

Harm Ratio: A Novel and Versatile Fairness Criterion

Soroush Ebadian
University of Toronto
soroush@cs.toronto.edu

Rupert Freeman
University of Virginia
FreemanR@darden.virginia.edu

Nisarg Shah
University of Toronto
nisarg@cs.toronto.edu

Abstract

Envy-freeness has become the cornerstone of fair division research. In settings where each individual is allocated a disjoint share of collective resources, it is a compelling fairness axiom which demands that no individual strictly prefer the allocation of another individual to their own. Unfortunately, in many real-life collective decision-making problems, the goal is to choose a (common) public outcome that is equally applicable to all individuals, and the notion of envy becomes vacuous. Consequently, this literature has avoided studying fairness criteria that focus on individuals feeling a sense of jealousy or resentment towards other individuals (rather than towards the system), missing out on a key aspect of fairness.

In this work, we propose a novel fairness criterion, *individual harm ratio*, which is inspired by envy-freeness but applies to a broad range of collective decision-making settings. Theoretically, we identify minimal conditions under which this criterion and its groupwise extensions can be guaranteed, and study the computational complexity of related problems. Empirically, we conduct experiments with real data to show that our fairness criterion is powerful enough to differentiate between prominent decision-making algorithms for a range of tasks from voting and fair division to participatory budgeting and peer review.

1 Introduction

How to make collective decisions while treating (groups of) individuals in a fair manner is a question that human societies have struggled to understand for centuries. Today, with the advent of algorithms making increasingly critical decisions, the field of algorithmic fairness has considered a plethora of novel fairness criteria (Pessach and Shmueli, 2023), albeit they often have limited applicability due to being handcrafted for specialized decision-making tasks. To surpass this limitation, researchers have looked towards computational social choice, a field at the intersection of economics and computer science (Finocchiaro et al., 2021; Gummadi and Heidari, 2019), where fairness criteria applicable to a broad range of domains have been proposed.

Envy-freeness (Tinbergen, 1930; Gamow and Stern, 1958; Foley, 1967) is arguably the most well-studied fairness criterion in computational social choice, defined for a general resource allocation setting in which a set of goods M is to be divided between a set of agents $N = \{1, \dots, n\}$. The goal is to find an allocation $A = (A_1, \dots, A_n)$, which is a partition of M , with A_i denoting the allocation to agent i . Each agent i has a utility function $u_i : 2^M \rightarrow \mathbb{R}_{\geq 0}$ and her utility under allocation A is given by $u_i(A_i)$. Envy-freeness demands that no agent strictly prefer the allocation to another agent over her own, i.e., $u_i(A_i) \geq u_i(A_j)$ for all $i, j \in N$. Envy-freeness is the canonical embodiment of two interesting features of a fairness criterion:

1. It is a pairwise individual fairness criterion, where an agent i , who feels being treated unfairly, points to another agent j as the reason for it.
2. It makes no interpersonal comparisons, i.e., the utility function u_i of agent i is never compared to the utility function of any other agent.

Significant literature has been devoted to exploring envy-freeness and its relaxations for a variety of resource allocation domains beyond the allocation of homogeneous divisible goods (Procaccia, 2016; Amanatidis et al., 2022).

A problem arises when one moves to more general collective decision-making domains. In the most general model, which we term the *public outcomes model*, there is a set N of n agents and a set \mathcal{O} of possible outcomes. Each agent i has a utility function $u_i : \mathcal{O} \rightarrow \mathbb{R}_{\geq 0}$ and the goal is to choose an outcome $o \in \mathcal{O}$. *What would envy-freeness mean here?* Even conceptually, can an agent i really envy another agent j , despite both experiencing the common outcome o ? From this vantage point, envy-freeness may seem vacuous for public outcomes, as several authors have pointed out.

“While envy is a compelling notion in the private goods setting, it makes less sense for public decisions.” (Conitzer et al., 2017)

“Those familiar with the literature on fair division of private goods will note the conspicuous absence of the envy-freeness property...Because we are considering public goods, envy freeness is only vacuously defined: the outcome in our setting is common to all agents.” (Fain et al., 2018)

Due to this, the literature on (various special cases of) the public outcomes model has ignored envy-freeness, which misses an essential consideration where an individual feels they have been unfairly treated due to another individual receiving disproportionate importance.

Note that the public outcomes model captures the resource allocation model as a special case, by setting \mathcal{O} as the set of all possible allocations A and defining $u_i(A) \triangleq u_i(A_i)$. But importantly, it also captures a host of other collective decision-making problems such as choosing the winner of an election based on voters’ preferences, choosing a company’s strategy based on its employees’ opinions, deciding which graduate applicants to admit based on the preferences of faculty members, or picking interesting posts to showcase on a social media website’s (non-personalized) front page based on its users’ interests.

1.1 Our Contributions

Our main contribution is to propose a novel fairness criterion, *individual harm ratio* (IHR), which is inspired by envy-freeness, satisfies its two features mentioned above — a pairwise individual fairness criterion that makes no interpersonal comparisons — and applies to the full range of the public outcomes model. Outcome $o \in \mathcal{O}$ has an individual harm ratio (IHR) of α if there are no agents $i, j \in N$ and outcome $o' \in \mathcal{O}$ such that $u_i(o') > \alpha \cdot u_i(o)$ and $u_k(o') \geq u_k(o)$ for all $k \in N \setminus \{i, j\}$. In words, agent i should not be able to point to a different outcome o' where, at the expense of only agent j (keeping all other agents at least as happy), agent i ’s utility could have improved by more than an α factor. This ratio quantifies, in a sense, the harm imposed by the presence of agent j on agent i (the lower the better).

Let us briefly compare this to envy-freeness. While envy-freeness of an allocation depends only on the allocation itself, calculating the individual harm ratio requires a search over the entire outcome space; this is unavoidable because simply considering the utilities of the agent for a given outcome provides too little to go on without making interpersonal comparisons.

In [Section 3](#), we formally define IHR and its groupwise extensions, *equal-sized group harm ratio* (EGHR) and *group harm ratio* (GHR), that are analogous to similar extensions of envy-freeness from resource allocation ([Varian, 1974](#); [Berliant et al., 1992](#); [Conitzer et al., 2019](#)). Then, we analyze their relations to each other and to other known fairness criteria. We observe that in resource allocation with additive utilities, an individual harm ratio of 1 logically implies (and is significantly stronger than) envy-freeness, and even in other contexts such as matching, individual harm ratio translates to an approximation of envy-freeness. We also observe that our fairness criteria provide a fairness hierarchy for the public outcomes model similar to one from resource allocation that utilizes envy-freeness ([Figure 2](#)).

In [Section 4](#), we show that despite the generality of the public outcomes model, even a group harm ratio (GHR) of 1 (which implies an individual harm ratio of 1) can be guaranteed by finding an outcome maximizing the Nash welfare (geometric mean of agent utilities), when the set of feasible utility vectors $\mathcal{U} = \{(u_1(o), \dots, u_n(o)) : o \in \mathcal{O}\}$ satisfies two mild conditions of compactness and upper convexity. We do so by establishing an equivalence between maximum Nash welfare and another well-studied criterion, proportional fairness (which implies a group harm ratio of 1), under these conditions. As we mention in [Section 1.2](#) (and elaborate in detail in [Appendix A](#)), this generalizes several results from prior work and, in some cases, provides a more elementary proof of a maximum Nash welfare outcome satisfying various fairness desiderata (or sometimes even of the existence of an outcome satisfying those desiderata). We also prove a partial converse: any outcome with a group harm ratio of 1 achieves a $\binom{n}{\lfloor n/2 \rfloor}^{1/n}$ -approximation — at least as good as a 2-approximation — of the maximum Nash welfare in any instance of the public outcomes model and that this guarantee is tight. This complements an existing characterization of maximum Nash welfare via group fairness due to [Freeman et al. \(2020\)](#).

In [Section 5](#), we turn to computational questions. For some special cases of the public outcomes model where an outcome with a group harm ratio of 2 is guaranteed to exist, we observe that an outcome with a group harm ratio of $2 + \epsilon$ can be computed in time polynomial in the instance size and $1/\epsilon$ using convex programming. Polynomial-time computation of an outcome with a group harm ratio of 2 (or even an individual harm ratio of 2) is left as an interesting open question. We also prove that for resource allocation (specifically, allocation of indivisible goods under additive utilities), while checking whether a given allocation is EF is trivially in P, verifying an individual harm ratio of 2 (which implies EF) is coNP-complete.

In [Section 6](#), we conduct experiments with simulated and real datasets across three real-world collective decision-making domains: allocation of indivisible goods ([Spliddit](#)), peer review (CVPR 2017, CVPR 2018, ICLR 2018), and participatory budgeting ([Pabulib](#)). Our findings suggest that individual and group harm ratios are powerful enough to distinguish between prominent rules in terms of the level of fairness they achieve, and correlate well with existing domain-specific fairness notions in each domain.

1.2 Summary of Significance and Implications

We argue the significance of our work in terms of providing a unifying framework to think of fairness in a broad range of settings and opening the doors to novel styles of fairness analyses in well-studied domains. In [Appendix A](#), we provide a technical comparison of our results to those from prior works, and identify exactly which prior results our work subsumes. Here, we briefly summarize four significant implications of our work.

1. Our work subsumes existing results from the resource allocation literature proving that MNW allocations are envy-free (EF) and Pareto optimal (PO), e.g., in cake-cutting ([Weller, 1985](#);

Segal-Halevi and Sziklai, 2019) (with our proof being more elementary than the existing ones), or that they are approximately EF and PO, e.g., in one-sided matching (Troebst and Vazirani, 2024). This is a result of 1-IHR implying exact or approximate EF in these domains.

2. Our work adds to a long line of literature identifying conditions guaranteeing the existence of EF+PO allocations (Varian, 1974; Svensson, 1983a; Diamantaras, 1992) and outcomes satisfying another fairness criterion, proportional fairness (Ray Chaudhury et al., 2022).
3. In some domains where our results hold, the fact that MNW allocations satisfy 1-IHR and PO is novel and interesting (Dall’Aglio, 2001; Cole and Tao, 2021; Caragiannis et al., 2022); in one case, we are able to establish (a previously unknown) existence of EF+PO allocations (Dall’Aglio, 2001).
4. Finally, our work significantly expands the possibility of conducting pairwise individual fairness analyses to a much larger set of collective decision-making domains than the resource allocation domains where envy-freeness is well-defined.

2 Model

For $t \in \mathbb{N}$, define $[t] \triangleq \{1, 2, \dots, t\}$. We consider a very general multi-agent decision-making setting, which we term the *public outcomes* model, where each instance is given by a set of outcomes \mathcal{O} , a finite set of agents $N = [n]$, and a utility profile $\mathbf{u} = (u_1, \dots, u_n)$ containing a utility function $u_i : \mathcal{O} \rightarrow \mathbb{R}_{\geq 0}$ for each agent $i \in N$. Without loss of generality, we assume that for each agent $i \in N$ there is some outcome $o \in \mathcal{O}$ such that $u_i(o) > 0$; otherwise, agent i can effectively be removed from consideration.

A priori, we impose no assumption on the set of outcomes \mathcal{O} (it could be finite, countably infinite, or uncountable) or the utility functions of the agents. As such, this general model subsumes various models studied in the literature as special cases, including cake-cutting (Steinhaus, 1948), allocation of homogeneous divisible or indivisible goods (Moulin, 2004), public decision-making (Conitzer et al., 2017), and allocation of public goods (Fain et al., 2018).

Utility set. Define the *utility set* $\mathcal{U} = \{(u_1(o), \dots, u_n(o)) : o \in \mathcal{O}\}$ to be the set of feasible utility vectors. Because we have placed no assumptions on the agent utilities, this is essentially the only object of interest as, given any \mathcal{U} , one can easily construct underlying outcome set \mathcal{O} and utility profile \mathbf{u} which induce \mathcal{U} . Crucially, the fairness notions we define below also depend only on \mathcal{U} .

2.1 Fairness

While many fairness notions defined in the literature for specialized models do not extend to this general model, the following ones do, and play a key role in our work. First, an outcome is proportionally fair if the agents are not happier (in an average multiplicative sense) when switching to any other outcome.

Definition 1 (Proportional Fairness (PF)). We say that an outcome $o \in \mathcal{O}$ is proportionally fair (PF) if $\frac{1}{n} \sum_{i \in N} \frac{u_i(o')}{u_i(o)} \leq 1$ for all $o' \in \mathcal{O}$.¹

A maximum Nash welfare outcome is one that maximizes the geometric mean of agent utilities.

Definition 2 (Maximum Nash Welfare (MNW)). Define the Nash welfare of an outcome $o \in \mathcal{O}$ as the geometric mean of agent utilities for it: $\text{NW}(o) = (\prod_{i \in N} u_i(o))^{1/n}$.² For $\alpha \geq 1$, we say that

¹If $u_i(o) = 0$, we let $u_i(o')/u_i(o) = \infty$ and PF is automatically violated.

²If $u_i(o) = 0$ for any agent $i \in N$, we say that $\text{NW}(o) = 0$.

an outcome $o \in \mathcal{O}$ is an α -approximate maximum Nash welfare (MNW) outcome if $\text{NW}(o) \geq (1/\alpha) \cdot \text{NW}(o')$ for all $o' \in \mathcal{O}$. When $\alpha = 1$, we simply refer to it as an MNW outcome.

An outcome is in the core (Foley, 1967) if no group can find an alternative outcome that provides all of them weakly higher utility, and at least one of them strictly higher utility, even after scaling their utilities by the size of the group in proportion to the set of all agents.

Definition 3 (Core). We say that an outcome $o \in \mathcal{O}$ is in the core if there is no group of agents $S \subseteq N$ and outcome $o' \in \mathcal{O}$ such that $\frac{|S|}{n} \cdot u_i(o') \geq u_i(o)$ for all $i \in S$ and at least one inequality is strict.

Proportionality (Steinhaus, 1948) says that each each should receive at least a $1/n$ fraction of the utility she can achieve in any outcome.

Definition 4 (Proportionality (Prop)). We say that an outcome $o \in \mathcal{O}$ is proportional (Prop) if $u_i(o) \geq \frac{1}{n} u_i(o')$ for all $o' \in \mathcal{O}$.

Finally, the following is traditionally considered a notion of efficiency, but we include it here because it connects well to the notions defined above. An outcome is Pareto optimal if there is no alternative outcome that makes each agent at least as happy and some agent strictly happier.

Definition 5 (Pareto Optimality (PO)). We say that an outcome $o \in \mathcal{O}$ is Pareto optimal (PO) if there is no outcome $o' \in \mathcal{O}$ such that $u_i(o') \geq u_i(o)$ for all $i \in N$ and at least one inequality is strict.

Note that these notions are *utilitarian*, i.e., they depend only on the utility vector induced by the outcome. Hence, one can equivalently speak of a utility vector satisfying these notions; if a utility vector satisfies a utilitarian notion, every outcome inducing it must satisfy the notion as well.

The following proposition describes well-known relationships between these notions. Its proof is in Appendix B.1.

Proposition 1 (PF \Rightarrow MNW, PF \Rightarrow Core \Rightarrow (Prop+PO)). *Any proportionally fair outcome is also a maximum Nash welfare outcome and lies in the core. Any outcome in the core satisfies proportionality and Pareto optimality.*

Remark 1. In many domains, every agent can possibly get zero utility under some outcome (i.e., $\forall i, \exists o \in \mathcal{O} : u_i(o) = 0$). However, when this is not the case, one can define stronger versions of the core and proportionality that take $u_i^{\min} \triangleq \inf_{o \in \mathcal{O}} u_i(o)$ into account.

In the definition of the core, we would write $\frac{|S|}{n} (u_i(o') - u_i^{\min}) \geq (u_i(o) - u_i^{\min})$, and for proportionality, we would write $u_i(o) \geq \frac{1}{n} u_i(o') + \frac{n-1}{n} u_i^{\min} \Leftrightarrow u_i(o) - u_i^{\min} \geq \frac{1}{n} (u_i(o') - u_i^{\min})$. One can check that this strengthens the respective definition. Aziz et al. (2023a) refer to the strengthened proportionality as general fair share (GFS).

Any method of achieving the regular versions can be modified to achieve the strengthened versions by feeding it translated utility functions $\{u'_i\}_{i \in N}$, where $u'_i(o) = u_i(o) - u_i^{\min}$ for all $i \in N$ and $o \in \mathcal{O}$. For maximum Nash welfare, we would then seek an outcome o that maximizes $\prod_{i \in N} (u_i(o) - u_i^{\min})$; this is akin to the Nash bargaining solution (Nash, 1950), if $u_1^{\min}, \dots, u_n^{\min}$ are treated as status quo utilities. Similarly, for proportional fairness, we would seek an outcome o for which $\frac{1}{n} \sum_{i \in N} \frac{u_i(o') - u_i^{\min}}{u_i(o) - u_i^{\min}} \leq 1$.

2.2 Private Goods Division

As demonstrated in the introduction, the public outcomes model subsumes a wide range of collective decision-making models studied in the literature. Below, we formally introduce the special case of private goods division, which we will refer to frequently in our results.

In a general *private goods division* setting, there is a set of goods M to be divided between a set of agents $N = [n]$. The outcome set \mathcal{O} is the set of allocations $A = (A_1, \dots, A_n)$ which are partitions of M into pairwise-disjoint measurable bundles. Each agent $i \in N$ has a measure u_i over M .³ Crucially, her utility for an allocation A is solely a function of the bundle A_i assigned to her, namely $u_i(A_i)$.

This model captures three prominent settings studied in the literature:

1. *Cake-cutting*: $M = [0, 1]$ and the utility function u_i of each agent i is a (countably additive) measure over M that is absolutely continuous with respect to the Lebesgue measure.
2. *Homogeneous divisible goods*: M is a set of m homogeneous divisible goods. An allocation can be described as $A = (A_{i,g})_{i \in N, g \in M}$ with $A_{i,g}$ being the fraction of good g allocated to agent i and $\sum_{i \in N} A_{i,g} = 1$ for all g . The utility of each agent i can be given by $u_i(A_i) = \sum_{g \in M} A_{i,g} \cdot v_{i,g}$, where $v_{i,g}$ is her value for receiving good g entirely.
3. *Indivisible goods*: This is identical to homogeneous divisible goods, except we further restrict $A_{i,g} \in \{0, 1\}$ for all $i \in N$ and $g \in M$.

For private goods division, additional fairness notions have been extensively studied, but they do not extend to the more general model of public outcomes; coming up with natural extensions of these notions is precisely the subject of our work. Envy-freeness (Foley, 1967) says that no agent should strictly prefer the allocation of another agent to her own.

Definition 6 (Envy-Freeness (EF)). We say that an allocation A is envy-free (EF) if $u_i(A_i) \geq u_i(A_j)$ for all $i, j \in N$.

Group fairness (Conitzer et al., 2019) requires that no group of agents should “envy” (in the sense of being able to Pareto improve by taking the resources allocated to the other group and redistributing amongst themselves, subject to an appropriate scaling factor to account for different-sized groups) any other group of agents.

Definition 7 (Group Fairness (GF)). We say that an allocation A is group fair (GF) if there is no pair of groups of agents $S, T \subseteq N$ and a division B of $\cup_{j \in T} A_j$ among agents in S such that $\frac{|S|}{|T|} \cdot u_i(B_i) \geq u_i(A_i)$ for all $i \in S$ and at least one inequality is strict.

For private goods division, envy-freeness implies proportionality, and group fairness implies the core (and therefore also proportionality). Additionally, restricting the group fairness definition to allow only pairs of groups with $|S| = |T|$ yields the group envy-freeness definition of Berliant et al. (1992).

Note that the difficulty in extending these notions is that in the public outcomes model, there is nothing “allocated” to individual agents. Instead, a common outcome is selected for all agents, making the concept of envy between individuals or groups vacuous.

³A measure assigns a non-negative value to each measurable subset of M (0 to the empty subset) and is countably additive.

3 Harm Ratio: A Novel Fairness Criterion

To motivate the definition of individual harm ratio, let us first revisit envy-freeness (EF) in private goods division with additive utilities (**Definition 6**): an allocation A is EF if $u_i(A_i) \geq u_i(A_j)$ for all $i, j \in N$. This is not a utilitarian notion: envy-freeness cannot be checked simply from the induced utility vector $(u_1(A), \dots, u_n(A))$, which makes it difficult to extend to the public outcomes model. However, if A violates EF, there exists a new allocation A' , given by $A'_i = A_i \cup A_j$, $A'_j = \emptyset$, and $A'_k = A_k$ for all $k \in N \setminus \{i, j\}$, which induces a utility vector $(u_1(A'), \dots, u_n(A'))$ satisfying the following: $u_i(A') > 2 \cdot u_i(A)$ and $u_k(A') \geq u_k(A)$ for all $k \in N \setminus \{i, j\}$. This is a purely utilitarian comparison, which can be generalized well beyond private goods division to the public outcomes model. As argued in the introduction, the utility improvement factor of 2 in this case measures the level of harm imposed on agent i due to the presence of agent j . We refer to it as the individual harm ratio (IHR).

The factor of 2 is important. Not only does it show up in connection to private goods division with additive utilities, we prove in **Theorem 2** that it remains the tightest achievable factor for the more general public outcomes model under certain conditions on the utility set \mathcal{U} . It also generalizes well to a factor of $\frac{|S \cup T|}{|S|}$ in the groupwise extension of IHR introduced later (**Definition 9**). However, a worse factor may be achieved when the utilities do not satisfy the requirements of **Theorem 2**, or a better factor may be achieved in practice on real-world instances. To account for these possibilities, we introduce an α -approximation of IHR.

Definition 8 (Individual Harm Ratio). We say that an outcome $o \in \mathcal{O}$ has an *individual harm ratio* of α , denoted α -IHR, if there are no agents $i, j \in N$ and outcome $o' \in \mathcal{O}$ such that $\frac{1}{2} \cdot u_i(o') > \alpha \cdot u_i(o)$ and $u_k(o') \geq u_k(o)$ for all $k \in N \setminus \{i, j\}$.

Definition 8 requires that no agent be able to find an outcome in which even half of her utility is higher (by a factor of α) and at most one other agent is hurt. Intuitively, a large individual harm ratio gives agent i a “justified claim” that the system is treating her unfairly due to its attempt to make agent j happy. The higher the α , the stronger the claim. Conversely, the lower the value of α for an outcome, the stronger its fairness guarantee.

3.1 Comparison to Envy-Freeness and Proportionality

Consider private goods division with additive utilities. From the argument above, it is clear that 1-IHR implies envy-freeness, but in fact, it is strictly stronger: even if moving the goods allocated to agent j to agent i does not suffice to more than double agent i 's utility, a complete reshuffling that also alters the allocations to the other agents may nonetheless be able to achieve this, while still keeping the other agents at least as happy. The next result shows that envy-freeness is much weaker than 1-IHR, as it implies only $(n/2)$ -IHR. All the proofs are in the appendix.

Theorem 1. *For private goods division with additive utilities:*

- (1-IHR \Rightarrow EF) *An individual harm ratio of 1 implies envy-freeness.*
- (Prop \Rightarrow $(n/2)$ -IHR) *Proportionality (and therefore envy-freeness) implies an individual harm ratio of $n/2$, and this is tight even for the allocation of homogeneous divisible goods.*

What happens once we go beyond private goods division? As noted earlier, envy-freeness is no longer well-defined for the public outcomes model, but proportionality still is. For private goods division with additive utilities, we noticed above that 1-IHR implies proportionality. This breaks down when we move to the public outcomes model, as the following simple example shows.

	g_1	g_2	g_3	\dots	g_{n-1}	g_n
u_1	1	0	0	\dots	0	0
u_2	1	1	0	\dots	0	0
u_3	1	1	1	\dots	0	0
\vdots				\ddots		\vdots
u_{n-1}	1	1	1	\dots	1	0
u_n	1	1	1	\dots	1	1

Figure 1: Instance for [Theorem 1](#). Blue and red rectangles show two different allocations; the rectangle in each good’s column covers the agents among which the good is equally divided. The blue allocation is EF, but $(n/2)$ -IHR, as witnessed by the existence of the red allocation.

Example 1. Consider an instance of the public outcomes model with a set of two outcomes $\mathcal{O} = \{o, o'\}$ and a set of three agents $N = [3]$. The utilities are $(u_1(o), u_2(o), u_3(o)) = (1, 1, 0)$ and $(u_1(o'), u_2(o'), u_3(o')) = (0.1, 0.1, 1)$. Note that outcome o is 1-IHR (in fact, it’s even better — $\frac{1}{2}$ -IHR): agents 1 and 2 are maximally happy, and while agent 3 receives zero utility, she cannot choose any other outcome where only one agent is hurt and thus cannot improve at all in any outcome that hurts at most one other agent. Note that o gives a zero approximation of proportionality for agent 3.

3.2 Groupwise Extension

[Example 1](#) points out a weakness of IHR in the public outcomes model. Finding an outcome that hurts at most one other agent can be quite limiting if there are no such alternative outcomes. What if we allow an agent (or a group of agents) to find outcomes that hurt multiple other agents? This must be accompanied with a greater requirement of utility improvement. It turns out we can strengthen IHR to group harm ratio (GHR) in the same way that envy-freeness is strengthened to group fairness (GF) in private goods division ([Conitzer et al., 2019](#)).

Definition 9 (Group Harm Ratio). We say that outcome $o \in \mathcal{O}$ has a *group harm ratio* of α , denoted α -GHR, if there are no non-empty groups of agents $S, T \subseteq N$ and outcome $o' \in \mathcal{O}$ such that the following two conditions hold:

1. For all agents $i \in S$, $\frac{|S|}{|S \cup T|} \cdot u_i(o') \geq \alpha \cdot u_i(o)$ with at least one inequality being strict.
2. For all agents $i \in N \setminus (S \cup T)$, $u_i(o') \geq u_i(o)$.

We remark that one can also define equal-sized group harm ratio (EGHR) by only imposing the GHR constraints when $|S| = |T|$, the same way GEF is defined for private goods division. The proof of the next result is in [Appendix C.2](#).

Proposition 2. *In the public outcomes model, $PF \Rightarrow 1\text{-GHR} \Rightarrow (PO + Core + 1\text{-IHR})$.*

Together with [Proposition 1](#) and [Theorem 1](#), [Proposition 2](#) paints a clear picture of the relations among the various fairness notions in the public outcomes model, depicted in [Figure 2\(b\)](#), mimicking a similar hierarchy for private goods division depicted in [Figure 2\(a\)](#) (except that 1-IHR no longer implies proportionality in the public outcomes model).

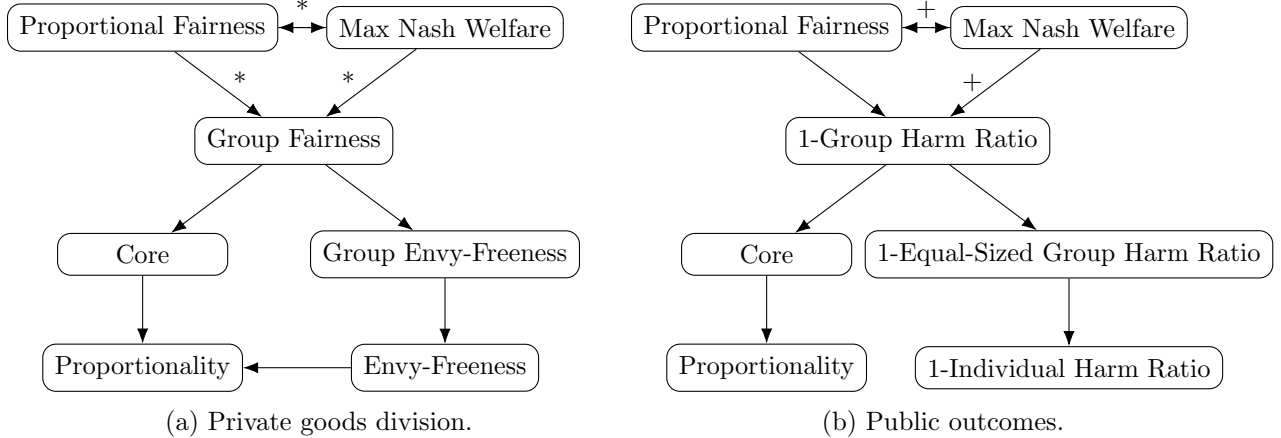


Figure 2: The figure depicts a hierarchy of fairness notions for private goods division (already known) and the public outcomes model (based on our novel fairness definitions and results). For private goods division, the implications marked with (*) hold for cake-cutting (Conitzer et al., 2019; Freeman et al., 2020). For the public outcomes model, we show the implications marked with (+) when the utility set \mathcal{U} is compact and upper convex (Theorem 2).

4 Maximum Nash Welfare Solution

Without any conditions on \mathcal{U} , an outcome satisfying any of the notions defined in Section 2.1 is not guaranteed to exist. Proportionality cannot be guaranteed in the standard example of allocating an indivisible item between two agents who value it at 1 ($\mathcal{U} = \{(1, 0), (0, 1)\}$). An MNW outcome is not guaranteed to exist if $x \in [0, 1)$ fraction of a single homogeneous divisible good must be allocated to a single agent with utility x for receiving x fraction of the good ($\mathcal{U} = \{x : x \in [0, 1)\}$); in particular, since the agent is not allowed to receive the whole good, a Nash welfare arbitrarily close to 1 can be attained but a Nash welfare of 1 cannot be attained. Nonexistence of the other properties follows from the implications in Proposition 1.

In this section, we show that even proportional fairness, which implies a group harm ratio of 1, can be guaranteed when the utility set \mathcal{U} satisfies two simple conditions:

- **Compactness:** \mathcal{U} is compact.
- **Upper convexity:** For any $\mathbf{x}, \mathbf{y} \in \mathcal{U}$ and $\alpha \in [0, 1]$, there exists $\mathbf{z} \in \mathcal{U}$ such that $z_i \geq \alpha \cdot x_i + (1 - \alpha) \cdot y_i$ for all $i \in N$.

Compactness and convexity (where even $\mathbf{z} = \alpha \cdot \mathbf{x} + (1 - \alpha) \cdot \mathbf{y} \in \mathcal{U}$ is guaranteed) of the utility set is a common feature of many decision-making models. Compactness of \mathcal{U} holds whenever the outcome set \mathcal{O} is compact and the utility functions u_i are continuous. Convexity of \mathcal{U} holds whenever the outcome set \mathcal{O} is the set of probability distributions (over some set) and one uses expected utilities (Young, 1995). In resource allocation contexts, it also holds for the richer class of concave utilities (Bertsimas et al., 2011, Proposition 1). The celebrated result of Dubins and Spanier (1961) establishes compactness and convexity of the utility set for cake-cutting.

Upper convexity is a slightly relaxed version of convexity, which allows convexity to fail as long as there is a utility vector that Pareto dominates the one obtained from a convex combination. Intuitively, this relaxed requirement should not cause any issues because one should always be able to switch to such a Pareto dominating outcome, although a priori it is not clear if such an outcome

would retain any fairness guarantees.⁴ Upper convexity is a useful relaxation.

Example 2. Consider allocating a single homogeneous divisible good between two agents. The outcome set is the simplex $\mathcal{O} = \{(x_1, x_2) : x_1, x_2 \geq 0, x_1 + x_2 = 1\}$. Suppose the agents have utility functions that are linearly increasing in the fraction of resource that they receive, but plateau when at least half of the resource is allocated to them. That is, $u_1((x_1, x_2)) = \min(0.5, x_1)$ and $u_2((x_1, x_2)) = \min(0.5, x_2)$. We have $(0.5, 0) \in \mathcal{U}$ and $(0, 0.5) \in \mathcal{U}$, but $(0.25, 0.25) \notin \mathcal{U}$. Upper convexity is still satisfied due to the fact that $(0.5, 0.5) \in \mathcal{U}$, which Pareto dominates $(0.25, 0.25)$.

Finally, it is worth noting that when \mathcal{U} is compact, an MNW outcome is guaranteed to exist because the product is a continuous function, which attains a maximum over a compact set via Weierstrass' extreme value theorem. Similarly, when \mathcal{U} is compact and upper convex, a proportional outcome is guaranteed to exist. For each agent i , let $\mathbf{x}^i \in \arg \max_{\mathbf{x} \in \mathcal{U}} x_i$ be a utility vector with the highest utility for agent i (compactness ensures that the maximum is attained in \mathcal{U}). Upper convexity guarantees the existence of a utility vector that (weakly) Pareto dominates $(1/n) \sum_{i \in N} \mathbf{x}^i$, and this is proportional by definition. Note that because this outcome has a strictly positive Nash welfare, so must any MNW outcome (so it must give a positive utility to each agent).

Proposition 3. *When \mathcal{U} is compact, an MNW outcome exists. When \mathcal{U} is compact and upper convex, a proportional outcome exists and every MNW outcome has a strictly positive Nash welfare.*

However, the existence of a proportionally fair (PF) outcome or even an outcome in the core is not trivial to establish even under compactness and upper convexity of \mathcal{U} . The next result shows that under these conditions, every MNW outcome (which exists by [Proposition 3](#)) is PF, making MNW and PF equivalent due to [Proposition 1](#) and yielding the existence of PF. Combining this with [Proposition 2](#), we get a 1-GHR (and thus a 1-IHR) implication as well.

Theorem 2. *Consider any instance of the public outcomes model where the utility set \mathcal{U} is compact and upper convex. A maximum Nash welfare (MNW) outcome exists, and every MNW outcome is proportionally fair (PF). Consequently, MNW and PF are equivalent, and every MNW outcome has a group harm ratio (and thus an individual harm ratio) of 1.*

For private goods division with additive utilities, 1-GHR implies PO and 1-IHR (which in turn implies EF). Hence, [Theorem 2](#) yields an alternative proof of the fact that every MNW allocation is EF and PO in cake-cutting, where the utility set is known to be compact and convex ([Dubins and Spanier, 1961](#)). We remark that our elementary proof is much simpler than the proof by [Segal-Halevi and Sziklai \(2019\)](#), who proved it by first establishing that an MNW allocation is characterized by a form of market equilibrium (s-CEEI), and the proof by [Weller \(1985\)](#), who showed only the existence of an EF+PO cake allocation using Kakutani's fixed point theorem.

The equivalence between PF and MNW is a well-known fact in many domains, going back to the work of [Kelly \(1997\)](#), who proved it in the context of rate control in communication networks. The proof provided above is very similar, except we strive to not make any assumptions on the nature of the decision-making (the outcome set \mathcal{O}) or the utility functions of the agents. The key to our proof is defining the function f in terms of the scalar α , which makes it differentiable without requiring any perturbations of utilities.

Once MNW is established as equivalent to PF, one may wonder how much stronger MNW is compared to the next property in the fairness hierarchy, namely 1-GHR. For private goods division, the equivalent question would be how strong MNW is compared to group fairness (GF) ([Conitzer](#)

⁴For example, in private goods division, it is well known that envy-freeness might not be preserved under Pareto improvements.

et al., 2019). Freeman et al. (2020) study allocation rules that map each instance to a set of (tied) allocations, and an allocation rule is said to satisfy GF if, on each instance, every allocation in the set returned by the rule satisfies GF. They prove that the rule returning the set of MNW allocations satisfies GF even for cake-cutting, and, subject to an axiom called replication invariance (Debreu and Scarf, 1963; Varian, 1974), it is the only rule satisfying GF.

Replication invariance informally requires that if one replicates an instance by creating k copies of each agent and good, then a replication of every allocation returned by the rule on the original instance must be returned by the rule on the new instance. It is difficult to extend to the public outcomes model.

Nonetheless, the next result shows that a group harm ratio of 1 implies a (tight) 2-approximation of maximum Nash welfare in the public outcomes model. This result nicely complements that of Freeman et al. (2020) and has two key advantages: first, it eliminates the need to impose any additional axioms, which significantly changes the technical arguments required, and second, the proof works on an instance-by-instance basis (rather than at the level of rules) in the very general public outcomes model.

Theorem 3. *Fix any instance of the public outcomes model. For every outcome $o \in \mathcal{O}$ with a group harm ratio of 1 (i.e., 1-GHR), we have that $\text{NW}(o) \geq \binom{n}{\lfloor n/2 \rfloor}^{-1/n} \cdot \text{NW}(o')$ for all $o' \in \mathcal{O}$, and this is tight. Finally, $\binom{n}{\lfloor n/2 \rfloor}^{1/n} \leq 2$ and approaches to 2 as $n \rightarrow \infty$.*

When \mathcal{U} is compact and upper convex, an MNW outcome exists (Proposition 3), and Theorem 3 shows that any 1-GHR outcome achieves a $\binom{n}{\lfloor n/2 \rfloor}^{1/n}$ -approximation of its Nash welfare.

5 Computation

Having established that MNW (equivalently, PF) outcomes are 1-GHR, and thus 1-IHR, under quite general conditions (Theorem 2), we now turn our attention to the question of computation. We are interested in two computational questions: checking whether a given outcome is 1-IHR and computing an outcome with a low individual harm ratio (ideally, 1).

For the special case of homogeneous divisible goods with additive utilities, an MNW/PF solution can be found in strongly polynomial time (Orlin, 2010). However, in general, exact polynomial time computation is not possible. We instead turn to approximate solutions.

Definition 10 (α -Proportional Fairness). An outcome $o \in \mathcal{C}$ is α -Proportionally Fair (α -PF) if $\frac{1}{n} \sum_{i \in N} \frac{u_i(o')}{u_i(o)} \leq \alpha$ for all $o' \in \mathcal{C}$.

When the utility space \mathcal{U} is convex and can be described by a polynomial number of linear constraints, as is the case for common settings such as public decision making (Conitzer et al., 2017), and budget-feasible participatory budgeting (Fain et al., 2016) (in all cases, assuming fractional solutions are allowed), a $(1 + \epsilon)$ -PF outcome can be computed in time polynomial in the input size and $1/\epsilon$ using standard convex programming techniques.⁵ However, approximate PF is interesting only insofar as it guarantees some approximation of other, normative, properties. Our most general of these is GHR.

By slightly adapting the proof of Proposition 2, it is easy to see that approximate PF implies approximate GHR. The proof of the next observation is in Appendix E.1.

⁵To see this, note that the function $\text{PF}(\mathbf{x}, \mathbf{y}) = \frac{1}{n} \sum_{i \in N} \frac{y_i}{x_i}$ is convex in \mathbf{x} . Therefore, since the supremum of convex functions is itself convex, $\text{PF}(\mathbf{x}) = \max_{\mathbf{y} \in \mathcal{U}} \frac{1}{n} \sum_{i \in N} \frac{y_i}{x_i}$ is convex in \mathbf{x} . Thus, $\min_{\mathbf{x} \in \mathcal{U}} \text{PF}(\mathbf{x})$ is a convex optimization problem with an optimal objective value of 1. It follows that a solution with objective value of $1 + \epsilon$ can be computed in time polynomial in the input size and $1/\epsilon$.

Proposition 4. *In the public outcomes model, $(1 + \epsilon)$ -PF \Rightarrow $(1 + n\epsilon)$ -GHR.*

Of course, our results have not ruled out the possibility of computing an exactly 1-IHR, or even a 1-GHR, outcome in polynomial time; perhaps this can be done more directly without having to appeal to (approximate) proportional fairness. We leave these as open questions.

With respect to checking whether a given outcome $o \in O$ is 1-IHR (i.e., has a individual harm ratio of 1), whenever the feasible utility space \mathcal{U} can be described by a polynomial number of linear constraints, the problem can be solved in polynomial time. To do so, we solve a linear feasibility program for every pair of agents i, j to check if there exists $\mathbf{x} \in \mathcal{U}$ for which $x_i > 2 \cdot u_i(o)$ and $x_k \geq u_k(o)$ for all $k \in N \setminus \{i, j\}$, which is a single linear program.

Indivisible goods. An interesting special case is that of indivisible goods allocation with additive utilities. Here, it is known that a 1-IHR outcome cannot be guaranteed. For the case of two agents, we know that EF and 1-IHR are equivalent ([Theorem 1](#)), and it is known via a standard reduction from the Partition problem that determining the existence of an EF outcome is NP-complete even with two agents.

Corollary 1. *For indivisible goods allocation with additive utilities, checking if a 1-IHR outcome exists is NP-complete, even for two agents.*

Lastly, we show that it is hard to check whether a given allocation of indivisible goods is 1-IHR. This is in sharp contrast with checking envy-freeness, which can be done easily by comparing the utility of every agent with their value for every other agent’s allocation. This result has an intricate proof, which is presented in [Appendix E.2](#).

Theorem 4. *For indivisible goods allocation with additive utilities, checking if a given outcome is 1-IHR is coNP-complete.*

Our proof uses a variable number of agents. For two agents, 1-IHR is equivalent to EF, so checking whether a given outcome is 1-IHR is in P. Could it perhaps be in P for any constant number of agents? We leave this as an open question.

6 Experiments

In this section, we validate our theoretical definitions by running numerical simulations on data from three real-world settings: private goods allocation, peer review, and participatory budgeting. Consistent with real-world constraints in these settings, we require that solutions be discrete. In this case, the utility sets are not upper convex and it is therefore no longer the case that an outcome with an IHR/GHR of at most 1—or even a finite IHR/GHR—always exists. Our approach will be to evaluate the individual and group harm ratios of prominent rules for these settings.

It is worth taking a moment to reflect on what one might like to see from these experiments. As a minimal requirement, our definition should differentiate between rules across a wide range of application settings, since the practical consequence of any fairness definition is limited to those cases where it can make a specific recommendation. Second, in settings where well-established and validated fairness notions already exist, we would like IHR to align well with those existing notions.⁶

6.1 Private Goods

Our first setting is that of indivisible private goods allocation. Every agent i has a value $v_{i,g}$ for every good g , with valuations normalized so that $\sum_g v_{i,g}$ is independent of i .

⁶Of course, the strength of IHR comes from the fact that it is broadly applicable to all these settings simultaneously.

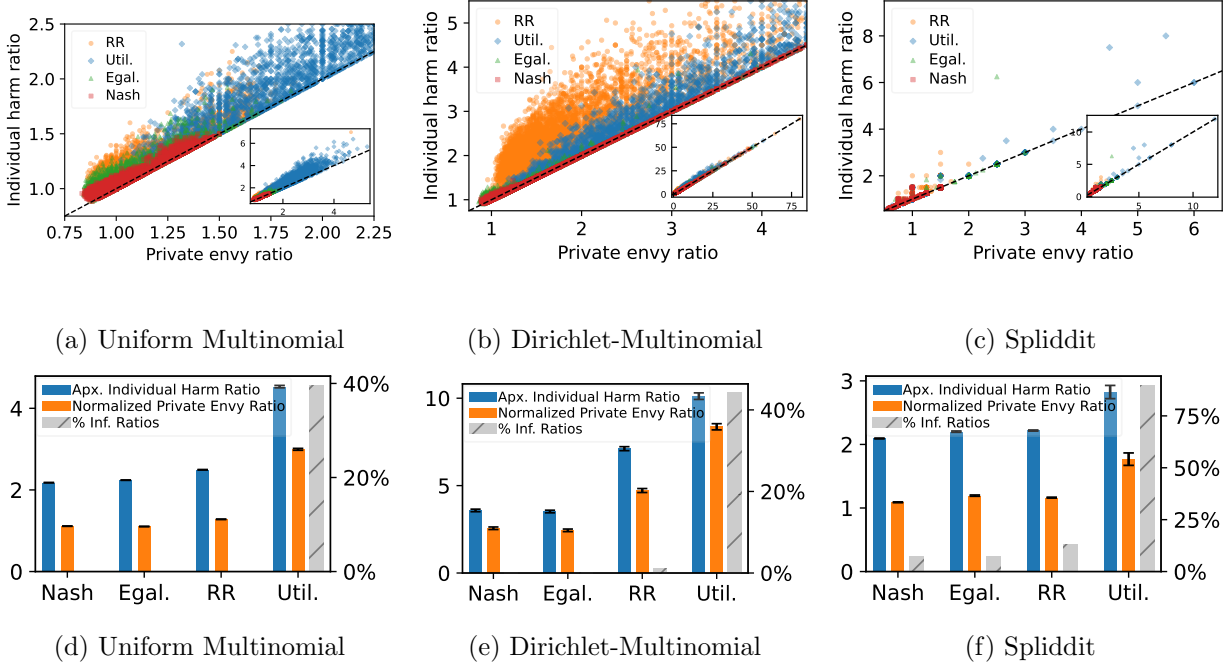


Figure 3: Scatter plots showing instance-wise individual harm ratio and private envy ratio, and bar charts displaying their averages as well as the fraction of instances with infinite ratios. In the scatter plots, the inset plot shows the full set of instances, while the main plot is zoomed into its lower left region that contains approximately 90% of the instances.

Datasets. We use two synthetic datasets and one real dataset. For the synthetic datasets, we first consider agent valuations drawn from a uniform multinomial distribution (i.e., randomly dividing 200 points between m goods) and, second, valuations drawn from a Dirichlet-multinomial distribution. The Dirichlet-multinomial can be thought of as first drawing an underlying “market value” for every good from a uniform Dirichlet distribution, and then drawing agent valuations from a multinomial distribution whose expectations equal to the market values, thus correlating agent valuations. For the synthetic experiments, we vary $n \in [3, 16]$ and $m \in [n, 5 \cdot n]$ and generate 10 samples for each combination of n and m . Finally, we also use real data from the popular fair division website [Spliddit.org](https://spliddit.org) (Goldman and Procaccia, 2014).

Rules. For every fair division instance, we compute the outcomes that maximize the egalitarian welfare (minimum utility), the Nash welfare (product of utilities), and the utilitarian welfare (sum of utilities) — we refer to these rules by Egal., Nash, Util. respectively. We additionally consider round robin (RR), which fixes an ordering of the agents and allows each to pick one good at a time in a cyclic fashion; this, like Nash, is guaranteed to achieve an approximate envy-freeness guarantee (Caragiannis et al., 2019).

Measurements. For each dataset, we display two graphs (Figure 3). The first (Figures 3(a) to 3(c)) is a scatter plot with each point in the plot being an outcome, color coded by the rule that produced it. The vertical axis measures the individual harm ratio (IHR), and the horizontal axis measures the private envy ratio, defined as $\text{per} \triangleq \frac{1}{2}(1 + \max_{i,j \in [n]} u_i(A_j)/u_i(A_i))$, which is the maximum multiplicative increase in utility that any agent could get by swapping their assigned bundle of goods with some other agent, linearly scaled so that an IHR of α guarantees $\text{per} \leq \alpha$. Hence, all points are guaranteed to lie on or above the 45 degree line $y = x$. Note that EF

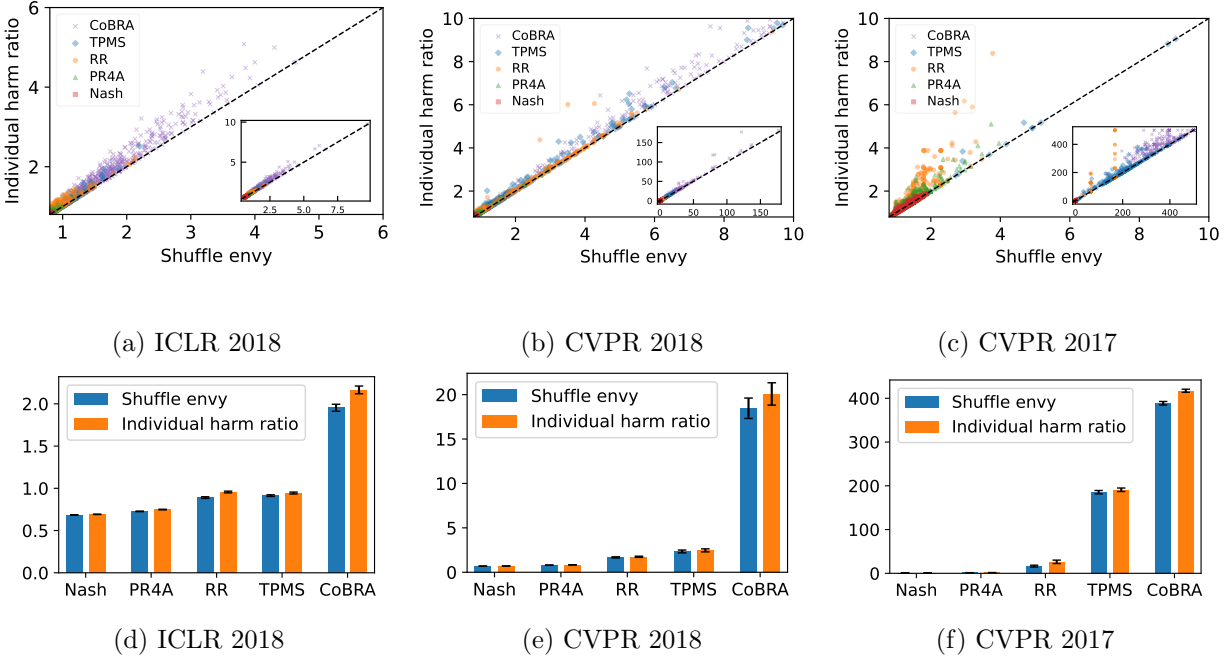


Figure 4: Scatter plots showing instance-wise public and shuffle envy ratios and bar charts displaying average public and shuffle envy ratios. In the scatter plots, the main plot is zoomed into the lower left region of the graph that contains more than 90% of all data points – except for CVPR’17 where CoBRA and TPMS lie far from the others – while the inset plot shows the full set of instances.

guarantees a private envy ratio of at most 1.

The second graph for each dataset is a bar graph that shows, for each rule, the average IHR, the average private envy ratio, and the fraction of instances with infinite IHR and private envy ratio. Infinite ratios occur whenever some agent receives zero utility from their allocation, and the average ratios are taken over instances with finite ratios.

Observations. In all datasets, we observe that individual harm ratio and private envy ratio are strongly correlated, with the points in the scatter plot lying somewhat close to the 45 degree line. The Nash and egalitarian rules give the lowest public and private envy ratios. Round robin is competitive on the uniform Dirichlet and Spliddit datasets, but less so for the Dirichlet-multinomial model. Meanwhile, the utilitarian rule has high individual harm ratio and private envy ratio, and a large fraction of instances with infinite ratios.

6.2 Peer Review

Datasets. In the peer review setting, we use real data from three conferences: ICLR 2018, CVPR 2017, and CVPR 2018. Each submission is an “agent” who has a value for being assigned a reviewer equal to a (system-generated) similarity score reflecting the reviewer’s expertise in the paper’s topic. Following the methodology of [Aziz et al. \(2023b\)](#), we assume that each submission has exactly one author, and that the set of reviewers coincides with the set of authors. The matching of reviewers to submissions is constrained in that no author can review their own submission, each submission must be assigned three distinct reviewers, and each reviewer must be assigned three submissions to review.

We also follow the methodology of Aziz et al. (2023b) to infer authorship information. The ICLR 2018 dataset provides a conflict matrix, indicating when reviewer j has a conflict with submission i . We find a maximum matching on the conflict matrix, and use this as the authorship matrix. We then subsample 300 instances of size 50 (that is, 50 submissions along with their 50 authors as the pool of reviewers). For the CVPR datasets, no conflict matrices are available. Instead, we again sample 300 instances of size 50 by randomly selecting one paper at a time. As its author, we select the author with the highest similarity score from those not already added to the pool of reviewers, and add that author to the pool of reviewers.

Rules. We compare five rules for peer review. The Toronto Paper Matching System (TPMS) (Charlin and Zemel, 2013) maximizes the utilitarian welfare, and we also include the Nash welfare-maximizing assignment and the PeerReview4All (PR4A) (Stelmakh et al., 2021) algorithm, which is a heuristic algorithm for the leximin solution, which maximizes the minimum utility (egalitarian welfare) and, subject to that, maximizes the second-minimum, and so on. We also compare the Core-Based Reviewer Assignment (CoBRA) algorithm of Aziz et al. (2023b), which computes a reviewer assignment that lies in the core, and an appropriate adaptation of the round robin rule. In the constrained setting of peer review, round robin works by letting agents choose a single reviewer in a cyclic fashion according to some fixed order, subject to their choice of reviewer leaving at least one feasible reviewer assignment remaining that is consistent with earlier choices.⁷

Measurements. For the peer review setting the standard definition of envy does not make sense, since the reviewers allocated to paper j may not be allowed to be allocated to paper i (in particular, if the author of paper i is a reviewer for paper j). Denoting j ’s “bundle” of assigned reviewers by A_j , it might be the case that $u_i(A_j)$ is undefined. To avoid this issue, we consider a modified notion of envy ratio that we term the *shuffle envy ratio*.⁸ For a pair of agents i, j , agent i ’s shuffle envy ratio towards j is half⁹ of the ratio between the maximum utility agent i could receive from a feasible allocation B such that $A_k = B_k$ for all $k \neq i, j$, and her utility for her assigned set of reviewers, $u_i(A_i)$. That is, agent i is allowed to improve the outcome (from her perspective) by shuffling her own allocation and agent j ’s allocation in whatever way she wishes, but cannot touch the allocation of any other agent. The shuffle envy ratio of a given allocation is the maximum shuffle envy ratio over all pairs of agents. Note that IHR must be at least as large as the shuffle envy ratio, since the modified allocations that shuffle envy allows are also allowed by the definition of IHR.

Observations. Examining the graphs, we again see that IHR is highly correlated with the shuffle envy ratio. Taken together with the private goods division results in the previous subsection, these results give us confidence that IHR is capturing an intuitive and meaningful aspect of fairness. In terms of the algorithms that we examine, Nash, PR4A, and round robin all perform well. TPMS does well on ICLR 2018 and CVPR 2018 data, but relatively badly on CVPR 2017. Finally, CoBRA performs much worse than all other algorithms in terms of both IHR and shuffle envy ratio.

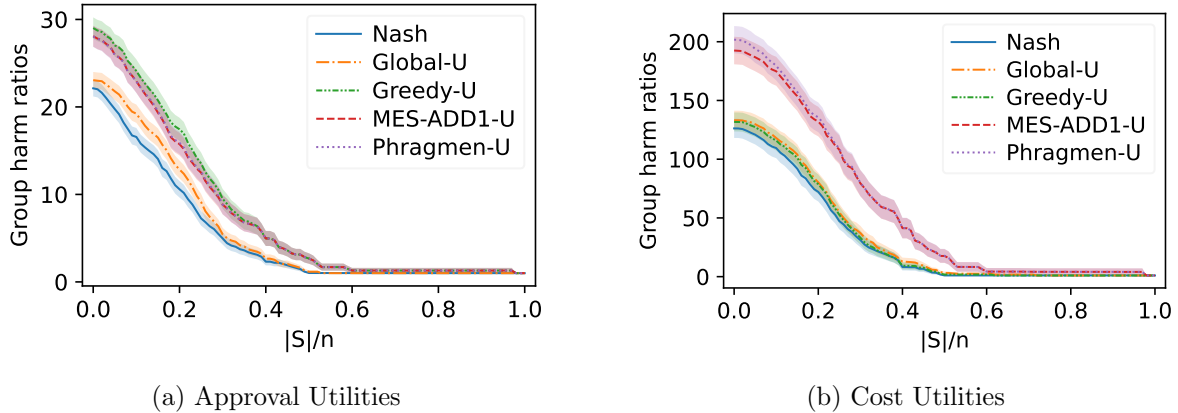


Figure 5: Plots show normalized public group envy ratio as a function of the minimum size of the envious group of agents S .

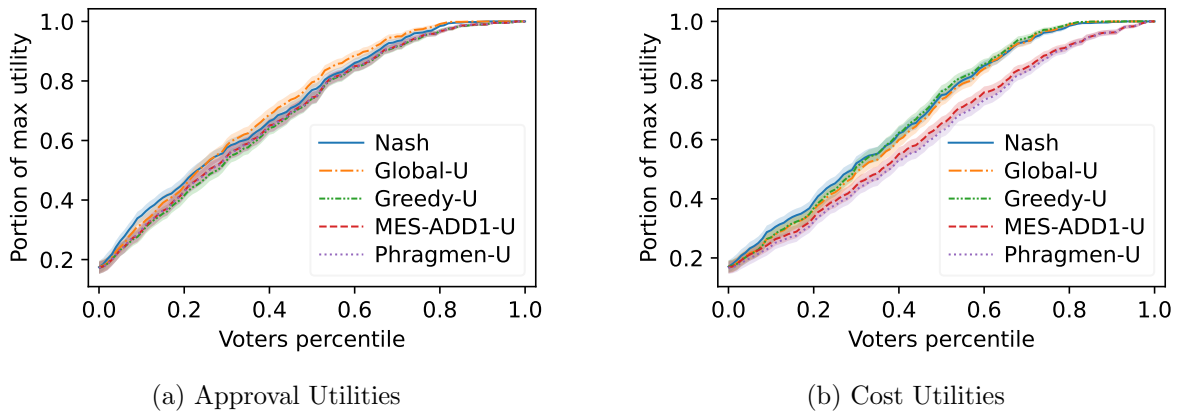


Figure 6: Plots show percentiles of voters proportional utilities (fraction of ones maximum feasible utility) averaged over all PB instances.

6.3 Participatory Budgeting

Our final setting is participatory budgeting, where the goal is for a city to select a subset of proposed projects to fund subject to a budget constraint, based on residents’ preferences. Specifically, there is a set of projects, each with an associated cost. Each voter casts an approval ballot, approving a subset of the projects. The outcome is a set of projects to fund whose total cost is at most a given budget.

Datasets. We use real participatory budgeting elections from Pabulib.org (Stolicki et al., 2020). We focus on five Polish cities: Warsaw, Gdynia, Lodz, Wroclaw, and Zabrze. For tractability, we

⁷We also assume that the minimum utility of any author for any reviewer is nonzero. Specifically, we set this minimum score to 10^{-3} while the maximum utility of each author for a reviewer is exactly 1. This avoids infinite ratios, and it is reasonable to assume receiving more reviews (with TPMS score of 0) is at least slightly better than receiving fewer.

⁸Similar results are obtained if we instead use the standard definition of envy, and only consider an agent’s envy towards papers that they themselves are not reviewing.

⁹This factor is chosen to align it with the definition of IHR, which also measures half of the ratio between the maximum utility of agent i in any outcome that hurts no agent but agent j and the current utility of agent i .

use the 321 elections with less than 15 projects on the ballot. We consider two common models of voter utility: approval utilities, in which the utility that a voter has for an outcome is the number of her approved projects that are funded, and cost utilities, in which her utility is the total cost of her approved projects that are funded. Like in peer review, we define a voter’s utility for having none of her approved projects funded to be ϵ instead of 0.¹⁰ This avoids infinite harm ratios and arguably better reflects reality, where we would expect voters to derive at least a small amount of utility even for projects that they did not vote for.

Measurements. When the number of agents is large, as in citywide elections, the individual harm ratio tends to be very small because it is hard for to find an alternative outcome that makes one agent better off, hurts at most one other agent, and keeps all the remaining (many) agents at least as happy. For this reason, we instead evaluate the group harm ratio in this section, which is always at least as large as the individual harm ratio. However, note that it is easy for a small group of agents to be harmed by the presence of a larger group. For example, a single agent who has no approved projects funded gets utility ϵ , but can increase her utility by a very large multiplicative factor if she is allowed to hurt all other agents and simply choose her most preferred budget-feasible set of projects to fund.

Recognizing this, we examine the average GHR as a function of the size of the harmed group S . By the above argument, small groups S will often produce large GHR, but these are violations that an election organizer might be comfortable with in a large election; unfairness towards a large group of agents, on the other hand, remains undesirable. Technically, we measure GHR as in [Definition 9](#), but only taking maximum over harmed groups S of a fixed size.

In [Figure 5](#), we plot the average (over elections) of the maximum group harm ratio across all (S, T) subject to the condition that $|S|/n \geq x$. For comparison, in [Figure 6](#), we consider the full distribution of agent utilities, normalized by their maximum possible utility.¹¹ For every instance, we order the agents by this normalized utility, and take the average over instances (we give voter percentile on the x axis to control for different instances having different numbers of agents).

Rules. We consider five commonly studied rules. Global-U is the utilitarian welfare-maximizing solution, Greedy-U is its greedy version that is more commonly used in practice ([Aziz and Shah, 2021](#)), and Nash denotes the MNW rule.¹² We also consider the method of equal shares (MES) ([Peters et al., 2021](#)) and Phragmen’s sequential rule (see, e.g., ([Brill et al., 2023](#))). Importantly, since MES and Phragmen might return an outcome that does not exhaust the allowed budget, we complete the outcomes of these rules into maximal solutions using the greedy utilitarian method (which adds budget-feasible unselected projects one by one, in the decreasing order of their utilitarian welfare). MES is considered a very attractive rule for participatory budgeting and is known to satisfy compelling proportionality properties, while the utilitarian rules are typically considered to be less fair than other rules due to their focus on satisfying a majority of voters, potentially at the expense of a minority.

Observations. In [Figure 5](#), we see that the group harm ratio is high when small harmed groups S are included, but decreases steadily as only larger harmed groups are considered. Under all rules and both utility models, only a very small group harm ratio remains for harmed groups of size larger than $0.5n$. Perhaps most interesting is the relative performance of the algorithms. For both

¹⁰We set $\epsilon = 10^{-2}$ for approval utilities, and $\epsilon = 10^{-3}B$ for cost utilities, where B is the total budget.

¹¹The ratio of an agent’s utility to her maximum possible utility is the approximation of proportionality for her.

¹²Greedy-U selects projects one by one in decreasing order of their utilitarian welfare, skipping over any project that cannot be added due to the budget constraint. In practice, instead of MNW, its two close approximations, namely proportional approval voting (PAV) and maximum smoothed Nash welfare, are more often used due to their better axiomatic properties. As [Figure 7](#) in the appendix shows, these rules have performance similar to that of MNW.

utility models, Nash yields the smallest harm ratios, closely followed by Global-U, the utilitarian solution. Somewhat surprisingly, MES and Phragmen, the two rules that are considered more fair in the literature, give outcomes with significantly higher GHR. The reason is clear when considering [Figure 6](#), which shows a clear dominance of the utility vectors achieved by Nash and Global-U over those achieved by MES and Phragmen, and it is this dominance that is being captured by the group harm ratio values in [Figure 5](#).

7 Discussion

Our proposal of novel fairness criteria opens the doors for a variety of extensions that have been popular in the fair division literature lately. For example, in integral settings where the utility set is discrete, can we guarantee “up to one” relaxations of our fairness criteria? Can the results be extended to weighted agents with unequal entitlements? [Theorem 2](#) goes through with the usual weighted extensions of MNW ($\arg \max_o \prod_{i \in N} u_i(o)^{w_i}$) and PF ($\max_{o'} \sum_{i \in N} w_i \cdot \frac{u_i(o')}{u_i(o)} \leq 1$), along with a natural weighted extension of group harm ratio, where the factor $\frac{|S|}{|S \cup T|}$ is replaced by $\frac{w(S)}{w(S \cup T)}$ (with $w(C) \triangleq \sum_{i \in C} w_i$). However, extending [Theorem 3](#) or considering the discrete case can be interesting. Finally, what can we say about the public bads model, where each agent has a non-positive utility for every outcome?

References

- Dana Pessach and Erez Shmueli. Algorithmic fairness. In *Machine Learning for Data Science Handbook: Data Mining and Knowledge Discovery Handbook*, pages 867–886. Springer, 2023.
- Jessie Finocchiaro, Roland Maio, Faidra Monachou, Gourab K Patro, Manish Raghavan, Ana-Andreea Stoica, and Stratis Tsirtsis. Bridging machine learning and mechanism design towards algorithmic fairness. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pages 489–503, 2021.
- Krishna P Gummadi and Hoda Heidari. Economic theories of distributive justice for fair machine learning. In *Companion Proceedings of the 2019 World Wide Web Conference*, pages 1301–1302, 2019.
- J. Tinbergen. *Mathematische psychologie*. *Mens en Maatschappij*, 6(4):342–352, 1930.
- George Gamow and Marvin Stern. *Puzzle-Math*. Viking, 1958.
- Duncan Karl Foley. Resource allocation and the public sector. *Yale Economics Essays*, 7:45–98, 1967.
- A. D. Procaccia. Cake cutting algorithms. In F. Brandt, V. Conitzer, U. Endress, J. Lang, and A. D. Procaccia, editors, *Handbook of Computational Social Choice*, chapter 13. Cambridge University Press, 2016.
- Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, and Alexandros A Voudouris. Fair division of indivisible goods: A survey. In *Proceedings of the 31st European Conference on Artificial Intelligence (ECAI)*, pages 5385–5393, 2022.
- Vincent Conitzer, Rupert Freeman, and Nisarg Shah. Fair public decision making. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC)*, pages 629–646, 2017.

- Brandon Fain, Kamesh Munagala, and Nisarg Shah. Fair allocation of indivisible public goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, pages 575–592, 2018.
- Hal R Varian. Equity, envy and efficiency. *Journal of Economic Theory*, 9:63–91, 1974.
- Marcus Berliant, William Thomson, and Karl Dunz. On the fair division of a heterogeneous commodity. *Journal of Mathematical Economics*, 21(3):201–216, 1992.
- Vincent Conitzer, Rupert Freeman, Nisarg Shah, and Jennifer Wortman Vaughan. Group fairness for the allocation of indivisible goods. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 1853–1860, 2019.
- Rupert Freeman, Nisarg Shah, and Rohit Vaish. Best of both worlds: Ex-ante and ex-post fairness in resource allocation. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, pages 21–22, 2020.
- Dietrich Weller. Fair division of a measurable space. *Journal of Mathematical Economics*, 14(1): 5–17, 1985.
- Erel Segal-Halevi and Balázs R Sziklai. Monotonicity and competitive equilibrium in cake-cutting. *Economic Theory*, 68(2):363–401, 2019.
- Thorben Troebst and Vijay V Vazirani. Cardinal-utility matching markets: The quest for envy-freeness, pareto-optimality, and efficient computability. *arXiv preprint arXiv:2402.08851*, 2024.
- L.-G. Svensson. Large indivisibles: An analysis with respect to price equilibrium and fairness. *Econometrica*, 51(4):939–954, 1983a.
- Dimitrios Diamantaras. On equity with public goods. *Social Choice and Welfare*, 9:141–157, 1992.
- Bhaskar Ray Chaudhury, Linyi Li, Mintong Kang, Bo Li, and Ruta Mehta. Fairness in federated learning via core-stability. *Advances in neural information processing systems*, 35:5738–5750, 2022.
- Marco Dall’Aglio. The dubins–spanier optimization problem in fair division theory. *Journal of Computational and Applied Mathematics*, 130(1-2):17–40, 2001.
- Richard Cole and Yixin Tao. On the existence of pareto efficient and envy-free allocations. *Journal of Economic Theory*, 193:105207, 2021.
- Ioannis Caragiannis, Vasilis Gkatzelis, Alexandros Psomas, and Daniel Schoepflin. Beyond cake cutting: Allocating homogeneous divisible goods. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems*, pages 208–216, 2022.
- Hugo Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.
- H. Moulin. *Fair Division and Collective Welfare*. MIT Press, 2004.
- Haris Aziz, Warut Suksompong, Zhaohong Sun, and Toby Walsh. Fairness concepts for indivisible items with externalities. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5472–5480, 2023a.
- J. Nash. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.

- H Peyton Young. *Equity: in theory and practice*. Princeton University Press, 1995.
- D. Bertsimas, V. F. Farias, and N. Trichakis. The price of fairness. *Operations Research*, 59(1): 17–31, 2011.
- L. E. Dubins and E. H. Spanier. How to cut a cake fairly. *American Mathematical Monthly*, 68(1): 1–17, 1961.
- Frank Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8:33–37, 1997.
- Gerard Debreu and Herbert Scarf. A limit theorem on the core of an economy. *International Economic Review*, 4(3):235–246, 1963.
- James B Orlin. Improved algorithms for computing fisher’s market clearing prices: Computing fisher’s market clearing prices. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 291–300, 2010.
- B. Fain, A. Goel, and K. Munagala. The core of the participatory budgeting problem. In *Proceedings of the 12th Conference on Web and Internet Economics (WINE)*, pages 384–399, 2016.
- J. Goldman and A. D. Procaccia. Spliddit: Unleashing fair division algorithms. *SIGecom Exchanges*, 13(2):41–46, 2014.
- Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation*, 7(3): Article 12, 2019.
- Haris Aziz, Evi Micha, and Nisarg Shah. Group fairness in peer review. In *Proceedings of the 37th Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2023b. Forthcoming.
- Laurent Charlin and Richard Zemel. The toronto paper matching system: an automated paper-reviewer assignment system. In *ICML 2013 Workshop on Peer Reviewing and Publishing Models*, 2013.
- Ivan Stelmakh, Nihar Shah, and Aarti Singh. Peerreview4all: Fair and accurate reviewer assignment in peer review. *The Journal of Machine Learning Research*, 22(1):7393–7458, 2021.
- Dariusz Stolicki, Stanisław Szufa, and Nimrod Talmon. Pabulib: A participatory budgeting library. *arXiv preprint arXiv:2012.06539*, 2020.
- Haris Aziz and Nisarg Shah. Participatory budgeting: Models and approaches. In Tamás Rudas and Gábor Péli, editors, *Pathways Between Social Science and Computational Social Science: Theories, Methods, and Interpretations*, pages 215–236. Springer, 2021.
- Dominik Peters, Grzegorz Pierczyński, and Piotr Skowron. Proportional participatory budgeting with additive utilities. *Advances in Neural Information Processing Systems*, 34:12726–12737, 2021.
- Markus Brill, Rupert Freeman, Svante Janson, and Martin Lackner. Phragmén’s voting methods and justified representation. *Mathematical Programming*, pages 1–30, 2023.
- Aanund Hylland and Richard Zeckhauser. The efficient allocation of individuals to positions. *The Journal of Political Economy*, 87(2):293–314, 1979.

- Lars-Gunnar Svensson. On the existence of fair allocations. *Zeitschrift für Nationalökonomie/Journal of Economics*, 43(3):301–308, 1983b.
- Xingyu Chen, Brandon Fain, Liang Lyu, and Kamesh Munagala. Proportionally fair clustering. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, pages 1032–1041, 2019.
- Evi Micha and Nisarg Shah. Proportionally fair clustering revisited. In *Proceedings of the 47th International Colloquium on Automata, Languages and Programming (ICALP)*, pages 85:1–85:16, 2020.
- Safwan Hossain, Andjela Mladenovic, and Nisarg Shah. Designing fairly fair classifiers via economic fairness notions. In *Proceedings of the International World Wide Web Conference (WWW)*, pages 1559–1569, 2020.

Appendix

A Detailed Significance and Implications

Simpler proof of envy-free and Pareto optimal cake-cutting. We prove that maximum Nash welfare (MNW) achieves group harm ratio (GHR) under compactness and upper convexity of the feasible utility set \mathcal{U} ([Theorem 2](#)). As we elaborate in [Section 4](#), for the prototypical fair division model of cake-cutting (formally defined in [Section 2.2](#)), these conditions hold ([Dubins and Spanier, 1961](#)) and GHR implies envy-freeness (EF) and Pareto optimality (PO). Thus, our result provides an alternative proof of the celebrated result that MNW achieves EF+PO in cake-cutting ([Weller, 1985](#); [Segal-Halevi and Sziklai, 2019](#)).

Beyond using the result of [Dubins and Spanier \(1961\)](#) that the feasible utility set \mathcal{U} is closed and convex for cake-cutting, our proof is fully elementary, e.g., teachable in a graduate course on the subject. Its domain-agnostic nature alleviates the need for defining domain-specific terminology (e.g., Borel sets and σ -algebra for cake-cutting), once closedness and convexity of \mathcal{U} (a subset of \mathbb{R}^n) has been established. Finally, it is much more straightforward than the proof of existence of an EF+PO allocation by [Weller \(1985\)](#), which uses Kakutani’s fixed point theorem, and the proof of MNW satisfying EF+PO by [Segal-Halevi and Sziklai \(2019\)](#), who first establish an equivalence between MNW and a market equilibrium concept called s-CEEI.

MNW allocations are 2-EF and PO when swaps are allowed. [Hylland and Zeckhauser \(1979\)](#) show existence of EF and PO outcomes in the one-sided matching setting, where n agents need to be (fractionally) matched to n items. This is equivalent to allocating n private goods to n agents but with the additional constraints that each agent receives a total fraction of 1 from all the goods. Very recently, [Troebst and Vazirani \(2024\)](#) show that finding a fractional matching that is EF and PO is PPAD-complete, and that the MNW solution – which can be approximated in polynomial time – achieves 2-EF. We recover this result as a direct corollary of [Theorem 2](#) (MNW \implies 1-IHR) and that 1-IHR implies 2-EF in the one-sided matching setting. To observe the latter implication, suppose by contradiction that a 1-IHR matching is not 2-EF, and the pair of agents i, j witness a violation of 2-EF. By swapping the bundles of i and j , we obtain another matching that the envious agent more than doubles their utilities without hurting anyone but the envied agent, thus failing 1-IHR. The key observation here is that while we are not allowed to take the union of the allocations to agents i and j , and give it entirely to agent i (as we did in cake-cutting), we are allowed to swap the allocations to agents i and j , and this is sufficient for 1-IHR to imply 2-EF.

Novel fair division implications. [Theorem 2](#) is connected to a series of works that establish the existence of EF+PO allocations in allocation of homogeneous divisible goods with strongly monotone preferences and a different additional condition: convex preferences ([Varian, 1974](#)), a unique allocation inducing any given (weakly) PO utility vector ([Varian, 1974](#)), the set of allocations inducing any given PO utility vector being convex ([Svensson, 1983b](#)), or the last set being a contractible space ([Diamantaras, 1992](#)). Our result provides the existence of a IHR +PO (in fact, a proportionally fair) allocation under a different condition of compactness and upper convexity of utility set \mathcal{U} .

[Ray Chaudhury et al. \(2022\)](#) show that a proportionally fair (PF) outcome exists under three conditions: \mathcal{O} is non-empty, compact and convex; each u_i is continuous; and the set of outcomes maximizing any non-negative weighted welfare, $\arg \max_{o \in \mathcal{O}} \sum_{i \in N} w_i u_i(o)$, is convex. This is one of the most general PF existence result known. [Theorem 2](#) is not subsumed by their result, and points to the possibility of an even more general result for the existence of a PF outcome.

Dall’Aglio (2001) studies a model that generalizes even cake-cutting, in which each point on the cake (represented as $[0, 1]$) can be “divided” between the agents, each agent receiving a fraction of the point. Formally, an allocation is given by $\phi = (\phi_i)_{i \in N}$ with $\phi_i : [0, 1] \rightarrow [0, 1]$ for each $i \in N$ and $\sum_{i \in N} \phi_i(x) = 1$ for all $x \in [0, 1]$; here, $\phi_i(x)$ denotes the fraction of point x allocated to agent i . The utility to agent i is given by $\int_x \phi_i(x) f_i(x) dx$, where f_i is a density function of agent i . Dall’Aglio (2001) proves that the utility set in this model is compact and convex, allowing Theorem 2 to go through to establish the existence of an EF+PO allocation.

Caragiannis et al. (2022) study a model of allocating homogeneous divisible goods, but where the utility function of an agent is weakly monotone over the fraction of a good allocated to her and only additive across the goods. They consider randomized allocations by evaluating the expected utility of the agents. Instead of EF+PO, they seek an allocation satisfying the weaker guarantee of EF and Pareto optimality within the set of EF allocations — a combination which is trivially guaranteed to be satisfiable. When the agent utilities are concave over the fraction of a good allocated, compactness and convexity of the utility set \mathcal{U} follows and Theorem 2 implies the existence of a deterministic IHR +PO allocation. IHR does not necessarily imply EF beyond fully additive utilities, so this is another relaxation of EF+PO, but one that is a priori not trivially satisfiable.

Cole and Tao (2021) study a resource allocation model where an allocation is a lottery over a given finite set of deterministic allocations. Agents have arbitrary utilities for deterministic allocations and calculate expected utilities for a randomized allocation. They show the existence of an EF+PO allocation under a swappability condition on the set of deterministic allocations. Theorem 2 implies the existence of a IHR +PO allocation (which is an incomparable guarantee to EF+PO) in their model without any swappability condition.

Envy in public outcomes. Finally, our work significantly expands the reach of envy-freeness (or more specifically, the idea of a pairwise notion of individual fairness) to a much larger set of collective decision making domains. In Section 5, we point out an open question of computing a IHR outcome in voting (where n voters express cardinal preferences over m candidates and a lottery over the candidates is selected) in polynomial time. One can also now explore IHR in domains where the concept of envy has not been explored (but other fairness notions have been), such as clustering (Chen et al., 2019; Micha and Shah, 2020), classification (Hossain et al., 2020), and federated learning (Ray Chaudhury et al., 2022).

B Proofs of Section 2.1

B.1 Proof of Proposition 1

Proof of Proposition 1. $PF \Rightarrow MNW$. Suppose $o \in \mathcal{O}$ is PF. Then, $u_i(o) > 0$ for all $i \in N$. For any $o' \in \mathcal{O}$, we have

$$\left(\prod_{i \in N} \frac{u_i(o')}{u_i(o)} \right)^{\frac{1}{n}} \leq \frac{1}{n} \sum_{i \in N} \frac{u_i(o')}{u_i(o)} \leq 1,$$

where we use the AM-GM inequality followed by the definition of PF. This shows that $NW(o') \leq NW(o)$, as desired.

$PF \Rightarrow Core$. Suppose $o \in \mathcal{O}$ is PF but violates the core. Then, there exists a group of agents $S \subseteq N$ and an outcome $o' \in \mathcal{O}$ such that $\frac{u_i(o')}{u_i(o)} \geq \frac{n}{|S|}$ for all $i \in S$ and at least one inequality is strict (we can place $u_i(o)$ in the denominator because o being PF ensures $u_i(o) > 0$). Thus, we have

$$\frac{1}{n} \cdot \sum_{i \in N} \frac{u_i(o')}{u_i(o)} \geq \frac{1}{n} \cdot \sum_{i \in S} \frac{u_i(o')}{u_i(o)} > \frac{1}{n} \cdot |S| \cdot \frac{n}{|S|} = 1,$$

which violates o being PF.

Core \Rightarrow (*Prop*+*PO*). *Core* \Rightarrow *Prop* follows from observing that proportionality imposes the same constraint as the core, but only for groups of agents S with $|S| = 1$. Similarly, Pareto optimality imposes the same constraint as the core but only for $S = N$. \square

C Proofs of Section 3

C.1 Proof of Theorem 1

Proof of Theorem 1. Fix any private goods division instance with additive utilities. We have already argued that 1-IHR implies EF. To see that *Prop* implies $(n/2)$ -IHR, consider any *Prop* allocation A . Hence, $u_i(A_i) \geq \frac{1}{n} \cdot u_i(M)$. Suppose for contradiction that A is not $(n/2)$ -IHR. Hence, there exist agents $i, j \in N$ and allocation A' such that $u_i(A'_i) > n \cdot u_i(A_i) \geq u_i(M)$, which is a contradiction.

For $n = 2$, this shows that 1-IHR is equivalent to EF, which is also evident from the definitions of IHR and EF. To show tightness for $n \geq 3$, consider the instance in Figure 1 where n homogeneous divisible goods need to be divided between n agents. Agent $i \in [n]$ values items $j \in \{1, \dots, i\}$ at 1 and the rest at 0.

Consider an allocation A (shown via blue rectangles in Figure 1) that divides g_1 equally among all agents and allocates good $j \in [n] \setminus \{1\}$ to agent j . Then, $u_1(A_1) = \frac{1}{n}$ and $u_i(A_i) = 1 + \frac{1}{n}$ for all $i \in [n] \setminus \{1\}$. It is easy to check that this is EF. However, we show that agent 1 is significantly harmed by the presence of agent 2 in the sense of IHR. Consider the allocation A' (shown via red rectangles in Figure 1) where each agent $i \in [n] \setminus \{2\}$ is allocated the i -th good, and the second good is equally divided among agents in $\{3, \dots, n\}$. Then, $u_1(A'_1) = 1$ and $u_i(A'_i) = 1 + \frac{1}{n-2} \geq 1 + \frac{1}{n} = u_i(A_i)$ for all $i \in [3, n]$. Each agent $i \in N \setminus \{1, 2\}$ is at least as happy under A' as under A , and for agent 1, $u_1(A'_1) = 1 = \frac{n}{2} \cdot 2 \cdot u_1(A_1)$. Hence, A is only $(n/2)$ -IHR. \square

C.2 Proof of Proposition 2

Proof of Proposition 2. PF \Rightarrow 1-GHR. Suppose $o \in \mathcal{O}$ is PF but violates 1-GHR. Then there exist non-empty $S, T \subseteq N$ and outcome $o' \in \mathcal{O}$ such that $\frac{|S|}{|S \cup T|} \cdot u_i(o') \geq u_i(o)$ for all $i \in S$, with at least one inequality being strict, and $u_i(o') \geq u_i(o)$ for all $i \in N \setminus (S \cup T)$. If $u_i(o) = 0$ for any $i \in N$, then o clearly violates PF, which is the desired contradiction.

Assume that $u_i(o) > 0$ for all $i \in N$. Then we have

$$\begin{aligned} \frac{1}{n} \sum_{i \in N} \frac{u_i(o')}{u_i(o)} &\geq \frac{1}{n} \left(\sum_{i \in S} \frac{u_i(o')}{u_i(o)} + \sum_{i \in N \setminus (S \cup T)} \frac{u_i(o')}{u_i(o)} \right) \\ &> \frac{1}{n} \left(\sum_{i \in S} \frac{|S \cup T|}{|S|} + \sum_{i \in N \setminus (S \cup T)} 1 \right) = \frac{1}{n} (|S \cup T| + |N \setminus (S \cup T)|) = 1, \end{aligned}$$

contradicting that o is PF.

1-GHR \Rightarrow (PO + Core + 1-IHR). Core, PO, and 1-IHR all impose the same constraints as 1-GHR, but only for some of possible pairs of groups S and T . PO imposes the constraint for $S = T = N$, the core imposes the constraints for $T = N$ (and any S), while 1-IHR imposes them for $|S| = |T| = 1$. \square

D Proofs of Section 4

D.1 Proof of Theorem 2

Proof of Theorem 2. The existence of an MNW outcome is due to Proposition 3. Let $o \in \mathcal{O}$ be any MNW outcome. Suppose for contradiction that it is not PF. Then, there exists an outcome $o' \in \mathcal{O}$ such that

$$\frac{1}{n} \sum_{i \in N} \frac{u_i(o')}{u_i(o)} > 1. \quad (1)$$

Define a function $f : (0, 1] \rightarrow \mathbb{R}$ such that for $\alpha \in (0, 1]$,

$$f(\alpha) = \frac{1}{n} \cdot \sum_{i \in N} \log(\alpha \cdot u_i(o) + (1 - \alpha) \cdot u_i(o')).$$

Note that this is well-defined because $u_i(o) > 0$ for all agents $i \in N$ (Proposition 3) and $\alpha > 0$, implying that the first term inside the logarithm is strictly positive (and the second term is weakly positive since $\alpha \leq 1$ and $u_i(o') \geq 0$). Note that $f(1) = \log \text{NW}(o)$.

Next, we prove that there exists $\alpha \in (0, 1)$ such that $f(\alpha) > f(1)$. Since f is differentiable in α , it suffices to show that the left-derivative of f at $\alpha = 1$ is strictly negative. It is easy to see that

$$\partial_- f(1) = \frac{1}{n} \sum_{i \in N} \frac{u_i(o) - u_i(o')}{u_i(o)} = 1 - \frac{1}{n} \cdot \sum_{i \in N} \frac{u_i(o')}{u_i(o)} < 0,$$

where the last inequality follows from Equation (1).

Thus, there exists $\alpha \in (0, 1)$ for which $f(\alpha) > f(1)$. Due to upper convexity of \mathcal{U} , there exists an outcome $\hat{o} \in \mathcal{O}$ such that $u_i(\hat{o}) \geq \alpha \cdot u_i(o) + (1 - \alpha) \cdot u_i(o')$ for all agents $i \in N$. Thus, $\log \text{NW}(\hat{o}) \geq f(\alpha) > f(1) = \log \text{NW}(o)$, contradicting the fact that o is an MNW outcome. Hence, we have proved that o must also be PF. Since PF implies MNW more generally (Proposition 1), this makes MNW and PF equivalent on every instance of the public outcomes model. \square

D.2 Proof of Theorem 3

Proof of Theorem 3. Fix any instance of the public outcomes model. Let $o, o' \in \mathcal{O}$ be any two outcomes satisfying 1-GHR. We want to show that $\text{NW}(o) \geq \binom{n}{\lfloor n/2 \rfloor}^{-1/n} \cdot \text{NW}(o')$. If $\text{NW}(o') = 0$, this holds trivially, so let us assume that $\text{NW}(o') > 0$, i.e., $u_i(o') > 0$ for all agents $i \in N$.

Sort the agents $i \in N$ in a non-decreasing order of $\frac{u_i(o)}{u_i(o')}$ (this is well-defined because $u_i(o') > 0$ for all $i \in N$); without loss of generality, let us rename the agents so that $\frac{u_1(o)}{u_1(o')} \leq \frac{u_2(o)}{u_2(o')} \leq \dots \leq \frac{u_n(o)}{u_n(o')}$. Let $L = \{i \in N : u_i(o') > u_i(o)\}$ and $H = \{i \in N : u_i(o') \leq u_i(o)\}$. By our renaming, we have that $L = \{1, 2, \dots, k\}$ and $H = \{k+1, k+2, \dots, n\}$, where $k = |L|$.

Fix any $i \in L$. Let $S = \{1, 2, \dots, i\} \subseteq L$ and $T = H$. Note that $S \cap T = \emptyset$. Since o is 1-GHR, there exists $j \in S$ for which $\frac{|S|}{|S \cup T|} \cdot u_j(o') \leq u_j(o)$, i.e., $\frac{u_j(o)}{u_j(o')} \geq \frac{i}{n-k+i}$. Due to the sorting of the agents, this implies $\frac{u_i(o)}{u_i(o')} \geq \frac{i}{n-k+i}$. For any $i \in H$, we have $\frac{u_i(o)}{u_i(o')} \geq 1$ by the definition of H .

Putting these together, we have

$$\frac{\text{NW}(o)}{\text{NW}(o')} = \left(\prod_{i \in N} \frac{u_i(o)}{u_i(o')} \right)^{\frac{1}{n}} \geq \left(\prod_{i=1}^k \frac{i}{n-k+i} \cdot \prod_{i=k+1}^n 1 \right)^{\frac{1}{n}} = \left(\frac{k!(n-k)!}{n!} \right)^{\frac{1}{n}} = \binom{n}{k}^{-\frac{1}{n}} \geq \binom{n}{\lfloor n/2 \rfloor}^{-\frac{1}{n}},$$

as required.

To prove tightness, fix an $\epsilon \in (0, 1/2)$ and consider an example with $\mathcal{O} = \{o, o^*\}$, where

$$u_i(o) = \begin{cases} \frac{i}{\lfloor n/2 \rfloor + i} + \epsilon, & i \leq \lfloor n/2 \rfloor \\ 1 + \epsilon, & i > \lfloor n/2 \rfloor, \end{cases}$$

and $u_i(o^*) = 1$ for all $i \in N$. We claim that o is 1-GHR. To see this, suppose for contradiction that there exist groups of agents S, T for which the two conditions from [Definition 9](#) hold (taking $o' = o^*$). Note that $T \supseteq \{i : i > \lfloor n/2 \rfloor\}$, since any agent who receives less utility in outcome o^* than in outcome o must be contained in T . In particular, $|T| \geq \lfloor n/2 \rfloor$. Next, observe that $S \subseteq \{i : i \leq \lfloor n/2 \rfloor\}$, since agents in S must receive more utility in outcome o^* than in outcome o . In particular, writing $|S| = k$, we have $k \leq \lfloor n/2 \rfloor$ and $|S \cup T| \geq k + \lfloor n/2 \rfloor$. Additionally, we have $\max_{i \in S} u_i(o) \geq \frac{k}{\lfloor n/2 \rfloor + k} + \epsilon$. Letting $i' = \arg \max_{i \in S} u_i(o)$,

$$\frac{|S|}{|S \cup T|} \cdot u_{i'}(o^*) \leq \frac{k}{k + \lfloor n/2 \rfloor} < \frac{k}{\lfloor n/2 \rfloor + k} + \epsilon \leq u_{i'}(o),$$

contradicting the violation of 1-GHR. To obtain the desired bound, note that

$$\frac{\text{NW}(o)}{\text{NW}(o^*)} \xrightarrow{\epsilon \rightarrow 0} \left(\prod_{i=1}^{\lfloor n/2 \rfloor} \frac{i}{\lfloor n/2 \rfloor + i} \right)^{1/n} = \left(\frac{\lfloor n/2 \rfloor! \lfloor n/2 \rfloor!}{n!} \right)^{1/n} = \binom{n}{\lfloor n/2 \rfloor}^{-1/n}.$$

Finally, $\binom{n}{\lfloor n/2 \rfloor}^{1/n} \leq 2$ holds because there are only 2^n subsets (of any size) of a set of size n , and the limiting value as $n \rightarrow \infty$ follows directly from applying Stirling's approximation. \square

E Proofs of Section 5

E.1 Proof of [Proposition 4](#)

Proof of [Proposition 4](#). Suppose $o \in \mathcal{O}$ is $(1 + \epsilon)$ -PF but violates $(1 + n\epsilon)$ -GHR. Then there exist non-empty $S, T \subseteq N$ and outcome $o' \in \mathcal{O}$ such that $\frac{|S|}{|S \cup T|} \cdot u_i(o') \geq (1 + n\epsilon)u_i(o)$ for all $i \in S$, and $u_i(o') \geq u_i(o)$ for all $i \in N \setminus (S \cup T)$. If $u_i(o) = 0$ for any $i \in N$, then o clearly violates $(1 + \epsilon)$ -PF, which is the desired contradiction.

Assume that $u_i(o) > 0$ for all $i \in N$. Then we have

$$\begin{aligned} \frac{1}{n} \sum_{i \in N} \frac{u_i(o')}{u_i(o)} &\geq \frac{1}{n} \left(\sum_{i \in S} \frac{u_i(o')}{u_i(o)} + \sum_{i \in N \setminus (S \cup T)} \frac{u_i(o')}{u_i(o)} \right) \\ &\geq \frac{1}{n} \left(\sum_{i \in S} \frac{(1 + n\epsilon)|S \cup T|}{|S|} + \sum_{i \in N \setminus (S \cup T)} 1 \right) \\ &= \frac{1 + n\epsilon}{n} \cdot |S \cup T| + \frac{1}{n} \cdot |N \setminus (S \cup T)| \geq 1 + \epsilon \cdot |S \cup T| \geq 1 + \epsilon, \end{aligned}$$

contradicting that o is $(1 + \epsilon)$ -PF. \square

		G^1			G^2			G^3			G^4			G^5			g^*
		g_1^1	...	g_{3d}^1	g_1^2	...	g_d^2	g_1^3	...	g_d^3	g_1^4	...	g_{3d}^4	g_1^5	...	g_{3d}^5	
Y	y_1	①	...	0	0	...	0	0	...	0	$1 - \epsilon$...	0	ϵ	...	0	0
	\vdots	\vdots	\ddots	\vdots	\vdots	\ddots	\vdots	\vdots	\ddots	\vdots	\vdots	\ddots	\vdots	\vdots	\ddots	\vdots	\vdots
	y_{3d}	0	...	①	0	...	0	0	...	0	0	...	$1 - \epsilon$	0	...	ϵ	0
Z	z_1	n_1	...	n_{3d}	①	...	0	①	...	0	0	...	0	0	...	0	0
	\vdots	\vdots	\ddots	\vdots	\vdots	\ddots	\vdots	\vdots	\ddots	\vdots	\vdots	\ddots	\vdots	\vdots	\ddots	\vdots	\vdots
	z_d	n_1	...	n_{3d}	0	...	①	0	...	①	0	...	0	0	...	0	0
W	w_1	0	...	0	2	...	2	0	...	0	①	...	①	0	...	0	ϵ
	w_2	0	...	0	0	...	0	0	...	0	0	...	0	①	...	①	①

Table 1: The indivisible goods allocation instance from the proof of Theorem 4. The outcome A is indicated by circles.

E.2 Proof of Theorem 4

Proof of Theorem 4. First, note that membership in coNP holds since an outcome o can be certified to not be 1-IHR by providing an alternative outcome o' and two agents i, j for which the violating conditions in Definition 8 hold.

To show coNP-completeness, we reduce from the 3-Partition problem: Given a multiset C of $3d$ numbers c_1, \dots, c_{3d} lying strictly between $1/4$ and $1/2$, can C be partitioned into d triplets C_1, \dots, C_d such that the sum of the members of each triplet is 1?

Given a 3-Partition instance, we define an indivisible goods instance (along with an outcome) with $m = 11d + 1$ goods arranged into the following subsets:

$$\begin{aligned}
G^1 &= \{g_1^1, \dots, g_{3d}^1\} \\
G^2 &= \{g_1^2, \dots, g_d^2\} \\
G^3 &= \{g_1^3, \dots, g_d^3\} \\
G^4 &= \{g_1^4, \dots, g_{3d}^4\} \\
G^5 &= \{g_1^5, \dots, g_{3d}^5\},
\end{aligned}$$

along with a single good g^* . There are $n = 4d + 2$ agents, arranged into the following subsets:

$$\begin{aligned}
Y &= \{y_1, \dots, y_{3d}\} \\
Z &= \{z_1, \dots, z_d\} \\
W &= \{w_1, w_2\}.
\end{aligned}$$

Let $0 < \epsilon < 1$. Valuations that the agents have for the goods are defined as follows (for ease of reading, we use the notation $v_i(g)$ rather than $v_{i,g}$; valuations that are not directly specified are equal to zero). For every $i \in Y$, $v_i(g_i^1) = 1$, $v_i(g_i^4) = 1 - \epsilon$, and $v_i(g_i^5) = \epsilon$. For every $i \in Z$ and every $j \in [3d]$, $v_i(g_j^1) = n_j$. Thus, these agents and items directly encode the values from the 3-Partition instance. Further, for every $i \in Z$, $v_i(g_i^2) = 1$, and $v_i(g_i^3) = d$. For agent w_1 we have $u_{w_1}(g_j^2) = 2$ for all $j \in [d]$, $u_{w_1}(g_j^4) = 1/3$ for all $j \in [3d]$, and $u_{w_1}(g^*) = \epsilon$. For agent w_2 , we have $u_{w_2}(g_j^5) = \epsilon$ for all $j \in [3d]$ and $u_{w_2}(g^*) = \epsilon$.

The allocation A is as follows, and is also depicted in Table 1. For every $j \in [3d]$, good g_j^1 is allocated to agent y_j . For every $j \in [d]$, goods g_j^2 and g_j^3 are allocated to agent z_j . Agent w_1 is

allocated the entire set of goods G^4 . Finally, agent w_2 is allocated the entire set of goods G^5 as well as good g^* . We have $u_{y_i}(A_{y_i}) = 1$ for all $i \in [3d]$, $u_{z_i}(A_{z_i}) = d + 1$ for all $i \in [d]$, $u_{w_1}(A_{w_1}) = d$, and $u_{w_2}(A_{w_2}) = (3d + 1)\epsilon$. Note that for all agents $i \neq w_1$, it is the case that $u_i(A_i) \geq u_i(M)/2$.

First suppose that there exists a solution to the 3-Partition instance. Then we can define an outcome A' as follows. For every $i \in [3d]$, agent $y_i \in Y$ is allocated g_i^4 and g_i^5 . For every $i \in [m]$, agent $z_i \in Z$ is allocated g_i^3 and a set of three goods from G^1 with total value exactly equal to 1 (this is possible by the assumption of a 3-Partition solution). Agent w_1 is allocated g^* and all goods in G^2 . We have $u_{y_i}(A'_{y_i}) = 1 = u_{y_i}(A_{y_i})$ for all $i \in [3d]$, $u_{z_i}(A'_{z_i}) = d + 1 = u_{z_i}(A_{z_i})$ for all $i \in [m]$, and $u_{w_1}(A'_{w_1}) = 2m + \epsilon > 2d = 2u_{w_1}(A_{w_1})$. Since only agent w_2 receives less utility under A' than under A , allocation A' is witness to a violation of 1-IHR.

Next, suppose that there does not exist a solution to the 3-Partition instance. We will show that outcome A is 1-IHR.


Since $u_i(A_i) \geq u_i(M)/2$ for all $i \neq w_1$, the only agent who can more than double her utility in some outcome $A' \neq A$ is w_1 . Consider such an outcome A' and suppose for contradiction that A' witnesses a violation of 1-IHR.

For it to be the case that $u_{w_1}(A'_{w_1}) > 2u_{w_1}(A_{w_1}) = 2m$, agent w_1 must be allocated some goods from G^2 (each of which she has utility 2 for). Suppose that she receives k such goods. We may assume that $m \geq 3$ and therefore receiving only $k = 1$ good from G^2 is not sufficient to double her utility; assume therefore that $k \geq 2$. We now follow a chain of implications about what this means for the outcome A' , given that there can be at most one agent i with $u_i(A'_i) < u_i(A_i)$.

At least k agents in Z must not be allocated any $g_j^2 \in G^2$ under allocation A' . Denote this set of agents by Z' . For A' to witness a violation of 1-IHR, it must be the case that for at least $k - 1 \geq 1$ of these agents in Z' , they receive utility at least $u_{z_i}(A'_{z_i}) \geq d + 1$, which implies that $A'_{z_i} = \{g_{z_i}^3\} \cup X$ for some $X \subseteq G^1$ with $|X| \geq 3$. The condition $X \subseteq G^1$ is required because, other than $g_{z_i}^3$ (for which agent z_i has utility d) and $g_{z_i}^2$ (which is not allocated to agent z_i , by definition), agents in Z' only get positive utility from G^1 , and $u_{z_i}(g_j^1) < 1/2$ for all $j \in [3d]$.

Now, there exist at least three agents $y_i \in Y$ for whom $g_i^1 \notin A'_{y_i}$. Since A' witnesses a violation of 1-IHR, it is necessary for at least two of these agents to have $u_{y_i}(A'_{y_i}) \geq u_{y_i}(A_{y_i}) = 1$, which implies that $A'_{y_i} = \{g_i^4, g_i^5\}$. In particular, if good g_i^5 is allocated to agent y_i under A' , then it is no longer allocated to agent w_2 , which means that $u_{w_2}(A'_{w_2}) < u_{w_2}(A_{w_2})$. Thus, for all agents $i \neq w_2$, it must be the case that $u_i(A'_i) \geq u_i(A_i)$, or else A' would not witness a violation of 1-IHR.

Having established this fact, the remaining analysis is simplified. Agent w_1 receives k goods from G^2 under A' . Therefore, at least $3k$ goods from G^1 must be allocated to agents in Z' , to compensate each of them for their lost good from G^2 . Accordingly, at least $3k$ goods from G^4 and $3k$ goods from G^5 must be allocated to agents in Y , to compensate those agents who lost a G^2 good. Since agent w_1 values each G^4 good at $1/3$, her utility from A' is $u_{w_1}(A'_{w_1}) \leq 2k + (3d - 3k)/3 + \epsilon = d + k + \epsilon$ (the additional ϵ comes from allocating g^* to w_1). Thus, unless $k = d$, $u_{w_1}(A'_{w_1}) < 2d$, contradicting the assumption that A' is a witness to a 1-IHR violation. So assume $k = d$. But in this case, all of the goods in G^1 need to be allocated to agents in Z . Since there does not exist a solution to the 3-Partition problem, it is impossible to assign them so that $u_{z_i}(A'_{z_i}) \geq d + 1 = u_{z_i}(A_{z_i})$ for all $i \in [d]$, since doing so would require each agent z_i receiving utility exactly 1 from a subset of the G^1 goods. Again, this contradicts A' witnessing a 1-IHR violation. Therefore, A is 1-IHR. This concludes the proof. \square



(a) Normalized public group envy ratio as a function of the minimum size of the envious group of agents S .

(b) percentiles of voters proportional utilities (fraction of ones maximum feasible utility) averaged over all PB instances.

Figure 7: Comparison among the Nash welfare-maximizing solution (Nash), Proportional Approval Voting (PAV), and the maximum smoothed Nash solution (Smooth-Nash), based on participatory budgeting datasets with approval utilities.

F Additional Experimental Plots

F.1 Nash Welfare Variants for Approval-Based Participatory Budgeting

For the experiments in the approval-based participatory budgeting setting, in [Figure 7](#), we compare the Nash welfare-maximizing solution to its two close approximations:

- Proportional Approval Voting Rule (PAV), which selects the set of projects O that maximizes $\sum_i H(u_i(O))$ where $H(0) = 0$ and $H(k) = \sum_{k'=1}^k \frac{1}{k'}$ for $k \in \mathbb{N}$ is the k th harmonic number.
- Maximum Smoothed Nash Welfare Rule (Smooth-Nash), which selects the set of projects that maximizes $\sum_i \log(u_i(O) + 1)$.

All three rules achieve similar results.