Ignorance is Often Bliss: Envy with Incomplete Information

Yiling Chen
Harvard University
yiling@seas.harvard.edu

Nisarg Shah
Harvard University
nisarg@g.harvard.edu

Abstract

We study the problem of fairly allocating indivisible goods to agents. Existing literature that focuses on the concept of envy inherently assumes that each agent can observe which goods the other agents are allocated before deciding if she envies them. In this paper, we propose a novel policy in which the principal hides from each agent the allocations made to the other agents. Each agent now attempts to infer these allocations, based on her knowledge of her own allocation, the mechanism used to divide the goods, and a prior over how the other agents value the goods, and then decides whether she envies them in expectation.

We propose eight measures for the total envy in the system, and show that under six of them, our policy is guaranteed to (weakly) reduce the total envy, for every mechanism on every instance. Our experiments show reduction of envy on synthetic as well as real-world data.

1 Introduction

In the mathematically rigorous study of fair division, which dates back to the work of Steinhaus [1948], perhaps the most difficult problem is to fairly divide a set of *indivisible goods* — goods which cannot be split or shared, such as houses, jewelry, or artwork — among a set of agents. The vast literature on fair division theory offers a slew of fairness desiderata such as envy-freeness [Foley, 1967], envy-freeness up to one good [Lipton *et al.*, 2004; Caragiannis *et al.*, 2016], proportionality, and the maximin share guarantee [Budish, 2011; Procaccia and Wang, 2014]. Perhaps the most prominent of them is envy-freeness, which informally requires that each agent value her own allocation at least as much as she values any other agent's allocation, i.e., that no agent envy any other agent. This definition inherently requires that every agent, in addition to knowing what she received, must know what every other agent received. We refer to this as the *full information case*. Often, the agents receive this knowledge directly from the principal making the division. What if the principal instead chooses to provide *incomplete information*, wherein each agent only knows her own allocation?

In some fair division applications, this can be simple to implement. For instance, in a cluster environment where heterogeneous computing cores with different CPU, RAM, and bandwidth configurations need to be divided among a set of processing jobs, agents submitting the jobs often observe the allocation of all cores through a digital interface, which can easily be modified to show to each agent only the cores that her own jobs received. In other applications, such as in estate division among heirs, this could be legally challenging because heirs have the legal right to obtain a copy of the will.

Nonetheless, it is interesting to consider the implications that such a policy may have. It is evident that incomplete information would lead to increased privacy, not only because agents are now oblivious to other agents' allocations, but also because they have less information to infer other agents' private valuations. But it is not apriori clear whether this policy would affect the degree of envy in the system.

Let us consider a simple example in which a house, a car, and a painting need to be divided among Alice, Bob, and Charlie. Let us assume that Alice has no distinctive information about how Bob or Charlie value the goods. In other words, from Alice's perspective every good that she does not receive is equally likely to be allocated to Bob or Charlie. Let us measure the total envy that Alice has by the number of agents she envies. In one scenario, imagine that Alice only wants the house, but does not receive it. With full information, she would envy only one agent who gets the house. But with incomplete information, she would envy both Bob and Charlie because according to her, they each receive the house with probability $1/2$, which is better than not receiving it at all. Thus, incomplete information *can increase* the total envy by Alice. On the other hand, imagine if Alice values the house at $\$100,000$ and the painting at $\$75,000$, does not care about the car, and receives only the painting. With full information, she would envy the agent who receives the house. But with incomplete information, she would think that Bob and Charlie each receive the house with probability $1/2$, have value $\$50,000$ for their allocations in expectation, and not envy either of them. Thus, incomplete information *can also reduce* the total envy by Alice.

In the above examples, we assumed specific allocations to demonstrate that incomplete information can lead to both increased and reduced envy. Can we always reduce (or increase) envy if we are allowed to choose the allocation care-fully? We also used a specific definition for the total envy by Alice. What happens if we measure it differently, say by paying attention to the difference in value observed by Alice?

1.1 Our Model and Results

We study fair division of goods in which there are n agents and m goods, each agent places a value on each good, and valuations are additive. In the standard framework, it is assumed that each agent i observes the entire allocation of goods to agents, and then envies an agent j if she values agent j's allocation more than she values her own allocation. We study a setting with incomplete information, in which agent i only knows her own allocation, infers about the allocations of other agents using a prior over their valuations, and decides her envy based on her expected value for their allocations.

In Section 2, we introduce eight measures for the total envy in the system. In Section 3, we first observe that under "symmetric beliefs", agents are envy-free with incomplete information if and only if they receive their proportional share. We also show that under six out of eight measures, total envy (weakly) reduces due to incomplete information, while for the other two measures, it could increase or decrease in the worst-case.

In Section 4, we present experiments on synthetic data in which we observe that in all our experiments, the average-case total envy in the system reduces due to incomplete information under all eight measures. For some cases, we also provide technical arguments establishing that the average-case envy in the full information setting is indeed qualitatively higher than in the incomplete information setting. We also present experiments on real-world goods division data from a fair division website `Spliddit.org`, in which we observe that the average-case total envy may slightly increase under two measures, but visibly reduces under the remaining six measures.

1.2 Related Work

To the best of our knowledge, there is no prior work on assessing envy in fair division when agents are oblivious to other agents' allocations, but have a prior over their valuations. However, a similar setup has been used previously to analyze other properties such as Bayesian incentive compatibility [Fujinaka, 2008].

The most closely related to ours is the work of Abebe *et al.* [2017], who studied the classic cake-cutting setting with a heterogeneous divisible good, but assume that the agents are connected via a social network structure, and each agent can only observe the allocations of her neighbors. They call an allocation locally envy-free if no agent envies her neighbors. As fully envy-free allocations (which are also locally envy-free) are guaranteed to exist in their setting, they focus on issues such as computational complexity and the price of fairness. Bouveret *et al.* [2010] also studied a setting with incomplete information, but in their setting it is the principal that does not have full information about agents' valuations, and attempts to argue whether a given allocation is necessarily or possibly envy-free.

We show that Bayesian envy-freeness coincides with proportionality for symmetric beliefs (Corollary 1). Hill [1987] demonstrated a lower bound on the minimum value each agent can be guaranteed as a fraction of their total value for all the goods, which can be viewed as an approximation to proportionality because proportionality requires that this fraction be at least $1/n$. Later, Markakis and Psomas [2011] provided a polynomial time algorithm for finding allocations that achieve this approximation, and Gourvès *et al.* [2014] provided improved guarantees in a more general setting. Clearly, these results directly inform approximation results for our Bayesian setting too. This also applies to results from the vast literature on the Santa Claus problem [Bansal and Sviridenko, 2006; Feige, 2008; Asadpour and Saberi, 2010], which also aims to approximate proportionality. We note that there is also a growing body of literature on approximating or relaxing envy-freeness [Lipton *et al.*, 2004; Caragiannis *et al.*, 2016; Budish, 2011; Nguyen and Rothe, 2013a], and that for symmetric beliefs Bayesian envy-freeness (i.e., proportionality) is also a relaxation of envy-freeness.

On a technical level, we use eight different measures for the total envy in the system, based on how envy is aggregated across different agents; this is inspired from similar definitions by Nguyen and Rothe [2013b].

2 Model

For $k \in \mathbb{N}$, define $[k] \triangleq \{1, \ldots, k\}$. Let $N = [n]$ denote a set of agents, and M denote a set of m goods to be divided among the agents. We assume that the goods are *indivisible*, i.e., a good cannot be split or shared between multiple agents. Each agent i is endowed with a valuation function $v_i : M \to \mathbb{R}$, where $v_i(g)$ denotes the value of agent i for good g. We assume additive valuations: slightly abusing the notation, we formally define $v_i(S) = \sum_{g \in S} v_i(g)$ for all $S \subseteq M$. We use $\mathbf{v} = (v_i)_{i \in N}$ to collectively denote the valuation functions of the agents.

An *allocation* $\mathbf{A}$ is a partition of the set of goods M among the set of agents N. We use A_i to denote the set of goods received by agent i under $\mathbf{A}$. The utility to agent i for her allocation is $v_i(A_i)$. An allocation is called *non-wasteful* if it allocates all the goods. We will assume non-wastefulness throughout the paper. An allocation $\mathbf{A}$ is said to satisfy *proportionality* if every agent receives her "proportional share", which is at least $1/n$ of her total value for the goods. In other words, proportionality requires $v_i(A_i) \geqslant v_i(M)/n$ for all $i \in N$. Note that an agent only needs to know her own allocation to assess whether she receives her proportional share. An allocation $\mathbf{A}$ is called *Pareto optimal* if no alternative allocation can make an agent happier without making some other agent strictly worse off, i.e., if for all alternative allocations $\mathbf{A}'$, $\exists i \in N, v_i(A'_i) > v_i(A_i)$ implies $\exists j \in N, v_j(A'_j) < v_j(A_j)$.

A mechanism takes as input the valuations of the agents and returns a (deterministic) allocation. We say that a mechanism satisfies a certain property (e.g., Pareto optimality, proportionality, or envy-freeness, which we define below) if it always returns an allocation satisfying the property.

Full information and envy-freeness. In the standard fair division setting, it is assumed that every agent can observe the entire allocation $\mathbf{A}$. In this case, $\mathbf{A}$ is called *envy-free* if

$v_i(A_i) \geqslant v_i(A_j)$ for all $i, j \in N$, i.e., if every agent values her own allocation at least as much as she values any other agent's allocation.

Incomplete information and Bayesian envy-freeness. The focus of this paper is to study the effect of the agents having incomplete information about other agents' allocation on envy. Specifically, each agent i knows her own valuation v_i and her allocation A_i (the mechanism used to divide the goods is also public knowledge), but does not have direct access to the allocations made to the other agents.

Instead, she has a prior over the valuations of the other agents: according to her, the valuation v_j of agent j is sampled (independently of other agents) from a distribution $\mathcal{D}_{i,j}$. Given her knowledge of the mechanism h used to divide the goods, and her own observed allocation A_i, she forms a posterior in which the probability of an allocation $\mathbf{A}'$ with $A'_i = A_i$ is given by

$$\Pr[\mathbf{A}'|A_i] = \frac{\Pr_{\{v_j \sim \mathcal{D}_{i,j}\}_{j \in N\setminus\{i\}}}[h(\mathbf{v}) = \mathbf{A}']}{\Pr_{\{v_j \sim \mathcal{D}_{i,j}\}_{j \in N\setminus\{i\}}}[h(\mathbf{v})_i = A_i]}.$$

Let us denote this posterior distribution by $\mathcal{D}_i$. To assess whether she envies another agent j, agent i now views her utility for agent j's allocation *in expectation*, where the expectation is taken over her uncertainty about the allocation. That is, agent i does not Bayesian-envy agent j if $v_i(A_i) \geqslant \mathbb{E}_{\mathbf{A}' \sim \mathcal{D}_i}[v_i(A'_j)]$. We say that an allocation $\mathbf{A}'$ is Bayesian envy-free if no agent envies another agent in expectation.

Quantitative measures of envy. It is easy to see that neither envy-freeness nor Bayesian envy-freeness can be guaranteed (imagine a single good divided between two agents). We therefore turn our attention to quantitative measures of the total envy in the system. To distinguish between envy in the full information and incomplete information cases, we call the former *full envy* (FE), and the latter *Bayesian envy* (BE).

Similarly to the three-step approach of Nguyen and Rothe [2013b], we take a four-step approach to define eight measures of the total envy in the system, starting from the envy of agent i for agent j, aggregating across j to find the total envy of agent i, and finally aggregating across i.

Step 1: The (basic) envy of agent i for agent j is given by

$$\text{FE}(i, j) = \max\big(v_i(A_j) - v_i(A_i), 0\big)$$
$$\text{BE}(i, j) = \max\big(\mathbb{E}_{\mathbf{A}' \sim \mathcal{D}_i}[v_i(A'_j)] - v_i(A_i), 0\big).$$

Step 2: Next, we either use this quantitative envy (i.e., use the identity operator $\mathcal{I}$ given by $\mathcal{I}(x) = x$), or convert it into a Boolean value (i.e., use the indicator operator $\mathbb{1}$ given by $\mathbb{1}(x) = 1$ if $x > 0$ and 0 otherwise). For $\text{E} \in \{\text{FE}, \text{BE}\}$,

$$\text{E}^{\mathcal{I}}(i, j) = \text{E}(i, j), \quad \text{E}^{\mathbb{1}}(i, j) = \mathbb{1}[\text{E}(i, j)].$$

Step 3: Next, to define the total envy of agent i, we aggregate her envy for all other agents by either adding (i.e., applying the sum operator Σ) or by taking the maximum (i.e., applying the max operator max). Formally, for $\text{E} \in \{\text{FE}, \text{BE}\}$ and $O_3 \in \{\mathcal{I}, \mathbb{1}\}$,

$$\text{E}^{\Sigma, O_3}(i) = \textstyle\sum_{j \in N\setminus\{i\}} \text{E}^{O_3}(i, j),$$
$$\text{E}^{\max, O_3}(i) = \max_{j \in N\setminus\{i\}} \text{E}^{O_3}(i, j).$$

Step 4: Finally, the total envy in the system is computed by aggregating the total envy of all agents, again, either by adding or by taking the maximum. Formally, for $\text{E} \in \{\text{FE}, \text{BE}\}$, $O_2 \in \{\Sigma, \max\}$, and $O_3 \in \{\mathcal{I}, \mathbb{1}\}$,

$$\text{E}^{\Sigma, O_2, O_3} = \textstyle\sum_{i \in N} \text{E}^{O_2, O_3}(i),$$
$$\text{E}^{\max, O_2, O_3} = \max_{i \in N} \text{E}^{O_2, O_3}(i).$$

For both the full information and the incomplete information case, this defines eight measures for the total envy in the system: E^{O_1, O_2, O_3} for $\text{E} \in \{\text{FE}, \text{BE}\}$, $O_1, O_2 \in \{\Sigma, \max\}$, and $O_3 \in \{\mathcal{I}, \mathbb{1}\}$. We remark that taking the sum of envies may be more meaningful when agents' valuations are normalized, i.e., if $\sum_{g \in M} v_i(g) = 1$ for all agents i. Our results in Section 3 hold both with and without such normalization.

3 Bayesian Envy versus Full Envy

In this section, we provide theoretical results comparing the total envy in the system in the full information case versus in the incomplete information case. For incomplete information, we focus on the special case of *symmetric beliefs*, where the prior that agent i has over the valuation function of agent j may depend on i, but not on j. Formally, we assume that $\mathcal{D}_{i,j} = \mathcal{D}_{i,j'}$ for all $i \in N$ and $j, j' \in N \setminus \{i\}$. We also assume that the mechanism is neutral, i.e., it does not distinguish between agents. With incomplete information, agent i now has no way to distinguish between the other agents. Thus, according to her, every good she does not receive might be allocated to each other agent with an equal probability. This leads to the following simple observation.

Proposition 1. *With incomplete information and symmetric beliefs, we have $v_i(A_i) \geqslant \mathbb{E}[v_i(A_j)]$ if and only if $v_i(A_i) \geqslant v_i(M)/n$, for all pairs of distinct agents $i, j \in N$.*

Proof. The goods the agent i does not receive are collectively valued $v_i(M) - v_i(A_i)$ by agent i, and each other agent j receives each such good with probability $1/(n-1)$ according to agent i. Hence, we have $\mathbb{E}[v_i(A_j)] = (v_i(M) - v_i(A_i))/(n-1)$. The desired result now follows by substituting this into $v_i(A_i) \geqslant \mathbb{E}[v_i(A_j)]$ and simplifying. ■

Note that $v_i(M)/n$ is the proportional share of agent i. Proposition 1 states that agent i does not Bayesian-envy other agents if and only if she receives her proportional share. As a direct corollary, Bayesian envy-freeness is equivalent to proportionality. Because envy-freeness with full information is known to imply proportionality, we find it interesting that with incomplete information it coincides with proportionality, which uncovers a stronger connection between the two classical fairness notions.

Corollary 1. *With incomplete information and symmetric beliefs, Bayesian envy-freeness is equivalent to proportionality, and is therefore implied by (full) envy-freeness.*

We do not focus on randomized mechanisms in this paper because the randomized Maximum Nash Welfare solution [Caragiannis *et al.*, 2016], or equivalently, the competitive equilibrium from equal incomes (CEEI) over randomized outcomes [Varian, 1974], is known to satisfy full envy-freeness, and from Corollary 1, it would satisfy Bayesian envy-freeness as well.

For deterministic mechanisms, on the other hand, neither full envy-freeness nor Bayesian envy-freeness can be guaranteed, as we observe in Section 2. We therefore ask whether the total amount of envy in the system can be reduced due to incomplete information as compared to the full information case. While *zero* full envy implies *zero* Bayesian envy (according to any of the eight measures) due to Corollary 1, the corollary makes no claim about the comparison of the total full envy and the total Bayesian envy in the case where the total full envy is *not zero*. In fact, in Section 1 we picked a particular measure — the total number of pairs (i, j) such that agent i envies agent j, i.e., $\mathrm{E}^{\Sigma,\Sigma,\mathbb{1}}$ for $\mathrm{E} \in \{\mathrm{FE}, \mathrm{BE}\}$ — and gave two instances of valuations and allocations such that in one case, the total full envy was higher, while in the other, the total Bayesian envy was higher. Our next result extends this incomparability observation to two of our eight measures and to (almost) all mechanisms, while showing that for the remaining six measures, the total Bayesian envy is always at most the total full envy, irrespective of the valuations or the allocation.

Theorem 1. *With symmetric beliefs and* $(O_1, O_2, O_3) \in \{\Sigma, \max\} \times \{\Sigma, \max\} \times \{\mathcal{I}, \mathbb{1}\}$*:*

1. *If* $(O_1, O_2, O_3) \in \{\Sigma, \max\} \times \{\Sigma\} \times \{\mathbb{1}\}$*, then* $\mathit{BE}^{O_1,O_2,O_3}$ *and* $\mathit{FE}^{O_1,O_2,O_3}$ *are incomparable. Formally, for every* $n \in \mathbb{N}$ *and every Pareto optimal mechanism, there exists an instance with* n *agents under which* $\mathit{BE}^{O_1,O_2,O_3} \leqslant \mathit{FE}^{O_1,O_2,O_3}$*, and for every* $n \in \mathbb{N}$ *and every mechanism, there exists an instance with* n *agents under which* $\mathit{BE}^{O_1,O_2,O_3} > \mathit{FE}^{O_1,O_2,O_3}$*.*
2. *Otherwise,* $\mathit{BE}^{O_1,O_2,O_3} \leqslant \mathit{FE}^{O_1,O_2,O_3}$ *for all mechanisms on all instances.*

Before we proceed to the proof, we remark that the Pareto optimality condition in part 1 of the theorem is weak and satisfied by most compelling mechanisms. In absence of this condition, the result may not hold. For example, consider the dictatorship mechanism that always allocates all the goods to agent 1. This violates Pareto optimality when there is a good that agent 1 has no value for, but another agent has a positive value for. Under this mechanism, agent 1 never envies another agent, while every other agent envies exactly one agent (agent 1) in the full information case and $n-1$ agents in the incomplete information case. Then, for every instance, we have $\mathrm{BE}^{O_1,\Sigma,\mathbb{1}} > \mathrm{FE}^{O_1,\Sigma,\mathbb{1}}$ for $O_1 \in \{\Sigma, \max\}$. Thus, part 1 of the theorem does not hold for this mechanism.

We also remark that our proof works both with and without assuming normalized agent valuations that sum to a fixed quantity; part 1 of the theorem is stronger when we insist on normalized valuations whereas part 2 is stronger when we do not assume normalized valuations.

Proof. Throughout this proof, let us denote BE^{O_2,O_3} by B and FE^{O_2,O_3} by E. We first prove the incomparability result. Fix $O_2 = \Sigma$ and $O_3 = \mathbb{1}$, and let $O_1 \in \{\Sigma, \max\}$.

To show that Bayesian envy can be strictly greater than full envy, consider an instance with n agents ($N = [n]$) and m goods ($M = \{g_1, \ldots, g_m\}$) in which all agents desire the same, single good. Formally, let the valuation of every agent i be given by $v_i(g_1) = 1$ and $v_i(g_j) = 0$ for $j \neq 1$. In any allocation, the agent i^* who receives good g_1 satisfies $B(i^*) = E(i^*) = 0$, whereas all other agents i satisfy $B(i) = n-1 > E(i) = 1$. It is easy to see that this implies $\sum_{i \in N} B(i) > \sum_{i \in N} E(i)$ as well as $\max_{i \in N} B(i) > \max_{i \in N} E(i)$, as required.

Next, we assume the mechanism is Pareto optimal, and show that full envy can be at least as much as Bayesian envy. Consider an instance with n agents ($N = [n]$) and $n+1$ goods ($M = \{g_1, \ldots, g_n, g_{n+1}\}$), in which the valuation of each agent i is given by $v_i(g_i) = 1/n$, $v_i(g_{n+1}) = (n-1)/n$, and $v_i(g_j) = 0$ for $j \notin \{i, n+1\}$.

The Pareto optimal mechanism must allocate good g_i to agent i for all $i \in [n]$. For the agent i^* that receives good g_{n+1}, we have $B(i^*) = E(i^*) = 0$, and for every other agent i, we have $B(i) = 0$ and $E(i) = 1$. One can now check that $\sum_{i \in N} B(i) < \sum_{i \in N} E(i)$ as well as $\max_{i \in N} B(i) < \max_{i \in N} E(i)$, as required.

We now show that for the remaining six measures, total Bayesian envy is no more than total full envy for every mechanism on every instance. To prove this, we show that for every $(O_2, O_3) \in \{(\max, \mathbb{1}), (\Sigma, \mathcal{I}), (\max, \mathcal{I})\}$, every instance, and every allocation, $B(i) \leqslant E(i)$ for every agent $i \in N$. Note that the inequality will be preserved when aggregation across the agents is performed by $O_1 \in \{\Sigma, \max\}$.

Fix an instance with valuations $\{v_i\}_{i \in N}$ and an allocation A. The argument for $O_2 = \max$ and $O_3 = \mathbb{1}$ is simple. For an agent i, we have $B(i), E(i) \in \{0, 1\}$. Further, $E(i) = 0$ implies $v_i(A_i) \geqslant v_i(A_j)$ for all $j \in N \setminus \{i\}$. Summing over j, we get $(n-1)v_i(A_i) \geqslant v_i(M) - v_i(A_i)$, i.e., $v_i(A_i) \geqslant v_i(M)/n$, which in turn implies $B(i) = 0$ (Proposition 1). Hence, $E(i) \geqslant B(i)$, as required.

For $O_2 = \max$ and $O_3 = \mathcal{I}$, we want to show

$$\max_{j \in N \setminus \{i\}} \max\{v_i(A_j) - v_i(A_i), 0\} \geqslant \max\left\{\frac{v_i(M) - n\, v_i(A_i)}{n-1}, 0\right\}, \tag{1}$$

where the RHS follows from the fact that in the incomplete information case, agent i has the same amount of envy for every other agent j. This is given by $(v_i(M) - v_i(A_i))/(n-1) - v_i(A_i)$, which simplifies to the expression on the RHS above. For $O_2 = \Sigma$ and $O_3 = \mathcal{I}$, we want to show

$$\sum_{j \in N \setminus \{i\}} \max\{v_i(A_j) - v_i(A_i), 0\} \geqslant (n-1) \cdot \max\left\{\frac{v_i(M) - n\, v_i(A_i)}{n-1}, 0\right\}, \tag{2}$$

which is equivalent to replacing $\max_{j \in N \setminus \{i\}}$ in Equation (1) with average over $j \in N \setminus \{i\}$. Because the maximum of a set of numbers is no less than their average, this implies Equation (1). For Equation (2), note that the LHS is trivially at least 0. We also have that

$$
\begin{aligned}
&\frac{1}{n-1}\sum_{j\in N\setminus\{i\}} \max\{v_i(A_j)-v_i(A_i),0\}\\
&\quad\geqslant \frac{1}{n-1}\sum_{j\in N\setminus\{i\}} v_i(A_j)-v_i(A_i)\\
&\quad= \frac{1}{n-1}\Big((v_i(M)-v_i(A_i))-(n-1)\cdot v_i(A_i)\Big)\\
&\quad= \frac{v_i(M)-n\,v_i(A_i)}{n-1},
\end{aligned}
$$

as required. ■

4 Experiments

We now present experiments on synthetic data as well as real-world data, in which we measure the total envy *on average*, with full and incomplete information. To make the envy by different agents comparable, we normalize valuations to sum to 1 before calculating envy. To make the different measures of total envy comparable, we replace summation by average; note that this simply amounts to dividing the total envy by a constant that depend on the number of agents, and does not change the comparison between envy with full and incomplete information.

As the mechanism to allocate goods to agents, we use the maximum Nash welfare (MNW) solution, which maximizes the product of agents' utilities. This mechanism is currently being used in practice on the fair division website `Spliddit.org` due to its attractive fairness guarantees [Caragiannis *et al.*, 2016]. Hence, the use of this mechanism tells us how incomplete information would affect the total envy in practice.

For the special case of symmetric beliefs, Bayesian envy has a simpler expression that only depends on the current allocation (and not on the mechanism used), which allows us to minimize this envy by solving an integer linear program.[1] In this case, for each combination of (O_1, O_2, O_3), we additionally compare the minimum possible full envy $\mathrm{FE}^{O_1,O_2,O_3}$ to the minimum possible Bayesian envy $\mathrm{BE}^{O_1,O_2,O_3}$, on average. In this case, we say the mechanisms used are the envy-minimizing (or MinEnvy) mechanisms. This compares the potential to reduce total envy under full and incomplete information, if we really sought to minimize it, but does not provide comparison for a specific mechanism.

Synthetic data. We use three methods for generating synthetic agent valuations. In each case, an agent's prior for other agents' valuations accurately reflects the distributions from which they were sampled. For each agent i and good g, we independently sample $v_i(g)$ from the normal distribution $\mathcal{N}(\mu_i(g), 1)$, truncated below 0, because negative valuations are not allowed. The methods differ in their choice of $\mu_i(g)$.

- Symmetric beliefs, homogeneous goods: $\mu_i(g) = 1/2$ for all i, g, i.e., all agents and goods are apriori identical.
- Symmetric beliefs, heterogeneous goods: $\mu_i(g) = \mu(g) \sim U[0,1]$ for all i, g, i.e., agents are apriori identical, but goods have different "market values" ($\mu_i(g)$).
- Asymmetric beliefs: $\mu_i(g) \sim U[0,1]$ for all i, g, i.e., all agents and goods are apriori different.

Due to space constraint, we only present graphs for the special case of $m = n$. Results for $m > n$ are similar, but the difference between Bayesian and full envy is smaller. We also omit the graphs for symmetric beliefs, heterogeneous goods, which are similar to the homogeneous goods case.

Each datapoint in our graphs is averaged over $1,000$ random valuations. For symmetric beliefs, we can quickly compute total Bayesian envy, allowing us to test $n = 10$ to $n = 50$. For asymmetric beliefs, we need a computationally intensive step, in which for each agent i, we repeatedly sample valuations of other agents from her prior, run the MNW solution until we find 100 samples in which agent i's allocation coincides with her actual allocation, and take her average value for other agents' allocations. This restricts our simulations to use $n = 5$ to $n = 25$.

Figures 1a and 1b show the comparison between full envy (solid lines) and Bayesian envy (dashed lines) on average under all eight measures, for MinEnvy and the MNW solution, respectively. Note that we observe qualitatively less Bayesian envy (which quickly drops to 0) than full envy (which either remains a constant or drops very slowly). This trend holds under asymmetric beliefs (Figure 1c) as well.

For MinEnvy under symmetric beliefs, the results can be explained theoretically. For $m = n$, Suksompong [2016] shows that under mild conditions, the probability that a proportional (i.e., Bayesian envy-free) allocation does not exist drops exponentially with n. Given the polynomial upper bound on the maximum amount of envy, we see that the expected (minimum possible) Bayesian envy must also drop exponentially with n. In contrast, there is at least a constant probability that no envy-free allocation exists [Dickerson *et al.*, 2014]. At least for some of the measures, it is possible to use this result to derive an at most polynomially decreasing lower bound on the expected (least possible) full envy.

Theorem 1 ensures that Bayesian envy can never be more than full envy for six measures under symmetric beliefs. For the remaining two measures, it was still lower in more than 95% of our simulations except for $n = 10$. Under asymmetric beliefs, we have no theoretical guarantee, but observe that Bayesian envy is again lower than full envy under all eight measures in more than 90% of our simulations, except for one measure at $n = 5$.

Real-world data. We use data from the fair division website `Spliddit.org` that allows fairly allocating a mix of divisible and indivisible goods. In particular, we use the $2,028$ instances created so far in which only indivisible goods were used. Unlike the synthetic data, this no longer has $m = n$. The number of agents n vary from 2 to 15, while the number of goods m vary from 2 to 93.

Because the valuations are already given, we only need to decide the priors. Both the "homogeneous goods" and "heterogeneous goods" priors (e.g., with $\mu_i(g) = \mathrm{avg}_i\ v_i(g)$ as the "market value" of good g) lead to symmetric beliefs, and

[1]For $m = n$, this can be accomplished in polynomial time solving maximum-weight bipartite matching when $O_1 = \Sigma$ or Linear Bottleneck Assignment Problem (LBAP) when $O_1 = \max$, but we omit the details due to lack of space.

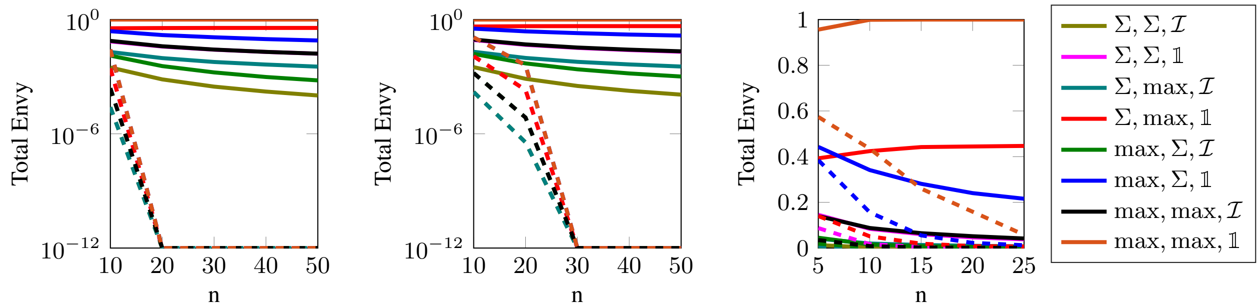

(a) MinEnvy, symmetric beliefs (b) MNW, symmetric beliefs (c) MNW, asymmetric beliefs

Fig. 1: Experiments on synthetic data with symmetric beliefs (homogeneous goods) and asymmetric beliefs. Solid and dashed lines represent total envy with full and incomplete information, respectively.

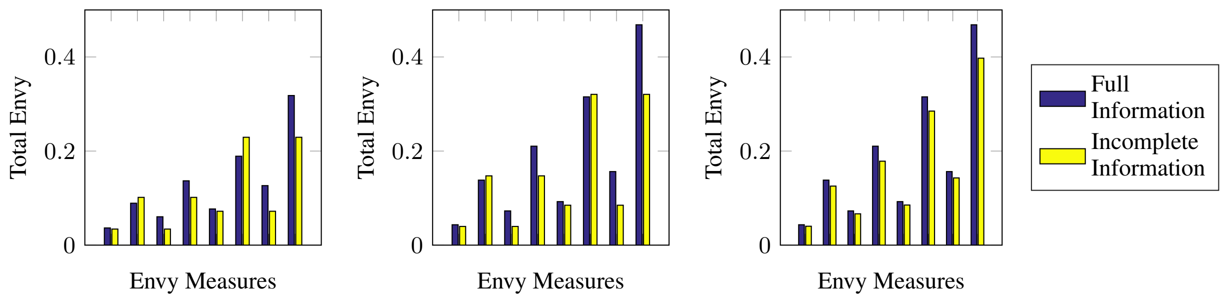

(a) MinEnvy, symmetric beliefs (b) MNW, symmetric beliefs (c) MNW, asymmetric beliefs

Fig. 2: Experiments on real-world data from `Spliddit.org` with symmetric beliefs and asymmetric beliefs. The eight envy measures, left to right, are in the same order as they appear in the legend of Figure 1, top to bottom.

with the valuations already fixed, lead to identical results. For asymmetric beliefs, we set $\mu_i(g) = v_j(g)$, i.e., priors are distributed around the true valuations. Figures 2a and 2b present the results for symmetric beliefs under the MinEnvy and the MNW solution, respectively, while Figure 2c presents the results for asymmetric beliefs under the MNW solution.

First, note that for asymmetric beliefs, Bayesian envy is still less than full envy on average under all eight measures (in fact, in more than 90% of simulations in each comparison). For symmetric beliefs, this comparison is guaranteed to always hold for six measures, and we observe a significant reduction under some of them. However, Bayesian envy is slightly higher on average than full envy under the two measures captured in part 1 of Theorem 1, for which envy in the two cases is incomparable in the worst case. Surprisingly, in more than 80% of the simulations, Bayesian envy is actually lower. A closer investigation revealed that data from Spliddit contains instances with "concentrated" valuations, in which more than one agent want the same good, and value other goods negligibly. The agent who does not receive this good envies only one agent in the full information case, but $n-1$ agents in the incomplete information case. Each such instance creates a dramatic increase in Bayesian envy when $O_2 = \Sigma$ and $O_3 = \mathbb{1}$. While such concentrated valuations may be common in some of the applications that Spliddit caters to (e.g., estate division), in other applications such as resource allocation in clusters, it is far less common because all cores with similar configuration are valued similarly.

5 Discussion

In this paper, we considered the problem of fair allocation of indivisible goods, proposed a novel policy wherein the principal hides from each agent the allocations of the other agents, and showed that it helps reduce the amount of total conflict (measured by aggregating envy of agents in different ways) both in theory and in practice.

We remark that such a policy could have additional benefits. For example, hiding the allocations of other agents can not only lead to increased privacy for the allocation, but can also make it harder for agents to infer the private valuations of other agents. An interesting direction for future work is to quantify the amount of privacy added by using notions of privacy such as differential privacy [Dwork *et al.*, 2006] or (information-theoretic) min-entropy leakage [Smith, 2009], which are closely related to each other [Alvim *et al.*, 2012].

We also foresee challenges in implementing this policy in the real world. First, effectively hiding information from the agents may require use of cryptographic schemes for secure multi-party computation [Yao, 1982] that provide privacy with respect to both inputs and outputs [Bresson *et al.*, 2006]. It would be interesting to study if outcomes under classic fair division mechanisms such as the MNW solution can be computed in a secure way.

References

R. Abebe, J. Kleinberg, and D. Parkes. Fair division via social comparison. In *Proceedings of the 16th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 2017. Forthcoming.

M. Alvim, M. Andrés, K. Chatzikokolakis, P. Degano, and C. Palamidessi. Differential privacy: On the trade-off between utility and information leakage. *Formal Aspects of Security and Trust*, pages 39–54, 2012.

A. Asadpour and A. Saberi. An approximation algorithm for max-min fair allocation of indivisible goods. *SIAM Journal on Computing*, 39(7):2970–2989, 2010.

N. Bansal and M. Sviridenko. The Santa Claus problem. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, pages 31–40, 2006.

S. Bouveret, U. Endriss, and J. Lang. Fair division under ordinal preferences: Computing envy-free allocations of indivisible goods. In *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI)*, pages 387–392, 2010.

E. Bresson, D. Catalano, N. Fazio, A. Nicolosi, and M. Yung. Output privacy in secure multiparty computation. Manuscript, 2006.

E. Budish. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes. *Journal of Political Economy*, 119(6):1061–1103, 2011.

I. Caragiannis, D. Kurokawa, H. Moulin, A. D. Procaccia, N. Shah, and J. Wang. The unreasonable fairness of maximum Nash welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, pages 305–322, 2016.

J. P. Dickerson, J. Goldman, J. Karp, A. D. Procaccia, and T. Sandholm. The computational rise and fall of fairness. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pages 1405–1411, 2014.

C. Dwork, F. McSherry, K. Nissim, and A. Smith. Calibrating noise to sensitivity in private data analysis. In *Theory of Cryptography Conference*, pages 265–284, 2006.

U. Feige. On allocations that maximize fairness. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 287–293, 2008.

D. Foley. Resource allocation and the public sector. *Yale Economics Essays*, 7:45–98, 1967.

Y. Fujinaka. A Bayesian incentive compatible mechanism for fair division. Manuscript, 2008.

L. Gourvès, J. Monnot, and L. Tlilane. Near fairness in matroids. In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI)*, pages 393–398, 2014.

T. Hill. Partitioning general probability measures. *Annals of Probability*, 15(2):804–813, 1987.

R. J. Lipton, E. Markakis, E. Mossel, and A. Saberi. On approximately fair allocations of indivisible goods. In *Proceedings of the 6th ACM Conference on Economics and Computation (EC)*, pages 125–131, 2004.

E. Markakis and C.-A. Psomas. On worst-case allocations in the presence of indivisible goods. In *Proceedings of the 5th Conference on Web and Internet Economics (WINE)*, pages 278–289, 2011.

T. T. Nguyen and J. Rothe. Envy-ratio and average-Nash social welfare optimization in multiagent resource allocation. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 1139–1140, 2013.

T. T. Nguyen and J. Rothe. How to decrease the degree of envy in allocations of indivisible goods. In *Proceedings of the 3rd International Conference on Algorithmic Decision Theory (ADT)*, pages 271–284, 2013.

A. D. Procaccia and J. Wang. Fair enough: Guaranteeing approximate maximin shares. In *Proceedings of the 14th ACM Conference on Economics and Computation (EC)*, pages 675–692, 2014.

G. Smith. On the foundations of quantitative information flow. *Foundations of Software Science and Computational Structures*, pages 288–302, 2009.

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

W. Suksompong. Asymptotic existence of proportionally fair allocations. *Mathematical Social Sciences*, 81:62–65, 2016.

H. Varian. Equity, envy and efficiency. *Journal of Economic Theory*, 9:63–91, 1974.

A. C. Yao. Protocols for secure computations. In *Proceedings of the 23rdAnnual Symposium on Foundations of Computer Science (SFCS)*, pages 160–164, 1982.