Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI-10)

Tolerable Manipulability in Dynamic Assignment without Money

James Zou
School of Engineering
and Applied Sciences
Harvard University
jzou@fas.harvard.edu

Sujit Gujar
Dept of Computer Science
and Automation
Indian Institute of Science
sujit@csa.iisc.ernet.in

David Parkes
School of Engineering
and Applied Sciences
Harvard University
parkes@eecs.harvard.edu

Abstract

We study a problem of dynamic allocation without money. Agents have arrivals and departures and strict preferences over items. Strategyproofness requires the use of an arrival-priority serial-dictatorship (APSD) mechanism, which is ex post Pareto efficient but has poor *ex ante* efficiency as measured through average rank efficiency. We introduce the *scoring-rule* (SR) mechanism, which biases in favor of allocating items that an agent values above the population consensus. The SR mechanism is not strategyproof but has tolerable manipulability in the sense that: (i) if every agent optimally manipulates, it reduces to APSD, and (ii) it significantly outperforms APSD for rank efficiency when only a fraction of agents are strategic. The performance of SR is also robust to mistakes by agents that manipulate on the basis of inaccurate information about the popularity of items.

Introduction

We are interested in assignment problems to a dynamic agent population where it is undesirable to use money; e.g., because of community norms, legal constraints, or inconvenience. Consider the allocation of tasks to agents that arrive and depart and have preferences on tasks and a time window within which they can be allocated a task. The problem can also be one of resource allocation, where agents arrive and demand access to a resource before departure. Motivating domains include those of car pooling in which the agents are commuters and the resources are seats in shared cars, or science collaborators in which the agents are people looking to perform useful work for the community.

For static assignment problems without money, the class of strategyproof mechanisms is more or less restricted to serial dictatorships in which each agent in turn selects its most preferred item from the unclaimed items. For the dynamic problem considered here, the additional requirement of strategyproofness with respect to arrival (i.e., preventing manipulation by reporting later arrival) necessitates the use of an *arrival-priority serial dictatorship* (APSD). Although APSD is *ex post* Pareto efficient it is unsatisfactory in another, more refined sense. If we consider *ex ante* efficiency, that is average utility received by agents under the mechanism, then its performance is quite poor. For risk-neutral agents for which the difference in utility for any pair of consecutive items in a preference ordering is constant, then the expected average rank (or *rank-efficiency*) measures the *ex ante* efficiency. APSD has poor rank efficiency because an early arrival may pick its most preferred item over its second most preferred item even if this leaves a later arrival with its least preferred item rather than its most preferred item.

This motivates our study of *tolerable manipulability*, an agenda for computational mechanism design first suggested by Feigenbaum and Shenker (2002). We introduce the *scoring-rule* (SR) mechanism, which biases in favor of allocating items that an agent values atypically highly. The SR mechanism is not strategyproof but has tolerable manipulability in the following sense: (i) if every agent optimally manipulates SR then it reduces to APSD, and thus the performance of the only strategyproof mechanism; and (ii) the SR mechanism significantly outperforms APSD for rank efficiency when only a fraction of agents are strategic. The performance of SR is also robust in the following sense: the optimal manipulation is a dominant strategy, and thus invariant to strategies of other agents, and SR continues to outperform APSD even when agents have inaccurate information about the distribution on preferences in the population and thus the rules of the SR mechanism.

Our simulation results on the SR mechanism demonstrate that for 10 agents, SR has 10% greater rank-efficiency than APSD when all agents are truthful and non-strategic. When 5 of the 10 agents are strategic, SR still maintains 5% greater rank-efficiency than APSD. Furthermore, the advantage of SR over APSD increases as the number of agents increase. With 25 agents, SR has 19% greater rank-efficiency over APSD. To further benchmark the performance of the SR mechanism we also compare against the rank efficiency of a sample-based stochastic optimization algorithm (Hentenryck and Bent 2006), namely, Consensus. When all the agents are truthful, SR outperforms Consensus by 4% when there are 10 agents. Even with up to four out of the ten agents acting strategically SR outperforms Consensus.

Related Work. For the house allocation problem, which is a static assignment problem in which agents and items are fixed and each agent has strict preferences over the assignment of one item, Svensson (1999) establishes that the only non-bossy, neutral and strategyproof mechanisms are serial dictatorships. Papai (2001) relaxes the requirement of neutrality and achieves a richer characterization.

Abdulkadiroglu and Loertscher (2005) study a dynamic house allocation problem that is quite different from our problem. It is a two period problem in which the agents and items are fixed and dynamics occur because agent preferences in period two are unknown in period one. Kurino (2009) considers a dynamic house allocation problem with fixed items and a dynamic (overlapping generations) agent population. His problem is again different from ours because each agent demands an item in every period. Another difference is that items are not consumed but rather returned to the market when an agent departs.

We are not aware of any prior work on tolerable manipulability in the context of dynamic mechanism design. Othman and Sandholm (2009) define the notion of *manipulation-optimal* mechanisms, in which: (i) the mechanism is undominated by any strategyproof mechanism when all agents are rational; and (ii) the performance is better than any strategyproof mechanism if any agent fails to be rational in any way. Their results are essentially negative, in that a manipulation-optimal mechanism is impossible whenever there is an agent and a joint type profile of other agents, such that there are two types of the agent for which its best-response is to be non-truthful and moreover its optimal misreport depends on its true type. This is a very weak condition and holds for our problem. Othman and Sandholm (2009) demonstrate positive results only in the case of an agent with two possible types.[1] Our definition of tolerable manipulability is weaker than that of manipulation-optimal mechanisms, in that we retain (i) but replace (ii) with (ii') the performance is better than any strategyproof mechanism if *a sufficient fraction of agents fail to be rational by being truthful.* With this approach we are able to achieve positive results. In a dynamic auction setting with money, Lavi and Nisan (2005), in studying the performance of mechanisms for a class of (set-Nash) equilibria, consider another form of tolerable manipulability in the sense that good properties are achieved as long as agents play strategies from a set that is closed under rational behavior.

The Model

There are $A = \{A_1, \ldots, A_n\}$ agents, $I = \{I_1, \ldots, I_m\}$ items, and each agent $A_k \in A$ has an arrival $\alpha_k \in T$, departure $\beta_k \in T$, demands one item and has preferences $\phi_k \in \Phi$ on items, where $T = \{1, 2, \ldots\}$ is the set of discrete time periods and Φ the set of preferences. Altogether this defines an agent's type $\theta_k = (\alpha_k, \beta_k, \phi_k) \in \Theta$ where Θ is the set of types. Preferences ϕ_k are strict and define a rank $r(k, j) \in \{1, \ldots, m\}$, where $r(k, 1)$ is the index of the most preferred item and so on. We write $\phi_k : I_{r(k,1)} \succ_k I_{r(k,2)} \succ_k \ldots \succ_k I_{r(k,m)}$ to denote an agent's preferences. Each agent A_k only cares about its allocated item in interval $\{\alpha_k, \ldots, \beta_k\}$, and is indifferent to the allocation to other agents. We consider a fixed set of items, all available from period one.

[1] See also Conitzer and Sandholm (2004) who first introduced an example to demonstrate the existence of a manipulation-optimal mechanism.

In a direct-revelation mechanism, the message space allows an agent A_k to report $\theta'_k = (\alpha'_k, \beta'_k, \phi'_k) \neq \theta_k$ in some period $t = \alpha'_k$. We assume $\alpha'_k \geq \alpha_k$ and $\beta'_k \leq \beta_k$, which together with $\beta'_k \geq \alpha_k$ implies that $t \in \{\alpha_k, \ldots, \beta_k\}$.[2] A mechanism is defined by a function $f : \Theta^n \rightarrow X$, where $x = (x_1, \ldots, x_n) \in X$ denotes an allocation of items $x_k \in I$ to each agent A_k and X is the set of *feasible* allocations, such that $x_k \neq x_\ell$ if $k \neq \ell$. We only consider deterministic mechanisms. Let $f_k(\theta) \in I$ denote the allocation to agent A_k in period $t \in \{\alpha_k, \ldots, \beta_k\}$, where $\theta = (\theta_1, \ldots, \theta_n)$ is the joint type profile. Let $\theta_{-k} = (\theta_1, \ldots, \theta_{k-1}, \theta_{k+1}, \ldots, \theta_n)$ and $f(\theta'_k, \theta_{-k})$ is shorthand for $f(\theta_1, \ldots, \theta_{k-1}, \theta'_k, \theta_{k+1}, \ldots, \theta_n)$.

To meet the requirements of an online mechanism, which cannot know future type reports, we require $f_k(\theta_k, \theta_{-k}) = f_k(\theta_k, \theta'_{-k})$ whenever $\theta_{\leq \beta_k} = \theta'_{\leq \beta_k}$, for all $k \in \{1, \ldots, n\}$, where $\theta_{\leq t}$ is the restriction of type profile θ to include only those agents that arrive no later than period t.

Some desiderata of online assignment mechanisms:

- *Strategyproof*: Mechanism f is strategyproof if $f_k(\theta_k, \theta_{-k}) \succ_k f_k(\theta'_k, \theta_{-k})$ for all $\theta'_k = (\alpha'_k, \beta'_k, \phi'_k)$ where $\alpha'_k \geq \alpha_k$ and $\beta'_k \leq \beta_k$, all k, and all θ. Truthful reporting is a dominant-strategy equilibrium.
- *Non-Bossy*. Mechanism f is non-bossy if $(f_k(\theta_k, \theta_{-k}) = f_k(\theta'_k, \theta_{-k})) \Rightarrow (f(\theta_k, \theta_{-k}) = f(\theta'_k, \theta_{-k}))$, for all $\theta'_k = (\alpha'_k, \beta'_k, \phi'_k)$ where $\alpha'_k \geq \alpha_k$ and $\beta'_k \leq \beta_k$, all k, and all θ.
- *Neutrality*. Mechanism f is neutral if $f_k(\theta) = \pi^{-1}(f_k(\pi_\theta(\theta)))$ for all agents A_k, type profiles θ, and item permutations π, where $\pi : \{1, \ldots, m\} \rightarrow \{1, \ldots, m\}$ is bijective, π^{-1} is the inverse, and $\pi_\theta(\theta) = \theta'$ is a type profile induced in the straightforward way by the permutation on items, so that rank $r'(k, j) = r(k, \pi(j))$ for all A_k and all items I_j, with arrival and departure times unchanged.
- *Pareto Efficient (PE)*. Mechanism f is *ex post* Pareto efficient if for all $\theta \in \Theta^n$, then there is no feasible allocation x' that is weakly preferred by all agents to $x = f(\theta)$ and strictly preferred by at least one agent.

Strategyproof Mechanisms

A serial dictatorship has a priority ranking $h : \{1, \ldots, n\} \rightarrow \{1, 2, \ldots, n\}$, such that agents are ordered $h(1), h(2), \ldots$, and assigned the most preferred item still available given the allocation to higher priority agents.

Definition 1. An online serial dictatorship is a serial dictatorship in which $(h(i) < h(j)) \Rightarrow \alpha_i \leq \beta_j$.

This follows from the online setting and what we term *schedulability*; i.e., we can only guarantee agent A_i a higher

[2] The arrival assumption ($\alpha'_k \geq \alpha_k$) is standard in problems of dynamic mechanism design (Parkes 2007), and can be motivated easily if the arrival is the period in which an agent realizes its demand or discovers the mechanism. The departure assumption ($\beta'_k \leq \beta_k$) is made for convenience; all our mechanisms allocate upon reported departure in any case and so an agent never has a useful strategy that involves reporting $\beta'_k > \beta_k$.

priority in the sense of serial dictatorship than agent A_j if it arrives before agent j departs.

Lemma 1. *Let f be a deterministic, strategyproof, neutral, non-bossy, online assignment mechanism. Then f must be an online serial dictatorship.*

Proof. Let f be any strategyproof, neutral, non-bossy online mechanism and given any reported schedule of agent arrival and departures, $\sigma = \{\alpha'_1, \beta'_1, ..., \alpha'_N, \beta'_N\}$. Now $f|_\sigma(\phi_1, ..., \phi_N) = \hat{f}(\theta_1, ..., \theta_N)$ defines an off-line mechanism $\hat{f}$, mapping a preference profile to an allocation of items. The restricted mechanism $f|_\sigma$ is Non-Bossy since $f_k|_\sigma(\phi_k, \phi_{-k}) = f_k|_\sigma(\phi'_k, \phi_{-k}) \Rightarrow f_k(\theta_k, \theta_{-k}) = f_k(\theta'_k, \theta_{-k})$. It is also neutral since any items permutation of $f|_\sigma$ is also an items permutation of f. Lastly, $f|_\sigma$ is strategyproof since any profitable misreport of preferences would carry over into a manipulation of f. Therefore by Svensson (1999), $f|_\sigma$ must be a offline serial dictatorship. For, f to be online feasible, any agent A_k should receive an object before departure. So, all the agents having higher priority than A_k must be allocated before A_k's departure. and the priority structure in f must be such that, if $h(i) < h(j)$, then $\alpha_i \leq \beta_j$. That is, f is online serial dictatorship. □

Definition 2. The Arrival-Priority Serial Dictatorship (APSD) or Greedy mechanism assigns priority by arrival, with an earlier arrival having a higher priority and with ties broken arbitrarily. An item is assigned upon an agent's arrival, and released (and thus allocated) upon departure.

Theorem 1. *A deterministic online mechanism is strategyproof, neutral and non-bossy if and only if it is the APSD mechanism.*

Proof. ($\Leftarrow$) It is immediate to check that APSD is neutral and non-bossy. To establish that it is strategyproof, note that since it is a serial dictatorship an agent A_k should report its true preference ϕ_k whatever its (α'_k, β'_k) report. Moreover, for any reported arrival α'_k, the outcome is invariant to its reported departure $\beta'_k \geq \alpha'_k$ (still with $\beta'_k \leq \beta_k$). Reporting a later arrival $\alpha'_k > \alpha_k$ only reduces its priority rank and reduces its choice set of items.

($\Rightarrow$) Consider a strategyproof mechanism that is neutral and non-bossy, and thus an online serial dictatorship by Lemma 1, but is not an arrival-priority serial dictatorship. In particular, there is some preference profile $\phi \in \Phi^n$ and some agent arrival/departure schedule such that $\alpha_i < \alpha_j$ but $h(i) > h(j)$ for some pair of agents A_i, A_j. To be schedulable, we must also have $\alpha_j \leq \beta_i$. But now, agent A_i can report $\beta'_i = \alpha_i < \alpha_j$, and force $h(i) < h(j)$, again by schedulability. Now suppose in particular that $\phi_i = \phi_j$, so that the item A_j receives when A_i is truthful is strictly preferred by A_i to its own allocation. Agent A_i will receive an item that is at least as good as that allocated to agent A_j when agent i is truthful, and thus this is a useful deviation. □

An online serial dictatorship is *ex post* Pareto efficient (PE) because agent preferences are strict. For agents except the top priority agent, a change in allocation to a more preferred item would require that an agent with higher priority be allocated a less preferred item than its current allocation.

Ex ante efficiency. A mechanism is *ex ante* PE if there is no function from types to allocations that induces greater expected utility than the distribution induced by a mechanism f and the type distribution. This is a stronger requirement than *ex post* PE. Following Budish and Cantillon (2009), one way to evaluate *ex ante* efficiency is further assume risk neutral agents, each with a constant difference in utility across items that are adjacent in their preference list. Let $R(x) = \frac{1}{n}\sum_k \mathbf{r}(k, x_k)$ define the average-rank score for allocation $x \in X$, where $\mathbf{r}(k, x_k)$ is agent $A'_k s$ rank for his allocated item x_k. Under this assumption, then the expected average-rank score measures *ex ante* efficiency.

Heuristic Allocation Methods

In this section, we compare the performance of APSD with two other algorithms, the scoring-rule (SR) and Consensus algorithms, considering for now only truthful inputs and without concern to manipulations. By assuming that ties are broken arbitrarily if multiple agents arrive together, we keep the presentation simple and consider the case of only one agent arriving in each period.

The Scoring-rule Algorithm. The general idea of the scoring-rule (SR) algorithm is to bias in favor of allocating an item to an agent that he values above the population consensus. For this, we define the score of item I_j, given that there are m items, as $S(I_j) = \sum_{k=1}^{m} \Pr(j, k)k$, where $\Pr(j, k)$ is the probability that item I_j is ranked in kth position by a random agent. For a uniform preference model, $\Pr(j, k) = 1/m$ for all items I_j, but this will typically be skewed in favor of some items over.

The SR algorithm works as follows:

1. Suppose that agent A_k arrives, with $\phi_k : I_{r(k,1)} \succ_k I_{r(k,2)} \succ_k \ldots \succ_k I_{r(k,m)}$.
2. Assign item I_j to agent A_k, where $j \in \arg\min_{\ell \in avail}[r^{-1}(k, \ell) - S(I_\ell)]$, where $avail$ is the set of available items given prior allocations and $r^{-1}(k, \ell)$ is the inverse rank function, giving agent $A'_k s$ rank for item I_ℓ. Release item I_j to A_k upon departure.

If agent A_k has a less popular I_ℓ item as one of his top choices (low $r^{-1}(k, \ell) - S(I_\ell)$), then SR will tend to allocate this item and save popular items for other agents; this is how the algorithm is designed to improve *ex ante* PE.

The Consensus algorithm. The Consensus algorithm is a sampled-based method of stochastic optimization (Hentenryck and Bent 2006).

1. When agent A_k arrives, generate L samples of the types of possible future agent arrivals; i.e., L samples of the types of $n - k$ agents where agent A_k is the kth agent to arrive. Let $\mathcal{A}_\ell$ denote the set of agents in the ℓth such sample.
2. For each sample $\ell \in \{1, \ldots, L\}$, compute an allocation x^*_l of the available items to agents $\{A_k\} \cup \mathcal{A}_\ell$ to minimize the average rank of allocated items.
3. Determine the item I_j that is allocated to agent A_k most often in the L solutions $\{x^*_1, \ldots, x^*_L\}$, breaking ties at random, and assign this item to the agent. Release the item to the agent upon its departure.

Rank Efficiency Analysis

To evaluate the rank efficiency of APSD, SR and the Consensus algorithm we adopt a simple model of the distribution on agent preferences that is parameterized by $(p_1, \ldots, p_m)$, with $p_j > 0$ to denote the *popularity* of item I_j. Given this, we have the following generative model for *weighted-popularity* preferences: For each agent A_k:

- Initialize $R_1 = \{1, 2, ..., m\}$. In the first round, an item is selected at random from R_1 with probability $p_j / \sum_{i \in R_1} p_i$. Let I_{k_1} denote the item selected in this round.
- Let $R_2 = R_1 \setminus \{k_1\}$ describe the remaining items. In this round Item I_j is selected from R_2 with probability $p_j / \sum_{i \in R_2} p_i$.
- Continue, to construct a preference rank $\phi_k : I_{k_1} \succ_k I_{k_2} \succ_k \ldots \succ_k I_{k_m}$ for the agent.

In our experiments, we adjust the popularity profile on items by adjusting a *similarity* parameter, z. For each item I_j, we set $p_j = \Psi_z(\frac{2j}{m})$, where there are m items and Ψ_z is the density function for Normal distribution $\mathcal{N}(1, z)$. High similarity corresponds to an environment in which all items are of similar popularity. Low similarity corresponds to an environment in which a few items are significantly more popular than other items. Rank-efficiency is evaluated over 1000 independent simulations, with the average rank for a run normalized to the rank of the optimal off-line allocation that minimizes the average rank based on true preferences. Smaller rank efficiency is better, and 1 is optimal.

Figure 1 provides representative results for 10 agents and 10 items.[3] SR outperforms APSD for all similarities, with improvement at least 10% for similarity $z \leq 0.3$. In absolute terms, an improvement of 10% in (normalized) rank efficiency for 10 agents and 10 items corresponds to an average absolute rank improvement of 0.4; so, roughly equivalent to an agent expecting to receive an average improvement of rank position of one every other time. For high similarity the performance of all algorithms becomes quite similar. Figure 2 considers increasing the number of agents, holding the number of items equal to the number of agents, and for similarity 0.3. The SR algorithm again has the best average rank-efficiency for all numbers of agents. For $n = 15$ and $n = 30$, the performance of APSD (or Greedy) is 40% and 60% worse, respectively, than the off-line solution, while SR is 20% and 35% worse.

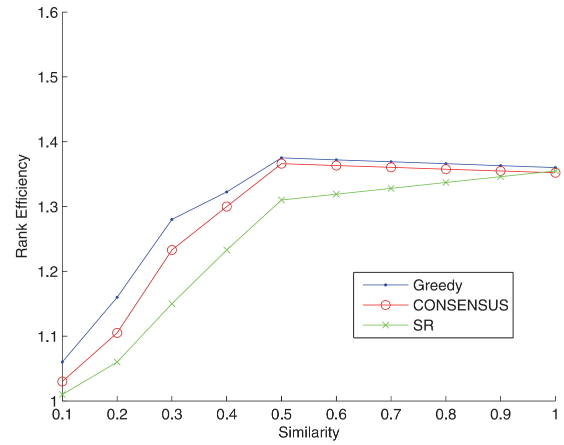

Figure 1: Rank efficiency under truthful agents as the similarity in item popularity is adjusted, for 10 items and 10 agents.

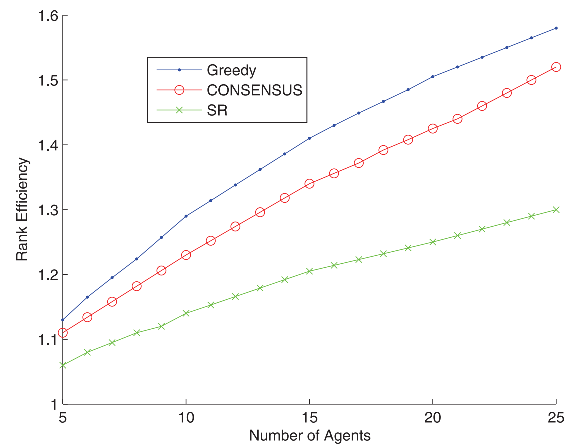

Figure 2: Rank efficiency under truthful agents as the population increases, for 10 items, 10 agents and similarity of 0.3.

What is The Scoring Rule Doing Right?

When all items have equal popularity, the SR algorithm agrees exactly with APSD. In this case, the score is equal for every item and $\arg\min_{\ell \in avail}[r^{-1}(k, \ell) - S(I_\ell)] = \arg\min_{\ell \in avail} r^{-1}(k, \ell)$, and selects the item for agent A_k with the smallest rank. Now consider a simple scenario where there are only two agents $\{A_1, A_2\}$ and two items $\{I_1, I_2\}$. A_1 arrives first and has preference ranks $\phi_1 : I_1 \succ_1 I_2$. Suppose that SR allocates $A_1 \rightarrow I_1$, so that $r(1,1) - S(I_2) < r(1,2) - S(I_1)$. Because the score of an item is precisely the expected rank for allocation that item, $S(I_1) = \mathbb{E}_{\phi_2} r(2,1)$ and $S(I_2) = \mathbb{E}_{\phi_2} r(2,2)$. Now, we see that SR allocates I_1 to A_1 exactly when $r(1,1) + S(I_1) = r(1,1) + \mathbb{E}_{\phi_2} r(2,1) < r(1,2) + \mathbb{E}_{\phi_2} r(2,2) = r(1,2) + S(I_2)$, and makes the allocation decision to minimize total expected rank. To test this intuition, we track the first occasion when SR and APSD make different allocations. let SR allocate I_1 while APSD allocates I_2. In being greedy, APSD achieves an average rank of 1.4 for I_2, but SR still achieves 1.62. On the other hand, while SR achieves an average rank of 3.54 on I_1, APSD struggles and achieves only 4.92 on this item. Looking at the score, we see that $S(I_1) = 6.43 > S(I_2) = 3.09$. SR made the right decision in allocating the less popular I_1 when a good opportunity arises while holding on to the more popular I_2.

Tolerable Manipulability

We focus now on the SR algorithm, which is effective in meeting the *ex ante* efficiency performance target but is manipulable by agents. The question that we ask in this section

[3] In all experiments, agents arrive in sequence, and this sequencing is sufficient to simulate; note that the performance of all algorithms are invariant to departure.

is whether the mechanism is tolerably manipulable?

Example 1. *Consider a simple example with three items* $\{I_1, I_2, I_3\}$. *Most agents agree that* I_1 *is the best,* I_2 *is the second best, and* I_3 *is the least desirable. Therefore,* $S(I_1) = 1+\epsilon$, $S(I_2) = 2$ *and* $S(I_3) = 3-\epsilon$, *for* $\epsilon > 0$ *small. Suppose the first agent to arrive,* A_1, *has preferences* $\phi_1 : I_1 \succ_1 I_3 \succ_1 I_2$. *The SR mechanism computes* $\{1-S(I_1), 2-S(I_3), 3-S(I,2)\} = \{-\epsilon, -1+\epsilon, 1\}$ *and allocates* I_3 *to* A_1. *This is the right decision. With high probability, the agents* A_2 *and* A_3 *will have the common preference* $\phi_2 (= \phi_3) : I_1 \succ I_2 \succ I_3$. *Then the SR mechanism would have allocation* $A_1 \rightarrow I_3, A_2 \rightarrow I_1, A_3 \rightarrow I_2$ *for an average rank of 1.67. The APSD mechanism would have allocated* $A_1 \rightarrow I_1, A_2 \rightarrow I_2$ *and* $A_3 \rightarrow I_3$ *for an average rank of 2. However,* A_1 *could misreport his preferences to be* $\phi'_1 : I_1 \succ'_1 I_2 \succ'_1 I_3$. *Then the SR mechanism would compute* $\{1-S(I_1), 2-S(I_2), 3-S(I_3)\} = \{-\epsilon, 0, \epsilon\}$ *and allocate* $A_1 \rightarrow I_1$, *which is his top choice.*

Having seen that SR is not strategyproof, we now analyze the optimal manipulation for an agent A_k in the SR mechanism. Suppose $\phi'_k = I_1 \succ'_k I_2 \succ'_k \ldots \succ'_k I_m$ yields the allocation of I_j, and that this is the best obtainable item for A_k under SR. Clearly, $\phi''_k : I_j \succ''_k I_1 \succ''_k I_2 \ldots \succ''_k I_{j-1} \succ''_k I_{j+1} \succ''_k \ldots I_m$ also leads to the same allocation. In fact, we show that if the agent can win item I_j with some misreport, then it can always win the item by placing it first, followed by the claimed items, followed by the other items in order of ascending score. Let $avail$ denote the set of available items, and $claimed$ the rest of the items. We propose the following manipulation algorithm for agent A_k:

1. Select the most preferred item $I_1 \in avail$.
2. Consider a preference profile ϕ'_k, with items ordered, from most preferred to least preferred, as
$$[I_1, claimed, sorted(avail \setminus \{I_1\})] \qquad (1)$$
This reports I_1 as the most preferred item, followed by the claimed items in any order, followed by the rest of the items sorted in ascending order of score.
3. Apply the SR calculation to ϕ'_k. If SR allocates I_1, then report ϕ'_k. Else, repeat steps 1-2 for the second preferred item, I_2, third preferred item and so on until an item is obtained, and report the corresponding ϕ'_k.

Lemma 2. *The preference report generated by this manipulation is the agent's best response to the current state of the SR mechanism and given a particular set of item scores.*

Proof. It is sufficient to show that for any I_j, the ordering of the remaining items I_l will is optimal in the sense of minimizing their adjusted score $r^{-1}(k, \ell) - S(I_\ell)$. To see this, consider two adjacent items $\{I_1, I_2\}$ (not equal to I_j) in a reported preference order, for which $S(I_2) < S(I_1)$. Let u denote the reported rank of the item in the first position. We claim that $\min(u - S(I_2), u+1-S(I_1)) > \min(u - S(I_1), u+1-S(I_2))$, and therefore it is best to report I_2 before I_1. This is by case analysis. Case (i): $S(I_1) - S(I_2) \leq 1$. Now we have $u - S(I_2) \leq u+1-S(I_1)$, and also $u - S(I_1) \leq u+1-S(I_2)$. Therefore, $\min(u-S(I_2), u+1-S(I_1)) = u-S(I_2) > u-S(I_1) = \min(u-S(I_1), u+1-S(I_2))$. Case (ii): $S(I_1)-S(I_2) > 1$. Now we have $u+1-S(I_1) \leq u-S(I_2)$ and $u-S(I_1) \leq u+1-S(I_2)$. Then, $\min(u-S(I_2), u+1-S(I_1)) = u+1-S(I_1) > u-S(I_1) = \min(u-S(I_1), u+1-S(I_2))$. Finally, it is easy to see that it is always just as good to list the unavailable items immediately after item I_j. □

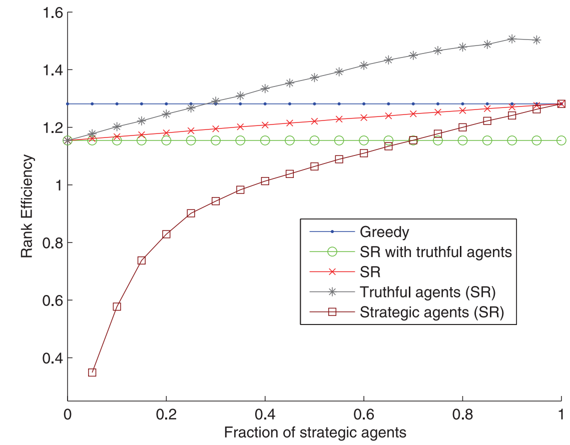

Figure 3: Rank efficiency of SR as the fraction of strategic agents varies, compared to SR with entirely truthful agents and to Greedy. 10 agents, 10 items, similarity=0.3.

Theorem 2. *If the difference in scores between successive items in SR is less than 1, and all agents are strategic, then the allocation under SR is identical to the APSD allocation.*

Proof. We show that agent A_k that follows its best response will receive its most preferred available item. The hardest case is when the agent's top ranked item is also the minimum score item. Label this item I_1. The optimal reported preference order is to sort the items in order of ascending score, e.g. $\phi'_k : I_1 \succ'_k I_2 \succ'_k \ldots \succ'_k I_m$. By assumption about consecutive scores, we have $1 - S(I_1) < 2 - S(I_2) < 3 - S(I_3) < \ldots < m - S(I_m)$ and SR would allocate the agent I_1. If the agent's top-ranked item is not the minimum score item, then the inequalities would still follow. □

This sufficient condition on adjacent scores can be interpreted within the weighted-popularity model. Sort the items, so that I_1 is the highest popularity, I_2 the second highest and so on. The condition on the gap between scores induces a simple requirement on popularity $\{p_1, p_2, \ldots, p_m\}$:

Proposition 1. *The scoring-gap condition is satisfied, and SR reduces to APSD under strategic behavior, in the weighted-popularity preference model if, for* $j = 1, 2, \ldots m-1$, *we have* $\frac{p_j}{\sum_{i=j}^m p_i} < \frac{p_{j+1}}{\sum_{i=j+1}^m p_i}$

The proof follows from series expansions and is omitted in the interest of space.

Experimental Results

Having established the first criteria that we introduce for tolerable manipulability (that the mechanism reduces to APSD when every agent is strategic) we now establish the second criteria: that the *ex ante* efficiency is better under SR than under ASPD when a sufficient fraction of agents fail to be

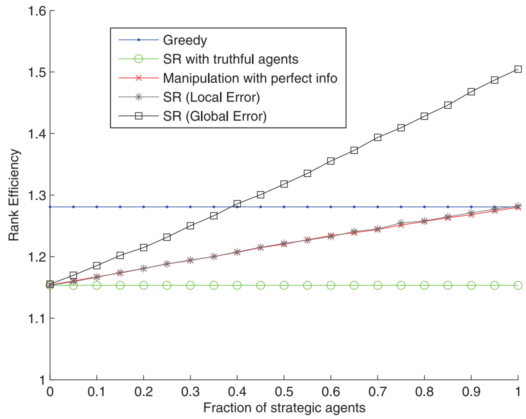

Figure 4: Rank efficiency of SR as the fraction of strategic agents varies and with errors in agent beliefs. 10 agents, 10 items, similarity =0.3.

rational and are instead truthful. In fact, the simulation results show that this holds *for any* fraction of strategic agents as long as they are well informed. We vary the fraction $q \in [0, 1]$ of agents that are strategic.

Figure 3 gives results for a simulation of 10 agents with 10 items and similarity of 0.3 (averaged over 5000 independent runs). The average rank-efficiency of SR worsens as the fraction of strategic agents increases, but is always better than APSD (i.e., Greedy) for any fraction $q < 1$ of strategic agents and equal to APSD when the fraction is one. Similar behavior is observed for other similarity values. We break down the performance of SR in terms of the average rank to strategic agents and to truthful agents. Strategic agents perform better than truthful agents and the gap is significant. However, when the fraction of strategic agents is less than a threshold (in this example 0.3), the truthful agents still do better than they would under APSD. Strategic agents receive better objects than they would receive in offline optimal.

We also consider the sensitivity of SR to strategic agents with imperfect information $\{\hat{p}_j\}$ about the relative popularity, and thus score, of items. Given a true profile of item popularities $\{p_1, \dots, p_m\}$ we consider two types of perturbations: *First experiment.* For each manipulating agent, for each item I_j, perturb the popularity so that $\hat{p_j}/p_j = \Delta_j \sim \mathcal{N}(1, z')$, where the variance, $z' \in \{0.1, 0.3\}$. This models "local" errors, in which the estimate for the popularity of each item is off by some random, relative amount. *Second experiment.* The true popularity profile, $\{p_j\}$, is constructed by evaluating the Normal pdf $\Psi_z(\frac{2j}{m})$ for each of $j \in \{1, \dots, m\}$, where Ψ_z is the density function for Normal distribution $\mathcal{N}(1, z)$. Here, we construct $\hat{p}_j = \Psi'_z(\frac{2j}{m})$, where Ψ'_z is the density function for Normal $\mathcal{N}(1 + \Delta, z)$ where $\Delta \sim U[-z''/2, z''/2]$. We consider $z'' \in \{0.1, 0.3\}$. This models "large-scale" errors, in which an agent has a systematic bias in popularity across all items.

Figure 4 presents the average rank-efficiency in the two experiments, for 10 agents, 10 items and similarity 0.3. We focus on $z' = 0.3$ and $z'' = 0.3$. The results are very similar for z' and z'' set to 0.1. The local perturbation turns out to not affect the performance much since $\{\hat{p}_i\}$ introduces a multiplicative error with zero expected bias. On the other hand, the large-scale perturbation can significantly reduce the rank-efficiency of SR under manipulative agents, introducing an additive error with bias; even here, SR still outperforms Greedy with less than 40% strategic agents.[4]

Conclusions

We have considered a dynamic mechanism design problem without money from the perspective of tolerable manipulability. This is an interesting domain because the unique strategyproof mechanism is easy to identify but has poor *ex ante* efficiency. Tolerable manipulability seems to be an interesting direction for practical, computational mechanism design, relaxing from the worst-case requirements of strategyproofness and providing better performance when only a fraction of agents are strategic and the rest are truthful. An appealing direction is to achieve stronger guarantees, finding a compromise between the framework we adopt and that of Othman and Sandholm (2009).

Acknowledgments

The second author would like to acknowledge Prof Y. Narahari and Infosys Technologies Pvt Ltd for financial support.

References

Abdulkadiroglu, A., and Loertscher, S. 2005. Dynamic house allocations. Technical report, Dept. of Economics, Columbia Univ.

Budish, E., and Cantillon, E. 2009. Strategic behavior in multi-unit assignment problems: Lessons for market design. Technical report, Harvard Business School.

Conitzer, V., and Sandholm, T. 2004. Computational criticisms of the revelation principle. In *Proceedings of the 5th ACM Conference on Electronic Commerce (EC-04)*, 262–263. (Short paper).

Feigenbaum, J., and Shenker, S. 2002. Distributed algorithmic mechanism design: Recent results and future directions. In *Proc. of the 6th Int. Workshop on Discrete Algorithms and Methods for Mobile Computing and Communication*, 1–13.

Hentenryck, P. V., and Bent, R. 2006. *Online Stochastic Combinatorial Optimization*. MIT Press.

Kurino, M. 2009. House allocation with overlapping agents: A dynamic mechanism design approach. Technical report, Department of Economics, University of Pittsburgh.

Lavi, R., and Nisan, N. 2005. Online ascending auctions for gradually expiring goods. In *SODA*.

Othman, A., and Sandholm, T. 2009. Better with Byzantine: Manipulation-Optimal mechanisms. In *Second International Symposium on Algorithmic Game Theory (SAGT)*.

Papai, S. 2001. Strategyproof and nonbossy multiple assignments. *Journal of Public Economic Theory* 3(3):257–71.

Parkes, D. C. 2007. Online mechanisms. In Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V., eds., *Algorithmic Game Theory*. Cambridge University Press. chapter 16.

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

[4]Although the analysis in this paper focuses on the case where the number of agents equals the number of items, simulation shows that SR continues to outperform APSD in the tight case when there are fewer items than agents.