

Efficient Resource Allocation with Secretive Agents

Soroush Ebadian¹, Rupert Freeman² and Nisarg Shah¹

¹University of Toronto

²University of Virginia

soroush@cs.toronto.edu, FreemanR@darden.virginia.edu, nisarg@cs.toronto.edu

Abstract

We consider the allocation of homogeneous divisible goods to agents with linear additive valuations. Our focus is on the case where some agents are secretive and reveal no preference information, while the remaining agents reveal full preference information. We study distortion, which is the worst-case approximation ratio when maximizing social welfare given such partial information about agent preferences. As a function of the number of secretive agents k relative to the overall number of agents n , we identify the exact distortion for every p -mean welfare function, which includes the utilitarian welfare ($p = 1$), the Nash welfare ($p \rightarrow 0$), and the egalitarian welfare ($p \rightarrow -\infty$).

1 Introduction

We study a resource allocation problem in which divisible goods must be allocated to agents with linear additive valuations.¹ Treating goods as divisible captures cases where they are inherently divisible (such as land or food), and where they are indivisible (such as jewelry or artwork) but can be allocated randomly or timeshared. Formally, an allocation is a matrix \mathbf{x} , where $x_{i,j} \in [0, 1]$ is the fraction of resource j given to agent i and $\sum_i x_{i,j} = 1$ for all j . The preferences of agent i are given by a valuation function v_i such that her utility from allocation \mathbf{x} is $v_i(\mathbf{x}_i) = \sum_j v_i(j) \cdot x_{i,j}$.

A classic solution is to allocate the resources in a way that maximizes some *social welfare function*, which maps the utilities of the agents to a single aggregate measure of allocation quality. Common examples include the utilitarian welfare ($\frac{1}{n} \sum_i v_i(\mathbf{x}_i)$), the Nash welfare ($(\prod_i v_i(\mathbf{x}_i))^{1/n}$), and the egalitarian welfare ($\min_i v_i(\mathbf{x}_i)$), where n is the number of agents. In fact, these are members of the broader class of p -mean welfare functions, given by $(\frac{1}{n} \sum_i v_i(\mathbf{x}_i)^p)^{1/p}$, with $p = 1$, $p \rightarrow 0$, and $p \rightarrow -\infty$ respectively.

When we have complete information about the valuation function of each agent, finding an allocation that maximizes social welfare is conceptually trivial (algorithmically, however, some welfare functions may be challenging to maxi-

mize [Lee, 2017; Garg *et al.*, 2021; Bezáková and Dani, 2005; Asadpour and Saberi, 2010]). But when we have only partial information, it is less clear what outcomes are prescribed by the social welfare maximization paradigm. One approach in the literature is to consider the *distortion*, which is the worst-case approximation ratio of the maximum social welfare that could be achieved with full information to the social welfare achieved by the allocation rule given partial information. Distortion can be viewed as the “price” of missing information, and minimizing distortion provably reduces the (worst-case) impact that the missing information has on the solution quality. Distortion was originally defined by Procaccia and Rosenschein [2006] in the context of voting, where it has led to an extensive literature of follow-up work; we point the reader to the recent survey by Anshelevich *et al.* [2021] for a summary. The approach has since been applied to other settings including matching [Amanatidis *et al.*, 2021; Ma *et al.*, 2021; Anshelevich and Zhu, 2021] and resource allocation [Halpern and Shah, 2021].

Traditionally, the distortion framework has been applied when every agent reports ordinal preferences [Boutilier *et al.*, 2015; Anshelevich *et al.*, 2018; Halpern and Shah, 2021]. In this paper, we introduce and study a different model, in which some agents provide complete cardinal valuation functions while others provide no information. We term the latter agents *secretive agents*. In practice, agents may be secretive because they do not want to disclose their valuations for privacy reasons, or because they are simply unresponsive to requests for information. For example, on a popular resource allocation website Spliddit.org, more than 10% of the goods division instances did not succeed because at least one user did not submit their valuation function.² Prior work in resource allocation has considered secretive agents [Asada *et al.*, 2018; Frick *et al.*, 2019; Chèze, 2019; Arunachaleswaran *et al.*, 2019], but these focus on guaranteeing certain fairness properties in the presence of secretive agents, not on welfare maximization or distortion. Further, unlike in our work, none of them allow more than a single agent to be secretive because guaranteeing the fairness properties they seek becomes trivially impossible in this case.

In the presence of one or more secretive agents, it is not a priori clear what a “good” allocation looks like. On the one

¹We discuss allocation of indivisible goods in the full version: <https://www.cs.toronto.edu/~nisarg/papers/secretive.pdf>

²For chore division instances, it was even higher at over 32%.

hand, if we assign any good to a secretive agent, she might turn out to have very low value for that good, resulting in the good effectively being wasted. On the other hand, if we allocate nothing to the secretive agents, we run the risk of facing high distortion due to instances where the secretive agents are the key to achieving high welfare. How do we balance these considerations? Should we allocate any resources to the secretive agents? If so, how do we determine how much of a resource should be allocated to the secretive agents? We answer these questions by identifying worst-case optimal allocation rules, which turn out to be surprisingly simple.

1.1 Our Results

Let n be the number of agents, k of whom are secretive. We present our results for divisible goods in the main body and defer the treatment of indivisible goods to the full version. For divisible goods, we provide a complete picture of the exact distortion for all p -mean social welfare functions. We introduce a family of allocation rules parameterized by $\alpha \in [0, 1]$ and show that all our upper bounds can be achieved by setting the right value of α as a function of p , n , and k . Given α , the corresponding rule allocates α fractions of all the goods to the non-secretive agents in such a way to maximize their social welfare and splits the remaining $1 - \alpha$ fraction of each good equally among the secretive agents. In each case, we can obtain an exactly matching lower bound. A summary of our results is presented in Table 1 and Figure 1 shows how the distortion varies with p , n , and k .

The distortion naturally increases as the number of secretive agents k increases; for every p , the distortion starts at 1 when $k = 0$ (full information) and increases to n at $k = n$ (no information). Interestingly, for $p = 1$ (the utilitarian welfare), $p \rightarrow 0$ (the Nash welfare), and $p \rightarrow -\infty$ (the egalitarian welfare), the distortion already becomes n at $k = n - 1$, meaning that knowing the valuation function of a single agent is not helpful for these welfare functions, but this is not the case for intermediate values of p . When $k = \Theta(n)$, the distortion is $\Theta(n)$ for $p \leq 1$ and $\Theta(n)^{1/p}$ for $p > 1$. When $k \ll n$, it is worth noting that the distortion for the Nash welfare is $\approx 1 + k \ln n/n$, which grows linearly in k like for the utilitarian and egalitarian welfare, but at a lower rate. More generally, the Nash welfare leads to a surprisingly low distortion; see Figure 1.

Finally, we conduct simulations on synthetic data and real data from Spliddit.org to evaluate the empirical performance of our algorithms with respect to the utilitarian social welfare. While every $\alpha \in [1/(k+1), 1]$ is optimal in the worst case, we find that higher values of α perform better empirically.

1.2 Related Work

In the voting literature, the idea of distortion has been analyzed under two primary frameworks, distinguished by what they assume the underlying expressive preference format to be: the utilitarian framework assumes that voters have utilities for candidates [Boutilier *et al.*, 2015; Caragiannis *et al.*, 2017; Benadè *et al.*, 2017], while the metric framework assumes that voters have costs for candidates satisfying the triangle inequality [Anshelevich *et al.*, 2018; Gkatzelis *et al.*, 2020]. Following Halpern and Shah [2021], our work follows

| Welfare | Distortion with $0 \leq k < n$ | Optimal α |
|----------|---|--|
| Egal. W. | $k + 1$ | $\alpha = \frac{1}{k+1}$ |
| | $(-\infty, 0)$ $n^{\frac{1}{p}} \left((n-k)^{\frac{1}{1-p}} + k \right)^{\frac{p-1}{p}}$ | $\alpha = \frac{(n-k)^{\frac{1}{1-p}}}{k + (n-k)^{\frac{1}{1-p}}}$ |
| Nash W. | $n(n-k)^{-\frac{n-k}{n}}$ | $\alpha = \frac{n-k}{n}$ |
| | $(0, 1)$ $n^{1-\frac{1}{p}} \left((n-k)^{1-p} + k \right)^{\frac{1}{p}}$ | $\alpha = \frac{n-k}{n}$ |
| Util. W. | $k + 1$ | $\alpha \in \left[\frac{1}{k+1}, 1 \right]$ |
| | $(1, \infty)$ $(k+1)^{\frac{1}{p}}$ | $\alpha = 1$ |

Table 1: Summary of results for divisible goods. For $k = 0$, all distortion values in the table evaluate to 1. However, for $k = n$ the correct distortion value is n for $p \leq 1$ and $n^{1/p}$ for $p > 1$.

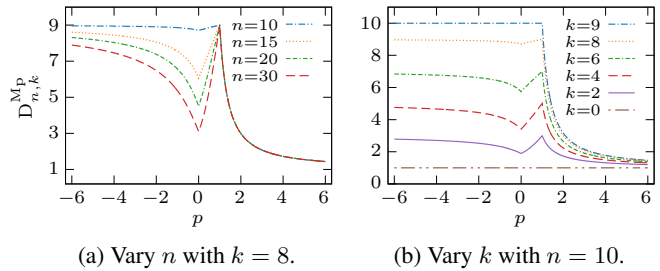


Figure 1: Distortion value with divisible items as a function of p .

the utilitarian framework as it is more applicable to allocating goods.

Halpern and Shah [2021], like us, assume that agents have additive cardinal valuations, but they study the case where every agent reports a ranking of her t most favorite goods. They analyze the best possible distortion with respect to the utilitarian social welfare as a function of t in relation to the number of goods m . In particular, when every agent ranks all the goods (i.e., $t = m$), they show that the best possible distortion (with a randomized rule) is n , which is what one can achieve with no preference information whatsoever. That is, they argue that having access to ordinal preference information is not helpful for welfare maximization. In contrast, our distortion bound is better when $k \leq n - 2$, i.e., even when we have access to the valuation functions of just *two* agents. In a sense, this shows the usefulness of eliciting cardinal preferences as opposed to ordinal preferences in resource allocation settings.

Finally, we note that the idea of secretive agents is also explored in the voting literature, albeit with very different motivations. Borodin *et al.* [2019, Lemma 4] show that constant metric distortion can be achieved in elections where any subset of voters that is at least a constant fraction of the electorate participate and submit ordinal preferences; such a strong guarantee is known to be impossible to achieve in the utilitarian framework, but may be possible if the participating subset of voters is assumed to be drawn at random. Micha and Shah [2020] study voting rules which have access to the votes of only a subset of voters, but instead of analyzing the distortion, their aim is to predict what popular voting rules would have returned given all the votes. One of their primary

motivations is to design voting rules to apply to polls in order to predict the outcome of an upcoming election.

2 Preliminaries

A resource allocation instance $(\mathcal{N}, \mathcal{M}, \mathbf{v})$ consists of a set of n agents \mathcal{N} , a set of m goods \mathcal{M} , and a utility profile $\mathbf{v} = (v_1, \dots, v_n)$, where $v_i: \mathcal{M} \rightarrow \mathbb{R}_{\geq 0}$ is the valuation function of agent i . We normalize the valuations of each agent i to sum up to one, i.e., $\sum_{j \in \mathcal{M}} v_i(j) = 1$. This unit-sum normalization is widely used in welfare maximization [Aziz, 2020].³

Allocations. An allocation is a division of the goods among the agents, denoted by $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$, where $x_{i,j}$ is the fraction of good j allocated to agent i and each good j is fully allocated (i.e., $\sum_{i \in \mathcal{N}} x_{i,j} = 1$ for each j). We consider the class of linear additive utilities, where the utility of agent i for her share \mathbf{x}_i is, with slight abuse of notation, defined as $v_i(\mathbf{x}_i) = \sum_{j \in \mathcal{M}} v_i(j) \cdot x_{i,j}$.

Welfare functions. A welfare function W aggregates the utilities to the agents under an allocation \mathbf{x} into a single non-negative real number measuring the efficiency of the allocation. Following Barman *et al.* [2020a], we consider the following class of welfare functions.

Definition 1 (p -Mean Welfare). For $p \in \mathbb{R}$, the p -mean welfare of allocation \mathbf{x} is defined as

$$M_p(\mathbf{x}) = \left(\frac{1}{n} \sum_{i \in \mathcal{N}} v_i(\mathbf{x}_i)^p \right)^{1/p}.$$

This class contains three popular welfare functions:

- Choosing $p = 1$ induces the *utilitarian welfare*, given by $UW(\mathbf{x}) = (1/n) \cdot \sum_{i \in \mathcal{N}} v_i(\mathbf{x}_i)$,
- The limit $p \rightarrow 0$ induces the *Nash welfare*, given by $NW(\mathbf{x}) = (\prod_{i \in \mathcal{N}} v_i(\mathbf{x}_i))^{1/n}$,
- The limit $p \rightarrow -\infty$ induces the *egalitarian welfare*, given by $EW(\mathbf{x}) = \min_{i \in \mathcal{N}} v_i(\mathbf{x}_i)$.

It is known that p -mean welfare functions are characterized by five natural axioms [Moulin, 2003, pp. 66-69], and further imposing the Pigou-Dalton principle induces $p \leq 1$. It is interesting that our result for the warm-up case of $k = n$ also differs depending on whether $p \leq 1$ or $p > 1$ (see Section 3.1).

2.1 Secretive Agents & Distortion

In our setting, we assume that we have no information about the valuation functions of k agents, whom we term *secretive agents*, while we have complete information on the valuation functions of the remaining agents, whom we term *non-secretive agents*. Our goal is to find an allocation that minimizes the worst-case multiplicative loss of efficiency measured by a p -mean welfare function.

More formally, let \mathcal{N}_{sec} and $\mathcal{N}_{\text{nonsec}}$ denote the sets of secretive and non-secretive agents, respectively. An instance

³Our results for the Nash welfare hold independently of any normalization. However, for every other p -mean welfare function, the optimal distortion is $k + 1$ in the absence of any normalization.

of resource allocation with secretive agents $(\mathcal{N}, \mathcal{M}, \mathbf{v}_{\text{nonsec}})$ consists of a set \mathcal{N} of agent, a set \mathcal{M} of items, and a valuation function v_i for each non-secretive agent (the valuation functions implicitly define the sets \mathcal{N}_{sec} and $\mathcal{N}_{\text{nonsec}}$). We aim to find an optimal strategy for the following game:

1. The adversary chooses the valuation functions of the non-secretive agents, denoted by $\mathbf{v}_{\text{nonsec}} = (v_i)_{i \in \mathcal{N}_{\text{nonsec}}}$.
2. The player chooses an allocation \mathbf{x} of the goods to all agents (secretive and non-secretive).
3. The adversary chooses the valuation functions of the secretive agents, denoted by $\mathbf{v}_{\text{sec}} = (v_i)_{i \in \mathcal{N}_{\text{sec}}}$, as well as an allocation \mathbf{x}^* .
4. The player incurs the (multiplicative) loss $W(\mathbf{x}^*)/W(\mathbf{x})$.

This game is formalized via the notion of *distortion*.

Definition 2 (Distortion with Secretive Agents). Given the number of agents n , the number of secretive agents k , and a welfare function W , the *distortion* is defined as

$$D_{n,k}^W = \sup_{\mathbf{v}_{\text{nonsec}}} \inf_{\mathbf{x}} \sup_{\mathbf{v}_{\text{sec}}, \mathbf{x}^*} \frac{W(\mathbf{x}^*)}{W(\mathbf{x})}.$$

Note that the distortion is always at least 1 as the adversary can always return the same allocation as the player returns, i.e., $\mathbf{x}^* = \mathbf{x}$.

A strategy for the player corresponds to an *allocation rule* that maps instances to allocations. Because we express distortion values that depend on n and m (that is, we are typically interested in varying $\mathbf{v}_{\text{nonsec}}$), we suppress the dependence on \mathcal{N} and \mathcal{M} and simply write $A(\mathbf{v}_{\text{nonsec}})$ to denote the output of allocation rule A on instance $(\mathcal{N}, \mathcal{M}, \mathbf{v}_{\text{nonsec}})$.

Definition 3 (Distortion of an Allocation Rule). Given n agents, k secretive agents, and a welfare function W , the *distortion of an allocation rule* A is defined as

$$D_{n,k}^W(A) = \sup_{\mathbf{v}_{\text{nonsec}}, \mathbf{v}_{\text{sec}}, \mathbf{x}^*} \frac{W(\mathbf{x}^*)}{W(A(\mathbf{v}_{\text{nonsec}}))}.$$

If $D_{n,k}^W(A) = D_{n,k}^W$ then we refer to A as an optimal strategy for the player.

3 Distortion Values

In this section, we present allocation rules that provide provable guarantees on the distortion with respect to p -mean welfare functions.

3.1 Warm-up: $k = 0$ and $k = n$

First, let us consider two extreme cases where $k = 0$ and $k = n$ which provide us some intuition for the general case.

Case $k = 0$. If there are no secretive agents, then we have full information of the utilities and we can return the allocation that maximizes the welfare for all agents. The adversary cannot obtain a welfare higher than us, therefore, the distortion value is 1. As $\mathcal{N} = \mathcal{N}_{\text{nonsec}}$ in this case, we may say our strategy was maximizing the welfare for the *non-secretive* agents. Denote this strategy by

$$\text{OPT}_{\text{nonsec}}(\mathbf{v}_{\text{nonsec}}) = \arg \max_{\mathbf{x}} \left(\frac{1}{n} \sum_{i \in \mathcal{N}_{\text{nonsec}}} v_i(\mathbf{x}_i)^p \right)^{1/p}.$$

Case $k = n$. Suppose all agents are secretive and we do not have any information from their utilities. To assist with intuition, assume $p \leq 1$. Our best response would be to return a uniform allocation, i.e. allocate $1/n$ of each item to each (secretive) agent. Intuitively speaking, this follows from the concavity of M_p for $p \leq 1$. If we act differently, the adversary can use the asymmetry in our allocation to incur a higher distortion (details provided in the full version). Denote this strategy by

$$\text{Uniform}_{\text{sec}}(\mathbf{v}_{\text{nonsec}}) = \{x_{i,j} = \frac{1}{|\mathcal{N}_{\text{sec}}|} \mid \forall j \in \mathcal{M}, i \in \mathcal{N}_{\text{sec}}\}.$$

Regardless of the utilities, for all agents we have $v_i(\mathbf{x}) = 1/n$. Hence, the welfare obtained is $1/n$. The adversary cannot achieve mean welfare of more than 1. Thus, we get an upper bound of $D_{n,k=n}^{\text{W}} \leq n$. In the full version, we also show a matching lower bound.

Lemma 1. *For all p -mean welfare functions with $p \in (-\infty, 1]$ (including NW and UW) and EW, the distortion with n secretive agents is $D_{n,n}^{\text{W}} = n$.*

The analysis presented does not hold for $p > 1$. By the convexity of M_p when $p > 1$, our best response is to allocate all items to one agent. Then, only one agent will have a utility of 1 while others get 0 utility. Therefore, $(\frac{1}{n} \sum_{i \in \mathcal{N}} v_i(\mathbf{x})^p)^{1/p} = (1/n)^{1/p}$ leading to an upper bound of $D_{n,n}^{\text{M}_p} \leq \frac{1}{n^{-1/p}} = n^{1/p}$.

Lemma 2. *For all p -mean welfare functions with $p \in (1, \infty)$, the distortion is $D_{n,n}^{\text{W}} = n^{1/p}$.*

3.2 Results for $1 \leq k \leq n - 1$

In general, our strategy for the general case is to mix the two strategies described for the extreme cases of $k \in \{0, n\}$. That is, our allocation rule is one from the following class of allocation rules,

$$\mathbf{A}_\alpha = \alpha \text{OPT}_{\text{nonsec}} + (1 - \alpha) \text{Uniform}_{\text{sec}}, \quad (1)$$

where we allocate $\alpha \in [0, 1]$ portion of each item according to the $\text{OPT}_{\text{nonsec}}$ rule, and the rest uniformly among the secretive agents. The proper choice of α however, depends on the chosen welfare function.

We begin with a lemma that provides an upper bound on the adversary's welfare. Throughout this paper, we often use $\beta = \sum_{i \in \mathcal{N}_{\text{nonsec}}} v_i(\text{OPT}_{\text{nonsec}}(\mathbf{v}_{\text{nonsec}}))^p$ to refer to the unnormalized welfare (i.e., the p -th power of the p -mean welfare) that the non-secretive agents receive under $\text{OPT}_{\text{nonsec}}$.

Lemma 3. *For any valuations of the non-secretive agents $\mathbf{v}_{\text{nonsec}}$ and $p \in \mathbb{R}$, it holds that*

$$\begin{aligned} M_p(\mathbf{x}^*) &\leq \left(\frac{1}{n}(\beta + k)\right)^{\frac{1}{p}} \\ &= (k + 1) \cdot M_p(\mathbf{A}_{1/(k+1)}(\mathbf{v}_{\text{nonsec}})), \end{aligned} \quad (2)$$

where $\beta = \sum_{i \in \mathcal{N}_{\text{nonsec}}} v_i(\text{OPT}_{\text{nonsec}}(\mathbf{v}_{\text{nonsec}}))^p$. That is, the welfare achieved by the adversary is at most $k + 1$ times higher than the welfare achieved by the allocation rule $\mathbf{A}_{1/(k+1)}$.

Proof. Let u_i denote the utility achieved by non-secretive agent i under $\text{OPT}_{\text{nonsec}}$. Note that each secretive agent receives utility at most 1 due to the unit-sum normalization. Hence,

$$\begin{aligned} M_p(\mathbf{x}^*) &= \left(\frac{1}{n} \cdot \sum_{i \in \mathcal{N}} v_i(\mathbf{x}^*)^p\right)^{\frac{1}{p}} \\ &\leq \left(\frac{1}{n} \left(\sum_{i \in \mathcal{N}_{\text{nonsec}}} u_i^p + \sum_{i \in \mathcal{N}_{\text{sec}}} 1\right)\right)^{\frac{1}{p}} = \left(\frac{\beta + k}{n}\right)^{\frac{1}{p}}. \end{aligned}$$

On the other hand, note that the utility vector under $\mathbf{A}_{1/(k+1)}(\mathbf{v}_{\text{nonsec}})$ is $(\frac{u_1}{k+1}, \dots, \frac{u_{n-k}}{k+1}, \frac{1}{k+1}, \dots, \frac{1}{k+1})$ because $\frac{1}{k+1}$ fraction of each item is allocated according to $\text{OPT}_{\text{nonsec}}$ and $\frac{1}{k+1}$ fraction of each item is allocated to each of k secretive agents. Hence, $M_p(\mathbf{A}_{1/(k+1)}(\mathbf{v}_{\text{nonsec}})) = \frac{1}{k+1} \cdot \left(\frac{\beta + k}{n}\right)^{\frac{1}{p}}$, which yields the desired relation. \square

Lemma 3 immediately implies the following Corollary.

Corollary 1. *For all p -mean welfare functions W , the allocation rule $\mathbf{A}_{1/(k+1)}$ has distortion $D_{n,k}^{\text{W}}(\mathbf{A}_{1/(k+1)}) \leq k + 1$ for all $n \geq k \geq 0$.*

It turns out that the upper bound of $k + 1$ is only tight for the egalitarian and utilitarian welfare functions. For other values of p , we can achieve lower distortion by tailoring our strategy to the particular welfare function. The next two lemmas contain common parts to the analysis that we will use to prove our guarantees.

Lemma 4. *Consider a resource allocation instance with secretive agents. For all $\alpha \in [0, 1]$ and any p -mean welfare M_p we have*

$$\frac{M_p(\mathbf{x}^*)}{M_p(\mathbf{A}_\alpha(\mathbf{v}_{\text{nonsec}}))} \leq \left(\frac{\beta + k}{\alpha^p \beta + \left(\frac{1-\alpha}{k}\right)^p k}\right)^{\frac{1}{p}} \triangleq f_p(\beta, \alpha), \quad (3)$$

where $\beta = \sum_{i \in \mathcal{N}_{\text{nonsec}}} v_i(\text{OPT}_{\text{nonsec}}(\mathbf{v}_{\text{nonsec}}))^p$.

Proof. By Lemma 3, the welfare achieved by the adversary is upper bounded by $M_p(\mathbf{x}^*) \leq \left(\frac{1}{n}(\beta + k)\right)^{1/p}$. Next, using the allocation rule \mathbf{A}_α , the player achieves a welfare of $M_p(\mathbf{A}_\alpha(\mathbf{v}_{\text{nonsec}})) = \left(\frac{1}{n}(\alpha^p \beta + \left(\frac{1-\alpha}{k}\right)^p k)\right)^{1/p}$. \square

Lemma 4 immediately implies that

$$D_{n,k}^{\text{M}_p} \leq D_{n,k}^{\text{M}_p}(\mathbf{A}_\alpha) \leq \max_{\mathbf{v}_{\text{nonsec}}} \left(\frac{\beta + k}{\alpha^p \beta + \left(\frac{1-\alpha}{k}\right)^p k}\right)^{\frac{1}{p}}.$$

Lemma 5. *Let $f_p(\beta, \alpha)$ be as defined in (3). Then, for a fixed $\alpha \geq \frac{1}{k+1}$, f_p is non-increasing over $\beta \geq 1$.*

Proof. As $f_p(\beta, \alpha) \geq 1$ and since \log preserves monotonicity, it is sufficient to show $\frac{d}{d\beta} \log f_p(\beta, \alpha) \leq 0$ for all $\beta \geq 1$.

$$\frac{d}{d\beta} \log f_p(\beta, \alpha) = \frac{1}{p} \left(\frac{1}{\beta + k} - \frac{1}{\beta + \left(\frac{1-\alpha}{\alpha k}\right)^p k}\right).$$

By $\alpha \geq \frac{1}{k+1}$, we have $\frac{1-\alpha}{\alpha k} \leq 1$. Then, we can check this expression is non-positive both for $p > 0$ and $p < 0$. \square

As f_p is non-increasing, to obtain an upper bound on the distortion, we need a lower bound on β . This value, as well as the proper choice of α , depends on p . In the rest of this section, we will find the proper choices for α and β based on the welfare function. We begin with $p \in (-\infty, 0)$.

Theorem 1. For $p \in (-\infty, 0)$, the allocation rule $A_{z/(k+z)}$ with $z = (n-k)^{\frac{1}{1-p}}$ achieves $D_{n,k}^{M_p}(A_\alpha) \leq n^{\frac{1}{p}}(z+k)^{\frac{p-1}{p}}$.

Taking the limit as $p \rightarrow 0$ and $p \rightarrow -\infty$ in Theorem 1 suggests upper bounds of $n(\frac{1}{n-k})^{\frac{n-k}{n}}$ and $k+1$ for the Nash and egalitarian welfare respectively. For the egalitarian welfare, we have already shown an upper bound of $k+1$ in Corollary 1, and the following lemma proves that this upper bound is achievable for the Nash welfare.

Theorem 2. For the Nash welfare, the allocation rule $A_{(n-k)/n}$ achieves $D_{n,k}^{NW}(A_\alpha) \leq n \left(\frac{1}{n-k}\right)^{\frac{n-k}{n}}$.

Now, we will focus on the range $p \in (0, 1]$.

Theorem 3. For $p \in (0, 1]$, the allocation rule $A_{(n-k)/n}$ achieves $D_{n,k}^{M_p}(A_\alpha) \leq n \left(\frac{(n-k)^{1-p} + k}{n}\right)^{\frac{1}{p}}$.

Proof. The requirement of Lemma 5 is met, as for $k < n$, $(n-k)(k+1) \geq n \Rightarrow \frac{n-k}{n} \geq \frac{1}{k+1}$.

By Lemma 4, distortion is bounded by (3), and by Lemma 5, this bound is maximized when β is minimized. For any given $\mathbf{v}_{\text{nonsec}}$, one suboptimal allocation is $\text{Uniform}_{\text{nonsec}}$. Each agent gets $v_i = \frac{1}{n-k}$ utility from this rule. Hence, $\beta \geq (n-k)\left(\frac{1}{n-k}\right)^p = (n-k)^{1-p}$.

By substituting β and α in (3), we have

$$\begin{aligned} D_{n,k}^{M_p}(A_\alpha) &\leq \left(\frac{(n-k)^{1-p} + k}{\left(\frac{n-k}{n}\right)^p (n-k)^{1-p} + \frac{k}{n^p}} \right)^{\frac{1}{p}} \\ &= n \left(\frac{(n-k)^{1-p} + k}{n-k+k} \right)^{\frac{1}{p}}. \quad \square \end{aligned}$$

Note that for $p = 1$, Theorem 3 implies an upper bound of $k+1$ for the utilitarian welfare, matching the upper bound from Corollary 1. Moreover, by taking the limit $p \rightarrow 0$ in Theorem 3, we get the same upper bound proven in Theorem 2.

In fact, for the case of utilitarian welfare, a range of strategies all yield a distortion of $k+1$.

Proposition 1. For the utilitarian welfare and for all $\alpha \in [\frac{1}{k+1}, 1]$, the allocation rule A_α achieves $D_{n,k}^{UW}(A_\alpha) \leq k+1$.

Proof. By Lemma 4 we have $D_{n,k}^{UW} \leq \max_\beta \frac{\beta+k}{\alpha\beta+(1-\alpha)}$. As the utilitarian welfare in any instance is at least 1, e.g. by giving all items to one agent, by Lemma 5 and setting $\beta = 1$ we have $D_{n,k}^{UW} \leq \frac{1+k}{\alpha+(1-\alpha)} = k+1$. \square

Lastly, the following theorem treats the case of $p > 1$.

Theorem 4. For $p \in (1, \infty)$, the allocation rule A_1 achieves $D_{n,k}^{M_p}(A_1) \leq (k+1)^{\frac{1}{p}}$.

In the full version, we present matching lower bounds for all of the upper bounds proven in this section.

4 Experiments

In this section, we measure the average-case approximation ratio of utilitarian welfare achieved by different rules based on synthetic and real-world data. In principle, one could conduct a similar analysis with other welfare measures, but we focus on utilitarianism for simplicity and conciseness.

Rules. We compare the following allocation rules motivated by Section 1.1: **Uniform** (allocate items uniformly to all agents), A_α with $\alpha = \frac{1}{k+1}$, $\alpha = \frac{n-k}{n}$, and $\alpha = 1$. Recall that A_1 returns a utilitarian welfare maximizing allocation for the nonsecretive agents, and all three of the A_α rules tested achieve the optimal distortion for utilitarian welfare.

Measurement. For a resource allocation instance with secretive agents, we measure the ratio between the maximum feasible welfare by full information and the welfare obtained by the rule, averaged over many instances. This provides us with an average-case analogue of the (worst-case) distortion.

4.1 Synthetic Data

Data Generation. We generate utilities for each agent, either secretive or nonsecretive, sampled i.i.d. from a Dirichlet distribution with m concentration parameters all set at 1, i.e. $\text{Dir}(1, \dots, 1)$. Each reported datum is the average of welfare ratios over 1000 randomly generated instances.

Experiments. We conduct three experiments each varying a parameter while fixing the others: vary k (Figure 2a), vary n with a fixed k (Figure 2b), vary n with a fixed ratio of k/n (Figure 2c), and vary m (in the full version).

Results. In all four figures we see a consistent relationship between the rules: rules with higher α outperform rules with lower α and all three of the A_α rules outperform **Uniform**. This is perhaps not surprising, since higher values of α more heavily exploit the information available to the rule from the non-secretive agents, with the **Uniform** rule being one example of an extreme case that ignores all available information about the utility functions.

In Figure 2a we see all three A_α rules achieve average welfare ratio 1 when $k = 0$, with the welfare ratio converging to that of **Uniform** when $k = n$, as expected. Of particular note is A_1 , which achieves an average welfare ratio close to 1 even for relatively large values of k (for example, the average welfare ratio is ~ 1.23 when $k = 10$) before rapidly increasing for large k . Of note is that all algorithms significantly outperform the worst-case bound of $k+1$ displayed with a dotted line in the figure.

Figure 2b reveals an interesting separation when A_1 and $A_{(n-k)/n}$ are compared to $A_{1/(k+1)}$ and **Uniform**. The average welfare ratio of the former rules decreases to 1 as n increases (with $k = 5$) while the average welfare ratio of the other rules actually increases with n . Figure 2c suggests that this increase persists even when the ratio k/n is held (approximately) constant.




Figure 2: Average welfare ratio achieved by different strategies. Error bands indicate the standard deviation. In all plots $m = 200$.

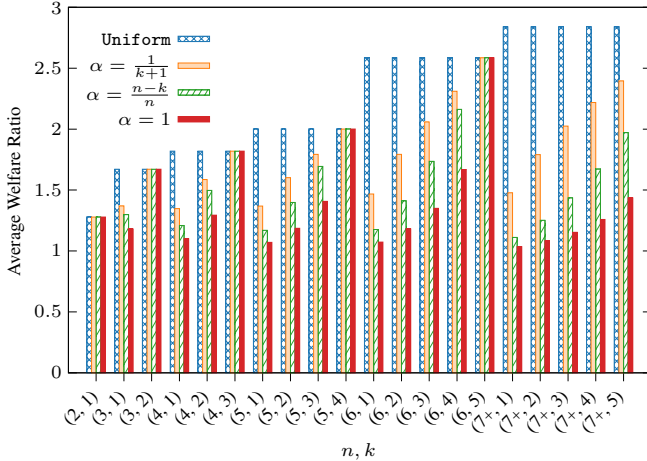


Figure 3: Average welfare ratio by different strategies on the Spliddit data. The x -axis is sorted by n and then k .

4.2 Spliddit Data

Data Generation. We also used the real-world goods division instances from Spliddit.org. For each instance with n agents and a fixed k , we randomly sampled k (secretive) agents, hid their utilities from the allocation rules and measured the welfare ratio based on the actual utilities. Similar to the simulated experiments, we report the average of 1000 such simulations.

Data Statistics. The report is based on 4679 Spliddit instances. The distribution of the number of agents n is $\{2 : 27.5\%, 3 : 67.3\%, 4 : 2.4\%, 5 : 1.7\%, 6 : 0.4\%, \text{ and } n \geq 7 : 0.7\%\}$. The number of goods m was in the range $[2, 96]$ with the mean and std. dev. of 31.1 ± 26.3 .

Experiments. We divided instances based on n and varied k from 1 to $\min(5, n - 1)$. The average welfare ratio is presented in Figure 3.

Results. In line with the results on synthetic data, we see higher α outperform lower α (and all outperform Uniform). The dependence on n and k also follows similar patterns as the synthetic case. Additionally, it is interesting to note the magnitude of the welfare ratio achieved by our rules. For Spliddit instances with 5 or fewer agents and at least 2 non-secretive agents ($k \leq n - 2$), the average welfare ratio is never higher than 1.5 for the rule A_1 . That is, on average, we

could achieve two-thirds of the maximum possible utilitarian welfare even if one or two agents do not respond to requests for their utility information.

5 Discussion

We studied distortion in resource allocation when k of the agents are secretive. For the utilitarian welfare, we identified a family of rules parametrized by $\alpha \in [1/k+1, 1]$ as worst-case optimal. Among this family, we find the rule with $\alpha = 1$ to be particularly interesting because it allocates no resources to the secretive agents and thus, unlike $\alpha < 1$, provides no incentive to an agent to be secretive. This can lead to fewer agents being secretive, which can further reduce distortion.

It is known that maximizing the Nash welfare yields desirable fairness guarantees [Caragiannis *et al.*, 2019; Conitzer *et al.*, 2019]. Fortunately, we find that the Nash welfare is the most approximable among all p -mean welfare functions.

Our work opens the door for interesting directions for future work. It would be interesting to study instance-wise optimal allocations, that is, allocations that minimize the worst-case approximation ratio on a given instance. It is likely that such allocations would more carefully decide which (and how much of) resources to allocate to the secretive agents depending on how highly they are valued by the non-secretive agents.

One may wish to reconcile distortion (welfare maximization) with fairness in the presence of secretive agents. If the goal is to only ensure fairness among the non-secretive agents, one can easily modify the rules proposed in this work by replacing $\text{OPT}_{\text{nonsec}}$ (the welfare-optimal allocation to the non-secretive agents) by an allocation to the non-secretive agents that maximizes welfare *subject to* the fairness guarantee. The additional loss in welfare incurred is precisely the *price of fairness*, which is well understood [Caragiannis *et al.*, 2012; Bertsimas *et al.*, 2011; Bei *et al.*, 2019; Barman *et al.*, 2020b].⁴ However, if the goal is to ensure fairness to *all* agents, it may be necessary that no more than a single agent is secretive, and even then, achieving fairness alone is already challenging [Arunachaleswaran *et al.*, 2019].

⁴For ensuring proportionality to the non-secretive agents, we would need $\alpha \geq (n - k)/n$, which can be set for $p \in [0, 1]$ while still using our analysis.

Acknowledgements

Shah was partially supported by an NSERC Discovery Grant.

References

- [Amanatidis *et al.*, 2021] Georgios Amanatidis, Georgios Birmpas, Aris Filos-Ratsikas, and Alexandros A Voudouris. A few queries go a long way: Information-distortion tradeoffs in matching. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence (AAAI)*, pages 5078–5085, 2021.
- [Anshelevich and Zhu, 2021] Elliot Anshelevich and Wennan Zhu. Ordinal approximation for social choice, matching, and facility location problems given candidate positions. *ACM Transactions on Economics and Computation (TEAC)*, 9(2):1–24, 2021.
- [Anshelevich *et al.*, 2018] Elliot Anshelevich, Onkar Bhardwaj, Edith Elkind, John Postl, and Piotr Skowron. Approximating optimal social choice under metric preferences. *Artificial Intelligence*, 264:27–51, 2018.
- [Anshelevich *et al.*, 2021] Elliot Anshelevich, Aris Filos-Ratsikas, Nisarg Shah, and Alexandros A Voudouris. Distortion in social choice problems: The first 15 years and beyond. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 4294–4301, 2021.
- [Arunachaleswaran *et al.*, 2019] Eshwar Ram Arunachaleswaran, Siddharth Barman, and Nidhi Rathi. Fair division with a secretive agent. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, volume 33, pages 1732–1739, 2019.
- [Asada *et al.*, 2018] Megumi Asada, Florian Frick, Vivek Pisharody, Maxwell Polevy, David Stoner, Ling Hei Tsang, and Zoe Wellner. Fair division and generalizations of sperner-and kkm-type results. *SIAM Journal on Discrete Mathematics*, 32(1):591–610, 2018.
- [Asadpour and Saberi, 2010] Arash Asadpour and Amin Saberi. An approximation algorithm for max-min fair allocation of indivisible goods. *SIAM Journal on Computing*, 39(7):2970–2989, 2010.
- [Aziz, 2020] Haris Aziz. Justifications of welfare guarantees under normalized utilities. *ACM SIGecom Exchanges*, 17(2):71–75, 2020.
- [Barman *et al.*, 2020a] Siddharth Barman, Umang Bhaskar, Anand Krishna, and Ranjani G. Sundaram. Tight approximation algorithms for p-mean welfare under subadditive valuations. In *Proceedings of the 28th Annual European Symposium on Algorithms (ESA)*, pages 11:1–11:17, 2020.
- [Barman *et al.*, 2020b] 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.
- [Bei *et al.*, 2019] X. Bei, X. Lu, P. Manurangsi, and W. Suksompong. The price of fairness for indivisible goods. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 81–87, 2019.
- [Benadè *et al.*, 2017] G. Benadè, S. Nath, A. D. Procaccia, and N. Shah. Preference elicitation for participatory budgeting. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pages 376–382, 2017.
- [Bertsimas *et al.*, 2011] D. Bertsimas, V. F. Farias, and N. Trichakis. The price of fairness. *Operations Research*, 59(1):17–31, 2011.
- [Bezáková and Dani, 2005] Ivona Bezáková and Varsha Dani. Allocating indivisible goods. *ACM SIGecom Exchanges*, 5(3):11–18, 2005.
- [Borodin *et al.*, 2019] A. Borodin, O. Lev, N. Shah, and T. Strangway. Primarily about primaries. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence (AAAI)*, pages 1804–1811, 2019.
- [Boutilier *et al.*, 2015] C. Boutilier, I. Caragiannis, S. Haber, T. Lu, A. D. Procaccia, and O. Sheffet. Optimal social choice functions: A utilitarian view. *Artificial Intelligence*, 227:190–213, 2015.
- [Caragiannis *et al.*, 2012] I. Caragiannis, C. Kaklamani, P. Kanellopoulos, and M. Kyropoulou. The efficiency of fair division. *Theory of Computing Systems*, 50(4):589–610, 2012.
- [Caragiannis *et al.*, 2017] I. Caragiannis, S. Nath, A. D. Procaccia, and N. Shah. Subset selection via implicit utilitarian voting. *Journal of Artificial Intelligence Research*, 58:123–152, 2017.
- [Caragiannis *et al.*, 2019] I. Caragiannis, D. Kurokawa, H. Moulin, A.D. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation (TEAC)*, 7(3):1–32, 2019.
- [Chèze, 2019] Guillaume Chèze. How to share a cake with a secret agent. *Mathematical Social Sciences*, 100:13–15, 2019.
- [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.
- [Frick *et al.*, 2019] Florian Frick, Kelsey Houston-Edwards, and Frédéric Meunier. Achieving rental harmony with a secretive roommate. *The American Mathematical Monthly*, 126(1):18–32, 2019.
- [Garg *et al.*, 2021] Jugal Garg, Edin Husić, Aniket Murhekar, and László Végh. Tractable fragments of the maximum nash welfare problem. arXiv:2112.10199, 2021.
- [Gkatzelis *et al.*, 2020] V. Gkatzelis, D. Halpern, and N. Shah. Resolving the optimal metric distortion conjecture. In *Proceedings of the 61st Symposium on Foundations of Computer Science (FOCS)*, pages 1427–1438, 2020.
- [Halpern and Shah, 2021] Daniel Halpern and Nisarg Shah. Fair and efficient resource allocation with partial information. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 224–230, 2021.
- [Lee, 2017] Euiwoong Lee. APX-hardness of maximizing nash social welfare with indivisible items. *Information Processing Letters*, 122:17–20, 2017.
- [Ma *et al.*, 2021] Thomas Ma, Vijay Menon, and Kate Larson. Improving welfare in one-sided matchings using simple threshold queries. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*, pages 5078–5085, 2021.
- [Micha and Shah, 2020] E. Micha and N. Shah. Can we predict the election outcome from sampled votes? In *Proceedings of the 44th AAAI Conference on Artificial Intelligence (AAAI)*, pages 2176–2183, 2020.
- [Moulin, 2003] H. Moulin. *Fair Division and Collective Welfare*. MIT Press, 2003.
- [Procaccia and Rosenschein, 2006] A. D. Procaccia and J. S. Rosenschein. The distortion of cardinal preferences in voting. In *Proceedings of the 10th International Workshop on Cooperative Information Agents (CIA)*, pages 317–331, 2006.