Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence

Matching with Dynamic Ordinal Preferences

Hadi Hosseini
Cheriton School of Computer Science
University of Waterloo
Waterloo, ON, Canada
h5hossei@uwaterloo.ca

Kate Larson
Cheriton School of Computer Science
University of Waterloo
Waterloo, ON, Canada
klarson@uwaterloo.ca

Robin Cohen
Cheriton School of Computer Science
University of Waterloo
Waterloo, ON, Canada
rcohen@uwaterloo.ca

Abstract

We consider the problem of repeatedly matching a set of alternatives to a set of agents with dynamic ordinal preferences. Despite a recent focus on designing one-shot matching mechanisms in the absence of monetary transfers, little study has been done on strategic behavior of agents in sequential assignment problems. We formulate a generic dynamic matching problem via a sequential stochastic matching process. We design a mechanism based on random serial dictatorship (RSD) that, given any history of preferences and matching decisions, guarantees global stochastic strategyproofness while satisfying desirable local properties. We further investigate the notion of envyfreeness in such sequential settings.

Introduction

One-sided matching problems have been extensively studied in the context of microeconomics, artificial intelligence, and mechanism design. These problems arise in various real-life application domains such as assigning dormitory rooms to college students, teaching load among faculty, college courses to students, medical residents to hospitals, scarce medical resources and organs to patients, *etc.* (Sönmez and Ünver 2010a; 2010b; Roth, Sönmez, and Ünver 2004; Ünver 2010; Dickerson, Procaccia, and Sandholm 2012; Ashlagi et al. 2013; Krysta et al. 2014).

Despite the interest in matching problems, little has been done in dynamic settings where agents' preferences evolve over time. In the real world, decisions do not exist in isolation, but rather are situated in a temporal context with other decisions and possibly stochastic events. Dynamic mechanism design (Parkes 2007) is a compelling research area that has attracted attention in recent years. In these settings, agents act to improve their outcomes over time, and decisions both in the present and in the past influence how the world and preferences look in the future. The dynamic pivot mechanism for dynamic auctions (Bergemann and Välimäki 2010), dynamic Groves mechanisms (Cavallo 2009), and many others (Vohra 2012; Athey and Segal 2013) are a few of myriad examples of mechanisms in dynamic settings that consider agents with private dynamic preferences. However, almost all of these works (excluding a recent study on dynamic social choice (Parkes and Procaccia 2013)) assume an underlying utility function with possible utility transfers.

We study dynamic matching problems in which a sequence of decisions must be made for agents whose private underlying preferences may change over time. In each period, the mechanism elicits ordinal preferences from agents, and each agent declares his preferences (truthfully or strategically) so as to improve his overall assignment, now or in the future. We propose a generic model to study the various properties of such environments including the strategic behavior of agents. Our model captures a diverse set of scenarios in real-life settings: assigning members to subcommittees each year, tasks among team members on various projects, nurses to various hospital shifts, teaching loads among faculty, and students to college housing each year.

Consider the problem of scheduling nurses to shifts (nurse rostering) in multiple planning periods.[1] Self rostering is one of the most advocated methods in nurse scheduling that caters to individual preferences (Siferd and Benton 1992; Burke et al. 2004). Each nurse has some internal preference over hospital shifts at various times. At each planning period (typically 4–6 weeks), an assignment decision is made based on the self-reported preferences. Although self rostering reduces administrative burden and improves nurse satisfaction, there is evidence that it encourages strategic behavior among nurses (Alsheddy and Tsang 2011; De Grano, Medeiros, and Eitel 2009). The preferences of nurses may change dynamically according to their internal desires and past assignments and "those who are savvy enough to game the system will always have an advantage over the procrastinators" (Bard and Purnomo 2005). This example and many other real-life applications raise several intriguing questions when designing matching mechanisms in dynamic and uncertain settings.

Our model and results We consider a setting where a sequence of assignments should be made for fixed number of agents and alternatives. Agents' preferences are represented as strict orderings over a set of alternatives, where these preferences may change over time. We formulate a general dynamic matching problem using a *history-dependent matching process*. The state of the matching process corresponds

[1] Assume skill category and hospital constraints are incorporated into shift schedules.

to a history of preference profiles and matching decisions.

We show that simply running a sequence of independent assignments induced by the random serial dictatorship (RSD) mechanism does not satisfy global strategyproofness. Subsequently, we design a stochastic matching policy that is not manipulable. Our key idea is to extend a notion first introduced for multi-period contracts (Townsend 1982), where future matching decisions are used to incentivize desirable behavior in the current time period. Our main result is a mechanism that, given any history of preferences and decisions and any preference dynamics, guarantees stochastic strategyproofness while satisfying desirable local properties of ex post efficiency, *sd*-strategyproofness, and equitability. We further formulate the notion of envy in the context of sequential matching by providing a systematic way of measuring the degree of envy towards individual agents, arguing that our mechanism provides a constant degree of individual envy by balancing priority orderings.

Related work In cardinal domains no matching mechanism exists that can satisfy ex ante Pareto efficiency, anonymity, and strategyproofness (Zhou 1990). Deterministic serial dictatorship mechanisms sacrifice fairness and anonymity for Pareto efficiency and strategyproofness. In fact, Svensson showed that under ordinal preferences the serial dictatorship mechanism is the only deterministic mechanism that is strategyproof, nonbossy, and neutral (Svensson 1999). It has been noted that randomization helps to restore strategyproofness and equity in assignment problems (Abdulkadiroğlu and Sönmez 1998; Gibbard 1977), several other measures such as efficiency and strategyproofness remain incompatible. With ordinal preferences, this incompatibility result persists (Bogomolnaia and Moulin 2001), and essentially states that no efficient solution can be trusted. That is, the outcome of any efficient mechanism may only be efficient with respect to the non-truthful preferences, and hence, be inefficient under true underlying preferences.

The random serial dictatorship (RSD) is an ordinal mechanism that is strategyproof, anonymous, and equitable but does not satisfy stochastic efficiency. Bogomolnaia and Moulin (2001) noted the inefficiency of RSD, and characterized an ordinal efficiency measure based on first-order stochastic dominance (*sd*). They proposed the probabilistic serial (PS) mechanism that is *sd*-efficient and *sd*-envyfree, but does not satisfy *sd*-strategyproofness. This incompatibility result immediately applies to majoritarian scheme in social choice literature (Aziz, Brandt, and Brill 2013). Therefore, in this paper we focus our attention only on RSD (or random priority) as a strategyproof matching mechanism.

In matching markets, dynamic arrival and departure of agents have been studied in various contexts such as campus housing and organ transplant, which assume time-invariant preferences (Kurino 2014; Bloch and Cantala 2013; Ünver 2010). Bade studied matching problems with endogenous information acquisition, and showed that simple serial dictatorship is the only ex ante Pareto optimal, strategyproof, and non-bossy mechanism when agents reveal their preferences only after acquiring information about those with higher priorities (Bade 2014).

A Matching Model for Dynamic Ordinal Preferences

In this section we introduce our model for dynamic preferences. We start by introducing key preference and matching terminologies for a single time step in the model. We then generalize to the dynamic setting studied in this paper.

There is a set $N = \{1, \dots, n\}$ of agents who have preferences over a finite set of alternatives $M = \{1, \dots, m\}$, where $n \geq m$.[2] Agents have preferences over alternatives, and we use the notation $a \succ_i^t b$ to mean that agent i strictly prefers alternative a to b at time t. We let $\mathcal{P}(M)$ or $\mathcal{P}$ denote the class of all strict linear preferences over M where $|\mathcal{P}| = m!$. Agent i's preference at time t is denoted by $\succ_i^t \in \mathcal{P}$, thus, $\succ^t = (\succ_1^t, \dots, \succ_n^t) \in \mathcal{P}^n$ denotes the *preference profile* of agents at time t. We write $\succ_{-i}^t$ to denote $(\succ_1^t, \dots, \succ_{i-1}^t, \succ_{i+1}^t, \dots, \succ_n^t)$, and thus $\succ^t = (\succ_i^t, \succ_{-i}^t)$.

A *matching* at time t, $\mu^t : N \rightarrow M$, is a bijective mapping, and we let $\mu^t(i)$ denote the alternative allocated to agent i under matching μ^t at time t. A matching is *feasible* at time t if and only if for all $i, j \in N$, $\mu^t(i) \neq \mu^t(j)$ when $i \neq j$. We let $\mathcal{M}$ denote the set of all feasible matchings over the set of alternatives M. We also allow for randomization where $\bar{\mu}^t$ denotes a probability distribution over the set of (deterministic) feasible matchings at time t . That is, $\bar{\mu}^t \in \Delta(\mathcal{M})$.

In this paper we are interested in settings where agents' preferences *evolve* over time. In particular, we assume that the preference held by an agent at time t depends on the preferences it held earlier along with allocations (*i.e.* matchings) made previously. We let h^t denote the joint history of the joint states defined by agents' preferences and realized matchings up to time $t-1$ and including the joint preferences at time t, that is, $h^t = (\succ^1, \mu^1, \dots, \succ^{t-1}, \mu^{t-1}, \succ^t)$. The set $\mathcal{H}^t$ contains all possible joint histories at time t. We assume there is an underlying stochastic kernel $P(\succ^{t+1} \mid \succ^t, \mu^t)$ which denotes the probability that agents will transition to a state where they have joint preference $\succ^{t+1}$ after matching decision μ^t in a state with joint preference $\succ^t$. In a more generic model, transitions could potentially be influenced by the complete history of preferences and decisions. Nonetheless, our analysis of the agents' reporting strategies considers generic strategic behaviors, and thus, holds under this assumption. Furthermore, we assume the stochastic kernel is common knowledge.

Since an agent's preferences can change, it may be the case that $a \succ_i^t b$ while $b \succ_i^{t+1} a$. For the sake of clarity, at each time t we rank alternatives under preference ordering $\succ_i^t$ such that $o^1 \succ_i^t o^2 \succ_i^t \dots \succ_i^t o^m$ denotes the ranking positions, where the index of each alternative indicates its ranking for agent i at time t.

We consider a discrete-time *sequential matching decision process* as a sequence of matchings prescribed by a matching policy. Given a history $h^t \in \mathcal{H}^t$, a *matching policy* $\pi(\mu | h^t)$ returns the probability of applying matching μ. Given a matching policy, π, and a history h^t the probability

[2]We accommodate the possibility of $n > m$ by adding dummy alternatives corresponding to a null assignment.

of agent i being allocated alternative x at time t is

$$p_i^t(x \mid h^t) = \sum_{\mu \in \mathcal{M}:\mu(i)=x} \pi(\mu \mid h^t). \tag{1}$$

where $\sum_{i \in N} p_i^t(x \mid h^t) = 1$. The definition of a matching policy, π, incorporates randomized or deterministic matching policies. When it is clear from the context, we will abuse notation and use $\pi(h^t)$ to also refer to the (random) matching prescribed by policy π given the history h^t.

Policy Evaluation To determine whether a particular matching policy, π, is a *good* policy, one must make comparisons between policies. In this paper we are interested in settings where agents have *ordinal* preferences and so do not rely on particular utility functions. Instead, to evaluate a policy π we look at weights which can be interpreted as expected probabilities of being allocated particular alternatives in the sequence of random matchings from time t onward. More concretely, let o^ℓ be any alternative ranked in position ℓ. Given h^t the expected probability that agent i receives alternatives with rankings as good as ℓ under matching policy π is defined recursively as

$$W_i^\pi(h^t, o^\ell) = \sum_{x=o^1}^{o^\ell} p_i^t(x|h^t) \times \tag{2}$$

$$\sum_{\mu \in \mathcal{M}} \sum_{\succ^{t+1} \in \mathcal{P}^n} \pi(\mu|h^t) P(\succ^{t+1} \mid \succ^t, \mu) W_i^\pi(h^{t+1}, o^\ell)$$

where $h^{t+1} = (h^t, \mu^t, \succ^{t+1})$ is the history at time $t+1$. Intuitively, $W_i^\pi(h^t, o^\ell)$ is the probability that agent i receives alternatives as good as ℓ in current period and all future sequences of matchings up to any desired planning horizon.

Example 1. *Consider assigning 3 objects in two decision periods to 3 agents with preference orderings, as shown in Figure 1, at periods 1 and 2. For simplicity assume deterministic transitions that are independent of the realized matching decisions, and let $h^1 = \succ^1$ so that the matching starts at $t = 1$. A matching $\mu = xyz$ denotes that agents $1, 2, 3$ receive objects x, y, z respectively. Consider a policy π that prescribes random matchings $\bar{\mu}^1 = (\frac{1}{2}abc, \frac{1}{2}cba)$ and $\bar{\mu}^2 = (\frac{1}{2}abc, \frac{1}{3}acb, \frac{1}{6}cab)$ at periods 1 and 2 respectively. Using Equation 2, the expected probability that agent 1 receives his first rank alternatives (i.e. o^1) in periods 1 and 2 given the above random decisions is calculated by (highlighted in pink in Figure 1):*

$$W_1^\pi(\succ^1, o^1) = (\tfrac{1}{2}) \times (\tfrac{1}{2} + \tfrac{1}{3}) = \tfrac{5}{12}$$

Similarly, we can compute the probability that agent 1 receives alternatives with rankings as good as rank 2 (i.e. o^1 and o^2) in periods 1 and 2:

$$W_1^\pi(\succ^1, o^2) = (\tfrac{1}{2} + \tfrac{1}{2}) \times (\tfrac{1}{2} + \tfrac{1}{3}) = \tfrac{5}{6}$$

Intuitively, $W_1^\pi(\succ^1, o^2)$ is the probability of all possible random sequences in which agent 1 receives alternatives that are ranked 1 or 2, that is, the following possible sequences $[\mu^1(1) = o^1, \mu^2(1) = o^1]$, $[\mu^1(1) = o^1, \mu^2(1) = o^2]$, $[\mu^1(1) = o^2, \mu^2(1) = o^1]$, $[\mu^1(1) = o^2, \mu^2(1) = o^2]$ (highlighted in Figure 1).

$$\begin{pmatrix} \succ_1^1: & a \succ c \succ b \\ \succ_2^1: & b \succ c \succ a \\ \succ_3^1: & a \succ c \succ b \end{pmatrix}$$

$$\begin{pmatrix} \succ_1^2: & a \succ b \succ c \\ \succ_2^2: & b \succ a \succ c \\ \succ_3^2: & b \succ c \succ a \end{pmatrix}$$

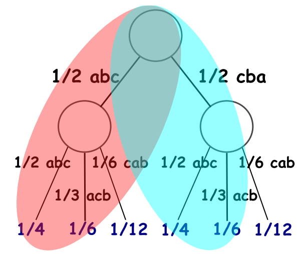

Figure 1: The preferences of agents and the probablity tree for the two-time matching problem in Example 1.

Agents' Strategies

We restrict attention to *dynamic mechanisms* in which each agent interacts with the mechanism simply by declaring his (perhaps untruthful) preference ordering at each decision period. In each period t, the mechanism elicits reports from agents regarding private preferences. Each agent observes his private preference $\succ_i^t$, and the history of past preferences and realized matchings h^{t-1}, and based on the underlying probabilistic kernel, takes an action by declaring a preference ordering for time t, $\hat{\succ}_i^t$. Then, the mechanism draws a matching $\mu^t \in \mathcal{M}$ according to matching policy π. Agent i's reporting *strategy* $\sigma_i(\succ_i^t)$ specifies his declared preference $\hat{\succ}_i^t = \sigma_i(\succ_i^t)$ when his true private preference is $\succ_i^t$.

Properties for the Model

Our goal is to implement stochastic matching policies so that agents truthfully reveal their preferences, no matter what other agents do, now or in the future. More specifically, we are interested in matching policies that satisfy global strategyproofness while inducing a sequence of locally strategyproof and ex post efficient random matchings. In this section, we formally define these local and global properties.

Stochastic dominance (sd) for two matching policies π and π' prescribes that given a transition model, for each rank ℓ, the expected probability that alternatives with rankings as good as ℓ get selected under π, is greater or equal to the expected probability that π' selects such alternatives.

Definition 1. *Given a transition model P, matching policy π **stochastically dominates (sd)** π', if at all states $h^t \in \mathcal{H}^t$, for all agents $i \in N$,*

$$\forall o^\ell \in M,\ W_i^\pi(h^t, o^\ell) \geq W_i^{\pi'}(h^t, o^\ell) \tag{3}$$

Strategyproofness in sequential settings (repeated assignments) prescribes that not only an agent cannot improve her current immediate outcome by misreporting her preferences but also the agent's current misreport will not make her better off in the future.

Definition 2. *A matching policy is **globally sd-strategyproof (gsd-strategyproof)** if and only if truthfulness is a stochastic dominant strategy for all possible realizations, i.e., for any transition kernel P, given any misreport $\hat{\succ}_i^t = \sigma_i(\succ_i^t)$ such that $\hat{h}_i^t = (\succ^1, \ldots, (\hat{\succ}_i^t, \succ_{-i}^t); \mu^1, \ldots, \mu^{t-1})$ at time t, for all agents $i \in N$,*

$$\forall o^\ell \in M,\ W_i^\pi(h^t, o^\ell) \geq W_i^\pi(\hat{h}^t, o^\ell) \tag{4}$$

Global *sd*-strategyproofness is an incentive requirement which states that under any possible transition of preference profiles, no agent can improve her sequence of random matchings (now or in future) by a strategic report.

A one-shot matching process is a special case of our model that coincides precisely with the random assignment problem (Abdulkadiroğlu and Sönmez 1998; Bogomolnaia and Moulin 2001). If there is a single decision period, agents become myopic. Thus, a random matching induced by matching mechanism π stochastically dominates another random matching induced by π' if for each item $x \in M$, the probability of selecting an outcome as good as x by π is greater than or equal to π'.

Definition 3. *Given a preference profile $\succ^t$, a random matching induced by π* ***stochastically dominates (sd)*** *another random matching prescribed by π', if for all agents $i \in N$,*

$$\forall y \in M, \sum_{\substack{x\in M:\\ x\succ_i^t y}} \sum_{\substack{\mu\in\mathcal{M}:\\ \mu(i)=x}} \pi(\mu|\succ^t) \geq \sum_{\substack{x\in M:\\ x\succ_i^t y}} \sum_{\substack{\mu\in\mathcal{M}:\\ \mu(i)=x}} \pi'(\mu|\succ^t)$$

In other words, in a single-shot setting, every agent with any utility model that is consistent with the private ordinal preferences will prefer π over π' if the random matching selected by π first-order stochastically dominates the random matching selected by π'.

A matching policy is *locally sd-strategyproof* in period t, if an agent's truthful report always results in a random matching that stochastically dominates his random matching under an untruthful misreport. In other words, at each period t no agent can improve his current random assignment by a strategic misreport.

Definition 4. *A matching policy π is* ***locally sd-strategyproof (lsd-strategyproof)*** *if and only if truthfulness is a stochastic dominant strategy for time t, that is, for any $\succ^t \in \mathcal{P}^n$, for all agents $i \in N$, $\forall y \in M$*

$$\sum_{\substack{x\in M:\\ x\succ_i^t y}} \sum_{\substack{\mu\in\mathcal{M}:\\ \mu(i)=x}} \pi(\mu|(\succ_i^t, \succ_{-i}^t)) \geq \sum_{\substack{x\in M:\\ x\succ_i^t y}} \sum_{\substack{\mu\in\mathcal{M}:\\ \mu(i)=x}} \pi(\mu|(\sigma_i(\succ_i^t), \succ_{-i}^t))$$

A matching is *Pareto efficient* if there is no other matching that makes all agents weakly better off and at least one agent strictly better off.

Definition 5. *A random matching is* ***ex post efficient*** *if it can be represented as a probability distribution over Pareto efficient deterministic matchings.*

Sequential RSD

In this section, we briefly describe the random serial dictatorship (RSD) (Abdulkadiroğlu and Sönmez 1998) mechanism for one-shot settings, and argue that a sequence of RSD-induced decisions is manipulable when preferences are dynamic.

To formally define the RSD mechanism, we first introduce priority orderings and serial dictatorships. A *priority ordering* $f : \{1, \ldots, n\} \rightarrow N$ is a one-to-one mapping that specifies an ordering of agents: agent $f(1)$ is ordered first, agent $f(2)$ is ordered second, and so on.

Given a priority ordering $f \in \mathcal{F}$ and a preference profile $\succ$, a **Serial Dictatorship**, $SD(f, \succ)$, is as follows: agent $f(1)$ receives her favorite object $m_1 \in M$ according to $\succ_{f(1)}$; $f(2)$ receives her favorite object $m_2 \in M \backslash \{m_1\}$; $f(n)$ receives her best object $m_n \in M \backslash \{m_1, \ldots, m_{n-1}\}$.

(a) Truthful		(b) Misreport	
$\succ_1$:	$a \succ c \succ b$	$\succ_1$:	$a \succ c \succ b$
$\succ_2$:	$b \succ c \succ a$	$\succ_2$:	$b \succ c \succ a$
$\succ_3$:	$a \succ c \succ b$	$\hat{\succ}_3$:	$a \succ b \succ c$

Table 1: Preferences revealed by three agents.

Random serial dictatorship is a convex combination of all feasible serial dictatorships, induced by a uniform distribution over all priority orderings, and it is formally defined as $RSD(\succ) = \frac{1}{n!} \sum_{f \in \mathcal{F}} SD(f, \succ)$.

In single-shot settings, RSD is *sd*-strategyproof, equitable (in terms of equal treatment of equals), and ex post efficient (Abdulkadiroğlu and Sönmez 1998). While RSD satisfies *lsd*-strategyproofness at each decision period, we argue that a stochastic policy consisting of a sequence of RSD induced random matchings (or *sequential RSD*) is prone to manipulation when agents have dynamic preferences. Sequential RSD selects a random matching at each round independent of the past history of decisions and preferences.

Theorem 1. *Sequential RSD (a sequence of RSD induced matchings) is not gsd-strategyproof.*

Proof. Consider 2 decision periods with deterministic preference dynamics known to agents. Let $\bar{\mu}^2_{\mu_1}$ denote the random matching at $t = 2$ after assignment μ_1 at $t = 1$. With truthful preferences (Table 1a), RSD induces the following random matching: $(\frac{1}{2}\mu_1, \frac{1}{2}\mu_2) = (\frac{1}{2}abc, \frac{1}{2}cba)$. If agent 3 misreports (Table 1b), the probability distribution would be $(\frac{2}{6}\mu_1, \frac{3}{6}\mu_2, \frac{1}{6}\mu_3) = (\frac{2}{6}abc, \frac{3}{6}cba, \frac{1}{6}acb)$. Assuming truthfulness in the second period, given a preference profile, identical decisions always result in identical next states. For each ranking position ℓ, we compute $W_3(\cdot, o^\ell)$:

	$W_3(\succ^t, o^\ell)$	$W_3((\hat{\succ}_i^t, \succ_{-i}^t), o^\ell)$
o^1	$\frac{3}{6} \times \frac{3}{6}(\bar{\mu}^2_{\mu_1})$	$\frac{3}{6} \times \frac{3}{6}(\bar{\mu}^2_{\mu_1})$
o^2	$1 \times \frac{3}{6}[\bar{\mu}^2_{\mu_1} + \bar{\mu}^2_{\mu_2}]$	$\frac{5}{6} \times [\frac{3}{6}(\bar{\mu}^2_{\mu_1}) + \frac{2}{6}(\bar{\mu}^2_{\mu_2})]$
o^3	$1 \times \frac{3}{6}[\bar{\mu}^2_{\mu_1} + \bar{\mu}^2_{\mu_2}]$	$1 \times [\frac{3}{6}(\bar{\mu}^2_{\mu_1}) + \frac{2}{6}(\bar{\mu}^2_{\mu_2}) + \frac{1}{6}(\bar{\mu}^2_{\mu_3})]$

Table 2: Evaluation of the matching policy for agent 3.

For o^1 and o^2, it is easy to see that the truthful revelation stochastically dominates the matching decisions when agent 3 is non-truthful. For strategyproofness we must show that for o^3, $1 \times \frac{3}{6}[\bar{\mu}^2_{\mu_1} + \bar{\mu}^2_{\mu_2}] \geq 1 \times [\frac{3}{6}(\bar{\mu}^2_{\mu_1}) + \frac{2}{6}(\bar{\mu}^2_{\mu_2}) + \frac{1}{6}(\bar{\mu}^2_{\mu_3})]$. By simple algebra, we see that for all preferences at the second period wherein agent 3's assignment under $\bar{\mu}^2_{\mu_3}$ stochastically dominates $\bar{\mu}^2_{\mu_2}$ the above inequality does not hold. Therefore, sequential RSD is not *gsd*-strategyproof. □

Sequential RSD with Adjusted Priorities

In the previous section we showed that sequential RSD is prone to manipulation when agents have dynamic preferences. In this section, we introduce a modification of RSD, which uses information contained in histories of agents to overcome strategic behavior. We start this section with some observations about relationships between agents.

A key property of RSD is that it prioritizes agents in each round, and the agent with higher priority gets to choose its more preferred item from the set of *remaining objects*. This gives rise to the concept of dictatorial domination.

Definition 6. *Given a preference profile and the realization of a matching decision, we say that agent i* ***dictatorially dominates*** *agent j at time t if and only if* $\mu^t(i) \succ_j^t \mu^t(j)$.

We can represent each agent's dictatorial dominance as a binary relation between every pair of agents. Each agent's dictatorial dominance on other agents is represented by $\boldsymbol{\omega}_i = (\omega_{i,1}, \ldots, \omega_{i,n})$, where $\omega_{i,j} = 1$ denotes that agent i has dominated agent j. A *dominance profile* is a matrix of agents' dictatorial dominances, $\boldsymbol{\omega} = (\boldsymbol{\omega}_i, \ldots, \boldsymbol{\omega}_n)$.

Given a random matching mechanism, the probability that agent i dominates agent j is equal to the sum of the probabilities of all deterministic matchings wherein agent j prefers the outcome of agent i to his own outcome.[3] Given a random matching policy π and h^t, the probability of i dominating j at period t is:

$$\omega_{ij}(\pi(h^t)) = \sum_{\mu \in \mathcal{M}: \mu(i) \succ_j^t \mu(j)} \pi(\mu \mid h^t) \tag{5}$$

Similarly the probability that agent j prefers his own outcome to agent i's outcome is $\bar{\omega}_{ij}(\pi(h^t)) = 1 - \omega_{ij}(\pi(h^t))$.

RSD ensures equal chance of dictatorships to agents, thus, an agent's strategic misreport can only increase his random dictatorial dominance on another agent to $\frac{1}{2}$.

Proposition 1. *Given RSD, for any* $\succ^t \in \mathcal{P}^n$ *we have* $\forall i, j \in N, \omega_{ij}(\pi(\succ^t)) \leq \frac{1}{2}$*, that is, the probability of agent i dominating another agent j is always bounded.*

The intuition comes from the fact that RSD is a uniform distribution over all priority orderings. Thus, when two agents have conflicting preferences over some alternatives, RSD assigns equal probability to all orderings that prioritize one agent lower than the other one.

Adjusted RSD

In this section, we provide a modification to the RSD mechanism based on adjusting agents' priority ordering. We introduce a simple structure to preserve the history of dominations throughout the matching process. Formally, let $\boldsymbol{d}^{t-1}$ be a matrix representing the complete *dominance history* of agents up to time t. As the dominance profile of agents is attained after realization of decisions (ex post), the state of the mechanism at t can be summarized by $h^t = (\succ^t, \boldsymbol{d}^{t-1})$. As shown in Algorithm 1, ARSD runs as follows:

- If for a pair of agents the dominance history is unbalanced, that is, $d_{ij}^{t-1} = 1$, then for each priority ordering $f \in \mathcal{F}$, if agent j is ranked before i, we change the priority ordering to f' such that j has higher priority than agent i, and add the new priority ordering to a new multiset $\mathcal{F}'$.
- The matching mechanism then draws a priority ordering from a uniform probability distribution over all priority orderings in the multiset $\mathcal{F}'$. Then, agents select alternatives according to the selected ordering.

[3]Deterministic matching is a special case where given μ^t, $\omega_{i,j} = 1$ iff $\mu^t(i) \succ_j^t \mu^t(j)$.

Algorithm 1: RSD with adjusted priorities (ARSD)

```
Input: preference profile ≻^t, dominance history d^{t-1}
Output: A probability distribution prescribed by π
foreach priority ordering f ∈ F do
    for i = 1 to n do
        for j = 1 to n do
            if f(i) < f(j) and d_ij^{t-1} = 1 then
                f' ← swap(f(i), f(j));
    foreach μ ∈ M do
        if SD(f', ≻^t) = μ then
            π(μ | (≻^t, d^{t-1})) ← π(μ | (≻^t, d^{t-1})) + 1/n! ;
// Updating d^t after realization of π
for i = 1 to n do
    for j = 1 to n do
        d_ij^t = ω_ij^t(μ^t) ⊕ d_ji^{t-1};
```

- After the realization of matching decision, and given the dominance profile $\boldsymbol{\omega}^t$ at time t, the mechanism updates the dominance history according to the following exclusive disjunction: $\boldsymbol{d}^t = \boldsymbol{\omega}^t \oplus (\boldsymbol{d}^{t-1})^{Tr}$, where $(\boldsymbol{d}^{t-1})^{Tr}$ is the transpose matrix of $\boldsymbol{d}^{t-1}$, *i.e.*, for each element d_{ij}^{t-1} the transpose element would be d_{ji}^{t-1}.

Proposition 2. *When the dominance history is balanced, i.e.* $d_{ij}^{t-1} = 0, \forall i, j \in N$*, ARSD is equivalent to RSD.*

Proof. According to Algorithm 1, for each priority ordering $f \in \mathcal{F}$ since for all $i, j \in N$ the dominance history is balanced, i.e. $d_{ij}^{t-1} = 0$, it is easy to see that the updated priority ordering $f' = f$. Thus at time t, the multiset $\mathcal{F}'$ is equivalent to the set of all priority orderings $\mathcal{F}$, and the a uniform distribution over $\mathcal{F}'$ is equivalent to the RSD mechanism. □

ARSD and Strategic Behavior At each period, an agent's strategic behavior is either through manipulating the immediate outcome (at the current period) or affecting the decision trajectory to gain advantage sometime in the future. We first focus on strategic behavior of agents at the current state and then show that no agent can obtain a more preferred outcome (now or in future) by misrepresenting his preferences.

Lemma 1. *Given any dominance history* $\boldsymbol{d}^{t-1}$*, ARSD is lsd-strategyproof, i.e., for each agent* $i \in N$*,*

$$\forall j \in N, \sum_{\substack{\mu \in \mathcal{M}: \\ \mu(i) \succ_i^t \mu(j)}} \pi(\mu | (\succ^t, \boldsymbol{d}^{t-1})) \geq \sum_{\substack{\mu \in \mathcal{M}: \\ \mu(i) \succ_i^t \mu(j)}} \pi(\mu | ((\hat{\succ}_i^t, \succ_{-i}^t), \boldsymbol{d}^{t-1}))$$

According to Lemma 1, an agent cannot improve the matching chosen at time t by misreporting at time t. Nor is an agent able to improve future matchings by increasing its dominance in the current period since this will only lead to the agent being given a lower priority in the matchings in the future. Therefore, the only possible strategy of interest to an agent is to try to decrease its dominance in the current time step, in the hope for improved priority (and thus, matchings) in the future. This potential strategizing is similar to dynamic house allocation for 2 periods of decisions where agents can opt out in the first period to get priority in the second period (Abdulkadiroğlu and Loerscher 2007).

We say that an agent's strategy is *dominance reducing* if it reports a preference that minimizes its current dominance.

Definition 7. *Agent i's reporting strategy $\sigma_i(\succ_i^t)$ at t is* ***dominance reducing*** *if given a matching policy π, $\forall j \in N$,*

$$\omega_{ij}(\pi(\sigma_i(\succ_i^t), \succ_{-i}^t)) \leq \omega_{ij}(\pi(\succ_i^t, \succ_{-i}^t)) \quad (6)$$

The next lemma states that given a random matching prescribed by ARSD, agent i's strategy to minimize her expected dictatorial dominance always results in reducing expected dominance caused by other agents on agent i.

Lemma 2. *Given ARSD, and strategy profile $\sigma_i(\succ_i^t)$, if $\omega_{ij}(\pi(\sigma_i(\succ_i^t), \succ_{-i}^t)) < \omega_{ij}(\pi(\succ_i^t, \succ_{-i}^t))$ for some $j \in N$, then*

$$\omega_{ji}(\pi(\sigma_i(\succ_i^t), \succ_{-i}^t)) < \omega_{ji}(\pi(\succ_i^t, \succ_{-i}^t)) \quad (7)$$

Proof. For the ease of notation, and since the history up to now would be fixed for any strategic report, we write $\omega_{ij}(\pi(h^t)) = \omega_{ij}(\pi(\succ))$ and drop the time indexes. By feasibility of the probability distribution induced by RSD, we can write $\forall \succ \in \mathcal{P}^n,\ \omega_{ij}(\pi(\succ)) + \bar{\omega}_{ij}(\pi(\succ)) = 1$. According to the dominance reducing strategy, for misreport $\hat{\succ}_i = \sigma_i(\succ_i)$ we have $\omega_{ij}(\pi((\hat{\succ}_i, \succ_{-i}))) < \omega_{ij}(\pi((\succ_i, \succ_{-i})))$. By applying the feasibility condition on both sides we have

$$1 - \bar{\omega}_{ij}(\pi((\hat{\succ}_i, \succ_{-i}))) < 1 - \bar{\omega}_{ij}(\pi((\succ_i, \succ_{-i}))) \quad (8)$$

$$\bar{\omega}_{ij}(\pi((\hat{\succ}_i, \succ_{-i}))) > \bar{\omega}_{ij}(\pi((\succ_i, \succ_{-i}))) \quad (9)$$

Thus, by reporting a dominance reducing preference, agent i's outcome gets improved, indicating that the probability of obtaining an object that is more preferred compared to $\mu(j)$ according to the report $\hat{\succ}_i$ will be higher than the same probability according to the truthful report $\succ_i$, i.e., $\bar{\omega}_{ji}(\pi((\hat{\succ}_i, \succ_{-i}))) > \bar{\omega}_{ji}(\pi((\succ_i, \succ_{-i})))$.

For contradiction assume that agent i increased the domination of agent j on himself by misreporting $\omega_{ji}(\pi(\hat{\succ}_i, \succ_{-i})) > \omega_{ji}(\pi(\succ_i, \succ_{-i}))$. By feasibility of the matching mechanism we have $\bar{\omega}_{ji}(\pi((\hat{\succ}_i, \succ_{-i}))) < \bar{\omega}_{ji}(\pi((\succ_i, \succ_{-i})))$, which contradicts the above finding. □

We show that ARSD is not manipulable in the global sense. Recall that our desired solution concept states that given any transition model, no agent can obtain a more preferred sequence of outcomes, no matter how other agents play now or in the future. Algorithm 1 ensures that after realization of the random decision, a dominating agent loses priority against a dominated agent in future steps. The next lemma follows directly from the ARSD algorithm.

Lemma 3. *Given two states $h^{t+1} = (\succ^{t+1}, \boldsymbol{d}^t)$ and $\hat{h}^{t+1} = (\succ^{t+1}, \hat{\boldsymbol{d}}^t)$ such that for agents $i, j \in N$, $\hat{d}_{ij}^t > d_{ij}^t$, we have*

$$\sum_{\substack{\mu \in \mathcal{M}: \\ \mu(i) \succ_i^{t+1} \mu(j)}} \pi(\mu \mid h^{t+1}) \geq \sum_{\substack{\mu \in \mathcal{M}: \\ \mu(i) \succ_i^{t+1} \mu(j)}} \pi(\mu \mid \hat{h}^{t+1}) \quad (10)$$

The implication of this lemma is that if an agent is dominated by another agent in some previous decisions, for all next steps where their dominance history is unbalanced, the agent always prefers his own outcome to the outcome of the dominating agent.

We now reach our main result in this section, which states that given any transition dynamics ARSD is not manipulable. Moreover, at each step ARSD prescribes a random matching that is ex post efficient and *lsd*-strategyproof.

Theorem 2. *ARSD is gsd-strategyproof while satisfying ex post efficiency and lsd-strategyproofness in each round.*

Proof. In each time, ARSD prescribes a random matching over deterministic Pareto efficient matchings induced by priority orderings, thus it satisfies ex post efficiency.

For strategyproofness, we must prove that agent i's misreport does not improve her overall allocation outcome, including immediate and future possible outcomes. That is, given any transition dynamic the following must hold, $\forall o^\ell \in M,\ W_i^\pi(h^t, o^\ell) \geq W_i^\pi(\hat{h}^t, o^\ell)$. Lemma 1 implies that no agent can immediately benefit from misreporting. For *gsd*-strategyproofness, we need to show that under all possible transition dynamics, there does not exist a period $\tau \geq t$ such that agent i will experience a better outcome by misreporting at time t. Let $\hat{\succ}_i^t \in \sigma_i(\succ_i^t)$ denote agent i's strategic report at time t. Let $\boldsymbol{d}^t$ and $\hat{\boldsymbol{d}}^t$ denote the dominance history reporting according to $\succ_i^t$ and $\hat{\succ}_i^t$ respectively. We analyze the following cases based on the possible strategic misreports:

Case 1: $\forall j \in N, \omega_{ij}(\pi(\hat{\succ}_i^t, \succ_{-i}^t)) \geq \omega_{ij}(\pi(\succ^t))$, meaning that agent i's misreport increases his dominance on all other agents. Since the dominance history up until t is the same, by the exclusive disjunction for the dominance profile we can immediately write $\hat{\boldsymbol{d}}_{ij}^t > \boldsymbol{d}_{ij}^t$. By Lemma 3 we have

$$\sum_{\substack{\mu \in \mathcal{M}: \\ \mu(i) \succ_i^{t+1} \mu(j)}} \pi(\mu \mid (\succ^{t+1}, \hat{\boldsymbol{d}}^t)) \leq \sum_{\substack{\mu \in \mathcal{M}: \\ \mu(i) \succ_i^{t+1} \mu(j)}} \pi(\mu \mid (\succ^{t+1}, \boldsymbol{d}^t))$$

Similarly for all future periods $\tau > t$ where the dominance history is unbalanced the above inequality holds, implying that agent i cannot achieve a better outcome by increasing her dominance at time t.

Case 2: $\forall j \in N, \omega_{ij}(\pi(\hat{\succ}_i^t, \succ_{-i}^t)) \leq \omega_{ij}(\pi(\succ^t))$, that is, agent i's strategic misreport reduces his dominance on all other agents. By Lemma 2, reducing dominance on an agent reduces the expected domination on agent i, thus, $\forall j \in N, \omega_{ji}(\pi(\hat{\succ}_i^t, \succ_{-i}^t)) \leq \omega_{ji}(\pi(\succ^t))$, and $\hat{\boldsymbol{d}}_{ji}^t < \boldsymbol{d}_{ji}^t$. By Lemma 1 we can assume that the agent is truthful in all future periods. By Lemma 3 we write $\forall j \in N, \bar{\omega}_{ij}(\pi((\succ^{t+1}, \hat{\boldsymbol{d}}^t))) \geq \bar{\omega}_{ij}(\pi((\succ^{t+1}, \boldsymbol{d}^t)))$, which holds $\forall \tau > t$.

For contradiction, assume that agent i has improved its overall assignment by misreporting. Therefore, $\sum_{\tau=t} \bar{\omega}_{ji}(\pi((\succ^\tau, \hat{\boldsymbol{d}}^{\tau-1}))) > \sum_{\tau=t} \bar{\omega}_{ji}(\pi((\succ^\tau, \boldsymbol{d}^{\tau-1})))$. Adding these inequalities yields that for all times,

$$\sum_{\tau=t}^{T} \sum_{i \in N} \bar{\omega}_{ji}(\pi((\succ^\tau, \hat{\boldsymbol{d}}^{\tau-1}))) > \sum_{\tau=t}^{T} \sum_{i \in N} \bar{\omega}_{ji}(\pi((\succ^\tau, \boldsymbol{d}^{\tau-1})))$$

which means that in a sequence of random matchings such that at some $\tau > t$ every agent's assignment (including agent i) strictly improves when agent i misreports, contradicting the feasibility of the ARSD mechanism.

Case 3: In this case, we assume that agent i's misreport has weakly increased his dominance on some agents while strictly reducing his dominance on a subset of agents. More formally, $\forall j \in N_d, \omega_{ij}(\pi(\hat{\succ}_i^t, \succ_{-i}^t)) < \omega_{ij}(\pi(\succ^t))$, and $\forall k \in N \backslash N_d, \omega_{ik}(\pi(\hat{\succ}_i^t, \succ_{-i}^t)) \geq \omega_{ik}(\pi(\succ^t))$, where $N_d \subseteq N$ denotes a subset of agents such that for some $j \in N_d$ agent i's dictatorial dominance changes to $\omega_{ij}(\pi(\hat{\succ}_i^t, \succ_{-i}^t))$.

For all $k \in N \backslash N_d$ since $\hat{\boldsymbol{d}}_{ik}^t \geq \boldsymbol{d}_{ik}^t$ by Lemma 3 we can write, $\bar{\omega}_{ki}(\pi(\succ^\tau, \hat{\boldsymbol{d}}^{\tau-1})) \leq \bar{\omega}_{ki}(\pi(\succ^\tau, \boldsymbol{d}^{\tau-1})), \forall \tau > t$. For the sake of our proof, assume the outcome for agent i after a dominance increasing report is as good as being truthful $\forall k \in N \backslash N_d$, i.e., $\bar{\omega}_{ki}(\pi(\succ^\tau, \hat{\boldsymbol{d}}^{\tau-1})) = \bar{\omega}_{ki}(\pi(\succ^\tau, \boldsymbol{d}^{\tau-1}))$.

For all $j \in N_d$, by Lemma 2, we have $\omega_{ji}(\pi(\hat{\succ}_i^t, \succ_{-i}^t)) < \omega_{ji}(\pi(\succ^t))$ yielding that $\hat{\boldsymbol{d}}_{ji}^t > \boldsymbol{d}_{ji}^t$. Thus for all $j \in N_d$ and all times $\tau > t$,

$$\bar{\omega}_{ij}(\pi(\succ^\tau, \hat{\boldsymbol{d}}^{\tau-1})) \leq \bar{\omega}_{ij}(\pi(\succ^\tau, \boldsymbol{d}^{\tau-1})) \quad (11)$$

$$\exists\, t' > t, \bar{\omega}_{ij}(\pi(\succ^{t'}, \hat{\boldsymbol{d}}^{t'-1})) < \bar{\omega}_{ij}(\pi(\succ^{t'}, \boldsymbol{d}^{t'-1})) \quad (12)$$

Adding the inequalities for all agents in N, we have $\forall j \in N \backslash i$, $\sum_{\tau=t}^T \bar{\omega}_{ij}(\pi(\succ^\tau, \hat{\boldsymbol{d}}^{\tau-1})) < \sum_{\tau=t}^T \bar{\omega}_{ij}(\pi(\succ^\tau, \boldsymbol{d}^{\tau-1}))$. Assume for contradiction that agent i has improved its overall outcome. This means that there exists at least one time step $\tau > t$ such that agent i is strictly better off: $\forall j \in N, \bar{\omega}_{ji}(\pi(\succ^\tau, \hat{\boldsymbol{d}}^{\tau-1})) > \bar{\omega}_{ji}(\pi(\succ^\tau, \boldsymbol{d}^{\tau-1}))$. This implies that misreporting has improved the expected assignments for agent i as well as for all other agents, which contradicts the feasibility of the random matching. Thus, agent i does not gain any expected benefit in the future by changing the trajectory of decisions, implying $W_i^\pi((\succ^t, \boldsymbol{d}^{t-1}), o^\ell) \geq W_i^\pi((\hat{\succ}_i^t, \succ_{-i}^t, \boldsymbol{d}^{t-1}), o^\ell), \forall o^\ell \in M$. □

Fairness in Sequential Matchings

We showed that a modification to the RSD mechanism in dynamic settings encourages truthful reporting by incorporating the history of past assignments. In this section, we argue that as a result of balancing priorities, ARSD also satisfies some desirable notions of local and global fairness.

We consider two notions of fairness in sequential matching problems; a local fairness notion of *equity* (equal treatment of equals), and a global notion of *ex post envyfreeness*.

Definition 8. *A sequential matching mechanism is **equitable** at each time step t if and only if $\forall i, j \in N$ with identical dominance histories $d_{ij}^{t-1} = d_{ji}^{t-1}$, if $\succ_i^t = \succ_j^t$ then*

$$\forall y \in M, \sum_{x \in M: x \succ_i^t y} p_i^t(x|(\succ^t, \boldsymbol{d}^{t-1})) = \sum_{x \in M: x \succ_j^t y} p_j^t(x|(\succ^t, \boldsymbol{d}^{t-1}))$$

In sequential mechanisms, envyfreeness relies on balancing priorities over the course of the assignment sequence. We consider ex post envyfreeness at each time based on the sequence of decisions up to and including the current period.

Definition 9. *Given a sequence of matchings $(\mu^1, \ldots \mu^t)$, a matching mechanism is **periodic ex post envyfree** (PEF) if at all times t, $\forall i, j \in N$, $\sum_{s=1}^t \omega_{ji}(\mu^s) = \sum_{s=1}^t \omega_{ij}(\mu^s)$.*

Proposition 3. *ARSD does not satisfy PEF, but does yield a sequence of equitable local matchings.*

PEF is a strict requirement that is not satisfied even with fixed preferences. Nevertheless, for two agents with fixed preferences there is a positive result.

Proposition 4. *When $n = 2$ with fixed preferences, ARSD satisfies PEF at least in all even periods.*

To see why this does not hold for $n > 2$, consider the following fixed preferences $\succ_1 = a \succ b \succ c$, $\succ_2 = b \succ a \succ c$, $\succ_3 = a \succ c \succ b$.

Degree of Envy The nonexistence of ex post envyfreeness for $n > 2$, raises a crucial question of whether our mechanism ensures some degree of envy. Among several plausible ways of defining envy (Chevaleyre and Endriss 2007), we consider a natural notion of envy; the envy of a single agent towards another agent. Given a history of assignments h^t, agent i's *degree of envy* with respect to agent j is

$$e_{ij}(h^t) = \sum_{s=1}^t \left[\omega_{ji}(\mu^s) - \omega_{ij}(\mu^s)\right] \quad (13)$$

Definition 10. *A sequential mechanism is c-envious if for all times t, $\forall i \in N$ agent i is envious to agent j for at most c assignments. That is, $c = \max_{ij} e_{ij}(h^t), \forall t$.*

Theorem 3. *An ARSD matching mechanism is* 1*-envious.*

In fact, ARSD interplays between random assignments in repeated decisions to maintain an approximately fair global policy. It is easy to see that the maximum envy of a society of agents with ARSD mechanism is $\frac{n(n-1)}{2}$.

Discussion

We studied the incentive and fairness properties of sequential matching problems with dynamic ordinal preferences. We showed that in contrast to one-shot settings, a sequence of RSD-induced matchings is prone to manipulation. Subsequently, we proposed a history-dependent matching policy (namely ARSD) that guarantees global strategyproofness while sustaining the local properties of the sequential RSD.

We restricted our analysis to problems where $n \geq m$. Although our main results immediately apply to problems with $n < m$ when the agent with the highest priority chooses $m - n + 1$ alternatives, in most situations RSD mechanism requires a careful method for picking the sequence at each serial dictatorship (Kalinowski et al. 2013; Bouveret and Lang 2014).

An important open question is how much efficiency is being sacrificed in order to guarantee strategyproofness. The incompatibility of ordinal efficiency and *sd*-strategyproofness in static settings (Bogomolnaia and Moulin 2001) prevents us from designing truthful optimal policies in dynamic settings. However, there may exist some approximately efficient random policies in the policy space that incentivizes truthfulness in sequential settings, perhaps, by renouncing the ex post efficiency in each round.

Finally, in static large markets where there are large number of copies of each object (such as assigning students to housing), the inefficiency of the RSD mechanism vanishes (Che and Kojima 2010). One potential direction would be to study the efficiency and envyfreeness of sequential matching problems in markets with multiple capacities and various agent/object ratios.

References

Abdulkadiroğlu, A., and Loerscher, S. 2007. Dynamic house allocations. *Working papar*.

Abdulkadiroğlu, A., and Sönmez, T. 1998. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica* 66(3):689–701.

Alsheddy, A., and Tsang, E. P. 2011. Empowerment scheduling for a field workforce. *Journal of Scheduling* 14(6):639–654.

Ashlagi, I.; Fischer, F.; Kash, I. A.; and Procaccia, A. D. 2013. Mix and match: A strategyproof mechanism for multi-hospital kidney exchange. *Games and Economic Behavior*.

Athey, S., and Segal, I. 2013. An efficient dynamic mechanism. *Econometrica* 81(6):2463–2485.

Aziz, H.; Brandt, F.; and Brill, M. 2013. On the trade-off between economic efficiency and strategy proofness in randomized social choice. In *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*, 455–462. International Foundation for Autonomous Agents and Multiagent Systems.

Bade, S. 2014. Serial dictatorship: the unique optimal allocation rule when information is endogenous. *Theoretical Economics* To appear.

Bard, J. F., and Purnomo, H. W. 2005. Preference scheduling for nurses using column generation. *European Journal of Operational Research* 164(2):510–534.

Bergemann, D., and Välimäki, J. 2010. The dynamic pivot mechanism. *Econometrica* 78(2):771–789.

Bloch, F., and Cantala, D. 2013. Markovian assignment rules. *Social Choice and Welfare* 40(1):1–25.

Bogomolnaia, A., and Moulin, H. 2001. A new solution to the random assignment problem. *Journal of Economic Theory* 100(2):295–328.

Bouveret, S., and Lang, J. 2014. Manipulating picking sequences. In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI'14)*. Prague, Czech Republic: IOS Press.

Burke, E. K.; De Causmaecker, P.; Berghe, G. V.; and Van Landeghem, H. 2004. The state of the art of nurse rostering. *Journal of scheduling* 7(6):441–499.

Cavallo, R. 2009. Mechanism design for dynamic settings. *ACM SIGecom Exchanges* 8(2):7.

Che, Y.-K., and Kojima, F. 2010. Asymptotic equivalence of probabilistic serial and random priority mechanisms. *Econometrica* 78(5):1625–1672.

Chevaleyre, Y., and Endriss, U. 2007. Reaching envy-free states in distributed negotiation settings. In *In Proceedings of IJCAI-2007*, 1239–1244.

De Grano, M. L.; Medeiros, D.; and Eitel, D. 2009. Accommodating individual preferences in nurse scheduling via auctions and optimization. *Health Care Management Science* 12(3):228–242.

Dickerson, J. P.; Procaccia, A. D.; and Sandholm, T. 2012. Dynamic matching via weighted myopia with application to kidney exchange. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1340–1346.

Gibbard, A. 1977. Manipulation of schemes that mix voting with chance. *Econometrica: Journal of the Econometric Society* 665–681.

Kalinowski, T.; Narodytska, N.; Walsh, T.; and Xia, L. 2013. Strategic behavior when allocating indivisible goods sequentially. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 452–458.

Krysta, P.; Manlove, D.; Rastegari, B.; and Zhang, J. 2014. Size versus truthfulness in the house allocation problem. In *Proceedings of the fifteenth ACM conference on Economics and computation*, 453–470. ACM.

Kurino, M. 2014. House allocation with overlapping generations. *American Economic Journal: Microeconomics* 6(1):258–289.

Parkes, D., and Procaccia, A. 2013. Dynamic social choice with evolving preferences. In *Proceedings of the AAAI Conference on Artificial Intelligence*, AAAI'13, 767–773.

Parkes, D. 2007. Online mechanisms. *Algorithmic Game Theory, ed. N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani* 411–439.

Roth, A. E.; Sönmez, T.; and Ünver, M. U. 2004. Kidney exchange. *The Quarterly Journal of Economics* 119(2):457–488.

Siferd, S. P., and Benton, W. 1992. Workforce staffing and scheduling: Hospital nursing specific models. *European Journal of Operational Research* 60(3):233–246.

Sönmez, T., and Ünver, M. U. 2010a. Course bidding at business schools. *International Economic Review* 51(1):99–123.

Sönmez, T., and Ünver, M. U. 2010b. House allocation with existing tenants: A characterization. *Games and Economic Behavior* 69(2):425–445.

Svensson, L.-G. 1999. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare* 16(4):557–567.

Townsend, R. M. 1982. Optimal multiperiod contracts and the gain from enduring relationships under private information. *The Journal of Political Economy* 1166–1186.

Ünver, M. U. 2010. Dynamic kidney exchange. *The Review of Economic Studies* 77(1):372–414.

Vohra, R. V. 2012. Dynamic mechanism design. *Surveys in Operations Research and Management Science* 17(1):60–68.

Zhou, L. 1990. On a conjecture by gale about one-sided matching problems. *Journal of Economic Theory* 52(1):123–135.