PMMS Allocations Need Not Exist for 3 Agents and Additive Valuations

Paul Gölz

September 5, 2026

Consider the setting of fair division with indivisible goods: Let $N = \{1, \dots, n\}$ denote the set of agents, M a set of m goods, and let each agent $i \in N$ have a value $u_i(\alpha) \geq 0$ for each good $\alpha \in M$. We assume that preferences are additive, which means that an agent i's valuation for a bundle of goods $B \subseteq M$ is simply $u_i(B) \coloneqq \sum_{\alpha \in B} u_i(\alpha)$. Our goal is to find an allocation of the goods, i.e., a partition of M into bundles $A_1, \dots, A_n$, each associated with its corresponding agent.

Caragiannis et al. [CKM+19] introduced the fairness axiom of *pairwise maximin share (PMMS)*. One reason to be interested in this axiom is that it is stronger (under mild conditions, for example that all agent–item values are positive) than *envy-freeness up to any good (EFX)*, whose existence is arguably the largest open question in fair division.

Definition 1. *An allocation $A_1, \dots, A_n$ satisfies* pairwise maximin share (PMMS) *if, for any two agents $i \neq j$, it holds that*

$$u_i(A_i) \geq \max_{S \subseteq (A_i \cup A_j)} \min \Big(u_i(S), u_i((A_i \cup A_j) \setminus S) \Big).$$

But even whether PMMS allocations always exist for additive preferences has since remained an open question, and was for example mentioned as Open Question 8 by Amanatidis et al. [ABF+24]. In this note, I give a counterexample, showing that PMMS allocations need not exist. That this impossibility only requires $n = 3$ agents contrasts with EFX, which is known to exist for three agents and additive preferences [CGM20].

Indeed, consider the following fair division instance with $n = 3$ agents and $m = 9$ goods $\alpha_1, \dots, \alpha_9$, with additive preferences, which has no PMMS allocation:

agent i	$u_i(\alpha_1)$	$u_i(\alpha_2)$	$u_i(\alpha_3)$	$u_i(\alpha_4)$	$u_i(\alpha_5)$	$u_i(\alpha_6)$	$u_i(\alpha_7)$	$u_i(\alpha_8)$	$u_i(\alpha_9)$
1	1	3	4	5	8	14	16	18	23
2	10	2	15	6	8	22	31	30	32
3	14	2	22	13	15	27	42	38	48

I have not yet found a conceptual proof showing that no PMMS allocation exists that would not devolve into a longish, uninformative case distinction. Computationally, it is easy to verify by complete enumeration that no PMMS allocation exists, for example with the following Python code, which can also be found at https://gist.github.com/pgoelz/e213cebb5ff6ee7ee00acd802246f88a:

```
from itertools import product

def bundle_value(utils, allocation, agent, bundle):
```

```python
    """Calculate u_i(A_j) for agent i and bundle j."""
    n = len(utils)
    m = len(allocation)
    return sum(utils[agent][j] for j in range(m) if allocation[j] == bundle)

def is_pmms(utils, allocation):
    """Check through exhaustive enumeration whether the allocation is PMMS for
    the given utilities.

    Represent an allocation as a vector of length m, where each entry gives
    the index (between 0 and n-1) of the agent the good is assigned to.
    """
    n = len(utils)
    m = len(allocation)
    for agent in range(n):
        for bundle in range(n):
            if bundle == agent:
                continue
            ownvalue = bundle_value(utils, allocation, agent, agent)
            reallocationindices = [j for j in range(m) if allocation[j]
                                   in [agent, bundle]]
            for split in product([agent, bundle],
                                 repeat=len(reallocationindices)):
                newallocation = list(allocation)
                for good, towhom in zip(reallocationindices, split):
                    newallocation[good] = towhom
                if (ownvalue < bundle_value(utils, tuple(newallocation),
                                            agent, agent)
                    and ownvalue < bundle_value(utils, tuple(newallocation),
                                                agent, bundle)):
                    return False  # PMMS-violating split found
    return True

def test_pmms_counterexample(utils):
    """Test through exhaustive enumeration whether there really is no PMMS
    allocation for the given additive utilities.

    Represent utils as a list of n lists (one per agent), each inner list of
    length m containing the agent's values for the goods in order.
    """
    n = len(utils)
    assert n >= 1
    m = len(utils[0])
    assert all(len(agentutils) == m for agentutils in utils)

    for allocation in product(range(n), repeat=m):
        if is_pmms(utils, allocation):
            print("Not a counterexample to PMMS, PMMS allocation exists:",
```

```
                allocation)
            return False
    print("No PMMS allocation exists, valid counterexample:", utils)
    return True


if __name__ == "__main__":
    candidate = [[1, 3, 4, 5, 8, 14, 16, 18, 23],
                 [10, 2, 15, 6, 8, 22, 31, 30, 32],
                 [14, 2, 22, 13, 15, 27, 42, 38, 48]]
    test_pmms_counterexample(candidate)
```

AI Usage

The text of this note and the above script for verifying the counterexample are entirely written by hand. But the counterexample was found by a large language model with a large degree of autonomy. I found an initial counterexample with $m = 12$ goods using OpenAI's GPT-6 Astra model in ultra mode, with a prompt that simply asked to resolve the existence of PMMS allocations. I subsequently tried a number of parallel attempts at simplifying the instance with ChatGPT, integer linear programming, and manual inspection. ChatGPT produced the above 9-good instance. I will expand on this note shortly.

References

[ABF+24] Georgios Amanatidis, Georgios Birmpas, Federico Fusco, Philip Lazos, Stefano Leonardi, and Rebecca Reiffenhäuser. "Allocating Indivisible Goods to Strategic Agents: Pure Nash Equilibria and Fairness". In: *Mathematics of Operations Research* 49.4 (2024), pp. 2425–2445.

[CGM20] Bhaskar Ray Chaudhury, Jugal Garg, and Kurt Mehlhorn. "EFX Exists for Three Agents". In: *Proceedings of the ACM Conference on Economics and Computation (EC)*. 2020, pp. 1–19.

[CKM+19] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. "The Unreasonable Fairness of Maximum Nash Welfare". In: *ACM Transactions on Economics and Computation* 7.3 (2019), pp. 1–32.