An Overview of Maximin Share Results under Various Constraints

CS 580: Algorithmic Game Theory

Albert Cao   Estelle Lee

November 30, 2022

Presentation Outline

1. Introduction and Definitions
2. Negative MMS results
3. MMS with Cardinality Constraints
4. MMS with binary, ternary, and bivalued valuations
5. Conclusion

Recap: Fair Division

Definition

In an instance of the fair division problem, there are

- N agents
- M items
- Each agent i has a valuation function v_i
- Items have non-negative value
- Valuations are additive: $v_i(S) = \sum_{j \in S} v_i(\{j\})$

We want to "fairly" allocate items to agents.

Recap: Maximin Share (MMS)

Maximin Share (MMS) is a notion of fairness.

Definition

Fix an agent i. Let agent i allocate items into N bundles, with the caveat that every other agent gets to choose a bundle before i. The optimal strategy is to maximize the minimum value of a bundle and this value is known as the Maximin Share (MMS).

Recap: Maximin Share (MMS)

Maximin Share (MMS) is a notion of fairness.

Definition

Fix an agent i. Let agent i allocate items into N bundles, with the caveat that every other agent gets to choose a bundle before i. The optimal strategy is to maximize the minimum value of a bundle and this value is known as the Maximin Share (MMS).

Definition

An agent's α-MMS value is α fraction of their MMS value.

Negative MMS results

Is it guaranteed that there is an allocation where every agent gets their MMS?

Negative MMS results

Is it guaranteed that there is an allocation where every agent gets their MMS? No!

Negative MMS results

Is it guaranteed that there is an allocation where every agent gets their MMS? No!

Example

Let there be 3 agents and 12 items. Items are arranged in a 3×4 matrix.
Define a base matrix B, oscillation matrix O, and 3 epsilon matrices $E^{(1)}, E^{(2)}, E^{(3)}$. Agent i's value for the item (j, k) is defined as

$$v_i(\{(j, k)\}) = 10^6 \cdot B_{jk} + 10^3 \cdot O_{jk} + E_{jk}^{(i)}$$

It is possible to choose values for those matrices such that there is no allocation where every agent gets their MMS.
(Procaccia and Wang [9])

Negative MMS results

Here's the value definition again:

$$v_i(\{(j,k)\}) = 10^6 \cdot B_{jk} + 10^3 \cdot O_{jk} + E_{jk}^{(i)}$$

On a high level, the purpose of each matrix is as follows:

- Base matrix: force every MMS allocation to contain 4 items in each bundle
- Oscillation matrix: constrain which bundles of 4 items can be in a MMS allocation
- Epsilon matrices: perturb each agent's valuations so that they prefer different partitions

Feige, Sapir, and Tauber [4] found an example with 3 agents and 9 items such that no allocation can guarantee 39/40-MMS for all agents.

Stronger negative results

Feige, Sapir, and Tauber [4] found an example with 3 agents and 9 items such that no allocation can guarantee 39/40-MMS for all agents.

For more agents: currently not known how to construct instances with a constant gap, the best gap is $\frac{1}{N^4}$

Cardinality Constraint Definition

What if we imposed a a limit on the number of items an agent could receive?

- Recall examples in class and on homework where an agent had unit demand, meaning that only the highest valued item they received contributed value to the agent.

Cardinality Constraint Definition

What if we imposed a a limit on the number of items an agent could receive?

- Recall examples in class and on homework where an agent had unit demand, meaning that only the highest valued item they received contributed value to the agent.

Definition

In fair division with cardinality constraint k, agents only receive the value of the k highest valued items in the bundle. An agent's MMS-k value is the value of the minimum bundle when they optimally allocate items into N bundles of size $\leq k$.

MMS with cardinality constraint $k = 1$ (Unit Demand)

Theorem

For any instance of the fair division problem, there exists an allocation that guarantees each agent their MMS-1 value.

MMS with cardinality constraint $k = 1$ (Unit Demand)

Theorem

For any instance of the fair division problem, there exists an allocation that guarantees each agent their MMS-1 value.

Proof.

We can run a "one round round-robin algorithm": Allow agents to pick an item from the unallocated set one at a time. An agent's MMS-1 value is the value of the N-th best item. They can always get the N-th best item (or better). $\square$

MMS with cardinality constraint $k = 2$

Theorem

For any instance of the fair division problem, there exists an allocation that guarantees each agent their MMS-2 value.

MMS with cardinality constraint $k = 2$

Theorem

For any instance of the fair division problem, there exists an allocation that guarantees each agent their MMS-2 value.

How? We can run a two round round-robin algorithm, with the second round proceeding in reverse order.

MMS with cardinality constraint $k = 2$

Example

Suppose there are 3 agents who all have the same valuation function and 6 items of values $\{10, 8, 7, 5, 3, 1\}$. The allocation would look like:

	Agent 1	Agent 2	Agent 3
Round 1	10		
Round 2			

MMS with cardinality constraint $k = 2$

Example

Suppose there are 3 agents who all have the same valuation function and 6 items of values $\{10, 8, 7, 5, 3, 1\}$. The allocation would look like:

	Agent 1	Agent 2	Agent 3
Round 1	10	8	
Round 2			

MMS with cardinality constraint $k = 2$

Example

Suppose there are 3 agents who all have the same valuation function and 6 items of values $\{10, 8, 7, 5, 3, 1\}$. The allocation would look like:

	Agent 1	Agent 2	Agent 3
Round 1	10	8	7
Round 2			

MMS with cardinality constraint $k = 2$

Example

Suppose there are 3 agents who all have the same valuation function and 6 items of values $\{10, 8, 7, 5, 3, 1\}$. The allocation would look like:

	Agent 1	Agent 2	Agent 3
Round 1	10	8	7
Round 2			5

MMS with cardinality constraint $k = 2$

Example

Suppose there are 3 agents who all have the same valuation function and 6 items of values $\{10, 8, 7, 5, 3, 1\}$. The allocation would look like:

	Agent 1	Agent 2	Agent 3
Round 1	10	8	7
Round 2		3	5

MMS with cardinality constraint $k = 2$

Example

Suppose there are 3 agents who all have the same valuation function and 6 items of values $\{10, 8, 7, 5, 3, 1\}$. The allocation would look like:

	Agent 1	Agent 2	Agent 3
Round 1	10	8	7
Round 2	1	3	5

MMS with cardinality constraint $k = 2$

Proof sketch.

Fix an agent i. Assume there are only $2N$ items (other items are irrelevant). To maximize the minimum bundle, agent i should pair the j-th best item with the j-th worst item. If there is a bundle with both items from the top N and a bundle with both items from the bottom N, we should swap items from those two bundles to increase the value of the minimum bundle, e.g.

$$\{10, 7\}, \{3, 5\} \to \{10, 3\}, \{7, 5\}$$

Similarly, if the j-th best item is not paired with the j-th worst item, then we should swap items to increase the value of the minimum bundle, e.g.

$$\{10, 3\}, \{8, 5\}, \{7, 1\} \to \{10, 1\}, \{8, 5\}, \{7, 3\} \to \{10, 1\}, \{8, 3\}, \{7, 5\}$$

2/3-MMS for cardinality constraint k = 3

Theorem

For any instance of the fair division problem, there exists an allocation that guarantees each agent 2/3 of their MMS-3 value.

2/3-MMS for cardinality constraint k = 3

Theorem

For any instance of the fair division problem, there exists an allocation that guarantees each agent 2/3 of their MMS-3 value.

This result uses the algorithm we just saw. An agent's MMS-3 value is at most $\frac{3}{2}$ of their MMS-2 value. (Proof omitted.)

2/3-MMS for cardinality constraint k = 3

Theorem

For any instance of the fair division problem, there exists an allocation that guarantees each agent 2/3 of their MMS-3 value.

This result uses the algorithm we just saw. An agent's MMS-3 value is at most $\frac{3}{2}$ of their MMS-2 value. (Proof omitted.)
Can we do better if we add a third round?

2/3-MMS for cardinality constraint k = 3

Theorem

For any instance of the fair division problem, there exists an allocation that guarantees each agent 2/3 of their MMS-3 value.

This result uses the algorithm we just saw. An agent's MMS-3 value is at most $\frac{3}{2}$ of their MMS-2 value. (Proof omitted.)
Can we do better if we add a third round? Unfortunately not.

Example

Suppose there are 2 agents with the same valuation function and 6 items $\{9, 3, 3, 3, 0, 0\}$. The first 2 rounds allocate the bundles $B_1 = \{9, 3\}, B_2 = \{3, 3\}$ and it doesn't matter what happens in the third round. Agent 2 only gets value 6, even though the optimal allocation is $B_1 = \{9, 0, 0\}, B_2 = \{3, 3, 3\}$.

NP-hardness of computing an agent's MMS-3 value

Theorem

Computing the MMS value with cardinality constraint k=3 for a single agent is NP-hard

NP-hardness of computing an agent's MMS-3 value

Theorem

Computing the MMS value with cardinality constraint k=3 for a single agent is NP-hard

Proof: Reduction from 3-partition.

NP-hardness of computing an agent's MMS-3 value

Theorem

Computing the MMS value with cardinality constraint k=3 for a single agent is NP-hard

Proof: Reduction from 3-partition.

Definition

The 3-partition problem is a decision problem, where, given a multiset S of $M = 3N$ positive integers such that the sum of all integers is NT, the output is the existence of a partition of N disjoint subsets of size 3 such that the sum of each subset is T.

Example

$S = \{20, 23, 25, 30, 49, 45, 27, 30, 30, 40, 22, 19\}$ can be partitioned into the four sets $\{20, 25, 45\}$,$\{23, 27, 40\}$,$\{49, 22, 19\}$,$\{30, 30, 30\}$ each of which sums to $T = 90$

NP-hardness of computing an agent's MMS-3 value

Theorem

Computing the MMS-3 value for a single agent is NP-hard

Proof sketch: Take an instance of the 3-partition problem and construct an instance of the MMS-3 problem.

- N agents
- Each $s \in S$ becomes an item where each agent has valuation for that item equal to s

An agent has MMS value T if and only if the 3-partition problem has a valid partition.

Existence of MMS-3

Theorem

For $N = 3$ agents, an MMS-3 allocation always exists.

Proof from Feige et al. [4]

Open Problem

For $N \geq 4$ agents, does an MMS-3 allocation always exist?

MMS-k existence for k $\geq$ 4

Theorem

An MMS-4 allocation does not always exist

MMS-k existence for k $\geq$ 4

Theorem

An MMS-4 allocation does not always exist

Proof.

Recall Procaccia and Wang [9]'s example for $N = 3$ agents and $M = 12$ items which showed that an MMS allocation does not always exist. Their example uses a base matrix to force every allocation to contain 4 items in each bundle, which exactly resolves to the MMS-4 case. $\square$

MMS with binary, ternary, and bivalued valuations

Previously we discussed the MMS problem within the lens of cardinality constraints on the number of items an agent can receive. Another variant of the MMS problem that we can consider is one where each agent has some fixed values they can have for each item.

Or in other words, for each agent i, their valuation function v_i satisfied

$$v_i \in \mathcal{F}$$

where $\mathcal{F}$ is a class of functions we can define.

MMS with binary valuations

The simplest constrained valuation function is one where each agent has binary valuations. This means that their valuation for each item is in $\{0, 1\}$. Intuitively, this can be thought of as a system where each agent either approves or disapproves of an item.

$$\mathcal{F} = \{v : \{1, ..., M\} \to \{0, 1\}\}$$

MMS with binary valuations

The simplest constrained valuation function is one where each agent has binary valuations. This means that their valuation for each item is in $\{0, 1\}$. Intuitively, this can be thought of as a system where each agent either approves or disapproves of an item.

$$\mathcal{F} = \{v : \{1, ..., M\} \to \{0, 1\}\}$$

This type of allocation problem can be easily be proved to have an MMS allocation through a simple round-robin allocation.

Theorem

For any instance of the fair allocation problem where agents have binary valuation functions, an MMS allocation exists.

MMS with binary valuations

Theorem

For any instance of the fair allocation problem where agents have binary valuation functions, an MMS allocation exists.

Proof.

Consider what the MMS of an agent is when they can only have either valuation 0 or 1 for an item. To maximize the minimum bundle, they will simply distribute all of the items with value 1 to the n bundles as evenly as possible. Remember that Bouveret and Lemiatre [2] showed that we can assume each agent's ranking of all the items is the same, even if their valuations may be different. Then if we follow the order of this ranking for round robin allocation, for each agent all items of value 1 will be distributed evenly. □

MMS with ternary valuations

The natural extension to binary valuations is ternary valuations, in which agents have possible valuations from $\{0, 1, 2\}$.

$$\mathcal{F} = \{v : \{1, ..., M\} \to \{0, 1, 2\}\}$$

MMS with ternary valuations

The natural extension to binary valuations is ternary valuations, in which agents have possible valuations from $\{0, 1, 2\}$.

$$\mathcal{F} = \{v : \{1, ..., M\} \to \{0, 1, 2\}\}$$

Amanatidis et al. [1] proved that this scenario always has an MMS allocation. The proof involves a variant of the greedy round robin allocation, where bundles are assigned items in a greedy fashion but are then modified by reversing rounds of the round robin so that they satisfy every agent's MMS requirement.

Theorem (Amanatadis et al. [1])

For any instance of the fair allocation problem where agents have ternary valuation functions, an MMS allocation exists.

MMS with bivalued valuations

Another extension to binary valuations is the idea of bivalued valutions. This means that instead of having possible valuations from $\{0, 1\}$, agents have possible valuations from $\{a, b\}$, where a and b are integer constants (and $a < b$ WLOG).

$$\mathcal{F} = \{v : \{1, ..., M\} \to \{a, b\}\}$$

MMS with bivalued valuations

Another extension to binary valuations is the idea of bivalued valutions. This means that instead of having possible valuations from $\{0, 1\}$, agents have possible valuations from $\{a, b\}$, where a and b are integer constants (and $a < b$ WLOG).

$$\mathcal{F} = \{v : \{1, ..., M\} \to \{a, b\}\}$$

Feige [3] proved that instances with bivalued valuations do have an MMS allocation.

Theorem (Feige [3])

For any instance of the fair allocation problem where agents have bivalued valuation functions, an MMS allocation exists.

References I

Amanatidis, Georgios, et al. ”Approximation algorithms for computing maximin share allocations.” ACM Transactions on Algorithms (TALG) 13.4 (2017): 1-28.

Bouveret, Sylvain; Lemaitre, Michel (2015). Characterizing conflicts in fair division of indivisible goods using a scale of criteria. Autonomous Agents and Multi-Agent Systems. 30 (2): 259

Feige, Uriel. Maximin Fair Allocations with Two Item Values. 22 Feb. 2022, https://www.wisdom.weizmann.ac.il/ feige/mypapers/MMSab.pdf.

Feige, Uriel, Ariel Sapir, and Laliv Tauber. ”A tight negative example for MMS fair allocations.” International Conference on Web and Internet Economics. Springer, Cham, 2021.

Garg, Jugal, McGlaughlin, Peter, and Taki, Setareh. Approximating Maximin Share Allocations. Retrieved from https://par.nsf.gov/biblio/10083883. Open access series in informatics 69. Web. 10.4230/OASIcs.SOSA.2019.20.

References II

Garg, Jugal, and Setareh Taki. ”An improved approximation algorithm for maximin shares.” Proceedings of the 21st ACM Conference on Economics and Computation. 2020.]

Ghodsi, Mohammad, et al. ”Fair allocation of indivisible goods: Improvements and generalizations.” Proceedings of the 2018 ACM Conference on Economics and Computation. 2018.

Hummel, Halvard, and Magnus Lie Hetland. ”Guaranteeing half-maximin shares under cardinality constraints.” arXiv preprint arXiv:2106.07300 (2021)

Procaccia, Ariel D., and Junxing Wang. ”Fair enough: Guaranteeing approximate maximin shares.” Proceedings of the fifteenth ACM conference on Economics and computation. 2014.

Thanks!

Questions?