Strategy-proof mechanisms for two-sided matching with minimum and maximum quotas

(Extended Abstract)

Suguru Ueda[1], Daniel Fragiadakis[2], Atsushi Iwasaki[1], Peter Troyan[2], and Makoto Yokoo[1]
1: Kyushu University, Japan, {ueda@agent., iwasaki@, yokoo@}inf.kyushu-u.ac.jp
2: Stanford University, California, {danielfragiadakis, petetroyan}@gmail.com

ABSTRACT

We consider the problem of allocating objects to agents when the objects have minimum quotas. There exist many real-world settings where minimum quotas are relevant. For example, in a hospital-resident matching problem, unconstrained matching may produce too few assignments to a rural hospital. Surprisingly, almost 50 years have passed after the seminal work by Gale and Shapley, no existing mechanism can guarantee minimum quotas so far; we did not know how to guarantee that a rural hospital has at least one resident.

In this paper, we propose mechanisms that can satisfy minimum quotas as well as standard maximum quotas. More specifically, we propose extended seat (ES) and multi-stage (MS) mechanisms modeled after the well-known deferred-acceptance (DA) and top trading cycles (TTC) mechanisms. Our proposed mechanisms are all strategy-proof, but a trade-off exists between the DA and TTC based mechanisms regarding Pareto efficiency and elimination of justified envy. In addition, there exist a tradeoff between ES and MS mechanisms depending on the size of minimum quotas.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Distributed Artificial Intelligence—*Multi-agent systems*; J.4 [**Social and Behavioral Sciences**]: Economics

General Terms

Algorithms, Economics, Theory

Keywords

Game theory, two-sided matching, deferred-acceptance, top trading cycle, minimum quotas

1. INTRODUCTION

The matching theory literature has developed numerous mechanisms to solve the problem of assigning objects to a group of agents when the agents have privately known preferences and the objects have priorities over the agents. Many problems fit broadly into this context, including assigning students to schools [1] and kidneys to patients [5]. Most of the previous literature considers only maximum quotas. However, in many real-world problems, minimum quotas may be imposed as well. This paper proposes several new mechanisms to deal with both quotas, starting from assigning students to labs at a university.

Two popular mechanisms that are often used with no minimum quotas are the deferred-acceptance (DA) mechanism [2] and the top-trading cycles (TTC) mechanism [4]. Both are strategy-proof, but TTC always produces a Pareto efficient assignment, while DA does not. However, DA eliminates justified envy, while TTC, which allows students to trade their priorities, does not.

The standard DA and TTC mechanisms may fail to produce feasible matchings that meet minimum quotas. Thus, we take the standard DA and TTC mechanisms and make two modifications to each: "extended seat" (ES) and "multi-stage" (MS). Thus, our four mechanisms are ES-DA, ES-TTC, MS-DA and MS-TTC. These mechanisms produce feasible matchings while preserving the good properties of DA and TTC.

Strategy-proofness is often taken as an important property in matching markets, and indeed, all four of our mechanisms are strategy-proof. However, there is a tradeoff between efficiency and fairness. The two TTC based mechanisms produce Pareto efficient assignments, but will lead to justified envy. It is known that there is no strategy-proof mechanism that completely eliminates justified envy when minimum quotas are imposed. Additionally, Hamada et al. showed that there may be no stable matching [3], extending the concept of justified envy. Even worse, they also showed that finding a matching with the minimum number of blocking pairs is hard to approximate even if labs have a master list (i.e., all labs use the same priority ordering).

Thus, we propose a slight strengthening of justified envy that eliminates some potential blocking pairs under the standard definition, and show that our two DA mechanisms eliminate all such justified envy. Thus, the choice between using a TTC mechanism or a DA mechanism depends on which goal, either Pareto efficiency or elimination of justified envy, policymakers consider more important.

2. MODEL

A market is a tuple $(S, L, p, q, \succ_S, \succ_L, \succ_{ML})$. $S = \{s_1, s_2, \ldots, s_n\}$ is a set of students, $L = \{l_1, l_2, \ldots, l_m\}$ a set of labs, and $p = (p_{l_1}, \ldots, p_{l_m})$ and $q = (q_{l_1}, \ldots, q_{l_m})$ are the minimum and maximum quotas, respectively, for

Appears in: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

each lab. We assume $0 \leq p_l \leq q_l$ for all $l \in L$ and $\sum_{l \in L} p_l \leq n \leq \sum_{l \in L} q_l$ to ensure a feasible matching exists. Define $e = n - \sum_{l \in L} p_l$ to be the number of "excess students".

Each student s has a strict preference relation $\succ_s$ over the labs, while each lab l has an idiosyncratic strict priority relation $\succ_l$ over the students. The vectors of all such relations are denoted $\succ_S = (\succ_s)_{s \in S}$ for the students and $\succ_L = (\succ_l)_{l \in L}$ for the labs. We assume that all labs are acceptable to all students and vice versa. In addition to the idiosyncratic lab preferences, there is a separate "master list (ML)" over all of the students. Without loss of generality, we let $s_1 \succ_{ML} s_2 \succ_{ML} \cdots \succ_{ML} s_n$.

A **matching** is a mapping $M : S \cup L \to S \cup L$ that satisfies: (i) $M(s) \in L$ for all $s \in S$, (ii) $M(l) \subseteq S$ for all $l \in L$, and (iii) for any s and l, we have $M(s) = l$ if and only if $s \in M(l)$. A matching is **feasible** if $p_l \leq |M(l)| \leq q_l$ for all $l \in L$. The labs are not strategic, i.e., $\succ_L$ and $\succ_{ML}$ are fixed and known to all students. Student s **envies** student s' at matching M if $M(s') \succ_s M(s)$ and if student s envies student s', this envy is **justified** if $s \succ_{M(s')} s'$. Also, we say s **strongly justifiably envies** s' at matching M if s envies s' and $s \succ_{ML} s'$ and $s \succ_{M(s')} s'$.

If student s justifiably envies student s', who is assigned laboratory $l = M(s')$, at matching M, then we say that student s and lab l form a blocking pair. A blocking pair can also be formed by a student and a lab with an empty seat, provided that moving the student to the lab with the empty seat results in a new matching that is feasible. Formally, given a matching M, student s **claims an empty seat at lab** l if (i) $l \succ_s M(s)$ (ii) $|M(l)| < q_l$ and (iii) $|M(M(s))| > p_{M(s)}$. Then, student s and lab l' form a **blocking pair** if either (i) s justifiably envies some student $s' \in M(l')$ or (ii) s claims an empty seat at l'.

3. EXTENDED-SEAT MECHANISMS

We consider an extended market $(S, \tilde{L}, \tilde{q}, \tilde{\succ}_S, \tilde{\succ}_L)$ where the set of students is unchanged, but for each "standard lab" l_j, we create an "extended lab" l_j^*. Thus, the set of labs is now $\tilde{L} = L \cup L^* = \{l_1, \ldots, l_m, l_1^*, \ldots, l_m^*\}$. In addition, we remove all minimum quotas, and define new maximum quotas $\tilde{q}_l$ for $l \in \tilde{L}$ as follows: if $l \in L$, we set $\tilde{q}_l = p_l$, while if $l^* \in L^*$, we set $\tilde{q}_{l^*} = q_l - p_l$.

For the lab priorities, if $l \in L$, then $\tilde{\succ}_l = \succ_l$; if $l^* \in L^*$, then $\tilde{\succ}_{l^*} = \succ_{ML}$. That is, the standard labs use the priorities from the original market, while all of the extended labs use the ML. For student s, the preferences over $L \cup L^*$ are created by taking the original preference relation $\succ_s$ and inserting lab l_j^* immediately after lab l_j. That is,

$$\text{preference relation } \succ_s\text{: } l_j \; l_k \; \cdots \text{ becomes } \tilde{\succ}_s : \; l_j \; l_j^* \; l_k \; l_k^* \; \cdots$$

Finally, no more than $e = n - \sum_{l \in L} p_l$ students can attain seats in extended labs. This restriction ensures that all quotas in the original matching problem will be satisfied.

Example 1. *[ES-DA] There are five students $s_1, \ldots, s_5$ and three labs l_1, l_2, l_3. For each lab, $p_l = 1$ and $q_l = 3$. The preferences and priorities are as follows:*

$$\begin{array}{llll} \succ_{s_1}: & l_1 & l_2 & l_3, \\ \succ_{s_2}: & l_2 & l_1 & l_3, \\ \succ_{s_3}: & l_2 & l_3 & l_1, \\ \succ_{s_4}: & l_2 & l_3 & l_1, \\ \succ_{s_5}: & l_2 & l_1 & l_3. \end{array} \qquad \begin{array}{llllll} \succ_{l_1}: & s_3 & s_5 & s_1 & s_2 & s_4, \\ \succ_{l_2}: & s_1 & s_4 & s_3 & s_5 & s_2, \\ \succ_{l_3}: & s_1 & s_2 & s_4 & s_5 & s_3. \end{array}$$

To run ES-DA, our extended market uses labs $L \cup L^* = \{l_1, l_2, l_3, l_1^*, l_2^*, l_3^*\}$, and maximum quotas $\tilde{q}_l = 1$ for $l \in L$ and $\tilde{q}_{l^*} = 3 - 1 = 2$ for $l^* \in L^*$. Note that there are no minimum quotas in this problem. We additionally modify all students' preferences by inserting lab l_j^* after lab l_j. For example, the modified preferences of student s_1 are as follows: $\tilde{\succ}_{s_1}$: $l_1 \; l_1^* \; l_2 \; l_2^* \; l_3 \; l_3^*$. For the lab priorities, we set $\tilde{\succ}_l = \succ_l$ for $l \in L$, while for $l^* \in L^*$, we set $\tilde{\succ}_{l^*} = \succ_{ML}$.

In round 1 of ES-DA, student s_1 applies to lab l_1 and students $s_2, \ldots s_5$ apply to lab l_2. Labs l_1 and l_2 tentatively accept s_1 and s_4, respectively. Lab l_2 rejects s_2, s_3 and s_5. In round 2, students s_2, s_3 and s_5 apply to l_2^*. Since only $e = 2$ students can be assigned to extended labs at the final matching, student s_5 is rejected. At the end, the following matching is produced:

$$l_1^* - \{s_1\}, \; l_2^* - \{s_2\}, \; l_3^* - \emptyset, l_1 - \{s_5\}, \; l_2 - \{s_4\}, \; l_3 - \{s_3\}.$$

Mapping this back to a matching in the original model:

$$l_1 - \{s_1, s_5\}, \; l_2 - \{s_2, s_4\}, \; l_3 - \{s_3\}.$$

The mechanism satisfies strategy-proofness and elimination of all strong justified envy. Furthermore, this idea can be applied to TTC and we obtain ES-TTC that is strategy-proof and efficient.

4. MULTI-STAGE MECHANISMS

The MS mechanisms proceed slightly differently. For these mechanisms, we first "reserve" a number of students equal to the sum of the minimum quotas across all labs. Then, we run the standard DA or TTC on the remaining set of students. This procedure is then repeated until all students are assigned. Because at each stage we reserve a number of students equal to the sum of the minimum quotas remaining, at the end of the mechanism, all minimum quotas will be satisfied. The MS mechanisms inherit the desirable properties from the ES mechanisms, i.e., MS-DA is strategy-proof and eliminates all strong justified envy, while MS-TTC is strategy-proof and efficient.

A tradeoff exists between the classes of ES and MS mechanisms, depending on the size of the minimum quotas. We empirically show that when the minimum quotas are small, the ES mechanisms tend to create many traditional blocking pairs compared to MS. When the minimum quotas are large, the reverse happens. Policymakers may find it advantageous to use the MS (ES) mechanisms when the minimum quotas are small (large).

5. REFERENCES

[1] A. Abdulkadiroglu and T. Sonmez. School choice: A mechanism design approach. *American Economic Review*, 93(3):729–747, 2003.

[2] D. Gale and L. S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.

[3] K. Hamada, K. Iwama, and S. Miyazaki. The hospitals/residents problem with quota lower bounds. In *19th Annual European Symposium on Algorithms*, pages 180–191, 2011.

[4] L. S. Shapley and H. Scarf. On cores and indivisibility. *Journal of Mathematical Economics*, 1:23–28, 1974.

[5] M. U. Unver. Dynamic kidney exchange. *Review of Economic Studies*, 77(1):372–414, 01 2010.