

Fair Division with Market Values

Siddharth Barman¹, Soroush Ebadian², Mohamad Latifian³, Nisarg Shah²

¹Indian Institute of Science Bangalore

²University of Toronto

³University of Edinburgh

barman@iisc.ac.in, soroush@cs.toronto.edu, mohamad.latifian@ed.ac.uk, nisarg@cs.toronto.edu

Abstract

We introduce a model of fair division with market values, where indivisible goods must be partitioned among agents with (additive) subjective valuations, and each good additionally has a market value. The market valuation can be viewed as a separate additive valuation that holds identically across all the agents. We seek allocations that are simultaneously fair with respect to the subjective valuations and with respect to the market valuation.

We show that an allocation that satisfies stochastically-dominant envy-freeness up to one good (SD-EF1) with respect to both the subjective valuations and the market valuation does not always exist, but the weaker guarantee of EF1 with respect to the subjective valuations along with SD-EF1 with respect to the market valuation can be guaranteed. We also study a number of other guarantees such as Pareto optimality, EFX, and MMS. In addition, we explore non-additive valuations and extend our model to cake-cutting. Along the way, we identify several tantalizing open questions.

1 Introduction

For centuries, human civilizations have pondered how to settle disputes centered around division of goods. Such resolution requirements come up often in cases such as divorce settlement and estate division (in the absence of a will). A common practice in the real world is not to liquidate all the goods but to assess their *fair market value* (FMV) and create an “equal division” whereby the total fair market value of the goods awarded to each party is (almost) equal. Such a division is sometimes even a legal requirement (Kagan 2021b). For instance, fair market values of land—as determined by assessors (Kagan 2021a)—are often used to settle land disputes (Shtechman, Gonen, and Segal-Halevi 2021). However, an FMV-based approach has two limitations: (i) parties may disagree on what the fair market value of a good is, and (ii) parties may have personal (i.e., subjective) cardinal preferences (over the goods) which may be quite different from the assessed market values.

Starting with the seminal work of Steinhaus (1948), the mathematical theory of fair division has made significant advances over the last century. The theory systematically addresses the issue (ii) mentioned above by developing models

wherein the participating agents (parties) can submit their own *heterogeneous* valuations over the goods. Here, the underlying goal is to find a division that each agent would consider fair *according to their own valuation*. Hence, the models provide a systematic way to avoid disagreements. In recent years, fair division literature has made impressive strides in the context of indivisible goods, resulting in allocation methods that achieve appealing fairness guarantees. One such notion is *envy-freeness up to one good* (EF1), which demands that no agent prefers the allocation of another agent over their own if we hypothetically exclude just one good from the envied agent’s allocation (Lipton et al. 2004; Budish 2011; Caragiannis et al. 2019).

Despite the mathematical and conceptual appeal of this framework, equal division of goods under market values still holds unshakable importance in real-world dispute resolutions. Hence, we study the framework of *fair division with market values*, where we seek an allocation of the goods that the agents find fair with respect to their heterogeneous personal preferences (“subjective utilities”), and that is also a near-equal division under the market values. Formally, a set \mathcal{M} of (indivisible) goods must be allocated among a set \mathcal{N} of n agents. Each agent $i \in \mathcal{N}$ submits an additive subjective utility function $u_i : 2^{\mathcal{M}} \rightarrow \mathbb{R}_{\geq 0}$ and there is an additive market value function $v : 2^{\mathcal{M}} \rightarrow \mathbb{R}_{\geq 0}$.¹ Our goal is to seek an allocation that is simultaneously fair (e.g., EF1) with respect to the subjective utilities $\{u_i\}_{i \in \mathcal{N}}$ and with respect to the market values v .

The significance of this formulation is further substantiated by the following observations: First, near-equal division under the market values can correspond to a feasibility constraint, imposed due to legal or normative reasons. Hence, one can view the requirement under market values as seeking an allocation that is fair (with respect to the agents’ subjective utilities) as well as *feasible* (with respect to the market values). Indeed, with this viewpoint, the current work contributes to the growing body of work on fair division under constraints (see, e.g., (Suksompong 2021)), and some of our results are derived using ideas from this line of work.

Second, one can adopt an epistemic perspective that the subjective utilities provided by the agents are in fact their

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹A set function $f : 2^{\mathcal{M}} \rightarrow \mathbb{R}_{\geq 0}$ is said to be additive if $f(S) = \sum_{r \in S} f(\{r\})$ for all $S \subseteq \mathcal{M}$ with $f(\emptyset) = 0$.

(personal) noisy estimates of the market values of the goods. Further, the “market values” that we take as input may be an expert’s estimate or algorithmic prediction. Achieving an equal division with respect to both agents’ personal estimates and the external estimate can provide greater robustness and prevent future disputes.

Our work is closely related to two very recent papers, which also study fair division in the presence of market values.² Dall’Aglío (2023) studies the same model, but with divisible goods and the subjective utility of each agent, for every good, modeled a function of the good’s market value. The focus of Dall’Aglío (2023) is on the behavior of prominent rules, rather than on concrete fairness guarantees. The work of Bu et al. (2023) is closer to ours, who consider a generalization of our model with *heterogeneous* market values $\{v_i\}_{i \in \mathcal{N}}$, motivated as an allocator’s preference over which agent should get which good. Their result for this general model establishes guarantees for proportionality, a notion weaker than the one studied in the current paper, namely, envy-freeness. In particular, there are two results where the current paper and the work of Bu et al. (2023) overlap: first, we show that a fairness guarantee for allocation of goods obtained by Bu et al. (2023) can be viewed as a corollary of a result due to Biswas and Barman (2018). Second, we show that an existential result by Bu et al. (2023) can be turned into a polynomial-time algorithm in our model (while Bu et al. (2023) provide a polynomial-time algorithm for a weaker notion). Our study of Pareto optimality, EFX, and MMS in this framework is entirely novel. We discuss these technical points at length in Section 1.2.

1.1 Our Contributions

In Section 3, we begin by asking whether there always exists an allocation of indivisible goods that satisfies *stochastic-dominance envy-freeness up to one good* (SD-EF1)—a strengthening of EF1—w.r.t.³ both additive subjective utilities and additive market values. We answer the question in the negative, using a carefully crafted instance with 2 agents and 7 goods. Complementing this result, we prove positive results for 2 agents and up to 6 goods, and for any number of agents with identical utilities. Then, relaxing the requirement slightly, we prove that an allocation that is EF1 w.r.t. subjective utilities and SD-EF1 w.r.t. market values always exists via a reduction to an algorithm by Biswas and Barman (2018). The complementary guarantee of SD-EF1 w.r.t. agents’ utilities and EF1 w.r.t. market valuation stands as an interesting open question.

Next, in Section 4, we consider other desiderata such as Pareto optimality (PO), maximin share fairness (MMS), and envy-freeness up to any good (EFX). In Section 4.1, our main result is that an allocation that is PO w.r.t. subjective utilities and EF1 w.r.t. market values always exists, but there are open questions surrounding several other combinations

²Our work was conducted concurrently and independently of these works. That said, there is very little overlap between our results and the ones obtained in (Dall’Aglío 2023) and (Bu et al. 2023), which we discuss at length in Section 1.2.

³We will, throughout, abbreviate ‘with respect to’ as w.r.t.

of desiderata. Section 4.2 shows that MMS or EFX cannot be attained w.r.t. one side (subjective utilities or market values) together with EF1, MMS, or EFX on the other side. Finally, in Section 4.3, we observe that the existence of an allocation that is EF1 w.r.t. even monotone subjective utilities and monotone market values remains open, but prove the existence of allocations that are $1/2$ -EF1 w.r.t. subadditive subjective utilities and SD-EF1 w.r.t. additive market values. These results are summarized in Table 1.

		Market	
		SD-EF1	EF1
Agents	SD-EF1	✗ (Theorem 1)	?
	EF1	✓ (Corollary 1)	
	PO	✗ (Theorem 4)	✓ (Corollary 2)
	MMS/EFX	✗ (Theorems 6 and 7)	

Table 1: Summary of our results. PO on the market side is always guaranteed. Positive result with PO on the agents’ side is conditioned on strictly positive utilities. MMS and EFX impossibilities hold even when swapping sides (i.e., EF1/SD-EF1 on the agents’ side and MMS/EFX on the market side), or demanding MMS/EFX on both sides.

Finally, in Section 5, we extend our setup of fair division with market values to cake-cutting. That is, we consider partitioning a heterogeneous divisible good (cake) among agents with subjective utilities in the presence of a market valuation, both subjective utilities and market valuation represented as countably additive measures over the cake. This model is formally introduced in Section 5. We observe that the existence of an allocation that is EF w.r.t. both subjective utilities and market values can be reduced to the known existence of a *perfect* cake division (Dubins and Spanier 1961). We prove bounds on the number of cuts (of the cake) and queries (of the subjective utility measures and the market value measure) required to achieve EF w.r.t. both sides. We also show that there may *not* exist any cake division that is EF+PO w.r.t. the subjective utilities along with EF w.r.t. the market values. Though, relaxing EF+PO to simply PO w.r.t. the subjective utilities yields a positive result.

Throughout, we identify a number of interesting and fundamental open questions.

1.2 Related Work

As mentioned previously, the works of Dall’Aglío (2023) and Bu et al. (2023) are the closest to ours. None of the results obtained in Dall’Aglío (2023) imply anything in our model, since this prior work addresses *homogeneous* divisible goods. Bu et al. (2023) study a generalization of our model with *heterogeneous* market valuations $\{v_i\}_{i \in \mathcal{N}}$. They view $\{v_i\}_{i \in \mathcal{N}}$ as the allocator’s preferences over which agent should get which item. Bu et al. (2023) provide four existential results:

1. EF1 on both sides, when there is a single market valuation (i.e., our model).
2. EF1 on both sides for the particular case of two agents.

3. PROP- $O(\log n)$ on both sides, where PROP (proportionality) is a notion weaker than envy-freeness.
4. PROP-2 on both sides under binary valuations.

We also study the first result in this list, but show that, in the case of goods (which is what Bu et al. address), the result is already implied by a result of Biswas and Barman (2018). In the list above, the second result is strong, but it does not come with an efficient construction. In fact, they show how to achieve EF2 on both sides in polynomial time, which we are able to improve to EF1, on both sides, in polynomial time. Their EF1 existential result extends to monotone subjective utilities and monotone market values, whereas our EF1 result extends only to monotone subjective utilities and additive market values (though still with polynomial value query complexity). The last two results in the list above address the weaker notion of proportionality, which we do not focus on. In addition, we establish negative results (e.g., Theorem 1) in a specialized version of the model considered in their work. Hence, the negative results are stronger. We also extend our results to the cake cutting setting, which is only mentioned in the appendix by Bu et al. (2023). In summary, the technical overlap between the work of Bu et al. (2023) and our work is little. Together, however, our works point to interesting open questions, which we highlight throughout.

A bit more broadly, fairness with respect to market values can be viewed as placing feasibility constraints (induced by the market values) on the allocation. Hence, fair division with market values can be considered as case of constrained fair division. Notably, many recent results address fair allocations subject to various feasibility constraints, such as cardinality constraints (Biswas and Barman 2018), matroid constraints (Dror, Feldman, and Segal-Halevi 2023), or budget constraints (Barman et al. 2023). As we show in this work, EF1 with respect to an additive market valuation can be achieved by placing appropriately defined cardinality constraints. However, it is not clear if EF1 with respect to a more general (e.g. submodular) market valuation or with respect to heterogeneous additive market valuations can be reduced to known constraints.

Our model is also related to recent works on fair division in two-sided markets (Patro et al. 2020; Freeman, Micha, and Shah 2021; Igarashi et al. 2023), where agents on two sides of a market are matched to each other, with agents on each side having preferences over those on the other side. Freeman, Micha, and Shah (2021) also seek EF1 on both sides, but in their case the second side has its own allocation, whereas in our case, the market side is also evaluated using the same allocation.

2 Preliminaries

An instance of fair division with market values is a quadruple, $(\mathcal{N}, \mathcal{M}, v, \vec{u} = (u_i)_{i \in \mathcal{N}})$, where \mathcal{N} is a set of n agents, \mathcal{M} is a set of m indivisible goods, $v : 2^{\mathcal{M}} \rightarrow \mathbb{R}_{\geq 0}$ is a monotone set function indicating the market valuation of the goods, and $u_i : 2^{\mathcal{M}} \rightarrow \mathbb{R}_{\geq 0}$ is the (subjective) monotone utility function of agent $i \in \mathcal{N}$. In particular, $u_i(S)$ and $v(S)$ denote agent i 's value and the market value for subset

of goods $S \subseteq \mathcal{M}$, respectively. Throughout most of the paper, we consider market valuation v and utilities \vec{u} that are *additive*: recall that a set function $f : 2^{\mathcal{M}} \rightarrow \mathbb{R}_{\geq 0}$ is additive if $f(S) = \sum_{g \in S} f(\{g\})$ for all $S \subseteq \mathcal{M}$. Section 4.3, in particular, goes beyond additive utilities and considers general, monotone ones. For notational convenience, we will write $u_i(g) := u_i(\{g\})$ and $v(g) := v(\{g\})$.

An allocation is n -partition $A = (A_1, \dots, A_n)$ of \mathcal{M} , where $A_i \subseteq \mathcal{M}$ denotes the bundle of goods given to agent i . Here, $A_i \cap A_j = \emptyset$ for all $i \neq j$ and $\cup_{i \in \mathcal{N}} A_i = \mathcal{M}$. Our goal is to find allocations that are fair with respect to both subjective utilities and market valuation. A prominent fairness criterion for indivisible goods is envy-freeness up to one good.

Definition 1 (Envy-free up to one good). For an instance $(\mathcal{N}, \mathcal{M}, v, (u_i)_{i \in \mathcal{N}})$, we say that an allocation $A = (A_1, \dots, A_n)$ is envy-free up to one good (EF1) with respect to subjective utilities if, for every pair of agents $i, j \in \mathcal{N}$, either $A_j = \emptyset$ or $u_i(A_i) \geq u_i(A_j \setminus \{g\})$ for some good $g \in A_j$. In addition, we say that A is EF1 with respect to market values if the above-mentioned inequalities hold with u_i replaced by v .

Next, we define a strengthening, SD-EF1, based on the stochastic dominance (SD) relation.

Definition 2 (SD-preference (Bogomolnaia and Moulin 2001)). For $X, Y \subseteq \mathcal{M}$, we say that agent i SD-prefers X to Y , denoted $X \succsim_i^{\text{SD}} Y$, if, for each good $g \in \mathcal{M}$, it holds that $|\{g' \in X : u_i(g') \geq u_i(g)\}| \geq |\{g' \in Y : u_i(g') \geq u_i(g)\}|$. That is, for each g , the subset X has at least as many goods with utility (to agent i) at least as much as g as Y does. Also, we define SD-preference under the market valuation, denoted $X \succsim_v^{\text{SD}} Y$, by replacing u_i with v .

Definition 3 (SD-EF1 (Aziz et al. 2024)). An allocation $A = (A_1, \dots, A_n)$ is said to be SD-envy-free up to one good (SD-EF1) if, for every pair of agents $i, j \in \mathcal{N}$, either $A_j = \emptyset$ or $A_i \succsim_i^{\text{SD}} A_j \setminus \{g\}$ for some good $g \in A_j$.

We will say that an additive function $u : \mathcal{M} \rightarrow \mathbb{R}_{\geq 0}$ induces a ranking (with ties) $\pi : \mathcal{M} \rightarrow [|\mathcal{M}|]$ if π ranks the goods in decreasing order of u , i.e., $\pi(g) \leq \pi(g')$ iff $u(g) \geq u(g')$ for all $g, g' \in \mathcal{M}$. Here, a weaker notion is that of consistency: u is said to be consistent with π if $\pi(g) \leq \pi(g')$ implies $u(g) \geq u(g')$ for all $g, g' \in \mathcal{M}$.

Note that SD-EF1 is a stronger notion than EF1: if an allocation is SD-EF1 with respect to utilities $(u_i)_{i \in \mathcal{N}}$, then it must be EF1 with respect to $(u_i)_{i \in \mathcal{N}}$, as well as with respect to any other utility profile $(u'_i)_{i \in \mathcal{N}}$ where, for each i , the utility u'_i is consistent with the ranking π_i induced by u_i .

The following lemma provides a useful observation about SD-EF1 with respect to a profile with identical valuations; recall that the market valuation can be seen as such a profile.

Lemma 1. *Let $\pi : \mathcal{M} \leftrightarrow [m]$ be a strict ranking over \mathcal{M} . If u_i induces π for every $i \in \mathcal{N}$, then A is SD-EF1 with respect to $(u_i)_{i \in \mathcal{N}}$ iff, for all agents $i \in \mathcal{N}$ and each index $\ell \in \{1, 2, \dots, \lceil m/n \rceil\}$, we have $|A_i \cap \{a \in \mathcal{M} : n\ell - n + 1 \leq \pi(a) \leq n\ell\}| \leq 1$. Further, if u'_i is consistent with π for every $i \in \mathcal{N}$, then such an allocation is still SD-EF1 with respect to $(u'_i)_{i \in \mathcal{N}}$.*

All the missing proofs can be found in the full version⁴.

3 Simultaneous Envy-Freeness

This section first shows that the strong notion of SD-EF1 cannot be guaranteed w.r.t. both the subjective utilities and the market valuation simultaneously.

Theorem 1. *There exists an instance with 2 agents and 7 goods in which no allocation is SD-EF1 w.r.t. both the subjective utilities and the market valuation.*

Proof. Let $\pi = (g_1 \succ g_2 \succ g_3 \succ g_4 \succ g_5 \succ g_6 \succ g_7)$ be the ranking induced by the market values, and those induced by the subjective utilities of the two agents be

$$\begin{aligned}\sigma_1 &= (g_1 \succ g_3 \succ g_2 \succ g_5 \succ g_4 \succ g_7 \succ g_6), \\ \sigma_2 &= (g_1 \succ g_5 \succ g_2 \succ g_3 \succ g_6 \succ g_7 \succ g_4).\end{aligned}$$

The impossibility result stated in the theorem will hold for any subjective utilities and market valuation that induce the above-mentioned rankings.

By Lemma 1, SD-EF1 w.r.t. the market valuation demands that each agent receive exactly one good from each of the sets $C_1 = \{g_1, g_2\}$, $C_2 = \{g_3, g_4\}$, and $C_3 = \{g_5, g_6\}$.

Suppose agent 1 receives good g_1 . Then, agent 2 must receive both good g_5 (for SD-EF1 w.r.t. the subjective utilities) and good g_2 (for SD-EF1 w.r.t. the market values). Then, considering the top four goods in σ_1 , agent 1 must receive g_3 to satisfy SD-EF1 w.r.t. its subjective utility. Next, recalling the constraints imposed by SD-EF1 w.r.t. market values through sets C_2 and C_3 above, we conclude that agent 2 must get g_4 and agent 1 must get g_6 . So far, we have $\{g_1, g_3, g_6\} \subseteq A_1$ and $\{g_2, g_4, g_5\} \subseteq A_2$. Since both agents have only two of their top six goods, both agents would need to have g_7 in order to satisfy SD-EF1 w.r.t. their subjective utilities, which is a contradiction.

The case of agent 2 receiving good g_1 leads to a similar argument, where we can deduce $\{g_2, g_3, g_6\} \subseteq A_1$ and $\{g_1, g_4, g_5\} \subseteq A_2$. Hence, in this case as well, both agents would require g_7 for SD-EF1 w.r.t. their subjective utilities, which is a contradiction. \square

We complete the picture by providing positive results for achieving SD-EF1 w.r.t. both the subjective utilities and the market values in two special cases.

Theorem 2. *In instances with two agents and at most six goods, there always exists an allocation that is SD-EF1 w.r.t. both the subjective utilities and the market valuation.*

Theorem 3. *When the subjective utilities of the agents induce the same ranking over the goods, an allocation that is SD-EF1 w.r.t. the subjective utilities and the market valuation always exists, and can be computed in polynomial time.*

Next, let us consider relaxing the goal of SD-EF1 to (the regular) EF1, at least on one side. Doing so on the market side leads to the following open problem, which seems to require considerably novel ideas towards a resolution.

Open Question 1: Does there always exist an allocation that is SD-EF1 w.r.t. the subjective utilities and EF1 w.r.t. the market valuation?

When we relax SD-EF1 to EF1 on the agents' side, a positive result emerges, even in the general case. We remark that Bu et al. (2023) prove the existence of an allocation that is EF1 w.r.t. both the subjective utilities and the market values. Upon careful observation, however, one finds that not only do they actually achieve SD-EF1 w.r.t. the market values, but in fact, their algorithm is precisely an instantiation of a more general algorithm designed previously by Biswas and Barman (2018). In other words, this problem can be reduced to the result of Biswas and Barman (2018); we lay out the simple reduction here, highlighting that there is no need to derive this result from scratch.

Biswas and Barman (2018) consider a setting in which there are no market values, but we are given a partition $\{\mathcal{M}_1, \dots, \mathcal{M}_\ell\}$ of the goods along with cardinality bounds $h_1, \dots, h_\ell \in \mathbb{Z}_+$. The goal here is to find a fair allocation A that satisfies cardinality constraints: $|A_i \cap \mathcal{M}_k| \leq h_k$ for each $k \in [\ell]$ and agent $i \in \mathcal{N}$. Note that, we need the bound $h_k \geq \lceil |\mathcal{M}_k|/n \rceil$, for each $k \in [\ell]$, to even ensure the existence of an allocation which satisfies the cardinality constraints. Biswas and Barman (2018) show that, in such a case, an EF1 allocation satisfying the cardinality constraints is guaranteed to exist.

For our reduction, we consider π_v the ranking induced by the market valuation v . Then, we set $\ell := \lceil m/n \rceil$, and, for each $k \in [\ell]$, define $\mathcal{M}_k := \{g \in \mathcal{M} : (k-1)n < \pi_v(g) \leq kn\}$. That is, \mathcal{M}_1 is the set of top n goods according to the market values, \mathcal{M}_2 is the set of next n goods according to the market values, and so on. Finally, we impose the cardinality constraints $|A_i \cap \mathcal{M}_k| \leq 1$ for all $i \in \mathcal{N}$ and $k \in [\ell]$. From Lemma 1, we see that this implies SD-EF1 w.r.t. the market values. The polynomial-time algorithm of Biswas and Barman (2018) then immediately yields an allocation that is EF1 w.r.t. the subjective utilities while satisfying the cardinality constraints (and, hence, SD-EF1 w.r.t. the market values).

Corollary 1. *An allocation that is EF1 w.r.t. the subjective utilities and SD-EF1 w.r.t. the market valuation always exists, and can be computed in polynomial time.*

4 Extensions to Other Criteria

Let us now extend our setup of fair division with market values to (i) consider other desiderata such as Pareto optimality, MMS, and EFX, and (ii) go beyond additive utilities.

4.1 Pareto Optimality

Definition 4 (Pareto Optimality). An allocation A is said to be *Pareto optimal* (PO) w.r.t. the subjective utilities if there is no allocation A' with the property that $u_i(A'_i) \geq u_i(A_i)$, for all $i \in \mathcal{N}$, with at least one inequality being strict.

In the full version, we address a strengthening called *fractional Pareto optimality* (fPO), which demands that not even a fractional allocation (where the goods may be fractionally assigned to agents) Pareto dominates the given integral one.

⁴Full version: <https://arxiv.org/pdf/2410.23137>

First, we note that every allocation is trivially PO w.r.t. the market values. Hence, the only interesting consideration is to demand PO w.r.t. the subjective utilities. Here, we show that PO cannot be achieved with SD-EF1 (Theorem 4). Interestingly, SD-EF1 implies both EF1 and balancedness ($|A_i| \in \{\lfloor m/n \rfloor, \lceil m/n \rceil\}$ for all i), and both these relaxations can individually be attained together with PO (Caragiannis et al. 2019; Conitzer, Freeman, and Shah 2017).

Theorem 4. *There exists an instance with 2 agents and 6 goods in which there is no allocation that is PO w.r.t. the subjective utilities and SD-EF1 w.r.t. the market valuation.*

Proof. Consider an instance with two agents with the following subjective utilities for goods g_1, \dots, g_6 :

	g_1	g_2	g_3	g_4	g_5	g_6
Agent 1	19	7	4	3	2	1
Agent 2	8	7	6	5	4	3

In addition, let the additive market valuation be identical to the utility of the first agent, i.e., $v = u_1$. Note that both agents must receive three goods, respectively, to ensure SD-EF1 (across the six goods) w.r.t. the market valuation.

Since agent 1 would be happy to give up any three goods to receive g_1 and agent 2 would be happy to accept any three goods in exchange for g_1 , PO demands that g_1 must be given to agent 1. Then, for SD-EF1, g_2 must be given to agent 2.

Among the remaining four goods $\{g_3, g_4, g_5, g_6\}$, each agent must receive two goods in such a way that no agent gets both g_3 and g_4 (or equivalently, both g_5 and g_6). However, it is easy to see that agent 1 would be happy to give up $A_1 \cap \{g_3, g_4, g_5, g_6\}$ to receive $\{g_2\}$, and agent 2 would also be happy to give up $\{g_2\}$ to receive $A_1 \cap \{g_3, g_4, g_5, g_6\}$, violating PO. \square

This leads us, in the PO context, to the following possibility: EF1+PO w.r.t. the subjective utilities and EF1 w.r.t. the market values. We leave this as an open question.

Open Question 2: Does there always exist an allocation that is EF1+PO w.r.t. the subjective utilities and EF1 w.r.t. the market values (or even balanced)?

Finally, relaxing SD-EF1 to EF1 w.r.t. the market values, we show that PO w.r.t. the subjective utilities can be attained together with EF1 w.r.t. the market values. Note that if the market values of all goods are set to be equal, then EF1 is equivalent to balancedness; thus, our result strictly strengthens the known result that a balanced PO allocation always exists Conitzer, Freeman, and Shah (2017).

In fact, we prove a slightly more general result, whereby we allow the market values to be heterogeneous (where each agent i has an additive market valuation v_i) and achieve *equitability up to one good* (EQ1) (Freeman et al. 2019) with respect to it.

Definition 5 (EQ1). An allocation A is called *equitable up to one good* (EQ1) w.r.t. a valuation profile $\vec{v} = (v_i)_{i \in \mathcal{N}}$ if, for every pair of agents $i, j \in \mathcal{N}$, either $A_j = \emptyset$ or $v_i(A_i) \geq v_j(A_j \setminus \{g\})$ for some good $g \in A_j$.

Freeman et al. (2019) prove that an allocation satisfying EQ1+PO does not always exist, but it does (and even EQ1+fPO can be guaranteed) when all values are nonzero (i.e., strictly positive). To obtain EQ1 and PO w.r.t. two different valuation profiles, respectively, it is clear that the profile w.r.t. which we want PO must have nonzero values.⁵ We show that PO (even fPO) w.r.t. a nonzero valuation profile and EQ1 w.r.t. a different valuation profile can be attained, strictly strengthening the result of Freeman et al. (2019).

Theorem 5. *An allocation that is fPO w.r.t. nonzero subjective utilities $(u_i)_{i \in \mathcal{N}}$ and EQ1 w.r.t. (heterogeneous) market valuations $(v_i)_{i \in \mathcal{N}}$ always exists, and can be computed in time polynomial in n , m , and V , where $V = \max_{i \in \mathcal{N}} |\{v_i(S) : S \subseteq \mathcal{M}\}|$ is the maximum number of distinct value levels for any i .*

Since both profiles are heterogeneous, one can swap them to achieve EQ1 w.r.t. subjective utilities and fPO w.r.t. heterogeneous nonzero market values as well. Since EQ1 and EF1 coincide when the market valuations are identical, $v_i = v$ for i , which is the case we consider throughout the paper, we obtain the following corollary.

Corollary 2. *There always exist an allocation that is fPO w.r.t. nonzero subjective utilities and EF1 w.r.t. the market valuation.*

4.2 MMS and EFX

In addition to EF1, two fairness criteria that have received significant attention in the discrete fair division are *maximin share fairness* (MMS) and *envy-freeness up to any good* (EFX). Here, we show that these notions lead to stark impossibilities in our setup of fair division with market values.

Definition 6 (α -MMS). Under utilities $(u_i)_{i \in \mathcal{N}}$, the maximin share of agent i is defined as $\mu_i := \max_{(B_1, \dots, B_n)} \min_{k \in [n]} u_i(B_k)$; here the max is considered over all n -partitions of \mathcal{M} . For $\alpha \in [0, 1]$, an allocation A is said to be α -MMS, w.r.t. subjective utilities, if $u_i(A_i) \geq \alpha \cdot \mu_i$ for all agents $i \in \mathcal{N}$.

We know that EF1 implies $(1/n)$ -MMS (Caragiannis et al. 2019). Hence, in the guarantee of Corollary 1 (EF1 w.r.t. the subjective utilities and SD-EF1 w.r.t. the market values), the fairness property on one or both sides can be replaced by $(1/n)$ -MMS. However, the following result shows that this is the best one can hope for, despite the fact that in the traditional fair division setup, even $(\frac{3}{4} + \frac{1}{12n})$ -MMS is known to exist (Akrami and Garg 2024).

Theorem 6. *For any $\alpha > 1/n$, there exists an instance with two identical valuation profiles (either one can be viewed as consisting of identical subjective utilities and the other as consisting of market values) in which no allocation is simultaneously α -MMS w.r.t. the first profile and either EF1 or α -MMS w.r.t. the second,*

⁵Otherwise, consider three goods positively valued only by agent 1, and a fourth good positively valued only by agent 2. The only PO allocation gives the first three goods to agent 1 and the fourth to agent 2, but this is not EQ1 under a valuation profile where all agents value all goods equally at 1.

Proof. Consider an instance with n agents and $2n-1$ goods, where the goods are partitioned into two sets, \mathcal{M}_1 with n goods and \mathcal{M}_2 with $n-1$ goods. Consider two identical valuation profiles with functions u and v as follows:

$$u(g) = \begin{cases} 1 & g \in \mathcal{M}_1 \\ n & g \in \mathcal{M}_2 \end{cases}, \quad v(g) = \begin{cases} 1 & g \in \mathcal{M}_1 \\ 0 & g \in \mathcal{M}_2 \end{cases}.$$

Note that EF1 or α -MMS (for any $\alpha > 0$) w.r.t. v requires that each agent receive exactly one good from \mathcal{M}_1 . Since $|\mathcal{M}_2| < n$, there exists an agent $i \in \mathcal{N}$ who does not receive any good from \mathcal{M}_2 . For this agent, $u_i(A_i) = 1$ whereas, under u , the maximin share $\mu_i = n$, causing a violation of α -MMS w.r.t. u , for any $\alpha > 1/n$. \square

Definition 7 (α -EFX). For $\alpha \in [0, 1]$, allocation A is said to be α -envy-free up to any good (α -EFX), w.r.t. the subjective utilities, if, for every pair of agents $i, j \in \mathcal{N}$, either $A_j = \emptyset$ or $u_i(A_i) \geq \alpha \cdot u_i(A_j \setminus \{g\})$ for every $g \in A_j$.

Theorem 7. For any $\alpha \in (0, 1]$, there exists an instance with two identical valuation profiles (either one can be viewed as consisting of identical subjective utilities and the other as market values) in which no allocation is simultaneously EF1 w.r.t. the first profile and α -EFX w.r.t. the second.

Proof. Consider an instance with 2 agents, 4 goods $\{g_1, g_2, g_3, g_4\}$, and two identical valuation profiles with valuation functions u and v as follows: $u(g_t) = 1$ for all $t \in [4]$, $v(g_1) = 2/\alpha + 1$, and $v(g_t) = 1$ for $t \in \{2, 3, 4\}$.

Note that EF1 w.r.t. u requires that each agent receive two goods. Consider any allocation A . Without loss of generality, assume that agent 1 receives g_1 and g_2 . Then, $v(A_2) = 2$ whereas $\alpha v(A_1 \setminus \{g_2\}) = 2 + \alpha > 2$, demonstrating a violation of α -EFX. \square

Remark 1. Since EFX implies EF1, we can replace EF1 with EFX in Theorem 6 to deduce that EFX w.r.t. one profile and α -MMS (for any $\alpha > \frac{1}{n}$) w.r.t. the other cannot be guaranteed.

4.3 Monotone Valuations

This section addresses whether one can extend the positive result of Corollary 1 to hold beyond additive valuations. That is, we consider whether EF1 can be achieved simultaneously for monotone (but not necessarily additive) subjective utilities and market valuation.⁶

We are interested in the broad classes of monotone, subadditive valuations. A valuation function $f: \mathcal{M} \rightarrow \mathbb{R}_{\geq 0}$ is said to be monotone if $f(X) \leq f(Y)$, for all $X \subseteq Y \subseteq \mathcal{M}$, and subadditive if $f(X \cup Y) \leq f(X) + f(Y)$ for all $X, Y \subseteq \mathcal{M}$.

An allocation that is EF1 w.r.t. monotone subjective utilities is guaranteed to exist (Lipton et al. 2004). Notably, with market values, the question remains open.

⁶SD-EF1 is not well-defined beyond additive valuations: there is no natural ordering over individual goods induced by a monotone valuation, and even if one were to induce it somehow, it would not suffice to guarantee EF1 w.r.t. every monotone valuation inducing it. Hence, we limit our attention to EF1.

Open Question 3: Does there always exist an allocation that is EF1 w.r.t. two heterogeneous monotone valuation profiles simultaneously (i.e., w.r.t. monotone subjective utilities and heterogeneous monotone market values)? What if one of the profiles consisted of additive valuations?

In fact, consider the special case of (homogeneous) additive market values where every good has the same market value. Then, EF1 w.r.t. such market values is simply balancedness, yielding the following tantalizing open question for the usual fair division setup without market values.

Open Question 4: Does there always exist a balanced EF1 allocation w.r.t. a (heterogeneous) monotone valuation profile?

While these questions remain open, we prove the existence of an allocation that is $1/2$ -EF1 w.r.t. subadditive subjective utilities and SD-EF1 w.r.t. additive market values. Formally, for $\alpha \in [0, 1]$, an allocation A is said to be α -EF1 w.r.t. subjective utilities $(u_i)_{i \in \mathcal{N}}$ if, for each pair of agents $i, j \in \mathcal{N}$ (with $A_j \neq \emptyset$), we have $u_i(A_i) \geq \alpha u_i(A_j \setminus \{g\})$ for some good $g \in A_j$. Here, we use the idea of minimally envied subsets from (Chaudhury et al. 2021).

Minimal Envied Swap (MES) Algorithm. As we did for establishing Corollary 1, we consider π_v the ranking induced by the (additive) market valuation v . Then, we partition the items into $(\mathcal{M}_1, \dots, \mathcal{M}_{\lceil m/n \rceil})$ such that $\mathcal{M}_k := \{g \in \mathcal{M} : (k-1)n < \pi_v(g) \leq kn\}$. As noted previously, for any allocation A , if $|A_i \cap \mathcal{M}_k| \leq 1$ for each $k \in [\lceil m/n \rceil]$ and all agents i , then A is SD-EF1 w.r.t. v . We now provide an algorithm that finds such an allocation that is $1/2$ -EF1 w.r.t. monotone, subadditive subjective utilities.

We start with an allocation A with empty bundles, $A_i = \emptyset$, for all agents $i \in \mathcal{N}$. Also, initialize ‘charity’ $C = \mathcal{M}$ as the set of all the unallocated goods. A subset $T \subseteq C$ is said to be an *envied feasible* subset of the charity if (1) there exists an agent $i \in \mathcal{N}$ envying it, i.e., $u_i(A_i) < u_i(T)$, and (2) $|T \cap \mathcal{M}_k| \leq 1$ for each $k \in [\lceil m/n \rceil]$. Further, we say that an envied feasible subset $T \subseteq P$ is *minimal* if no agent envies any strict subset of T . It is easy to verify that while there exists an envied feasible subset of the charity, there also exists a minimal envied feasible subset of the charity. The algorithm now works in two phases.

Phase 1: While there exists a minimal envied feasible subset T of the charity C , select an agent i who envies T , and update $A_i = T$ and charity (unallocated good) $C = \mathcal{M} \setminus (\cup_{i \in \mathcal{N}} A_i)$.

Phase 2: Once Phase 1 has terminated, allocate goods in C arbitrarily subject to maintaining $|A_i \cap \mathcal{M}_k| \leq 1$ for all agents $i \in \mathcal{N}$ and each $k \in [\lceil m/n \rceil]$.

Theorem 8. Every fair division instance with (monotone) subadditive subjective utilities and additive market values admits an allocation that is $1/2$ -EF1 w.r.t. subjective utilities and SD-EF1 w.r.t. market valuation. Further, such an allocation is returned by the MES algorithm detailed above.

Proof. First, we argue that the algorithm terminates and returns an allocation that is SD-EF1 w.r.t. the market values. Note that each iteration of Phase 1 strictly increases the utility of the selected agent, and the utilities of the remaining agents remain unchanged. Hence, Phase 1 must terminate. Further, throughout the execution of Phase 1, the selected subset of goods T is feasible, i.e., throughout, we maintain $|A_i \cap \mathcal{M}_k| \leq 1$ for each $i \in \mathcal{N}$ and $k \in \lceil \lceil m/n \rceil \rceil$. Therefore, at the end of Phase 1 and for each k , the number of agents with exactly one good from \mathcal{M}_k is equal to $|\mathcal{M}_k \setminus C|$. This observation implies that there remain $n - |\mathcal{M}_k \setminus C|$ agents who can each still receive a good from \mathcal{M}_k , while maintaining the cardinality constraint. Since $|\mathcal{M}_k| \leq n$, we obtain the following bound for the unassigned goods $|\mathcal{M}_k \cap C| \leq n - |\mathcal{M}_k \setminus C|$. Hence, we can allocate the unassigned goods in $\mathcal{M}_k \cap C$, one each, to the remaining agents while maintaining the feasibility constraints. As noted above, this ensures SD-EF1 w.r.t. the market values.

To show that the returned allocation is $1/2$ -EF1 w.r.t. the utilities, we note that the initial allocation (with empty bundles) is EF1. During Phase 1, assigning $A_i = T$ ensures no agent envies any strict subset of A_i . Therefore, the allocation remains EF1 even after the swap and, hence, the EF1 property is maintained until the end of Phase 1. During Phase 2, each agent i receives a subset of goods $B_i \subseteq C$ that is feasible ($|B_i \cap \mathcal{M}_k| \leq 1$ for all k) and, hence, is not envied by any other agent j . This follows from the fact that there are no envied feasible subsets of the charity C left after Phase 1. Finally, the subadditivity of the utilities ensures that the final allocation is $1/2$ -EF1. The theorem stands proved. \square

Next, we provide an exact version of Theorem 8 for the special case of two agents. The proof of this theorem follows closely from a result of Kyropoulou, Suksompong, and Voudouris (2020) and can be found in the full version.

Theorem 9. *For the case of two agents, there always exists an allocation that is EF1 w.r.t. monotone subjective utilities and SD-EF1 w.r.t. additive market values, and it can be computed using polynomially many value queries.*

5 Cake Division

This section extends our fair division with market valuation setup to heterogeneous divisible goods, i.e., to the canonical setting of cake cutting. We mention the cake division results here; the technical details can be found in the full version.

The cake is represented by the interval $[0, 1]$ and a piece of the cake is a finite union of intervals of $[0, 1]$. Here, among n agents, an allocation $A = (A_1, \dots, A_n)$ is a partition of the cake $[0, 1]$ into n pairwise-disjoint pieces, wherein piece A_i is allocated to agent i . The preferences of the agents i , over the cake, are represented by (integrable) densities $u_i : [0, 1] \rightarrow \mathbb{R}_+$. Overloading notation, we will use u_i s to also denote the subjective utilities of the agents and v to denote the market valuation. That is, $u_i(X)$ denotes the utility of i , for piece X , and $v(X)$ is its market value.

An allocation $A = (A_1, \dots, A_n)$ is envy-free (EF) w.r.t. utilities if $u_i(A_i) \geq u_i(A_j)$ for all agents $i, j \in \mathcal{N}$, and EF with respect to the market valuation if $v(A_i) = v(A_j)$ for all $i, j \in \mathcal{N}$. The existence of an EF cake division is rather

easy to establish, hence, we focus on two questions: (a) how many cuts are required in such an allocation, and (b) how many queries must be made in the Robertson-Webb model to find such an allocation.

For question (a), we show that $2n - 2$ cuts are necessary and, towards an upper bound, $n(n - 1)$ cuts suffice. The gap here leaves open the following tantalizing question.

Open Question 5: Does there always exist a cake division with $O(n)$ cuts that is envy-free w.r.t. the subjective utilities and the market valuation?

We believe that the flexibility of only having to achieve EF with respect to the utilities (rather than a perfect division) should permit using $o(n^2)$ cuts, perhaps even just $\Theta(n)$ cuts.

For the above-mentioned question (b), we note that, in the standard Robertson-Webb query model for cake division, the work of Aziz and Mackenzie (2016) provides a protocol that finds envy-free division—under subjective utilities—with a finite (albeit hyper-exponential) number of queries. In the full version we show that, by contrast, a protocol with finite query complexity does not exist when EF is required w.r.t. both subjective utilities and market valuation.

For cake division, we also consider Pareto Efficiency as an additional desideratum. It is known that, under agents' subjective utilities, there always exists an allocation that is EF and Pareto optimal (PO) (Weller 1985). Contrasting this result, we show that EF and PO w.r.t. the subjective utilities and EF w.r.t. the market valuation cannot be guaranteed in general (this result can be found in the full version). For the specific case of two agents this combination—EF+PO w.r.t. utilities and EF w.r.t. market valuation—remains open.

6 Conclusion and Future Work

The current work identifies a conceptually important frontier of fair division that requires fairness under market values, along with the standard desiderata under subjective utilities. This framework raises many interesting questions and admits technically rich connections with settings, such as fair division with constraints. Several of the raised questions are, in fact, fundamental even in the standard fair division setting (i.e., without an explicit invocation of the market values), e.g., the open question regarding the existence of balanced EF1 allocations under monotone utilities.

For multiple criteria—including SD-EF1, EF1, PO, MMS, and EFX—we obtain positive and negative results on simultaneous fairness under subjective utilities and the market valuation. With this groundwork, interesting directions for future work in the current framework include the allocation of chores, division of homogeneous goods, and price of fairness. The open questions explicitly identified throughout the paper are also relevant avenues for exploration.

Acknowledgments

Siddharth Barman gratefully acknowledges the support of the Walmart Center for Tech Excellence (CSR WMGT-23-0001) and a SERB Core research grant (CRG/2021/006165).

Nisarg Shah gratefully acknowledges support from an NSERC Discovery Grant.

References

- Akrami, H.; and Garg, J. 2024. Breaking the $3/4$ barrier for approximate maximin share. In *Proceedings of the 35th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 74–91.
- Aziz, H.; Freeman, R.; Shah, N.; and Vaish, R. 2024. Best of both worlds: Ex ante and ex post fairness in resource allocation. *Operations Research*, 72(4): 1674–1688.
- Aziz, H.; and Mackenzie, S. 2016. A discrete and bounded envy-free cake cutting protocol for any number of agents. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, 416–427. IEEE.
- Barman, S.; Khan, A.; Shyam, S.; and Sreenivas, K. 2023. Finding fair allocations under budget constraints. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 5481–5489.
- Biswas, A.; and Barman, S. 2018. Fair Division Under Cardinality Constraints. In *Proceedings of the 37th International Joint Conference on Artificial Intelligence (IJCAI)*, 91–97.
- Bogomolnaia, A.; and Moulin, H. 2001. A new solution to the random assignment problem. *Journal of Economic theory*, 100(2): 295–328.
- Bu, X.; Li, Z.; Liu, S.; Song, J.; and Tao, B. 2023. Fair division with allocator’s preference. In *Proceedings of the 19th Conference on Web and Internet Economics (WINE)*, 77–94.
- Budish, E. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6): 1061–1103.
- Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2019. The Unreasonable Fairness of Maximum Nash Welfare. *ACM Transactions on Economics and Computation*, 7(3): Article 12.
- Chaudhury, B. R.; Kavitha, T.; Mehlhorn, K.; and Sgouritsa, A. 2021. A little charity guarantees almost envy-freeness. *SIAM Journal on Computing*, 50(4): 1336–1358.
- Conitzer, V.; Freeman, R.; and Shah, N. 2017. Fair Public Decision Making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, 629–646.
- Dall’Aglio, M. 2023. Fair division of goods in the shadow of market values. *European Journal of Operational Research*, 307(2): 785–801.
- Dror, A.; Feldman, M.; and Segal-Halevi, E. 2023. On fair division under heterogeneous matroid constraints. *Journal of Artificial Intelligence Research*, 76: 567–611.
- Dubins, L. E.; and Spanier, E. H. 1961. How to cut a cake fairly. *American Mathematical Monthly*, 68(1): 1–17.
- Freeman, R.; Micha, E.; and Shah, N. 2021. Two-Sided Matching Meets Fair Division. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, 203–209.
- Freeman, R.; Sikdar, S.; Vaish, R.; and Xia, L. 2019. Equitable allocations of indivisible goods. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 280–286.
- Igarashi, A.; Kawase, Y.; Suksompong, W.; and Sumita, H. 2023. Fair Division with Two-Sided Preferences. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, 2756–2764.
- Kagan, J. 2021a. Assessor: Meaning, Certification, Why They Matter. <https://www.investopedia.com/terms/a/assessor.asp>.
- Kagan, J. 2021b. Equitable Distribution: Definition, State Laws, Exempt Property. <https://www.investopedia.com/terms/e/equitable-division.asp>.
- Kyropoulou, M.; Suksompong, W.; and Voudouris, A. A. 2020. Almost envy-freeness in group resource allocation. *Theoretical Computer Science*, 841: 110–123.
- Lipton, R. J.; Markakis, E.; Mossel, E.; and Saberi, A. 2004. On Approximately Fair Allocations of Indivisible Goods. In *Proceedings of the 6th ACM Conference on Economics and Computation (EC)*, 125–131.
- Patro, G. K.; Biswas, A.; Ganguly, N.; Gummadi, K. P.; and Chakraborty, A. 2020. Fairrec: Two-sided fairness for personalized recommendations in two-sided platforms. In *Proceedings of the 29th International World Wide Web Conference (TheWebConf)*, 1194–1204.
- Shtechman, I.; Gonen, R.; and Segal-Halevi, E. 2021. Fair cake-cutting algorithms with real land-value data. *Autonomous Agents and Multi-Agent Systems*, 35: 1–28.
- Steinhaus, H. 1948. The Problem of Fair Division. *Econometrica*, 16: 101–104.
- Suksompong, W. 2021. Constraints in fair division. *ACM SIGecom Exchanges*, 19(2): 46–61.
- Weller, D. 1985. Fair division of a measurable space. *Journal of Mathematical Economics*, 14(1): 5–17.