# Pushing the Limits of Fairness in Algorithmic Decision-Making

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

## Abstract

Designing provably fair decision-making algorithms is a task of growing interest and importance. In this article, I argue that preference-based notions of fairness proposed decades ago in the economics literature and subsequently explored in-depth within computer science (specifically, within the field of computational social choice) are aptly suited for a wide range of modern decision-making systems, from conference peer review to recommender systems to participatory budgeting.

## 1 Introduction

With machine learning (ML) deployed worldwide to make critical decisions, there is exploding interest in ensuring that the ML models treat (groups of) people fairly. While algorithmic fairness is a nascent subject of study in machine learning, it has deep roots in social choice theory from economics. The more recent field of computational social choice at the intersection of computer science and economics has explored the applicability of these economic fairness notions to algorithmic paradigms. The purpose of this article is to demonstrate that these notions are well-suited to a broad range of algorithmic decision-making settings and deserve an in-depth study.

Take, for example, the classical economic notion of the core [Gillies, 1953], which demands that when collective resources are divided amongst a group of people, no subset of them be able to find a better[1] allocation of their "entitled share" of the collective resources. This applies to participatory budgeting [Fain *et al.*, 2018], where a city allocates a public budget to fund infrastructure projects and the "entitled share" of any group of residents is defined in proportion to the size of the group. The same notion also applies to conference peer review [Aziz *et al.*, 2023], where the "resource" being divided is the reviewing capacity and the "entitled share" of any group of authors is the reviewing capacity they contribute by also serving as potential reviewers.

These types of fairness notions are appealing because they are well-defined for a wide range of domains, including ones with highly complex decision spaces, and they pay explicit attention to the preferences of the stakeholders involved. In this article, I will survey their applications to various algorithmic decision-making paradigms, based partly on my own work, and argue that these notions can play a key role in realizing the grand vision of an overarching theory for algorithmic fairness that spans a diverse set of applications.

In Section 2, I will first define several fairness definitions, some proposed in my own work, using an example application: the allocation of homogeneous divisible goods. Then, in Section 3, I will demonstrate their applicability to a wide range of real-world domains, which my work has explored extensively. Finally, in Section 4, I will conclude with a call to arms for exploring the applicability of such preference-based fairness definitions in novel domains in an attempt to develop an overarching theory for algorithmic fairness.

## 2 Fairness Definitions

Let us review the basic model for allocating homogeneous divisible goods, which will serve as our reference context. There is a set of agents $N$ and a set of divisible goods $M$. Each agent $i \in N$ values each good $g \in M$ at $v_{i,g}$, and agents have additive linear preferences: for receiving $X_g \in [0,1]$ fraction of each good $g$, the utility of agent $i$ is given by $v_i(X) = \sum_{g \in M} X_g \cdot v_{i,g}$, where $X = (X_g)_{g \in M}$. Let $v_i(M) \triangleq \sum_{g \in M} v_{i,g}$ be the total value of agent $i$.

An allocation $A \in [0,1]^{N \times M}$ allocates $A_{i,g}$ fraction of each good $g$ to each agent $i$; a valid allocation must satisfy $\sum_{g \in M} A_{i,g} \leq 1$ for all $i \in N$. Denoting $A_i = (A_{i,g})_{g \in M}$, each agent $i$ derives utility $v_i(A_i)$ under this allocation. This models many real-world applications where divisible resources must be allocated; an example is the division of food items (such as milk or rice) donated to a food bank.

Next, let us review some prominent fairness definitions.

### 2.1 Entitlement-Based Notions

Some definitions focus on a notion of *entitlement* of an agent or a group of agents and aim to treat each agent or group of agents no worse than their entitlement. I refer to them as *entitlement-based notions*. One example is the classical notion of *proportionality* [Steinhaus, 1948].

**Definition 1** (Proportionality)**.** An allocation $A$ is proportional if $v_i(A_i) \geq (1/n) \cdot v_i(M)$ for each $i \in N$. Here, agent

[1] Here, "better" means a Pareto improvement, which makes at least one person happier without hurting anyone.

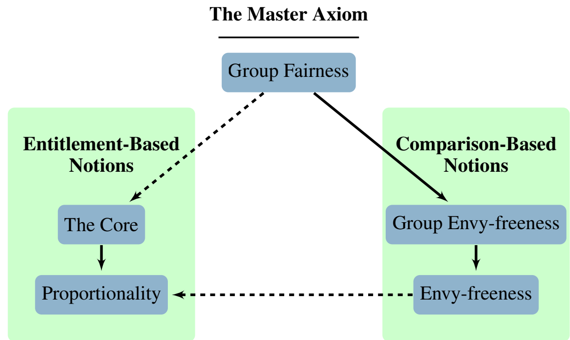


Figure 1: Logical relations between fairness axioms for the allocation of homogeneous divisible goods. Solid arrows indicate implications that hold for arbitrary monotone preferences, while dashed arrows indicate implications that hold under additive linear preferences.

$i$ is viewed as entitled to receiving $1/n$-th of her total value.

The notion of the *core* [Gillies, 1953], imported from the economic theory of non-transferable utility games, lifts this idea to all ($2^n - 1$ many) non-empty groups of agents.

**Definition 2** (The Core)**.** An allocation $A$ is in the core if there is no non-empty group of agents $S \subseteq N$ and allocation $B$ of the goods to agents in $S$ such that $v_i(A_i) \leq (|S|/n) \cdot v_i(B_i)$ for all $i \in S$ and at least one inequality is strict. Here, each group $S$ is viewed as entitled to receiving $|S|/n$-th of any combination of values they can achieve by dividing the goods amongst themselves.

Note that the core logically implies proportionality, as the latter basically imposes the core condition on singleton groups. Let me remark that instead of letting group $S$ divide all the goods amongst themselves and scaling their utility down by a factor of $|S|/n$ (*utility-scaling version*), one may allow them to divide only $|S|/n$ fraction of each good and not perform any utility scaling (*endowment-scaling version*). For our reference application, the two are equivalent due to additive linear preferences, but in general they can differ significantly (e.g., see Section 3.2). The utility-scaling version is more broadly defined because in some applications it may not be clear what resources to scale or how. But when there is a compelling way to scale resources, the endowment-scaling version can often make more sense.

## 2.2 Comparison-Based Notions

Other notions focus on *comparing* the treatment provided to two agents or two groups of agents. I refer to them as *comparison-based notions*. A widely known example is *envy-freeness* [Gamow and Stern, 1958; Foley, 1967].

**Definition 3** (Envy-freeness)**.** An allocation $A$ is envy-free if $v_i(A_i) \geq v_i(A_j)$ for all $i, j \in N$. Here, the allocations to agents $i$ and $j$ are being compared using the preferences of agent $i$ (and, due to the quantifier, also those of agent $j$).

Crucially, the definition uses the same valuation function on both sides of the equation, thus avoiding any *interpersonal comparison of utilities*. *Group envy-freeness* [Varian, 1974; Berliant *et al.*, 1992] extends this to all pairs of groups of agents of equal size.

**Definition 4** (Group Envy-freeness)**.** An allocation $A$ is group envy-free if there are no non-empty groups $S, T \subseteq N$ with $|S| = |T|$ and allocation $B$ of the resources $\cup_{j \in T} A_j$ allocated to group $T$ to agents in $S$ such that $v_i(A_i) \leq v_i(B_i)$ for all $i \in S$ and at least one inequality is strict. This ensures that group $S$ does not envy group $T$ collectively.

## 2.3 The Master Axiom

In recent work [Conitzer *et al.*, 2019], we formulated a novel fairness definition that is logically stronger than both the core and group envy-freeness, thus sitting at the apex of both entitlement-based and comparison-based notions defined above. These logical implications are depicted in Figure 1. Group fairness extends group envy-freeness to all pairs of groups $S$ and $T$, but with a careful adjustment to the notion of envy based on the difference in the sizes of the groups.

**Definition 5** (Group Fairness)**.** An allocation $A$ is group fair if there are no non-empty groups $S, T \subseteq N$ and allocation $B$ of the resources $\cup_{j \in T} A_j$ allocated to $T$ to agents in $S$ such that $v_i(A_i) \leq (|S|/|T|) \cdot v_i(B_i)$ for all $i \in S$ and at least one inequality is strict. Here, after allowing $S$ to reallocate $T$'s resources amongst themselves, their utilities are scaled down by $|S|/|T|$, reflecting the difference in group sizes.

At this point, one wonders whether any algorithm can provably satisfy these notions of fairness. Proportionality and envy-freeness are rather easy to satisfy, e.g., by dividing each good equally among the agents. Prior work in economics shows that an allocation maximizing the Nash social welfare (i.e., product of utilities) of the agents, which coincides with the competitive equilibrium from equal incomes (CEEI), satisfies the core [Varian, 1974] and group envy-freeness [Varian, 1974; Berliant *et al.*, 1992]. We prove that it actually satisfies group fairness [Conitzer *et al.*, 2019] and, together with a mild axiom, it is in fact characterized by group fairness [Freeman *et al.*, 2020].

# 3 Expansive Applicability

I will now review several applications for which these fairness notions have been shown to be appealing in recent work, including in some of my own. In addition to the applications listed below, I have also explored them for public good allocation [Conitzer *et al.*, 2017a; Fain *et al.*, 2018; Banerjee *et* *al.*, 2023], ad allocation [Hosseini *et al.*, 2023], multi-armed bandits [Hossain *et al.*, 2021], land division [Caragiannis *et al.*, 2022], and team formation [Li *et al.*, 2023].

### 3.1 Allocating Indivisible Goods

In the reference setting of Section 2, if the goods being allocated are *indivisible* (i.e., cannot be split fractionally between agents), none of the fairness notions can be guaranteed. Nonetheless, our work has shown that maximizing the Nash social welfare[2] satisfies "up to one good"-style relaxations of these notions,[3] all the way up to group fairness [Caragiannis *et al.*, 2019; Conitzer *et al.*, 2019].

One may find the violation of fairness, even up to one good, troubling, especially if a high-valued good causes a large violation. There are multiple ways to address this. On the one end, our work has shown that the violation can be removed ex-ante by using randomization: specifically, there always exists a randomization over envy-free up to one good allocations that is (exactly) envy-free in expectation [Freeman *et al.*, 2020]. On the other end, we have also identified a stronger ex-post guarantee of *envy-freeness up to any good* (EFX), which informally demands that the fairness violation be caused by the least-valued good. Resolving whether an EFX allocation always exists is "fair division's biggest problem" [Procaccia, 2020], and there are promising recent developments towards a (positive) resolution [Chaudhury *et al.*, 2020; Amanatidis *et al.*, 2021; Akrami *et al.*, 2023].

Our work has also shown that some of these fairness guarantees remain achievable when allocating (negatively-valued) *chores* instead of goods [Ebadian *et al.*, 2022b; Freeman *et al.*, 2020], but many important questions remain open; see the recent survey by Amanatidis *et al.* [2022].

### 3.2 Participatory Budgeting

Participatory budgeting (PB) is a process whereby residents of a geographical region vote over how a portion of the public budget should be allocated to fund some of the proposed public projects, and hundreds of millions of dollars have been allocated worldwide via PB. See our recent book chapter for a detailed account of research on PB [Aziz and Shah, 2021].

Here, (the endowment-scaling version of) the core especially makes natural sense: it demands that no group of residents be able to find an allocation of their proportional share of the budget[4] that they prefer to chosen allocation of the whole budget. We show that $\widetilde{O}(\log m)$-approximate core is achievable with $m$ proposed projects [Fain *et al.*, 2018].

Stronger notions such as group fairness or proportional fairness [Kelly, 1997] remain unexplored for PB, although our recent work explores proportional fairness for randomized single-winner selection [Ebadian *et al.*, 2022a], which can be viewed as a special case of PB.

[2]The exact rule is a subtle refinement of this.

[3]Informally, violation of fairness must be due to a single good.

[4]It is reasonable to demand that tax dollars be spent equitably, so each resident is entitled an equal share of the available budget and any group of residents can pool their entitlements together.

### 3.3 Conference Peer Review

Large conferences such as IJCAI, AAAI, and NeurIPS invite submissions from many subcommunities, and the authors of the submissions often serve as reviewers too. An advantage touted by such conferences is that their large reviewer pool can enable finding suitable reviewers with diverse expertise for many submissions. But the algorithms used to match reviewers to submissions can also mistreat a community by assigning its submissions to less qualified external reviewers, incentivizing the community to leave and set up its own conference, in which its submissions can be reviewed by more qualified reviewers from within the community.

The core turns out to be aptly-suited to addressing this issue. Defining a group of researcher's "entitled share of resources" as the reviewing capacity they offer, the core demands that no group be able to find a better reviewing assignment for their submissions using their entitled share of resources, thus preventing groups from breaking off and setting up their own conferences. Here, in addition to guaranteeing fairness, the core also contributes stability to the system.

We prove that, at least in a limited model with single-author submissions and mild conditions on authors' preferences over reviewers, a reviewing assignment in the core that satisfies load and conflict avoidance constraints always exists, and can be computed efficiently [Aziz *et al.*, 2023]. It remains to be seen whether this can be extended to more realistic models.

### 3.4 Locating Public Facilities

Another interesting application of the core is public facility location via clustering, where the goal is to choose $k$ locations for building public facilities that would fairly serve $n$ people in a geographical region (represented as points in metric space). This application admits a natural definition of entitlements: each group of $n/k$ people is entitled to one public facility, each group of $2n/k$ people is entitled to two public facilities, and so on. Then, the core demands that, for any $\ell$, no group of at least $\ell n/k$ people be able to find $\ell$ locations such that each member is closer to some one of the $\ell$ new locations than to any of the $k$ locations chosen by the algorithm.

Chen *et al.* [2019] prove that there are instances with no clustering in the core, but a $(1+\sqrt{2})$-approximate core clustering exists for any metric. We prove that for the common case of $L^2$ distance over a Euclidean space $\mathbb{R}^t$, the approximation factor can be improved to 2 [Micha and Shah, 2020].

### 3.5 Allocating Educational/Computing Resources

While the maximum Nash social welfare solution works well for additive preferences, our prior work shows that another solution called the leximin solution, a refinement of Rawls' egalitarian criterion, is more appealing for other preferences.

We consider the problem of allocating unused classrooms in public schools to local charter schools [Kurokawa *et al.*, 2018], a process mandated by California's Proposition 39. After extensive discussions with public school districts in California, we observe that a particular style of dichotomous preferences best models the needs of the charter schools, and prove that for such preferences the (randomized) leximin solution satisfies proportionality and envy-freeness, while also ensuring that no group of charter schools can manipulate the process to their advantage.

In other work, we consider the problem of allocating computing resources such as CPU, RAM, and network bandwidth among processes in a cluster environment [Parkes *et al.*, 2015], where the class of Leontief preferences is more suitable. Noticing that the *Dominant Resource Fairness* (DRF) algorithm implemented in the popular distributed computing framework Apache Hadoop essentially implements the egalitarian criterion, we prove that using its leximin refinement instead would again satisfy proportionality, envy-freeness, and resistance to strategic manipulations by groups of users.

### 3.6 Recommender Systems

While the above applications model one-sided markets in which resources are allocated to agents with preferences, these versatile fairness notions can also be extended to *two-sided markets*, in which agents on two sides of a market are matched to each other and agents on each side have preferences over those on the other side. This models, e.g., matching consumers to products (and, in turn, to producers) in recommender systems, matching students to schools or medical residents to hospitals, or matching refugees to pro bono service providers. A natural goal in this case is to ensure fairness among agents on each side of the market simultaneously.

In recent work [Freeman *et al.*, 2021], we propose such an extension of envy-freeness, dubbed *double envy-freeness*, for many-to-many two-sided matching markets, and prove that, at least when agents on each side agree over a ranking of the agents on the other side (but may disagree in intensities), an "up to one"-style relaxation of double envy-freeness can be guaranteed in polynomial time. A natural direction for the future is to design algorithms satisfying two-sided versions of other fairness notions from Section 2.

### 3.7 Classification

The groupwise fairness notions of the core and group envy-freeness *strengthen* the individual fairness notions of proportionality and envy-freeness, respectively. But in some applications, individual fairness may already be a bar too high. For example, when using a classification algorithm to determine which defendants to grant bail or which loan applications to approve, certain individuals would inevitably be left empty-handed. Traditional ML fairness notions such as statistical parity and equalized odds thus impose fairness with respect to specific groups only *on average over their members*.

Building on prior work [Balcan *et al.*, 2019], we show that envy-freeness can also be extended in this fashion to prevent envy between groups on average, this new fairness notion subsumes notions like statistical parity and equalized odds as special cases, and it generalizes well from a small sample to an underlying population [Hossain *et al.*, 2020].

The key advantage of such preference-based fairness notions is that they naturally apply to non-binary decisions. Thus, they prevent classifiers from "gerrymandering" fairness guarantees by, e.g., granting bail to black and white defendants at the same rate, thus satisfying statistical parity, but discriminating in the average bail amounts assigned to the two groups [Arnold *et al.*, 2018].

## 4 The Quest for an Overarching Theory

The grand vision here is to develop an overarching theory for algorithmic fairness, which can be used to pick or formulate the best-suited notions of fairness for any application at hand and design efficient algorithms provably satisfying them.

One should not mistake this for a quest for an "ultimate notion of fairness" that supersedes all others. Such an *assertive* notion, which certifies that no unfairness would remain upon its satisfaction, does not yet exist as this requires a deep understanding of long-term impacts of algorithmic decisions. Instead, we have *preventive* notions, each of which prevent the algorithm from imposing a specific type of harm. Since there exist many types of harms, we need to aim for (approximately) satisfying multiple fairness notions simultaneously.

Several milestones must be conquered along the long road to achieving this grand vision. Surely, we need to identify novel types of harms and formulate fairness notions which prevent them. As we design provably fair algorithms, we also need to understand their generalization guarantees [Balcan *et al.*, 2019; Micha and Shah, 2020; Hossain *et al.*, 2020] and the price of imposing fairness in terms of other objectives of interest [Barman *et al.*, 2020; Hossain *et al.*, 2020].

But most importantly, we need to expand the study of fairness to new domains. For example, these notions are applicable to online problems where agents and/or resources arrive and depart over time [Kash *et al.*, 2014; Hosseini *et al.*, 2023; Banerjee *et al.*, 2023]. Another domain of significant recent interest is political redistricting or gerrymandering [Borodin *et al.*, 2018; Borodin *et al.*, 2022], for which the core seems aptly-suited to define a fair redistricting [Benade *et al.*, 2023]. Finally, the core can also be used to automate moral decision-making [Noothigattu *et al.*, 2018; Lee *et al.*, 2019; Conitzer *et al.*, 2017b]. For example, in the classical trolley problem, where one must decide whether to divert a trolley to save five people at the expense of one, the core demands that the diversion happen with probability exactly $5/6$, the collective entitlement of the five people. This is an ethical position worth exploring, but the real strength of the core lies in generalizing this basic idea to arbitrarily complex ethical scenarios.

To that end, let me reiterate that notions like the core are appealing because they simultaneously provide guarantees for all possible groups without needing to prespecify them based on fixed attributes such as race, gender, or sexual orientation. Most other fairness notions apply either only to prespecified groups (e.g., statistical parity, equalized odds [Hardt *et al.*, 2016], and calibration [Dawid, 1982; Pleiss *et al.*, 2017]) or to exponentially many groups (e.g., multicalibration [Hébert-Johnson *et al.*, 2018] and subgroup fairness [Kearns *et al.*, 2018]), but not to *all possible* groups.

I conclude with a call to arms for exploring preference-based fairness notions in a wide range of domains and understanding how well they align with human perceptions of fairness [Saxena *et al.*, 2019; Lee *et al.*, 2019; Gal *et al.*, 2017].

## Acknowledgments

I acknowledge the support of an NSERC Discovery Grant. I am grateful to all my collaborators, who were the driving force behind the work reported in this article.

## References


[Akrami *et al.*, 2023] Hannaneh Akrami, Noga Alon, Bhaskar Ray Chaudhury, Jugal Garg, Kurt Mehlhorn, and Ruta Mehta. EFX: A simpler approach and an (almost) optimal guarantee via rainbow cycle number. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*, 2023. Forthcoming.

[Amanatidis *et al.*, 2021] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, Alexandros Hollender, and Alexandros A Voudouris. Maximum Nash welfare and other stories about EFX. *Theoretical Computer Science*, 863:69–85, 2021.

[Amanatidis *et al.*, 2022] Georgios Amanatidis, Haris Aziz, Georgios Birmpas, Aris Filos-Ratsikas, Bo Li, Hervé Moulin, Alexandros A Voudouris, and Xiaowei Wu. Fair division of indivisible goods: A survey. arXiv:2208.08782, 2022.

[Arnold *et al.*, 2018] David Arnold, Will Dobbie, and Crystal S Yang. Racial bias in bail decisions. *The Quarterly Journal of Economics*, 133(4):1885–1932, 2018.

[Aziz and Shah, 2021] 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.

[Aziz *et al.*, 2023] Haris Aziz, Evi Micha, and Nisarg Shah. Group fairness in peer review. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2023. Forthcoming.

[Balcan *et al.*, 2019] Maria-Florina Balcan, Travis Dick, Ritesh Noothigattu, and Ariel D Procaccia. Envy-free classification. In *Proceedings of the 33rd Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 1238–1248, 2019.

[Banerjee *et al.*, 2023] Siddhartha Banerjee, Vasilis Gkatzelis, Safwan Hossain, Billy Jin, Evi Micha, and Nisarg Shah. Proportionally fair online allocation of public goods with predictions. In *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI)*, 2023. Forthcoming.

[Barman *et al.*, 2020] Siddharth Barman, Umang Bhaskar, and Nisarg Shah. Optimal bounds on the price of fairness for indivisible goods. In *Proceedings of the 16th Conference on Web and Internet Economics (WINE)*, pages 356–369, 2020.

[Benade *et al.*, 2023] Gerdus Benade, Ariel D Procaccia, and Jamie Tucker-Foltz. You can have your cake and redistrict it too. In *Proceedings of the 24th ACM Conference on Economics and Computation (EC)*, 2023. Forthcoming.

[Berliant *et al.*, 1992] Marcus Berliant, William Thomson, and Karl Dunz. On the fair division of a heterogeneous commodity. *Journal of Mathematical Economics*, 21(3):201–216, 1992.

[Borodin *et al.*, 2018] Allan Borodin, Omer Lev, Nisarg Shah, and Tyrone Strangway. Big city vs. the great outdoors: Voter distribution and how it affects gerrymandering. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 98–104, 2018.

[Borodin *et al.*, 2022] Allan Borodin, Omer Lev, Nisarg Shah, and Tyrone Strangway. Little house (seat) on the prairie: Compactness, gerrymandering, and population distribution. In *Proceedings of the 21st International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 154–162, 2022.

[Caragiannis *et al.*, 2019] 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.

[Caragiannis *et al.*, 2022] Ioannis Caragiannis, Evi Micha, and Nisarg Shah. A little charity guarantees fair connected graph partitioning. In *Proceedings of the 36th AAAI Conference on Artificial Intelligence (AAAI)*, pages 4908–4916, 2022.

[Chaudhury *et al.*, 2020] Bhaskar Ray Chaudhury, Jugal Garg, and Kurt Mehlhorn. EFX exists for three agents. In *Proceedings of the 21st ACM Conference on Economics and Computation (EC)*, pages 1–19, 2020.

[Chen *et al.*, 2019] 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.

[Conitzer *et al.*, 2017a] 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.

[Conitzer *et al.*, 2017b] Vincent Conitzer, Walter Sinnott-Armstrong, Jana Schaich Borg, Yuan Deng, and Max Kramer. Moral decision making frameworks for artificial intelligence. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 4831–4835, 2017.

[Conitzer *et al.*, 2019] 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.

[Dawid, 1982] A. Philip Dawid. The well-calibrated Bayesian. *Journal of the American Statistical Association*, 77(379):605–610, 1982.

[Ebadian *et al.*, 2022a] Soroush Ebadian, Anson Kahng, Dominik Peters, and Nisarg Shah. Optimized distortion and proportional fairness in voting. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*, pages 563–600, 2022.

[Ebadian *et al.*, 2022b] Soroush Ebadian, Dominik Peters, and Nisarg Shah. How to fairly allocate easy and difficult chores. In *Proceedings of the 21st International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 372–380, 2022.

[Fain *et al.*, 2018] 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.

[Foley, 1967] Duncan Karl Foley. Resource allocation and the public sector. *Yale Economics Essays*, 7:45–98, 1967.

[Freeman *et al.*, 2020] 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.

[Freeman *et al.*, 2021] Rupert Freeman, Evi Micha, and Nisarg Shah. Two-sided matching meets fair division. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 203–209, 2021.

[Gal *et al.*, 2017] Ya'akov Gal, Moshe Mash, Ariel D Procaccia, and Yair Zick. Which is the fairest (rent division) of them all? *Journal of the ACM*, 64(6): article 39, 2017.

[Gamow and Stern, 1958] George Gamow and Marvin Stern. *Puzzle-Math*. Viking, 1958.

[Gillies, 1953] Donald Bruce Gillies. *Some theorems on n-person games*. Princeton University, 1953.

[Hardt *et al.*, 2016] Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. In *Proceedings of the 30th Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 3315–3323, 2016.

[Hébert-Johnson *et al.*, 2018] Ursula Hébert-Johnson, Michael Kim, Omer Reingold, and Guy Rothblum. Multicalibration: Calibration for the (computationally-identifiable) masses. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, pages 1939–1948, 2018.

[Hossain *et al.*, 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 (TheWebConf)*, pages 1559–1569, 2020.

[Hossain *et al.*, 2021] Safwan Hossain, Evi Micha, and Nisarg Shah. Fair algorithms for multi-agent multi-armed bandits. In *Proceedings of the 34th Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 24005–24017, 2021.

[Hosseini *et al.*, 2023] Hadi Hosseini, Zhiyi Huang, Ayumi Igarashi, and Nisarg Shah. Class fairness in online matching. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 2023. Forthcoming.

[Kash *et al.*, 2014] Ian Kash, Ariel D Procaccia, and Nisarg Shah. No agent left behind: Dynamic fair division of multiple resources. *Journal of Artificial Intelligence Research*, 51:579–603, 2014.

[Kearns *et al.*, 2018] Michael Kearns, Seth Neel, Aaron Roth, and Zhiwei Steven Wu. Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, pages 2564–2572, 2018.

[Kelly, 1997] Frank Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8:33–37, 1997.

[Kurokawa *et al.*, 2018] David Kurokawa, Ariel D. Procaccia, and Nisarg Shah. Leximin allocations in the real world. *ACM Transactions on Economics and Computation*, 6(3–4): Article 11, 2018.

[Lee *et al.*, 2019] Min Kyung Lee, Daniel Kusbit, Anson Kahng, Ji Tae Kim, Xinran Yuan, Allissa Chan, Daniel See, Ritesh Noothigattu, Siheon Lee, Alexandros Psomas, and Ariel D Procaccia. WeBuildAI: Participatory framework for fair and efficient algorithmic governance. In *Proceedings of the 22nd ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*, article 181, 2019.

[Li *et al.*, 2023] Lily Li, Evi Micha, Aleksandar Nikolov, and Nisarg Shah. Partitioning friends fairly. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI)*, 2023. Forthcoming.

[Micha and Shah, 2020] 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.

[Noothigattu *et al.*, 2018] Ritesh Noothigattu, Snehalkumar Gaikwad, Edmond Awad, Sohan Dsouza, Iyad Rahwan, Pradeep Ravikumar, and Ariel Procaccia. A voting-based system for ethical decision making. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, pages 1587–1594, 2018.

[Parkes *et al.*, 2015] David C. Parkes, Ariel D. Procaccia, and Nisarg Shah. Beyond Dominant Resource Fairness: Extensions, limitations, and indivisibilities. *ACM Transactions on Economics and Computation*, 3(1): Article 3, 2015.

[Pleiss *et al.*, 2017] Geoff Pleiss, Manish Raghavan, Felix Wu, Jon Kleinberg, and Kilian Q Weinberger. On fairness and calibration. In *Proceedings of the 31st Annual Conference on Neural Information Processing Systems (NeurIPS)*, pages 5680–5689, 2017.

[Procaccia, 2020] Ariel D Procaccia. Technical perspective: An answer to fair division's most enigmatic question. *Communications of the ACM*, 63(4):118–118, 2020.

[Saxena *et al.*, 2019] Nripsuta Ani Saxena, Karen Huang, Evan DeFilippis, Goran Radanovic, David C Parkes, and Yang Liu. How do fairness definitions fare? Examining public attitudes towards algorithmic definitions of fairness. In *Proceedings of the 2nd AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society (AIES)*, pages 99–106, 2019.

[Steinhaus, 1948] Hugo Steinhaus. The problem of fair division. *Econometrica*, 16:101–104, 1948.

[Varian, 1974] Hal R Varian. Equity, envy and efficiency. *Journal of Economic Theory*, 9:63–91, 1974.