t^{q+1} + 2$ and $A_{t+1}^q = A_t^q - 2$ and $A_{t+1}^{q'} = A_t^{q'}$ for $q' \in L \setminus \{q, q+1\}$. Thus the contribution of the couple to Y_t drops by 10. If c applies to an unvisited hospital and one visited hospital then for every $q \in L$, $A_{t+1}^q \leq A_t^q + 2$ since at most one other couple lost a temporary assignment. If it applies to two visited hospitals then for every $q \in L$, $A_{t+1}^q \leq A_t^q + 4$ since at most two additional couples lose a temporary assignment. For singles similar bounds can be used. For each $q = 0, 1, 2$ let Q_t^q be the event that at the beginning of step t a couple with q temporary assignments is chosen, and by W_t^q the event that a single with q temporary assignments is chosen. Therefore for every $q \in L \setminus \{3\}$

$$\begin{aligned} E[X_{t+1} | X_1, \dots, X_t, Q_{t+1}^q] &= E[X_{t+1} | X_t, Q_{t+1}^q] \\ &\leq \frac{(\lambda n - V_t)^2}{(\lambda n)^2} (V_t + 2 + 15A_t^0 + 10A_t^1 + 5A_t^2 - 10) \\ &\quad + 2 \cdot \frac{(\lambda n - V_t)V_t}{(\lambda n)^2} (V_t + 1 + 15A_t^0 + 10A_t^1 + 5A_t^2 + 10) \\ &\quad + \frac{V_t^2}{(\lambda n)^2} (V_t + 15A_t^0 + 10A_t^1 + 5A_t^2 + 20) + t + 1 \\ &\leq V_t + 15A_t^0 + 10A_t^1 + 5A_t^2 + t, \end{aligned}$$

where the last inequality holds for any $V_t \leq (\lambda n)/10$. Similarly,

$$\begin{aligned} E[X_{t+1} | X_1, \dots, X_t, W_{t+1}^q] &= E[X_{t+1} | X_t, W_{t+1}^q] \\ &\leq \frac{(\lambda n - V_t)}{\lambda n} (V_t + 1 + 15A_t^0 + 10A_t^1 + 5A_t^2 - 5) \\ &\quad + \frac{V_t}{\lambda n} (V_t + 15A_t^0 + 10A_t^1 + 5A_t^2 + 10) + t + 1 \\ &\leq V_t + 15A_t^0 + 10A_t^1 + 5A_t^2 + t. \end{aligned}$$

Therefore, since either a couple or a single is chosen at each step, we obtain that $E[X_{t+1} | X_1, \dots, X_t] \leq V_t + 15A_t^0 + 10A_t^1 + 5A_t^2 + t$. \square

As argued in the claim $|X_{t+1} - X_t| < 22$ for every $t > 1$. Therefore, by Azuma-Hoeffding's inequality, we obtain that for any $T \geq 1$

$$\begin{aligned} \Pr\left(V_T - V_0 \geq \frac{\lambda n}{10}\right) &\leq \Pr\left(X_T - X_0 \geq \frac{\lambda n}{10} - 45n + T\right) \\ &\leq e^{-(\lambda n/10 - 45n + T)^2 / (968T)} < 1 - \beta, \end{aligned}$$

for some constant $\beta > 0$ and a sufficiently large λ , i.e., with constant probability the process will never reach $(\lambda n)/10$ visited hospitals. \square

Lemma 7 provides that in the 3-Pessimistic DA process described above, the number of hospitals visited is with constant probability only a fraction of the total number of hospitals, which implies that the doctors in the process (all but c' and s') will never visit h_1 and h_2 .

By Lemma 7 and the definition of 3-Pessimistic DA, each player i , single or couple, obtains three different temporary assignments, p_i^1 , p_i^2 , and p_i^3 (thus if i is a couple, p_i^j is a pair of hospitals) and observe that $p_i^1 \succ_i p_i^2 \succ_i p_i^3$.

To finish the proof we argue that in every stable matching, no agent i will be assigned to a pair of hospitals less preferred to p_i^3 . Call a player i (a single or a couple) that gets a hospital less preferred to p_i^3 *poor*, and let U be the set of poor players. Suppose that $|U| = k > 0$. For a player i to be poor, at least one hospital in each p_i^1 , p_i^2 and p_i^3 should be taken (if i is a single then all p_i^j are single hospitals and all should be taken). Since for each two players j, l , $\{p_j^1, p_j^2, p_j^3\} \cap \{p_l^1, p_l^2, p_l^3\} = \emptyset$ there are at least $3k$ hospitals that need to be assigned. These hospitals cannot be assigned to players that are not poor (since they get better choices for themselves) since p_i^1, p_i^2, p_i^3 are assigned exclusively to player i , i.e., no other doctor in the process ever applied to either of these three hospitals; indeed for each $j \neq i$ these hospitals are lower on player j 's list than p_j^1, p_j^2, p_j^3 , and the only way player j could occupy any of these hospitals is if he is poor.

Since there are only k poor players, with a total of up to $2k$ doctors, they cannot be assigned to all $3k$ hospitals—a contradiction.

The second part of the theorem follows since no doctor other than s, m_c and f_c will ever apply to h_1, h_2, h_3 and h_4 with constant probability since the preferences lists are bounded by a constant $k > 0$ (observe that the probability that a given hospital h will not be ranked by any doctor or couple is at least $(1 - 1/n)^{nk}$, which approaches e^{-k} as n tends to infinity). \square

Endnotes

1. In fact there were approximately 40,000 doctors, but only 16,000 of them were from American institutions and most couples were from American institutions. As we later discuss, since doctors who graduated in the United States are usually considered superior by the hospitals to doctors who graduated outside of the United States, the correct comparison is to the number of U.S. graduates and not to the number of participants.
2. Preferences that can be represented as a ranking list over the other side of the market.
3. Ronn (1990) showed that even determining whether a stable matching exists may be computationally intractable (NP-complete).
4. Formally, the number of doctors grows at a rate $n^{1-\epsilon}$. In fact ϵ may be even a slow decreasing function converging to zero.
5. In a recent study, Biró et al. (2011) investigate the matching problem with couples and compare various matching algorithms on simulated data.
6. Similar results have been obtained by Kelso and Crawford (1982) for such markets when firms can pay salaries to workers, by Gul and Stacchetti (1999) in auction settings when one considers the existence of a competitive equilibrium, and by Hatfield and Milgrom (2005) in a setting in which firms and workers sign contracts as well as many others. Interestingly, Echenique (2011) showed an equivalence result between contracts and salaries when preferences are substitutable.

7. We believe our techniques can be adapted to prove directly that the RP algorithm also succeeds with high probability in large random markets. See the last section for further discussion.
8. There are λn positions for some $\lambda > 1$.
9. All our positive results also hold (with similar proofs similar to the ones presented here) in the exact same model used by Kojima et al. (2013). The differences between our models are further discussed in §3.
10. This assumption is necessary to prevent too many hospitals to artificially be closed. Furthermore, in practice doctors can list hospitals they interviewed with, and therefore it is natural that only doctors can have unacceptable hospitals in a random model.
11. Our results can also extend to allow a blocking pair to involve a couple c and a hospital h such that one of the members of the couple is matched to h while the other is unassigned.
12. Step 3(b21) in the formal definition of SoDA.
13. Steps 2(a11), 2(a21), and 3(b22) in the formal definition.
14. In the formal definition of SoDA, all doctors in the set R in steps 2(a1), 2(a2), and 3(b2) are evicted by the applying couple c .
15. If preference lists for doctors are of constant length, our results hold also if for each couple $c = (f, m)$ after f and m draw their lists independently, their lists can be arbitrarily aggregated in a similar way as described in Kojima et al. (2013) (any joint preference list for c can be generated from m 's and f 's lists as long as no $h \in H$ that does not appear in f 's (h 's) list appears in the joint list for f (m)).
16. For any two functions f and g we write $f = \Omega(g)$ if $g = O(f)$.
17. This follows from the fact that in markets without couples, the set of stable matching forms a lattice. However, in a matching market with couples even if a stable matching exists, the set of stable matchings do not necessarily form a lattice.
18. A weakly connected component in a directed graph is a connected component in the graph obtained by removing the directions of the edges.
19. A topological sort π is an order over the couples such that no couple has an edge to a couple ahead of him in the order.
20. We do not rule out here that h was inserted to the influence tree by two different doctors. However, we will later show that the probability of this event is negligible.
21. In particular one will need to define influence trees for hospitals, show that with high probability a hospital does not encounter any couple, and with a bit of effort apply Lemma 10 in Kojima and Pathak (2009), which asserts the desired result for hospitals in markets without couples.
22. The results can be slightly improved by randomizing a new insert order each time the algorithm fails (doing this a small arbitrary number of times).
23. In fact in the NRMP more than 20,000 doctors participate, but 16,000 are from the United States and are ranked higher in the match.
24. The result is true also for αn couples for any constant $\alpha > 0$.
25. This question was raised by Federico Echenique.
26. In fact the existence result holds even without limiting the number of workers that can be hired in each firm.
27. The condition on P only assures that an eviction chain ends in a worker who is yet to be assigned with positive probability. To limit the size of influence trees a similar union bound method (as in the proof of Theorem 1) can be used.
28. Kojima et al. (2013) allowed to aggregate the preference lists almost arbitrarily (see footnote 15). However, since their preference lists are bounded by a constant, the probability that the members of

- a given couple will draw the same hospital is very small implying that in their joint list they will never apply to the same hospital (in our model, they may apply to the same hospital in different steps of the algorithm, but that will happen with low probability).
29. It is possible that if two influence trees intersect they will have other nodes (\tilde{h}, \tilde{d}) in common, since there might be common paths that continue from the point they intersect.
30. We write $\text{polylog } n$ for a polynomial in $\log n$. In particular $\text{polylog } n/n^{2\epsilon}$ tends to zero as n tends to infinity.
31. A set of nodes in an undirected graph is called a connected component if there exists a path between each to nodes in the set.
32. The presence of a cycle does not necessarily imply that there is no stable matching. In fact the SoDA will often find stable matchings even when there are cycles in the couples graph.
33. A backward edge is an edge from a newly inserted couple to a previously inserted one.
34. If capacities can be more than one, reduce capacity of h by the number of doctors that are members of a couples that assigned h .
35. It can be shown that the SoDA algorithm without this modification will run with at most $(3/\epsilon)^{3/\epsilon} \cdot n^{1-\epsilon}$ iterations.
36. As usual if a member of a couple is rejected from some hospital, its other member is also rejected.

References

- Biró P, Irving RW, Schlotter I (2011) Stable matching with couples: An empirical study. *J. Experiment. Algorithmics* 16:Article 1.2.
- Che Y, Kojima F (2010) Asymptotic equivalence of probabilistic serial and random priority mechanisms. *Econometrica* 78:1625–1672.
- Cripps MW, Swinkels JM (2006) Efficiency in large double auctions. *Econometrica* 74:47–92.
- Echenique F (2011) Contracts vs. salaries in matching. *Amer. Econom. Rev.* 102(1):594–601.
- Gale D, Shapley LL (1962) College admissions and the stability of marriage. *Amer. Math. Monthly* 69:9–15.
- Gul F, Stacchetti E (1999) Walrasian equilibrium with gross substitutes. *J. Econom. Theory* 87:95–124.
- Hatfield JW, Kojima F (2010) Substitutes and stability for matching with contracts. *J. Econom. Theory*. 145:1704–1723.
- Hatfield JW, Kominers SD (2014) Multilateral matching. *J. Econom. Theory*. Forthcoming.
- Hatfield JW, Milgrom PR (2005) Matching with contracts. *Amer. Econom. Rev.* 95:913–935.
- Immorlica N, Mahdian M (2005) Marriage, honesty, and stability. *Proc. Sixteenth Annual ACM-SIAM Sympos. Discrete Algorithms* (SIAM, Philadelphia), 53–62.
- Jackson MO, Manelli AM (1997) Approximately competitive equilibria in large finite economies. *J. Econom. Theory* 77:354–376.
- Jackson MO, Swinkels JM (2005) Existence of equilibrium in single and double private value auctions. *Econometrica* 73:93–139.
- Klaus B, Klijn F (2005) Stable matchings and preferences of couples. *J. Econom. Theory* 121:75–106.
- Kelso AS, Crawford VP (1982) Coalition formulation, and gross substitutes.. *Econometrica* 50:1483–1504.
- Kojima F, Manea M (2010) Incentives in the probabilistic serial mechanism. *J. Econom. Theory* 145:106–123.
- Kojima F, Pathak PA (2009) Incentives and stability in large two-sided matching markets. *Amer. Econom. Rev.* 99:608–627.
- Kojima F, Pathak P, Roth AE (2013) Matching with couples: Stability and incentives in large markets. *Quart. J. Econom.* 128(4):1585–1632.
- Lahaie S, Parkes DC (2009) Fair package assignment. *Proc. First Conf. Auctions, Market Mechanisms and Their Appl.*

- Lee S (2011) Incentive compatibility of large centralized matching markets. Working paper.
- Milgrom PR (2004) *Putting Auction Theory to Work* (Cambridge University Press, Cambridge, UK).
- Ostrovsky M (2008) Stability in supply chain networks. *Amer. Econom. Rev.* 98(3):897–923.
- Pycia M (2012) Stability and preference alignment in matching and coalition formation. *Econometrica* 80(1):323–362.
- Roberts DJ, Postlewaite A (1976) The incentives for price-taking behavior in large exchange economies. *Econometrica* 44:115–127.
- Ronn E (1990) NP-complete stable matching problems. *J. Algorithms* 11:285–304.
- Roth AE (1982) The economics of matching: Stability and incentives. *Math. Oper. Res.* 7(4):617–628.
- Roth AE (1984) The evolution of the labor market for medical interns and residents: A case study in game theory. *J. Political Econom.* 92:991–1016.
- Roth AE (2009) The origins, history, and design of the resident match. The origins, history, and design of the resident match. *J. Amer. Medical Assoc.* 289:909–912.
- Roth AE, Peranson E (1999) The redesign of the matching market for American physicians: Some engineering aspects of economic design. *Amer. Econom. Rev.* 89:748–780.
- Roth AE, Sotomayor MAO (1990) *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis* (Cambridge University Press, Cambridge, UK).
- Rustichini A, Satterthwaite M, Williams S (1994) Convergence to efficiency in a simple market with incomplete information. *Econometrica* 62:1041–1064.
- Satterthwaite MA, Williams SR (1989) The rate of convergence to efficiency in the buyer's bid double auction as the market becomes large. *Rev. Econom. Stud.* 56:477–498.
- Sun N, Yang Z (2006) Equilibria and indivisibilities: Gross substitutes and complements. *Econometrica* 74:1385–1402.

Itai Ashlagi is an assistant professor of operations management at the Sloan School of Management of Massachusetts Institute of Technology. He is interested in topics on the border of microeconomic theory and operations research and in particular focuses on topics in market design and game theory.

Mark Braverman is an assistant professor of computer science at Princeton University. His research interests include complexity theory, information theory, algorithms, and applications of computer science and game theory in healthcare and medicine.

Avinatan Hassidim is an assistant professor in the computer science department of Bar Ilan University. His research interests lie in the intersection of computer science and economics, with a focus on market design, labor markets, and online marketplaces.