Strategyproof Multi-Item Exchange Under Single-Minded Dichotomous Preferences

JAAMAS Track

Haris Aziz
UNSW Sydney and Data61, CSIRO
Sydney, Australia
haris.aziz@unsw.edu.au

ABSTRACT

We consider multi-item exchange markets in which agents want to receive one of their target bundles of resources. The model encompasses well-studied markets for kidney exchange, lung exchange, and multi-organ exchange. We identify a general and sufficient condition called weak consistency for the exchange mechanisms to be strategyproof even if we impose any kind of distributional, diversity, or exchange cycle constraints. Within the class of weakly consistent and strategyproof mechanisms, we highlight two important ones that satisfy constrained Pareto optimality and strong individual rationality. Several results in the literature follow from our insights. We also derive impossibility results when constrained Pareto optimality is defined with respect to more permissive individual rationality requirements.

KEYWORDS

Housing markets, Organ exchange, and Strategyproofness

ACM Reference Format:
Haris Aziz. 2020. Strategyproof Multi-Item Exchange Under Single-Minded Dichotomous Preferences. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020,* IFAAMAS, 3 pages.

1 INTRODUCTION

Clearing complex exchange markets are one of the successful applications of algorithmic economics and multi-agent systems (see e.g., [5]). Key domains for which algorithms have been designed include housing markets and kidney exchange markets. In recent years, market designers are turning their attention to liver and lung exchange markets. There are proposals to explore the efficiency gains via multi-organ exchange markets (see e.g., [3]). Apart from exchanging organs or housing, several new digital platforms have come up that facilitate bartering of goods. Exchange markets are also useful to model time-bank scenarios in which people exchange services rather than items.

In most of the organ exchange markets, it is supposed that agents have single-minded dichotomous preferences over outcomes: they are either satisfied with an allocation or they are not. We consider a general model of item or organ exchange that captures all such organ markets. Each agent is endowed with a set of items and each agent has a set of target sets of items. An agent is *satisfied* if she gets any one of those target sets of items.

Our primary contribution is formalizing a general property of mechanisms called *weak consistency* and proving that any mechanism satisfying weak consistency is strategyproof (SP). Within the class of weakly consistent mechanisms, we highlight two subclasses of mechanisms which have proved useful in restricted domains. The first one is called *constrained priority (CP)* and the second one is called *constrained utilitarian priority (CUP)*. Both mechanisms run a serial dictatorship type priority mechanism on the set of feasible and individually rational allocations. We show that the mechanisms satisfy strategyproofness and constrained Pareto optimality. We then show how a subtle difference in imposing a different version of individual rationality results in impossibility results even for single-minded agents. We summarize the results in the paper [1].

2 MODEL AND CONCEPTS

An exchange market is a tuple $I = (N, O, e, D)$ where $N = \{1, \dots, n\}$ be a set of n agents and O be the set of items. The vector $e = (e_1, \dots, e_n)$ specifies the endowment $e_i \subset O$ of each agent $i \in N$. We suppose that $\bigcup_{i \in N} e_i = O$ and $e_i \cap e_j = \emptyset$ for all $i, j \in N$ such that $i \neq j$. . Each agent has a demand set $D_i \subset 2^{2^O}$. Each element of D_i is a bundle of items that is acceptable to agent i and meets her goal of being in the exchange market. We say that $I' = (N, O, e', D')$ is *more constrained* than $I = (N, O, e, D)$ is $e'_i \subseteq e_i$ for all $i \in N$ and $d'_i \supseteq d_i$ for all $i \in N$. We will write $I' \leq I$ if I' is more constrained than I. Just like the endowment, any allocation $x = (x_1, \dots, x_n)$ specifies the allocated bundle $x_i \subset O$ of each agent $i \in N$. In any allocation x, $x_i \cap x_j = \emptyset$ for all $i, j \in N$ such that $i \neq j$. Where the context is clear, we refer to the allocation bundle x_i as the allocation of agent i. We say that allocation x *satisfies* agent i if $x_i \supseteq d$ for some $d \in D_i$. Our model captures any kind of exchange market in which agents are interested in getting one of the acceptable bundles.

For any two allocations $x_i, y_i \subseteq O$, one can define the preference relation $\succsim_i$ of agent i where $x_i \succsim_i y_i$ if and only if y_i satisfies i implies that x_i satisfies i. The weak preference relation gives rise to the indifference relation $\sim_i$ which holds if $x_i \succsim_i y_i$ and $y_i \succsim_i x_i$. It also gives rise to the strict part $\succ_i$ of the relation where $x_i \succ_i y_i$ if $x_i \succsim_i y_i$ and $x_i \not\succsim_i y_i$.

An allocation x is *strongly individually rational (S-IR)* if $x_i \supseteq e_i$ or $x_i \supseteq d$ for some $d \in D_i$. Informally, if an agent's goal is not met, she is not interested in using up any of her endowed resources. This is a standard assumption is settings such as kidney exchange. Enforcing the S-IR requirement can also be viewed as a special type of feasibility constraint.

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand.

Our model can also model altruistic donors who are happy with any allocation. We are interested in maximizing the number of satisfied agents. We say that an allocation x is *constrained Pareto optimal* within the set of allocations satisfying set τ of properties if there exists no allocation y satisfying τ such that $y_i \succsim_i x_i$ for all $i \in N$ and $y_i \succ_i x_i$ for some $i \in N$.

Mechanisms. A mechanism M is a function that maps each problem instance $I = (N, O, e, D)$ into an allocation. We say that a mechanism is S-IR if it returns an S-IR allocation for each instance.

Let $\rho(I)$ denote the set of allocations that are feasible with respect to some feasibility constraints denoted by ρ. Depending on what the feasibility constraints ρ capture, the term $\rho(I)$ could denote the set of allocations based on pairwise exchanges or exchanges cycles of size at most 3, or satisfying some distribution or diversity constraint. In principle, ρ could also capture the S-IR requirement. Whatever feasibility conditions we assume on the model, we will assume that they allow the endowment allocation to be a feasible allocation.

Since we are interested in mechanisms that satisfy any given feasibility constraints represented by ρ, we can view a mechanism as a function that maps the set of ρ-feasible allocations to some ρ-feasible allocation. Without loss of generality, we can think of mechanisms as choice functions that choose one of the 'best' allocations from the set of feasible allocations. More generally, we will refer to $M(I, Y)$ as forcing M to return an allocation from the set of allocations $Y \subseteq \rho(I)$.

3 RESULTS

We first define weak consistency. Recall that in our setting, for any agent i with demand set D_i, and for any two allocations x_i and y_i, it is the case that $x_i \sim_i y_i$ if both x_i and y_i satisfy i with respect to D_i or neither satisfy her. Consider $\rho(I)$ a set of feasible allocations defined with respect to some feasibility constraints ρ. A mechanism M is *weakly consistent* if for all $I = (N, O, e, D)$, for any $I' \leq I$, for any Y, Y' such that $Y' \subseteq Y \subseteq \rho(I)$, and $x = M(I, Y)$, if there exists some feasible allocation $y \in Y' \subset Y$ under instance I' such that $x_i \sim_i y_i$ for all $i \in N$, then $M(I', Y') = z$ where $z_i \sim_i x_i$ for all $i \in N$. The preference relation $\sim_i$ used in the definition is with respect to instance I. Note that weak consistency is weaker than the consistency condition considered by Hatfield [4].

Theorem 3.1. *Any weakly consistent and S-IR mechanism is SP.*

Next, we focus on two natural weakly consistent mechanisms. The mechanisms are adaptations of the idea of applying serial dictatorship and priority over the set of all feasible outcomes (see e.g. [2]). The mechanisms are parametrized with respect to ρ (a set of feasibility constraints) and π which is a priority ordering over the agent set in which the agent in j-th turn is denoted by $\pi(j)$. Both rules are based on lexicographic comparisons.

For any permutation π of N, the CP mechanism is defined as follows. $CP(I, \rho, \pi) = \arg\max_{x \in \rho(I)}(u_{\pi(1)}(x), \ldots, u_{\pi(n)}(x))$. For any permutation π of N, the CUP mechanism is defined as follows. $CUP(I, \rho, \pi) = \arg\max_{x \in \rho(I)}(\sum_{i \in N} u_i(x), u_{\pi(1)}(x), \ldots, u_{\pi(n)}(x))$.

CP starts from the set of feasible allocations and then refines this set by using a priority ordering over the agents. CUP starts from the set of feasible allocations satisfying the maximum number of agents and then refines this set by using a priority ordering over the agents. Both CP and CUP are flexible enough to enforce or not enforce S-IR as part of the feasibility constraints.

The following remark points out that CUP is a general mechanism that has been used in restricted domains when S-IR is enforced.

Next, it is shown that CUP and CP are weakly consistent.

Lemma 3.2. *CUP and CP are weakly consistent.*

Theorem 3.3. *For any feasibility restriction ρ on the set of S-IR allocations, CUP is SP.*

Theorem 3.4. *For any feasibility restriction ρ on the set of S-IR allocations, CP is SP.*

If we include S-IR in the set of feasibility requirements ρ, then both CP and CUP return an allocation that is Pareto optimal allocation within the set of allocations satisfying S-IR and ρ.

We can rephrase our result in the form of the following theorem.

Theorem 3.5. *Under dichotomous preferences, for any restriction on allocations ρ, there exists a SP mechanism that returns an allocation that satisfies constrained Pareto optimality among the set of all allocations that satisfy ρ and S-IR.*

Since imposing S-IR is similar to imposing IR when agents have trichotomous preferences, we can rephrase the theorem above as follows.

Theorem 3.6. *Under trichotomous preferences where each agent only has her own endowment as the second preferred outcome, for any restriction on allocations ρ, there exists a strategyproof mechanism that returns an allocation that satisfies constrained Pareto optimality among the set of all allocations that satisfy ρ and IR.*

In the previous part, we limited our attention to allocations that satisfy S-IR (strong individual rationality). We now explore the consequence of dropping the S-IR requirement. An allocation x is *individually rational (IR)* if $e_i \supseteq d$ for some $d \in D_i$ implies that $x_i \supseteq d'$ for some $d' \in D_i$.

Whereas IR is a less stringent requirement than S-IR, constrained Pareto optimality with respect to allocations satisfying IR is a stronger property than constrained Pareto optimality with respect to allocations satisfying S-IR. In fact, the property is stronger enough that our central results in the previous sections collapse.

We say that an allocation is a *result of pairwise exchanges* if it is a result of pairs of agents make one-for-one exchange for items and with each item changing ownership at most once. Considering the exchange cycles view of allocations as mentioned in Section 3, an allocation as a result of pairwise exchange is characterized by exchange cycles in which each cycle has at most two agents in it.

Theorem 3.7. *Consider ρ as the restriction of allowing allocations that are a result of pairwise item exchanges in which agents get desirable items. For $|N| \geq 3$, there exists no SP mechanism that returns an allocation that is constrained Pareto optimal among the set of all allocations that satisfy ρ.*

ACKNOWLEDGMENTS

Aziz gratefully acknowledges the UNSW Scientia Fellowship and Defence Science and Technology (DST). He thanks Barton Lee, Bahar Rastegari, and the reviewers of JAAMAS for useful feedback.

REFERENCES

[1] H. Aziz. 2020. Strategyproof multi-item exchange under single-minded dichotomous preferences. *Autonomous Agents and Multi-Agent Systems* 34, 1 (2020), 3.

[2] H. Aziz, F. Brandt, and M. Brill. 2013. The Computational Complexity of Random Serial Dictatorship. *Economics Letters* 121, 3 (2013), 341–345.

[3] J. P. Dickerson and T. Sandholm. 2017. Multi-Organ Exchange. *J. Artif. Intell. Res.* 60 (2017), 639–679.

[4] J. W. Hatfield. 2005. Pairwise kidney exchange: Comment. *Journal of Economic Theory* 125 (2005), 189–193.

[5] T. Sönmez and M. U. Ünver. 2011. Matching, Allocation, and Exchange of Discrete Resources. In *Handbook of Social Economics*, J. Benhabib, M. O. Jackson, and A. Bisin (Eds.). Vol. 1. Elsevier, Chapter 17, 781–852.