
Robust Allocations with Diversity Constraints

Zeyu Shen
Duke University
Durham NC 27708-0129
zeyu.shen@duke.edu

Lodewijk Gelauff **Ashish Goel**
Management Science and Engineering
Stanford University, Stanford CA 94305
{lodewijk,ashishg}@stanford.edu

Aleksandra Korolova
Department of Computer Science
University of Southern California
korolova@usc.edu

Kamesh Munagala
Department of Computer Science
Duke University, Durham NC 27708-0129
kamesh@cs.duke.edu

Abstract

We consider the problem of allocating divisible items among multiple agents, and consider the setting where any agent is allowed to introduce *diversity constraints* on the items they are allocated. We motivate this via settings where the items themselves correspond to user ad slots or task workers with attributes such as race and gender on which the principal seeks to achieve demographic parity. We consider the following question: When an agent expresses diversity constraints into an allocation rule, is the allocation of other agents hurt significantly? If this happens, the cost of introducing such constraints is disproportionately borne by agents who do not benefit from diversity. We codify this via two desiderata capturing *robustness*. These are *no negative externality* – other agents are not hurt – and *monotonicity* – the agent enforcing the constraint does not see a large increase in value. We show in a formal sense that the Nash Welfare rule that maximizes product of agent values is *uniquely* positioned to be robust when diversity constraints are introduced, while almost all other natural allocation rules fail this criterion. We also show that the guarantees achieved by Nash Welfare are nearly optimal within a widely studied class of allocation rules. We finally perform an empirical simulation on real-world data that models ad allocations to show that this gap between Nash Welfare and other rules persists in the wild.

1 Introduction

Allocating heterogeneous divisible items (or resources) among agents with different values over these items is a central problem in economics, with literature dating back many decades [48, 10, 42, 49, 5]. These problems have gained importance in computer science and machine learning due to their applications in, among other things, large-scale online advertising [38, 7, 31, 6, 26, 8], resource allocation in datacenters [30] and sharing economy platforms.

Given their wide applicability, there has been a large body of work on elucidating desirable properties of such allocations. One desirable property is *fairness* or *equity* in the value the allocation provides to different agents. Though there is no universally agreed-upon notion of fairness, one appealing notion [49] defines fair allocations as those that are both Pareto-optimal on the value they provide to the agents as well as *envy-free* [48], meaning that no agent should derive more value from another agent’s allocation. A less restrictive notion is the so-called Pigou-Dalton principle [19, 40] or *welfarism*, which states that given any fixed total value of the agents, an allocation should prefer distributing these values so that there is no potential transfer of value from an agent with larger value (the “rich”) to one with smaller value (the “poor”). This loosely translates to allocations that optimize

a (weakly) concave social welfare function over agents’ values. We call rules that optimize separable, symmetric, concave functions of agent values as *welfarist* allocation rules¹

Due to their simplicity of implementation and ease of understanding, most allocation rules implemented in practice are welfarist rules. Indeed, even when allocations are supported by prices, for instance in revenue-maximizing ad auctions, by Myerson’s celebrated transformation [41, 24] and its generalizations, these can be viewed as optimizing a linear “virtual welfare” function over allocations [12, 14]. Further, envy-freeness can also be implemented by the welfarist *Nash Welfare* rule [42, 23, 49] that optimizes the product of agents’ values, and which we will discuss extensively.

1.1 Allocations with Diversity

In this paper, we consider a different desirable facet of such allocations – diversity. The notion of fairness presented above attempts to spread value fairly across agents. Suppose, in addition, the items being allocated were also associated with individuals. For instance, consider the following scenario motivated by online advertising – the items are ad view slots in an advertising ecosystem, each item labeled with the attributes, such as race, age, gender, of the viewing individual on a social media site. The agent represents an advertiser who has different preferences over the viewing individuals. The agent may not only seek an allocation of slots that maximizes its own value, but may also want these slots to not be all views by any particular race. In other words, the agent seeks an allocation that is also diverse on the attributes of the corresponding slots. Similarly, in a team formation platform, an agent seeking such a team may value diversity in attributes such as gender and race of the team members in addition to competence. Diversity has become an increasingly important consideration for eliminating bias in allocation platforms where the items map to humans. Allowing agents to specify some form of diversity constraints on allocations are one way for achieving this goal [50, 4, 13, 43].

The simplest diversity constraint we consider is the case where the agent wants items with different attributes (such as race and gender) in fixed proportions, which models the diversity desiderata in the advertising and team formation settings considered above. We term these “proportionality constraints”; see Section 2.2 for a formal definition. More generally, such constraints capture *complementarity* in the allocation, where an agent can say “I want at least as much of this attribute as that attribute.” A different application for such constraints is in machine allocation in data centers, where a client could request machines in different geographic regions in certain proportions [46]. For simplicity, we call these constraints *diversity constraints*, though it should be kept in mind that they model complementarity in general and can easily arise in settings beyond achieving diversity.

In Section 2.2 we model diversity constraints imposed by agent i as simply a convex polytope \mathcal{P}_i such that the empty allocation $\vec{0} \in \mathcal{P}_i$. First, this means such constraints preserve feasibility of the allocation. More importantly, these constraints, including proportionality constraints considered above, can have both positive and negative coefficients, and this leads to interesting and non-obvious behavior of welfarist allocation rules. This forms the focus of this paper.

Externality in Diverse Allocations. One natural solution to incorporating diversity is for the allocation platform to ask agents for their constraints and add them to the welfarist optimization problem. Indeed, recent work on large scale online allocation problems [20, 21, 26] shows that even in the presence of *arbitrary* convex constraints on the allocations that can run across time, the resulting problem can be viewed as online stochastic convex programming and admits efficient approximation algorithms [3, 8]. Similarly, recent works on auctions with diversity constraints [14, 16, 28] also follow this route and simply add these constraints to an optimization procedure.

In this paper, our main focus is to understand how adding such constraints affects the *outcome* of the allocation. More precisely, for any welfarist allocation rule, when an agent imposes diversity constraints, this changes the allocation of not just this agent, but all other agents as well. Naturally, seeking a diverse allocation will incur a cost in terms of diminished total value of all the agents. This cost is, of course, counterbalanced by the benefit of diversity to the agent seeking diversity (which does not explicitly appear in our model). Since the agent seeking diversity is the one benefiting from the constraints, it would be desirable from both the other agents’ and the platform’s perspective that this agent should be the one bearing the majority of this cost, while the values for other agents should

¹We borrow this term from [44], though the concept itself is classic [37, 39] (See Section 2.1 for formal definitions.)

not be significantly impacted. In other words, there should be *little negative externality* on other agents. We present this desideratum formally in Section 2.3.

It is not *a priori* clear whether the welfarist allocation rules are robust in the above sense and thus the extent to which imposing diversity constraints, even by one agent, may impact other agents. In fact, this lack of clarity may be a major reason for a platform to hesitate to implement functionality enabling agents to specify such diversity constraints, a significant shortcoming for applications such as team selection and online advertising. Thus the focus of our work is on quantifying the potential negative externality and on developing recommendations for finding robust allocations.

1.2 Our Results

To start with, one could reasonably hope that welfarist allocation rules that attempt to find equitable or fair allocations across agents would be naturally robust to the addition of diversity constraints, in the sense that to a good approximation, the negative externality to other agents' values is bounded. *Our first analytical result shows that the above intuition is false.* In Section 3, we show that *almost all* welfarist allocation rules – regardless of how equitably or fairly they distribute value among agents – are not robust in the sense that the negative externality on other agents that stems from imposing even a single diversity constraint is unbounded (see Theorem 1.) On the positive side, we show that the Nash Welfare objective [42, 23], which optimizes the product of agents' values, is the exception in that it achieves negative externality that is bounded above by an absolute constant factor even when a constant number of agents impose arbitrary diversity constraints. Further, this constant bound is nearly optimal for any number of agents enforcing such constraints since any allocation rule satisfying the Pigou-Dalton principle also suffers similar negative externality. (Theorems 2 and 3.)

We then consider how an agent's imposing diversity constraints affects its own value. We define an allocation rule to be *monotone* if an agent's imposing diversity constraints weakly reduces her own value, and is monotone to a constant factor if the value increases by at most a constant factor. In Section 4, we show such monotonicity is correlated with how egalitarian the allocation rule is – the more egalitarian (or concave) a rule is, the closer to monotone it is (see Theorem 4.) Further, the Nash Welfare allocation rule is monotone to an absolute constant factor.

In Section 5, we use real-world datasets to model ad allocation. For a budget-capped valuation function that is widely studied in this setting, we compare the negative externality induced by various allocation rules. In this setting, the social welfare maximizing rule models first price auctions with budget constrained advertisers, and we empirically show that it indeed suffers large externality and non-monotonicity. However, the Nash Welfare rule suffers very small externality and is monotone, hence making a case for that our theoretical insights will apply in the wild as well.

We summarize the upper and lower bounds on the amount of negative externality and non-monotonicity for various allocation rules in Section 2.3. Taken together, our main contribution is therefore to show that the Nash Welfare objective is uniquely positioned among a large class of welfarist allocation rules to be robust, achieving both no negative externality and monotonicity to within small constant factors. Our results also suggest that care must be exercised by the allocation platform when allowing agents to express diversity or complementarity constraints. If these are directly encoded into the optimization like in [14, 21, 26, 3, 6], they may create second-order unfairness where the resulting decrease in value is borne by agents not directly benefiting from the diversity.

1.3 Related Work

Our work studies robustness of allocations when users can express diversity constraints. This can be equivalently viewed as an agent specifying a different utility function. For instance, a proportionality constraint corresponds to a user expressing a *Leontief* utility function over the allocation instead of linear utility. In that sense, our results can be equivalently viewed as studying the robustness to agents reporting utility functions that capture diversity. Robustness in our sense has been widely studied when agents can arbitrarily misreport utility functions. In this context, monotonicity is simply strategy-proofness, and non-negative externality maps to non-bossiness [47]. The main result in [47] shows that the only allocation rules that satisfy these properties are variants of serial dictatorships, where agents go one after the other and choose their optimal allocation, and these rules are incompatible with Pareto-optimality and welfarist rules. The only exception is settings resembling matchings [45, 34], which are much more restrictive than the setting we study.

Note however that a diversity constraint is a very specific type of “misreport” – given any allocation, the new utility function is necessarily weakly smaller in value, and this is crucial to all our results. Monotonicity can now be viewed as strategy proofness with such misreports. But now, rules that are not strategyproof in general can easily be monotone. Indeed, we show in Theorem 4 that the MMF rule is actually monotone for general constraints, Pareto-optimal, and welfarist, but is trivially very far from being strategyproof if arbitrary misreports are allowed. In the same vein, the non-negative externality property is much more specific than non-bossiness with general misreports, and it is a priori not obvious that is incompatible with welfarist rules (Theorem 3). One of our main results is to show that even with the specific type of utility “misreports” that diversity constraints imply, this property is incompatible to any factor with welfarist rules, except for the NW rule.

Several recent works have considered the question of computing solutions to discrete optimization problems, such as b -matchings [4], stable matching [43], clustering [17], ranking [15, 29], voting [13] and packing [18], when the items belong to groups that must be allocated fairly in the resulting solution. For instance, each cluster in the clustering solution must be assigned a balanced cohort of red and blue points [17]. Our work tackles a normative question that arises in this space – assuming only a subset of agents care about diversity, what do these constraints mean for the allocations to other agents? Extending our work to discrete optimization problems is an interesting open question.

For lack of space, we present a detailed comparison with other notions in Economics and related work on ad auctions in the Supplementary material (Appendix A).

2 The Diverse Allocation Model

There are n agents and m divisible items, each with unit supply.² The goal of the platform is to compute an allocation \vec{x} , where x_{ij} represents the fraction of item j that is assigned to agent i . Each item is assigned at most once fractionally, that is, $\sum_i x_{ij} \leq 1 \forall \text{ items } j$.

In the context of ad auctions, an item represents a type of viewer or keyword, available in a certain quantity. A fractional allocation would allocate the corresponding fraction of the ad slots of that type to the advertiser (or agent).

We assume that any agent i 's value for an allocation \vec{x}_i is linear. We denote the value as $V_i(\vec{x}_i) = \sum_j v_{ij}x_{ij}$, where v_{ij} is the value of the agent for item j . Our positive results extend to the more general setting where V_i are arbitrary non-decreasing, continuous, concave functions, while our negative results apply even to the linear case.

2.1 Welfarist Allocation Rules

We will consider a general subclass of allocation rules called *welfarist* allocation rules. These optimize $\sum_i f(V_i(\vec{x}_i))$, where f is a monotonically non-decreasing, continuous and concave function.

Formally, given a choice of f , the welfarist optimization problem is as follows:

$$\text{Max } \sum_i f(V_i) \quad \text{s.t.} \quad \vec{V}_i \in \left\{ V_i = \sum_j v_{ij}x_{ij} \forall i; \sum_i x_{ij} \leq 1 \forall j; x_{ij} \geq 0 \forall i, j \right\}$$

Typically, the more concave f is, the more egalitarian across agents is the resulting allocation. The main class of rules we consider are γ -**fair** rules [39], defined as:

$$f(y) = \frac{y^\gamma}{\gamma}, \quad \gamma \in (-\infty, 1]$$

At $\gamma = 0$, this rule is defined by its limit as $\gamma \rightarrow 0$, so that $f(y) = \ln y$.

As special cases, this yields well-known rules in increasing order of how concave f is, in a sense capturing increasing fairness in the resulting allocations.

Social Welfare (SW). $f(y) = y$, corresponding to $\gamma = 1$.

Nash Welfare (NW). As $\gamma \rightarrow 0$, we obtain $f(y) = \ln y$.

²The model trivially generalizes to the setting with arbitrary but known supply by scaling the values.

Max-min Fairness (MMF). This computes an allocation with \vec{V} such that if $V_i < V_j$, there is no other allocation where agent i 's value is larger than V_i and j 's value is smaller than V_j . This corresponds to $\gamma \rightarrow -\infty$.

Among these functions, SW is the least egalitarian, while MMF is the most. We now note some features of welfarist allocation rules that we will use later.

Pareto-optimality (PO). The optimal allocation \vec{x} is such that there is no other allocation \vec{y} such that $V_i(\vec{y}_i) \geq V_i(\vec{x}_i)$ for all agents i , with one inequality strict.

Pigou-Dalton (PD) Principle [19][40]. Consider any two allocations with values \vec{V} and \vec{V}' where $\sum_i V_i = \sum_i V'_i$. Fix two agents k and ℓ and suppose $V_j = V'_j$ for all $j \neq k, \ell$, while $V_k < V'_k < V_\ell$ and $V_\ell < V'_\ell < V_k$. Then the allocation rule weakly prefers \vec{V}' to \vec{V} .

Anonymity. We finally say that a rule is *anonymous* if the allocation only depends on the revealed vector of values, and not on the identity of the agents. In particular, agents that reveal identical values receive identical allocations.

Note that all welfarist rules described above satisfy (PO), (PD), and anonymity. The NW rule has additional desirable properties. First, it is *scale invariant*: If an agent i scales all her values v_i by any factor α with other agents remaining the same, the optimal allocation does not change. Next, the allocation is *proportional*: Given n agents, the value obtained by any agent in the NW allocation is always at least $1/n$ of the value it could have obtained had it been the only agent in the system.

2.2 Specifying Diversity

We assume an agent i can be interested in receiving a diverse mix of items, and can specify diversity as a *constraint set* \mathcal{P}_i on the allocation \vec{x}_i it receives. This constraint set is added to the welfarist optimization problem described above. We assume \mathcal{P}_i is convex and that $\vec{0} \in \mathcal{P}_i$. As mentioned before, such constraints are motivated by settings where the items have attributes related to gender, race, income, etc., and an agent can be interested in obtaining a balanced mix of items, as opposed to a welfare maximizing set of items. We consider two types of constraints:

Proportionality Constraint. Here, agent i specifies a partition $S_{i1}, S_{i2}, \dots, S_{ik_i}$ of a subset T of the items, along with proportions $\alpha_{i1}, \alpha_{i2}, \dots, \alpha_{ik_i}$ where $\sum_{r=1}^{k_i} \alpha_{ir} \leq 1$. It seeks allocations from each S_{ir} in proportion to α_{ir} . In other words, the constraint set \mathcal{P}_i is:

$$\sum_{j \in S_{ir}} x_{ij} = \alpha_{ir} \sum_j x_{ij} \quad \forall r \in \{1, 2, \dots, k_i\}$$

We can also define an ϵ -approximate notion of proportionality, where the right hand side of the above equality is constrained to be within $(1 \pm \epsilon)$ factor of the left hand side. In our constructions, we will only consider exact proportionality.

General Constraint. There is an arbitrary convex set \mathcal{P}_i , with $\vec{0} \in \mathcal{P}_i$. The agent needs $\vec{x}_i \in \mathcal{P}_i$. This is more general than the proportionality setting, and could for instance capture multiple proportionality constraints over different partitions of the items. Further, the unconstrained setting where \mathcal{P}_i is all possible allocations is also a special case.

We call an agent who does not specify any constraints an *unconstrained* agent. A constrained agent can specify a general constraint set, and this includes being unconstrained. Our positive results (Theorems 2 and 4) hold in the most general setting with any number of agents who are initially constrained, while the specific agent switches from being unconstrained to expressing a general constraint. Note that this captures as a special case, the setting mentioned before where all agents are initially unconstrained, and one agent switches to expressing a general constraint. Our impossibility results (Theorems 1, 3 and 4, and Corollaries 1 and 3) on the other hand hold even in the special case where there are two agents who are unconstrained and one agent switches to expressing a single proportionality constraint.

2.3 Robustness Desiderata for Allocations with Diversity

The high level question we consider is: For welfarist allocation rules, does providing the ability for an agent to express diversity constraints lead to undesirable externality in the resulting allocations?

We study two kinds of externalities: Whether this additional functionality can hurt the agents who do not use this functionality; and whether an agent can gain by misrepresenting themselves as diversity constrained when they are unconstrained in reality. Formally, we consider two natural desiderata for the allocation rule.

Non-negative Externality. (NNE) Suppose \vec{x} is the allocation when agent i is unconstrained. Suppose the agent expresses a diversity constraint \mathcal{P}_i and the new allocation with \mathcal{P}_i included in the optimization is \vec{y} . Then for any $\ell \neq i$, we should have $V_\ell(\vec{y}_\ell) \geq V_\ell(\vec{x}_\ell)$. In other words, if agent i expresses a diversity constraint, it should not create negative externality to other agents. For a parameter $q \in [0, 1]$, we say that an allocation rule is q -NNE if $V_\ell(\vec{y}_\ell) \geq qV_\ell(\vec{x}_\ell)$ for all $\ell \neq i$. For some of our results, this can be naturally generalized to k agents are initially unconstrained and simultaneously express arbitrary diversity constraints.

Monotonicity. (MON) Consider the same setting as above. We should have $V_i(\vec{x}_i) \geq V_i(\vec{y}_i)$. In other words, expressing a diversity constraint should not increase the linear value of the agent, since such an increase has the undesirable effect of an agent having the incentive to misreport their constraints. For a parameter $p \in [0, 1]$, we say that an allocation rule is p -MON if $V_i(\vec{x}_i) \geq pV_i(\vec{y}_i)$.

An allocation rule is (p, q) -robust if it is p -MON, and q -NNE, where $p, q \in (0, 1)$ are constants. As discussed before, robustness is desirable for practical implementation of diversity constraints.

Summary of Results. In the sequel, we will bound the (p, q) -robustness achievable for various welfarist allocation rules. In particular, we show that among the class of γ -Fair rules, only the Nash Welfare rule achieves constant $q \geq \frac{1}{4}$, that is, has bounded negative externality. All other rules have $q = 0$. Further, the bound attained by NW cannot be significantly improved: No rule satisfying (PO) and (PD) as defined in Section 2.1 can achieve $q > \frac{1}{2}$.

For monotonicity, we show that p increases with decreasing γ , that is, as the function f becomes more concave or fair. At one extreme, for MMF, we have $p = 1$, and at the other, for SW, we have $p = 0$. The Nash Welfare objective achieves $p = \frac{1}{2}$, and the bounds we obtain are tight for all $\gamma \in (-\infty, 1]$.

3 The Non-negative Externality (NNE) Condition

In this part, we will assume f is differentiable and study which functions f satisfy the q -NNE condition for some absolute constant $q > 0$. At first glance, one might intuit that the more concave or fair f is, the more likely it is to satisfy q -NNE. Surprisingly, we show this intuition is false. Our main result in this section is that f needs to have a very specific form that is unrelated to its fairness or concavity in order for the allocation rule to be q -NNE. In particular, we show that most functions that are not scale-invariant fail the NNE property regardless of how concave they are, while the scale-invariant NW allocation satisfies q -NNE for an absolute constant q .

3.1 An Impossibility Result

Define $g(y) = yf'(y)$. We will assume g is continuous. We need the following technical definition.

Definition 1. For $\delta \leq 1$, we say that f is δ -scaled if g is continuous and $\min_{x, y \geq 0} \frac{g(y)}{g(x)} = \delta$.

We show an impossibility result for q -NNE in the following theorem³.

Theorem 1. If f is δ -scaled, then it is not q -NNE for any $q > \delta$.

The above theorem rules out achieving q -NNE for any constant q for a wide range functions. In particular, we have:

Corollary 1. The γ -Fairness rule for any fixed $\gamma \neq 0$ does not satisfy q -NNE for any constant $q > 0$.

This rules out q -NNE for any constant $q > 0$ for SW and MMF. A similar corollary holds also for additional functions commonly considered in the literature. For the *Exponential* function $f(y) = 1 - e^{-\lambda y}$, note that $g(y) = \lambda y e^{-\lambda y}$, and $g(0) = 0$. For the *Smooth Nash Welfare* [27, 25] function $f(y) = \ln(1 + y)$, we have $g(y) = \frac{y}{y+1}$, which is increasing with $g(0) = 0$. Similarly, for $f(y) = \ln \ln(1 + y)$, we have $g(y) = \frac{y}{y+1} \frac{1}{\log(y+1)}$ is decreasing with $g(y) \rightarrow 0$ as $y \rightarrow \infty$. Therefore, these functions are only δ -scaled for $\delta \rightarrow 0$, and are hence not q -NNE for any $q > 0$.

³All omitted proofs are in the Supplementary material (Appendix B).

3.2 Externality of Nash Welfare

We now show quite surprisingly that in a sense, the converse of Theorem 1 is also true. Specifically, this theorem does not rule out achieving q -NNE for functions f that are δ -scaled for constant $\delta > 0$. We show that a subclass of these functions satisfy q -NNE for constant $q > 0$, and this subclass includes the NW allocation rule. Further, we show that any allocation rule that is Pareto-optimal and satisfies the Pigou-Dalton principle incurs an almost matching negative externality.

Theorem 2. *Allocation rules where $g(y) = yf'(y)$ is non-increasing and f is δ -scaled satisfy q -NNE for $q = \frac{\delta}{\delta+3}$.*

The next corollary follows by directly applying the above theorem on NW rule. The proof follows by observing that for NW, $f(y) = \ln y$, so that $g(y) = 1$ is a constant, so that it is δ -scaled for $\delta = 1$.

Corollary 2. *The NW allocation rule satisfies q -NNE for $q = \frac{1}{4}$.*

We note that there could be other functions f satisfying the preconditions of Theorem 2. For instance, $f(x) = 2 \ln x - \ln(1+x)$, which is a combination of Nash Welfare and Smooth Nash Welfare, is δ -scaled for $\delta = \frac{1}{2}$, so that it is q -NNE for $q = \frac{1}{7}$.

To complete the picture, we now show that no rule that satisfies Pareto-optimality (PO) and the Pigou-Dalton principle (PD) as defined in Section 2.1 can achieve q -NNE for $q > 1/2$.

Theorem 3. *No allocation rule that satisfies (PO) and (PD) satisfies q -NNE for any constant $q > \frac{1}{2}$.*

Observations and Extensions. Theorem 2 does not require $V_i(\vec{x}_i)$ to be a linear function of x_i . It holds as long as $V_i(\vec{x}_i)$ is any concave non-decreasing function as long as $V_i(\vec{0}) = 0$. The same holds for all the positive results we subsequently present, e.g., Theorem 4.

As an extension of Theorem 2, suppose k agents express diversity constraints. A naive application yields q -NNE for $q = \frac{\delta^k}{(\delta+3)^k}$, since each of the k agents expressing a diversity constraint decreases the value of another agent by a factor of at most $\frac{\delta}{\delta+3}$. However, we show that as long as a small subset of agents expresses diversity constraints, the externality for δ -scaled functions is bounded.

Corollary 3. *For any k , suppose a set S of k agents switch from being unconstrained to expressing general diversity constraints. Then, allocation rules where g is non-increasing and f is δ -scaled satisfy q -NNE for $q = \frac{\delta}{2k+\delta+1}$.*

Therefore, the NW allocation is q -NNE for $q = \frac{1}{2(k+1)}$. We show a matching lower bound below.

Corollary 4. *There are instances where any allocation satisfying (PO), (PD), and anonymity is not q -NNE for $q = \frac{1+\epsilon}{k+1+\epsilon}$, where $\epsilon > 0$ is any constant when each of k agents switches to expressing a proportionality constraint.*

4 The Monotonicity (MON) Condition

Unlike the case of NNE where almost all rules other than NW fail, for monotonicity, we show that increasing fairness or concavity of f is correlated with achieving p -MON for constant $p > 0$. This is surprising since one would expect that if a diversity constraint hurts another agent by an arbitrary amount, it should help the agent enforcing the constraint also by an arbitrary amount. However, we show that this intuition is false and MON and NNE are indeed very different properties.

For γ -Fairness, we now show nearly matching upper and lower bounds on p -MON as a function of γ , and show that p decreases with increasing γ .

Theorem 4. *For $\gamma \in (-\infty, 1]$, the γ -Fairness rule satisfies p -MON where p is the solution to $p^{1-\gamma} + p = 1$. Further, this bound on p is tight.*

The above theorem shows that p increases with decreasing γ , confirming the intuition that achieving p -MON is easier as the function becomes more concave. In particular, we have the following corollary by setting γ appropriately:

Corollary 5. *The following bounds on p -MON are tight: $p = 0$ for social welfare ($\gamma = 1$); $p = \frac{1}{2}$ for Nash welfare ($\gamma \rightarrow 0$), and $p = 1$ for max-min fairness ($\gamma \rightarrow -\infty$).*

Therefore, the MMF allocation rule satisfies monotonicity, combined with (PO) and (PD), while Theorem 3 rules out a corresponding result for 1-NNE.

5 Empirical Simulation

We now augment our theoretical analysis with a quantitative empirical demonstration of externality and monotonicity achieved by various allocation rules. We consider a budget-constrained valuation function that is widely used in ad allocations [38, 22, 20]. We generate instances from real-world datasets to mimic bidding data in ad allocation, and empirically compare the externality and monotonicity properties of various allocation rules. We show that even empirically, rules other than Nash Welfare incur significant negative externality, thus providing evidence that our results are not merely a worst-case analysis, but correspond to what may also happen in practice.

Datasets and Setup. We use two datasets. The UCI Adult dataset [1] tabulates census information such as geographic area, job, gender, and race for around 50,000 individuals. The Yahoo A3 dataset [2] contains bid information for advertisers for a set of ads.

For the Adult dataset, we consider the two genders as items, and the 14 job categories as agents. The value of a gender (item) to a job (agent) is set to be linearly and positively correlated with the number of people of that gender working in that job. We take the first 500 data points. For the Yahoo dataset, we consider the advertisements as items, and the advertisers as agents. The value of an advertisement to an advertiser is set as its bid on the advertisement.⁴ We take the first 10,000,000 data points. We arbitrarily take 20 advertisements, together with 6 advertisers who have bid on most of these advertisements. In both cases, the diversity constraint for an agent is the proportionality constraint that equalizes the allocation across items where this agent has non-zero values. (We exclude items with zero value since the agent is clearly not interested in this item.)

In order to make the valuation function mirror advertising applications, we use the budget-capped valuation function [38, 22], where for each agent i , its value is

$$V_i(\vec{x}_i) = \min \left(B_i, \sum_j v_{ij} x_{ij} \right). \quad (1)$$

where B_i is the budget of the agent. This function has also been termed *liquid welfare*, since it is an upper bound on the amount of welfare the platform can generate given budget constraints of the advertisers. In our simulation, we draw B_i independently and uniformly from $(0, T)$. For the Adult dataset, we set $T = 20$, and for the Yahoo dataset, we set $T = 10$. These are chosen commensurate with the magnitude of the values of the agents for items.

Simulation. We conduct three sets of simulations on each dataset and compare the allocation rules for γ -Fairness with $\gamma = 1$ (Social Welfare), $\gamma = 0.5$, $\gamma = 0.1$, $\gamma \rightarrow 0$ (Nash Welfare), and $\gamma = -1$ (approximating MMF).

In the first *single agent* simulation, a specific agent expresses the diversity constraint, requiring equal allocation of all items with non-zero values. For each allocation rule and each agent i , we record the largest q_i such that this rule satisfies q_i -NNE when this agent expresses a diversity constraint. We report $q_{\min} = \min_i q_i$ for each rule. In the second *double agent* simulation, two agents simultaneously express diversity constraints, each requiring equal allocations of all items with non-zero value. As before, we compute the minimum q value achieved over all pairs of agents.

For both these simulations, to prevent the negative externality from being dominated by tiny values, when computing the negative externality, we ignore agents whose values in the unconstrained case are less than 10% of their budget.⁵ We repeat these simulations 10 times with random seeds.

In the third *monotonicity* simulation, we again consider the case where a single agent expresses a diversity constraint, requiring its allocation of all items to be equal. For each allocation rule and agent i , we record the largest p_i such that this rule satisfies p_i -MON when this agent expresses a diversity constraint, and report $p_{\min} = \min_i p_i$ for each rule.

⁴In reality, the value is the bid times the CTR. Though our dataset does not have CTR information, we believe our results will extend to that case.

⁵Including these preserves performance of NW and γ -Fairness for $\gamma = 0.1$, while other rules only do worse.




Figure 1: The plots of $(1 - q_{\min})$ for Social Welfare (SW), Nash Welfare (NW), and γ -Fairness for $\gamma = 0.5, 0.1, -1$.

Results. In Figures [1a](#) and [1b](#), we present the box plots of $(1 - q_{\min})$ for the single and double agent simulations respectively. For both datasets, Nash Welfare and γ -Fairness for $\gamma = 0.1$ give the best performance with a loss in value below 10% in most cases, while $\gamma = 0.5$, SW, and $\gamma = -1$ are significantly worse, with negative externality at least twice as large. The exception is the good performance of $\gamma = -1$ on the Adult dataset. Moreover, in the double agent simulation, NW and γ -Fairness for $\gamma = 0.1$ perform much better than the worst-case bounds presented before.

For the monotonicity simulation, we find that $p_{\min} = 1$ for all rules except SW. In other words, these rules are monotone. For SW, we have $p_{\min} = 0.36$ for Yahoo and $p_{\min} = 0.94$ for Adult.⁶

Relevance to Ad Auctions. The Social Welfare maximizing rule under the valuation function from Eq [\(1\)](#) closely mirrors allocation rules used in ad auctions [\[38, 20, 26, 51\]](#). In particular, via standard LP duality, the SW allocation has the following structure: It is a first price auction where bid (or value) of an advertiser is scaled down by an advertiser-dependent parameter. Such an allocation rule mirrors the widely used smooth delivery allocation rules [\[51\]](#) for budget constrained advertisers. Further, rules that allocate adwords in an online fashion essentially compute an approximately SW allocation [\[38, 20\]](#). Our empirical results show that such a rule suffers large negative externality when advertisers are allowed to express diversity constraints, and this can be mitigated if the platform instead uses Nash Welfare. This complements a recent line of work in advertising on implementing Nash Welfare [\[8\]](#) and regularization [\[6\]](#) as approaches to mitigate unfairness; we show that these approaches mitigate negative impacts of diversity constraints as well.

6 Conclusion

The conceptual message of this paper is that incorporating diversity constraints into an allocation platform requires careful selection of the underlying optimization algorithm to prevent negative externality. One could of course wonder whether negative externality is necessarily a bad thing since it forces other agents to consider diversity themselves. However, our results should be viewed as saying that many allocation rules change values of other agents in an unpredictable or counterintuitive fashion. Since the agents typically are not symmetric in either values or what constraints they desire, the resulting externalities will also be asymmetric. This argues for minimizing such negative externality in the first place instead of using it as a disincentive tool.

Our work is just a first step towards understanding the robustness of allocation rules with diversity constraints, and we now present several open questions. In addition to tightening our bounds, it would be interesting to study pricing rules and auctions. Though we do not present details, these allocation rules suffer from similar drawbacks. For instance, truthful auctions [\[32\]](#) or online allocations [\[38\]](#) with budgets are not Pareto-optimal with diversity constraints, while rules that compute optimal auctions [\[41, 14\]](#) or market-clearing solutions [\[5\]](#) are not robust to any approximation even with quasi-linear utilities. We leave a deeper examination of rules with prices and budgets as an interesting open question. Finally, it would be interesting to study such externality in contexts other than allocation problems, for instance, discrete ML problems such as ranking or clustering.

⁶The performance of SW is unstable under different sets of random budgets, while the other four rules are stable. Thus, for the sake of exposition, we select a bad instance for SW to present here.

Acknowledgments and Disclosure of Funding

This work is supported by NSF awards CCF-1637397, CNS-1943584, CNS-1956435, CNS-1916153, CCF-2113798, ONR award N00014-19-1-2268, DARPA award FA8650-18-C-7880, and gifts from Google and Microsoft.

References

- [1] <https://archive.ics.uci.edu/ml/datasets/Adult>.
- [2] <https://webscope.sandbox.yahoo.com/catalog.php?datatype=a>.
- [3] S. Agrawal and N. R. Devanur. Fast algorithms for online stochastic convex programming. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, page 1405–1424, 2015.
- [4] F. Ahmed, J. P. Dickerson, and M. Fuge. Diverse weighted bipartite b-matching. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence, IJCAI’17*, page 35–41. AAAI Press, 2017.
- [5] K. J. Arrow and G. Debreu. Existence of an equilibrium for a competitive economy. *Econometrica*, 22(3):265–290, 1954.
- [6] S. Balseiro, H. Lu, and V. Mirrokni. Regularized online allocation problems: Fairness and beyond. *arXiv 2007.00514*, 2020.
- [7] S. R. Balseiro, J. Feldman, V. S. Mirrokni, and S. Muthukrishnan. Yield optimization of display advertising with ad exchange. *Management Science*, 60(12):2886–2907, 2014.
- [8] M. H. Bateni, Y. Chen, D. Ciocan, and V. Mirrokni. Fair resource allocation in a volatile marketplace. *Available at SSRN*, 2018.
- [9] W. C. Brainard and H. E. Scarf. How to compute equilibrium prices in 1891. *American Journal of Economics and Sociology*, 64(1):57–83, 2005.
- [10] S. J. Brams and A. D. Taylor. *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, 1996.
- [11] E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061 – 1103, 2011.
- [12] Y. Cai, C. Daskalakis, and S. M. Weinberg. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, pages 130–139, 2012.
- [13] L. E. Celis, L. Huang, and N. K. Vishnoi. Multiwinner voting with fairness constraints. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pages 144–151. International Joint Conferences on Artificial Intelligence Organization, 7 2018.
- [14] L. E. Celis, A. Mehta, and N. K. Vishnoi. Toward controlling discrimination in online ad auctions. In *Proceedings of the 36th International Conference on Machine Learning, ICML*, pages 4456–4465, 2019.
- [15] L. E. Celis, D. Straszak, and N. K. Vishnoi. Ranking with fairness constraints. *CoRR*, abs/1704.06840, 2017.
- [16] S. Chawla and M. Jagadeesan. Fairness in ad auctions through inverse proportionality. *CoRR*, abs/2003.13966, 2020.
- [17] F. Chierichetti, R. Kumar, S. Lattanzi, and S. Vassilvitskii. Fair clustering through fairlets. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017.
- [18] F. Chierichetti, R. Kumar, S. Lattanzi, and S. Vassilvitskii. Matroids, matchings, and fairness. In K. Chaudhuri and M. Sugiyama, editors, *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, volume 89 of *Proceedings of Machine Learning Research*, pages 2212–2220. PMLR, 16–18 Apr 2019.

- [19] H. Dalton. The measurement of the inequality of incomes. *The Economic Journal*, 30(119):348–361, 1920.
- [20] N. R. Devanur and T. P. Hayes. The adwords problem: Online keyword matching with budgeted bidders under random permutations. In *Proceedings of the 10th ACM Conference on Electronic Commerce*, EC '09, pages 71–78, 2009.
- [21] N. R. Devanur, K. Jain, B. Sivan, and C. A. Wilkens. Near optimal online algorithms and fast approximation algorithms for resource allocation problems. In *Proc. 12th ACM EC*, pages 29–38, 2011.
- [22] S. Dobzinski and R. P. Leme. Efficiency guarantees in auctions with budgets. In *Proc. Intl. Conf. Automata, Languages, and Programming (ICALP)*, pages 392–404, Berlin, Heidelberg, 2014. Springer Berlin Heidelberg.
- [23] E. Eisenberg and D. Gale. Consensus of subjective probabilities: The pari-mutuel method. *The Annals of Mathematical Statistics*, 30(1):165–168, 1959.
- [24] E. Elkind. Designing and learning optimal finite support auctions. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 736–745. SIAM, 2007.
- [25] B. Fain, K. Munagala, and N. Shah. Fair allocation of indivisible public goods. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, EC '18, page 575–592, 2018.
- [26] J. Feldman, M. Henzinger, N. Korula, V. S. Mirrokni, and C. Stein. Online stochastic packing applied to display ad allocation. In *Proceedings of the 18th Annual European Conference on Algorithms*, page 182–194, Berlin, Heidelberg, 2010. Springer-Verlag.
- [27] T. Fluschnik, P. Skowron, M. Triphaus, and K. Wilker. Fair knapsack. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):1941–1948, Jul. 2019.
- [28] L. Gelauff, A. Goel, A. Korolova, K. Munagala, and S. Yandamuri. Advertising for demographically fair outcomes. *CoRR/2006.03983*, 2020.
- [29] S. C. Geyik, S. Ambler, and K. Kenthapadi. Fairness-aware ranking in search & recommendation systems with application to linkedin talent search. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '19*, page 2221–2231, New York, NY, USA, 2019. Association for Computing Machinery.
- [30] A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica. Dominant resource fairness: Fair allocation of multiple resource types. In *8th USENIX Symposium on Networked Systems Design and Implementation (NSDI 11)*, Boston, MA, Mar. 2011. USENIX Association.
- [31] G. Goel, S. Leonardi, V. S. Mirrokni, A. Nikzad, and R. P. Leme. Reservation exchange markets for internet advertising. In *43rd Intl. Colloq. Automata, Languages, and Programming, ICALP 2016*, pages 142:1–142:13, 2016.
- [32] G. Goel, V. Mirrokni, and R. P. Leme. Polyhedral clinching auctions and the adwords polytope. *J. ACM*, 62(3), June 2015.
- [33] F. Gul and E. Stacchetti. Walrasian equilibrium with gross substitutes. *Journal of Economic Theory*, 87(1):95–124, 1999.
- [34] H. Hosseini and K. Larson. Multiple assignment problems under lexicographic preferences. In *AAMAS*, 2019.
- [35] A. Hylland and R. Zeckhauser. The efficient allocation of individuals to positions. *Journal of Political Economy*, 87(2):293–314, 1979.
- [36] K. Jain and V. V. Vazirani. Eisenberg Gale markets: Algorithms and game-theoretic properties. *Games and Economic Behavior*, 70(1):84–106, 2010.
- [37] F. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, 8(1):33–37, 1997.
- [38] A. Mehta, A. Saberi, U. Vazirani, and V. Vazirani. Adwords and generalized online matching. *J. ACM*, 54(5), 2007.
- [39] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking*, 8(5):556–567, 2000.
- [40] H. Moulin. *Fair Division and Collective Welfare*. The MIT Press, 2003.

- [41] R. B. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- [42] J. F. Nash. The bargaining problem. *Econometrica*, 18(2):155–162, 1950.
- [43] T. Nguyen and R. Vohra. Stable matching with proportionality constraints. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, EC ’17, page 675–676, 2017.
- [44] D. Peters and P. Skowron. Proportionality and the limits of welfarism. In *Proceedings of the 21st ACM Conference on Economics and Computation*, EC ’20, page 793–794, 2020.
- [45] S. Pápai. Strategyproof assignment by hierarchical exchange. *Econometrica*, 68(6):1403–1433, 2000.
- [46] A. Qureshi, R. Weber, H. Balakrishnan, J. V. Guttag, and B. M. Maggs. Cutting the electric bill for internet-scale systems. In *Proceedings of the ACM SIGCOMM 2009 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, Barcelona, Spain, August 16-21, 2009*, pages 123–134. ACM, 2009.
- [47] M. A. Satterthwaite and H. Sonnenschein. Strategy-proof allocation mechanisms at differentiable points. *Review of Economic Studies*, 48(4):587–597, 1981.
- [48] H. Steinhaus. The problem of fair division. *Econometrica*, 17:315–319, 1949.
- [49] H. R. Varian. Two problems in the theory of fairness. *Journal of Public Economics*, 5(3):249–260, 1976.
- [50] K. Yang, V. Gkatzelis, and J. Stoyanovich. Balanced ranking with diversity constraints. In S. Kraus, editor, *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 6035–6042. ijcai.org, 2019.
- [51] W. Zhang, S. Yuan, and J. Wang. Optimal real-time bidding for display advertising. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, page 1077–1086, New York, NY, USA, 2014. Association for Computing Machinery.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes] We have described some limitations in the introduction.
 - (c) Did you discuss any potential negative societal impacts of your work? [N/A]
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
 - (b) Did you include complete proofs of all theoretical results? [Yes] All proofs are in the Supplementary material.
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] The code is uploaded as part of the submission. The datasets are publicly available.
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [Yes]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

Supplementary Material

A Other Related Work

The Nash Welfare mechanism is classic approach to finding equitable allocations, and achieves fairness via the notion of *market clearing* [5, 9, 42]. It was shown in [23] that this objective finds the solution to the special case of the *Fisher market* [9]. Given linear valuation functions of agents and equal initial budgets, this market computes equilibrium prices for the items, such that when each agent buys their value maximizing allocation subject to exhausting their budget, then each item with positive price is fully allocated. This concept is also called *competitive equilibrium with equal incomes* (CEEI) and is widely studied as a fair allocation rule [35, 11, 49]. The Fisher market with linear valuation functions satisfies the *Gross Substitutes* property [33], which states that if the prices of some items increase, the demand for the other items cannot go down. This implies adding (resp. removing) an agent weakly reduces (resp. increases) the value obtained by the other agents. This property is also called competition monotonicity [36].

Though our results about the externality induced by Nash Welfare appear superficially similar to Gross Substitutes, we cannot find a formal connection. This is because our setting considers the same set of agents but adds constraints, and these constraints do not preserve the Gross Substitutes property. Furthermore, in contrast to the Gross Substitutes property that holds in a strict sense for linear valuation functions, in Section 3 we present a lower bound showing any welfarist allocation rule must result in some negative externality.

Our work considers the pure allocations problem with additional diversity constraints. In some settings such as online advertising, it is also possible to consider allocations with prices and budgets. For instance, if the utility of an agent is its value minus the price it pays (called *quasi-linear*), then the dual of the social welfare maximizing allocation also yields an equilibrium [5, 33]. Similarly, the online algorithms literature [38, 20, 21, 3, 6] considers the model where the total value derived by an agent is constrained by its budget, while auction literature [41, 12, 14, 16, 32] finds allocations and prices where agents do not have incentive to misreport. As mentioned before, though some of these mechanisms suffer from the same drawbacks as welfarist allocation rules, we leave a deeper examination of rules with prices and budgets (from the perspective of robustness) as an interesting open question.

B Omitted Proofs

Our proofs will extensively use the following property of the optimal allocation under welfarist rule f . Let \vec{V}^* denote the values of the agents in the optimal solution, and let \vec{V} be the values in any other feasible allocation. Since the valuation function is concave and continuous and so is f , the space of feasible \vec{V} is convex. This implies the following gradient optimality property:

$$\nabla f(\vec{V}^*) \cdot (\vec{V} - \vec{V}^*) \leq 0 \quad (2)$$

B.1 Proof of Theorem 1

Proof. Consider agents 1 and 2, and items a and b . Set $v_{1a} = \alpha$, $v_{2b} = \beta$, and $v_{1b} = v_{2a} = 0$. (We can set the latter values to be any small $\epsilon > 0$ as well.) In the absence of diversity constraint, we clearly have $x_{1a} = x_{2b} = 1$. So $V_2 = \beta$.

Now, agent 1 expresses a diversity constraint and requires $x_{1a} = x_{1b}$. Set $x_{1a} = x_{1b} = x$. Then, we optimize:

$$\max \ell(x) = f(\alpha x) + f(\beta(1 - x)), \quad x \in [0, 1].$$

We will show that $\ell'(1 - \delta) > 0$. Since ℓ is concave, this implies that the optimal solution $\tilde{x} > 1 - \delta$, which in turn implies that $V_2 < \delta\beta$, which thus shows that the allocation rule is not q -NNE for $q = \delta$.

To show this, we observe:

$$\ell'(1 - \delta) = \alpha f'(\alpha(1 - \delta)) - \beta f'(\delta\beta) \geq g(\alpha) - \frac{g(\beta\delta)}{\delta},$$

where the final inequality follows since $f'(\alpha(1 - \delta)) < f'(\alpha)$ by the concavity of f . We now set $\alpha = \operatorname{argmax}_y g(y)$ and $\beta = \frac{y_\delta}{\delta}$, where $g(y_\delta) = \delta \max_y g(y)$. This yields $\ell'(1 - \delta) > 0$, completing the proof. \square

B.2 Proof of Corollary 1

Proof. For γ -Fairness, note that $g(y) = x^\gamma$ and is monotonically increasing and unbounded if $\gamma \in (0, 1]$ and is monotonically decreasing and unbounded when $\gamma < 0$. Therefore, in either case, it is not δ -scaled for any $\delta > 0$, which shows that there is no constant $q > 0$ for which the allocation is q -NNE. \square

B.3 Proof of Theorem 2

Proof. Suppose agent i imposes a diversity constraint. Let \vec{A} denote the values of the agents in the optimal allocation if the constraint is not enforced, and \vec{B} the vector of values in the allocation if the constraint is enforced. Let $r_k = \frac{A_k}{B_k}$. Consider an agent $\ell \neq i$.

Since \vec{B} corresponds to the values in a feasible allocation even without agent i 's constraint, by the gradient optimality condition (Eq (2)), we have:

$$\sum_k \left(g(A_k) \frac{1}{r_k} - g(A_k) \right) = \sum_k f'(A_k)(B_k - A_k) \leq 0,$$

which can be rewritten as:

$$g(A_\ell) \frac{1}{r_\ell} + \sum_{k \neq \ell} g(A_k) \frac{1}{r_k} \leq \sum_k g(A_k). \quad (3)$$

Similarly, suppose we take the allocation without agent i 's constraint and remove agent i 's allocation from it, the resulting allocation is feasible for the problem where agent i has a constraint. This is because the empty allocation is feasible for agent i 's constraints. Applying Eq (2) again, we have:

$$-g(B_i) + \sum_{k \neq i} (g(B_k)r_k - g(B_k)) = f'(B_i)(0 - B_i) + \sum_{k \neq i} f'(B_k)(A_k - B_k) \leq 0,$$

which can be rewritten as:

$$g(B_\ell)r_\ell + \sum_{k \neq i, \ell} g(B_k)r_k \leq \sum_k g(B_k). \quad (4)$$

Since g is non-increasing, so is $g(x)/x$. Therefore, by the Rearrangement inequality, we have for all $k \neq i, \ell$:

$$g(A_k) \frac{1}{r_k} + g(B_k)r_k = \frac{g(A_k)}{A_k} B_k + \frac{g(B_k)}{B_k} A_k \geq \frac{g(A_k)}{A_k} A_k + \frac{g(B_k)}{B_k} B_k = g(A_k) + g(B_k). \quad (5)$$

We now have the following, where the first inequality follows from Eq (5), and the final inequality follows by adding Equations (3) and (4):

$$\begin{aligned} & g(A_\ell) \frac{1}{r_\ell} + g(B_\ell)r_\ell + \sum_{k \neq i, \ell} (g(A_k) + g(B_k)) \\ & \leq \left(g(A_\ell) \frac{1}{r_\ell} + \sum_{k \neq \ell} g(A_k) \frac{1}{r_k} \right) + \left(g(B_\ell)r_\ell + \sum_{k \neq i, \ell} g(B_k)r_k \right) \\ & \leq \sum_k (g(A_k) + g(B_k)). \end{aligned}$$

Simplifying, this further gives:

$$g(A_\ell) \frac{1}{r_\ell} + g(B_\ell)r_\ell \leq g(A_\ell) + g(B_\ell) + g(A_i) + g(B_i).$$

Dividing both sides by $g(B_\ell)$ gives:

$$r_\ell \leq 1 + \frac{g(A_\ell)}{g(B_\ell)} + \frac{g(A_i)}{g(B_\ell)} + \frac{g(B_i)}{g(B_\ell)} \leq 1 + \frac{3}{\delta}.$$

i.e. $\frac{A_\ell}{B_\ell} = r_\ell \leq 1 + \frac{3}{\delta}$. Therefore, $A_\ell \leq B_\ell (1 + \frac{3}{\delta})$, showing q -NNE for $q \geq \frac{\delta}{\delta+3}$. \square

B.4 Proof of Theorem 3

Proof. Consider the following example: There are 2 items a, b and two agents 1, 2. The values are $v_{1a} = v_{2a} = v_{2b} = 1$, and the rest of the values are zero. Agent 2 enforces the proportionality constraint $\epsilon x_{2a} = x_{2b}$. First consider the setting with no proportionality constraint. Suppose a rule allocates $x_{1a} = 1 - x$, so that her value is $V_1 = 1 - x$. Then $x_{2a} = x$ and $x_{2b} = 1$ so that the value of agent 2 is $V_2 = 1 + x$. Since the allocation rule satisfies (PD), this forces $x = 0$, so that $V_1 = 1$. Now suppose agent 2 enforces the proportionality constraint and as before, let $x_{2a} = x$. This forces $x_{2b} = \epsilon x$, so that $V_2 = (1 + \epsilon)x$. As before $V_1 = 1 - x$. If $x < \frac{1}{2+\epsilon}$, this allocation cannot satisfy (PD). Therefore $V_1 \leq \frac{1+\epsilon}{2+\epsilon}$ for any allocation satisfying (PO) and (PD). Now taking $\epsilon \rightarrow 0$ shows that the allocation cannot be q -NNE for any constant $q > \frac{1}{2}$. \square

B.5 Proof of Corollary 3

Proof. Let $\ell \notin S$ denote the agent whose value we are bounding. The inequalities obtained by generalizing Eq (3) and (4) to omit the set S instead of a single agent i now yields:

$$g(A_\ell) \frac{1}{r_\ell} + \sum_{k \neq \ell} g(A_k) \frac{1}{r_k} \leq \sum_k g(A_k), \quad g(B_\ell) r_\ell + \sum_{k \notin S \cup \{\ell\}} g(B_k) r_k \leq \sum_k g(B_k).$$

The same line of reasoning gives:

$$g(A_\ell) \frac{1}{r_\ell} + g(B_\ell) r_\ell \leq \sum_{k \in S \cup \{\ell\}} (g(A_k) + g(B_k)).$$

Dividing both sides by $g(B_\ell)$ gives:

$$r_\ell \leq 1 + \frac{g(A_\ell)}{g(B_\ell)} + \sum_{k \in S} \left(\frac{g(A_k)}{g(B_\ell)} + \frac{g(B_k)}{g(B_\ell)} \right) \leq 1 + \frac{2k+1}{\delta},$$

thus showing q -NNE for $q \geq \frac{\delta}{2k+\delta+1}$. \square

B.6 Proof of Corollary 4

Proof. The lower bound follows by extending Theorem 3. There is an item a such that for agent 1, $v_{1a} = 1$. For $i \in \{2, 3, \dots, k+1\}$, there is an item i such that agent i has $v_{ii} = 1$. For $i \in \{2, 3, \dots, k+1\}$, we also have $v_{ia} = 1$. All other values are zero. Without the diversity constraint, suppose $x_{2a} = x_{3a} = \dots = x$ and $x_{1a} = 1 - kx$, then $V_2 = V_3 = \dots = 1 + x$ and $V_1 = 1 - kx$. Then (PD) implies $x = 0$ so that $V_1 = 1$.

Now suppose each agent $i \in \{2, 3, \dots, k+1\}$ express the proportionality constraint $\epsilon x_{ia} = x_{ii}$. If $x_{1a} = 1 - ky$, then by anonymity, we have $x_{ia} = y$ and $x_{ii} = \epsilon y$ for all $i \in \{2, 3, \dots, k+1\}$. Therefore, $V_1 = 1 - ky$ and $V_2 = V_3 = \dots = (1 + \epsilon)y$. Now, any allocation with $y < \frac{1}{k+1+\epsilon}$ does not satisfy (PD). Therefore, any allocation satisfying (PO), anonymity, and (PD) has $V_1 \leq \frac{1+\epsilon}{k+1+\epsilon}$. This completes the proof. \square

B.7 Proof of Theorem 4

Proof. We first present the upper bound. The allocation maximizes $\sum_i \frac{1}{\gamma} V_i^\gamma$. As in the proof of Theorem 2, let \vec{A} denote the vector of values of the agents if agent i did not have a diversity constraint, and let \vec{B} denote the vector of values if the constraint is enforced. Assume by scaling all values by

the same amount that $A_i = 1$, and denote $B_i = x$. Our goal is to upper bound x , which will yield the value of p . We will assume below that $\gamma \neq 0$, and present the proof for $\gamma \rightarrow 0$ separately.

The gradient optimality condition Equation (2) now simplifies to:

$$\sum_i V_i(V_i^*)^{\gamma-1} \leq \sum_i (V_i^*)^\gamma. \quad (6)$$

Consider the unconstrained allocation, but set agent i 's allocation to zero. This is feasible for the constrained version whose optimal solution is \vec{B} . Applying Eq (6), we have:

$$x^\gamma + \sum_{k \neq i} B_k^\gamma \geq \sum_{k \neq i} \frac{A_k}{B_k^{1-\gamma}}. \quad (7)$$

Similarly, the constrained allocation \vec{B} is feasible for the unconstrained problem. Applying Eq (6), we have:

$$1 + \sum_{k \neq i} A_k^\gamma \geq x + \sum_{k \neq i} \frac{B_k}{A_k^{1-\gamma}}. \quad (8)$$

Combining Equations (8) and (7), we have

$$x - x^\gamma \leq 1 + \sum_{k \neq i} \left(A_k^\gamma + B_k^\gamma - \frac{A_k}{B_k^{1-\gamma}} - \frac{B_k}{A_k^{1-\gamma}} \right). \quad (9)$$

We will now show that

$$A_k^\gamma + B_k^\gamma - \frac{A_k}{B_k^{1-\gamma}} - \frac{B_k}{A_k^{1-\gamma}} \leq 0.$$

This is equivalent to:

$$\left(\frac{1}{A_k^{1-\gamma}} - \frac{1}{B_k^{1-\gamma}} \right) (A_k - B_k) \leq 0.$$

It is easy to check that for any $\gamma \in (-\infty, 1]$, we have $A_k \geq B_k$ iff $A_k^{1-\gamma} \geq B_k^{1-\gamma}$. This proves the inequality. Plugging it into Eq (9), we have

$$x - x^\gamma \leq 1.$$

It is easy to check that for $\gamma \in (-\infty, 1]$, x is maximized when $x = 1 + x^\gamma$, completing the proof.

Nash Welfare. When $\gamma \rightarrow 0$, the above proof does not directly apply. Nevertheless, we can obtain the bound $p = \frac{1}{2}$ as follows. We follow the same outline as the proof of Theorem 2 but do not fix another agent ℓ . For NW rule, Eq (3) can be rewritten as:

$$\frac{1}{r_i} + \sum_{k \neq i} \frac{1}{r_k} \leq n.$$

and Eq (4) implies

$$\sum_{k \neq i} r_k \leq n.$$

Our goal now is to upper bound $\frac{1}{r_i}$. Therefore, we solve:

$$\min \sum_{k \neq i} \frac{1}{r_k} \quad \text{s.t.} \quad \sum_{k \neq i} r_k \leq n, \quad r_k \geq 0.$$

This implies $r_k = \frac{n}{n-1}$ in the optimal solution, so that $\sum_{k \neq i} \frac{1}{r_k} \geq \frac{(n-1)^2}{n}$. Therefore

$$\frac{1}{r_i} \leq n - \sum_{k \neq i} \frac{1}{r_k} \leq n - \frac{(n-1)^2}{n} \leq 2.$$

Therefore, $A_i \geq \frac{1}{2} B_i$, which shows p -MON for $p \geq \frac{1}{2}$.

Tight Lower Bound. There are $n + 1$ of agents $\{0, 1, 2, \dots, n\}$ where $n \rightarrow \infty$ and let $c = \beta n$, where $\beta \in (0, 1)$ is a constant to be determined later. There are $2n + 1$ items $\{0, 1, 2, \dots, 2n\}$. We have $v_{00} = c$; for each $i \in \{1, 2, \dots, n\}$ we have $v_{0i} = v_{ii} = 1$; and $v_{i(i+n)} = c - 1$. All other values are zero. Agent 0 expresses the proportionality constraint $x_{00} = x_{01} = \dots = x_{0n}$.

We always have $x_{i(i+n)} = 1$ for $i \in \{1, 2, \dots, n\}$. In the absence of the diversity constraint, we have $x_{00} = 1$. Suppose $x_{01} = x_{02} = \dots = x_{0n} = x$ and $x_{11} = x_{22} = \dots = x_{nn} = 1 - x$. The optimal allocation solves

$$\max_x n \frac{(c-x)^\gamma}{\gamma} + \frac{(nx+c)^\gamma}{\gamma} \quad x \in [0, 1].$$

It is easy to check that $x = 0$ in this solution, so that $V_0 = nx + c = c = \beta n$.

When agent 0 adds the diversity constraint, we have $x_{00} = x_{01} = \dots = x_{0n} = x$ and $x_{11} = x_{22} = \dots = x_{nn} = 1 - x$. Again, the optimization problem is:

$$\max_x n \frac{(c-x)^\gamma}{\gamma} + \frac{((n+c)x)^\gamma}{\gamma} \quad x \in [0, 1].$$

This yields

$$x = \min \left(1, c \frac{(n+c)^{\frac{\gamma}{1-\gamma}}}{n^{\frac{1}{1-\gamma}} + (n+c)^{\frac{\gamma}{1-\gamma}}} \right).$$

We will constrain c so that

$$(c-1)(n+c)^{\frac{\gamma}{1-\gamma}} \geq n^{\frac{1}{1-\gamma}}, \quad (10)$$

so that at the optimal solution, we have $x = 1$ implying $V_0 = (n+c)x = n+c$. This will show $p = \frac{c}{n+c} = \frac{\beta}{1+\beta}$.

The constraint Eq (10) can be written as $(c-1)^{1-\gamma}(n+c)^\gamma \geq n$. Dividing by n and observing that when $\beta > 0$ and $n \rightarrow \infty$, $\frac{c-1}{n} \rightarrow \beta$, this constraint reduces to:

$$\beta^{1-\gamma}(1+\beta)^\gamma \geq 1.$$

Setting $\theta = \frac{\beta}{1+\beta}$, this implies that this instance is not p -MON for $p > \theta$, where θ is constrained by $\theta^{1-\gamma} + \theta \geq 1$, completing the proof. \square