Robust Rent Division

Dominik Peters
University of Toronto
dominik@cs.toronto.edu

Ariel D. Procaccia
Harvard University
arielpro@seas.harvard.edu

David Zhu
Harvard University
david.zhu@gmail.com

Abstract

In fair rent division, the problem is to assign rooms to roommates and fairly split the rent based on roommates' reported valuations for the rooms. Envy-free rent division is the most popular application on the fair division website Spliddit. The standard model assumes that agents can correctly report their valuations for each room. In practice, agents may be unsure about their valuations, for example because they have had only limited time to inspect the rooms. Our goal is to find a robust rent division that remains fair even if agent valuations are slightly different from the reported ones. We introduce the lexislack solution, which selects a rent division that remains envy-free for valuations within as large a radius as possible of the reported valuations. We also consider robustness notions for valuations that come from a probability distribution, and use results from learning theory to show how we can find rent divisions that (almost) maximize the probability of being envy-free, or that minimize the expected envy. In particular, we show that we can identify an almost optimal allocation based on polynomially many samples from the valuation distribution. We show that finding the best allocation given these samples is NP-hard, but in practice such an allocation can be computed using integer linear programming.

1 Introduction

The literature on fair division of resources has produced allocation mechanisms for many domains, such as course allocation, indivisible goods, chores, house assignment, and the selection of citizens' assemblies (Budish 2011; Caragiannis et al. 2019; Moulin 2019; Flanigan et al. 2021). But arguably the most widely *used* example is *rent division*: this is the most popular application on the fair division website spliddit.org (Goldman and Procaccia 2014), where it has been used more than 30,000 times since its launch in 2014.

Rent division deals with the common situation where a group of n future roommates are planning to move into a house or apartment which has n rooms, one for each roommate. They will split the rent payments among themselves. The roommates may differ in how much they are willing to pay for different rooms. Given the room valuations of each roommate, our task is to assign the rooms, and to decide how to split the rent. We wish to do this fairly, and so we will choose an allocation that is *envy-free*: no roommate would strictly prefer to get another room, given the prices we have assigned to those rooms. Such an allocation is guaranteed to exist (Svensson 1983).

	Room 1	Room 2	Room 3
Alice	**300**	400	300
Bob	300	**700**	0
Charlie	300	100	**600**
Spliddit	100	500	400
Lexislack	200	450	350

Table 1: Example of valuations. In any envy-free allocation, Alice gets room 1, Bob gets room 2, and Charlie gets room 3. The lower rows display the price vectors selected by Spliddit's rule (maximin) and by our lexislack rule.

Let us consider an example with $n = 3$ roommates, and let the total rent be \$1000. Table 1 shows the valuation that each agent assigns to each room. Given this information, the algorithm in use on Spliddit will assign Room 1 to Alice, Room 2 to Bob, and Room 3 to Charlie, charging them \$100, \$500, and \$400 respectively. This allocation is envy-free under the assumption (which we will make throughout) that agents have *quasilinear utilities*: their utility under an allocation is the value of their room minus its price. For example, Alice has utility $300 - 100 = 200$. She does not envy the others: Bob's room would give her only utility $400 - 500 < 200$, and Charlie's room $300 - 400 < 200$.

On a typical instance, there are infinitely many allocations that are envy-free. Spliddit's algorithm chooses the one that maximizes the utility of the worst-off agent, subject to envy-freeness (Gal et al. 2017). This is known as the *maximin* rule. In optimizing this objective, Spliddit might choose an outcome that is only barely envy-free. In the example, Bob has utility $700 - 500 = 200$, but he would gain the same utility from having Alice's room: $300 - 100$. If, upon moving in, Bob discovers a defect in his room and now only values it at 600, say, then he would envy Alice. Thus, the envy-freeness of Spliddit's allocation is not robust.

We study the rent division problem with the goal of finding allocations that are robustly envy-free, in the sense that they remain envy-free even if valuations change slightly. For this, we introduce the *lexislack rule*, which selects an envy-free allocation where the minimum "slack" (the amount by which agent i prefers her allocation to agent j's) is maximized lexicographically. This produces an allocation that remains envy-free for all valuation profiles that are within a maximally large ℓ_1-radius of the reported profile. In the example of Table 1, the lexislack rule assigns the rooms in the same way as does Spliddit, but charges the roommates \$200, \$450, and \$350. One can verify that with these prices, each agent prefers their allocation to any other agent's by at least 150. This means that even after Bob's adjustment to 600, he does not envy Alice. We show that the lexislack rule always selects an essentially unique outcome, which can be found in polynomial time by linear programming.

This notion of robustness based on ℓ_1-distance may not always be appropriate. A more flexible model would assume that valuations are drawn from some probability distribution $\mathcal{D}$. For example, we could add some Gaussian noise around the reported valuations. In this setting, it is natural to look for allocations that are envy-free with maximum probability. Another target could be an allocation that minimizes the expected amount by which one agent envies another. To find such allocations, we propose an approach based on sampling: we sample a number of valuation profiles from $\mathcal{D}$, and then find an allocation that is optimal on this sample, for example using integer linear programming (ILP). We give polynomial sample complexity bounds which show how many samples are sufficient so that this approach identifies an almost optimal allocation with high probability. To prove these bounds, we draw on a somewhat unexpected connection to PAC learning. We also consider the computational complexity of finding the best allocation on a given sample, and prove that these optimization problems are NP-complete. The ILP approach to these problems appears to be feasible in practice, however.

We end with some experiments on data taken from Spliddit. They suggest that our three new rules significantly outperform the Spliddit maximin rule on robustness metrics. Interestingly, the lexislack solution does comparably well to the rules based on sampling. Given its conceptual simplicity and easy computation, this suggests lexislack as a good rule for applications in which robustness is desired.

Related Work

The rent division model is well-studied in the economics literature (Svensson 1983; Alkan, Demange, and Gale 1991; Aragones 1995; Su 1999; Velez 2018), often without assuming quasilinear utilities. That literature includes results on the structure of the envy-free set and about strategic aspects. Computer scientists have studied the computation of allocation rules (Gal et al. 2017; Procaccia, Velez, and Yu 2018). Bei et al. (2021) study a generalization of the rent division problem.

Robustness has been studied in several areas of computational social choice, such as in voting (Shiryaev, Yu, and Elkind 2013), in committee elections (Bredereck et al. 2021; Gawron and Faliszewski 2019; Misra and Sonar 2019), and in stable matching (Chen, Skowron, and Sorge 2019; Mai and Vazirani 2018). We are not aware of such work for fair division, though Menon and Larson (2020) study a related problem of "stability" which requires that the allocation should not change much if valuations change slightly. For rent division, a blog post by Critch (2015) argues in favor of aiming for robustness in the rent division problem. Critch (2015) implemented an algorithm for robust rent division that appears in experiments to maximize the slack, but it does not coincide with the lexislack rule, and a theoretical analysis of this algorithm is not available.

2 Preliminaries: Rent Division

Let $n \in \mathbb{N}$ and write $[n] = \{1, \dots, n\}$. Let $N = [n]$ be a set of n agents, and let $R = [n]$ be a set of n rooms. Without loss of generality, we let the total rent be 1. A *(valuation) profile* $v = (v_{ir})_{i \in N, r \in R}$ is a collection of values $v_{ir} \in \mathbb{Q}_+$, one for each agent $i \in N$ and each room $r \in R$.

A *room assignment* is a bijection $\sigma : N \to R$, so that agent i is assigned room $\sigma(i)$. Given a valuation profile v, we say that σ is *optimal* if it maximizes utilitarian social welfare $\sum_{i \in N} v_{i\sigma(i)}$. An *allocation* (σ, p) is a room assignment σ together with a payment vector $p = (p_1, \dots, p_n) \in \mathbb{R}^n$ with $\sum_{r \in R} p_r = 1$, where p_r is the rent of room r. (The value p_r is usually but not necessarily non-negative.)

We assume that agents have *quasilinear utilities*. This means that if v is a valuation profile and (σ, p) is an allocation, then agent i's utility in this allocation is $v_{i\sigma(i)} - p_{\sigma(i)}$, i.e., the valuation of i for her room $\sigma(i)$ minus the room's rent. An allocation (σ, p) is *envy-free* if $v_{i\sigma(i)} - p_{\sigma(i)} \geqslant v_{ir} - p_r$ for all $i \in N$ and $r \in R$, so that each agent i weakly prefers her allocation to receiving any other room.

The following facts about rent division are well-known (see, e.g., Velez 2018). We include proofs for convenience.

Theorem 1. *(a) For every optimal room assignment σ, there are prices p such that (σ, p) is envy-free.*
(b) If (σ, p) is envy-free then σ is optimal.
(c) Let σ_1 and σ_2 be optimal room assignments, and let (σ_1, p) be an envy-free allocation. Then (σ_2, p) is also an envy-free allocation, with all agents indifferent between the two: $v_{i\sigma_1(i)} - p_{\sigma_1(i)} = v_{i\sigma_2(i)} - p_{\sigma_2(i)}$ for all $i \in N$.

Proof. (a) An optimal room assignment σ forms a solution to the standard assignment problem (see, e.g., Wolsey 1998, Section 4.3). The dual of the assignment problem LP is

$$\min \sum_{i \in N} q_i + \sum_{r \in R} p_r \ \text{ s.t. } \ q_i + p_r \geqslant v_{ir} \ \text{ for } i \in N, r \in R.$$

Since σ is an optimal room assignment, by complementary slackness there exists a solution $(q_i), (p_r)$ to the dual program where $q_i + p_{\sigma(i)} = v_{i\sigma(i)}$ for each $i \in N$. Thus

$$v_{i\sigma(i)} - p_{\sigma(i)} = q_i \geqslant v_{ir} - p_r \quad \text{for all } i \in N, r \in R,$$

using dual feasibility. Thus (p_r) is an envy-free price vector, but we must ensure that $\sum_{r \in R} p_r = 1$, which we can do by adding a constant to each p_r. This preserves envy-freeness.

(b) Suppose (σ, p) is envy-free and σ' is any room assignment. Then $\sum_{i \in N} v_{i\sigma(i)} \geqslant \sum_{i \in N} (v_{i\sigma'(i)} - p_{\sigma'(i)} + p_{\sigma(i)}) = (\sum_{i \in N} v_{i\sigma'(i)}) - 1 + 1 = \sum_{i \in N} v_{i\sigma'(i)}$, where the inequality follows from envy-freeness. Thus σ has at least the welfare of σ'. Since σ' was arbitrary, σ is an optimal assignment.

(c) We show $v_{i\sigma_1(i)} - p_{\sigma_1(i)} = v_{i\sigma_2(i)} - p_{\sigma_2(i)}$ for $i \in N$. From this, envy-freeness of (σ_2, p) follows immediately. We have $v_{i\sigma_1(i)} - p_{\sigma_1(i)} \geqslant v_{i\sigma_2(i)} - p_{\sigma_2(i)}$ for all $i \in N$ since (σ_1, p) is envy-free. Sum these inequalities to get

$$(\textstyle\sum_{i\in N} v_{i\sigma_1(i)}) - 1 \geqslant (\sum_{i\in N} v_{i\sigma_2(i)}) - 1.$$

But the two sides of this inequality are equal, since both σ_1 and σ_2 are optimal. Hence each inequality is satisfied with equality, as required. □

Theorem 1(a) implies that an envy-free allocation exists for all valuation profiles. It also implies that we can find one in polynomial time: find an optimal room assignment σ using bipartite matching, then use linear programming to find prices p that make the allocation envy-free (Gal et al. 2017).

Theorem 1(c) implies that when selecting among envy-free allocations, we can restrict attention to any fixed σ and only consider allocations of the form (σ, p). By Theorem 1(c), all utility vectors achievable in an envy-free allocation are achieved by allocations of this form.

3 The Lexislack Solution

We start by considering a common form of robustness: we look for allocations that remain fair for all valuations that are within some radius of input valuations, for as large a radius as possible. Thus, unlike in later sections, we do not assume that valuations come from a probability distribution $\mathcal{D}$.

Let v be a valuation profile, fixed throughout. Let (σ, p) be an allocation. For $i \in N$ and $r \in R$, define

$$\Delta_{ir}(\sigma, p) = (v_{i\sigma(i)} - p_{\sigma(i)}) - (v_{ir} - p_r).$$

Then define the *slack* of this allocation as

$$\mathsf{slack}(\sigma, p) = \min_{i\in N} \min_{r\neq\sigma(i)} \Delta_{ir}.$$

Thus, an allocation has positive slack if every agent strictly prefers their allocation to all other agents' allocations. An allocation (σ, p) is envy-free if and only if $\mathsf{slack}(\sigma, p) \geqslant 0$.

Slack is a measure of how robustly fair an allocation is, which we can make precise in the following result.

Proposition 2. *Let (σ, p) be an envy-free allocation with $\mathsf{slack}(\sigma, p) = s \geqslant 0$. If v' is a valuation profile that is s-close to v in the sense that*

$$\|v_i - v'_i\|_1 = \textstyle\sum_{r\in R} |v_{ir} - v'_{ir}| \leqslant s$$

for all $i \in N$, then (σ, p) is also envy-free under v'.

Proof. Let $i, j \in N$. Then $\sum_{r\in R} |v_{ir} - v'_{ir}| \leqslant s$ implies

$$(v_{i\sigma(i)} - v'_{i\sigma(i)}) + (v_{i\sigma(j)} - v'_{i\sigma(j)}) \leqslant s \qquad (1)$$

Adding $p_{\sigma(j)} - p_{\sigma(i)}$ to both sides and rearranging, we get

$$\begin{aligned}&(v'_{i\sigma(i)} - p_{\sigma(i)}) - (v'_{i\sigma(j)} - p_{\sigma(j)})\\ &\geqslant (v_{i\sigma(i)} - p_{\sigma(i)}) - (v_{i\sigma(j)} - p_{\sigma(j)}) - s \geqslant 0\end{aligned}$$

by definition of slack. Thus, i does not envy j under v'. □

One can also prove alternatives version of Proposition 2. For example, $\|v_i - v'_i\|_\infty \leqslant s/2$ is sufficient to imply (1).

If we wish to ensure robustness in a sense like in Proposition 2, this suggests the following allocation rule:

$$\mathsf{maxislack}(v) = \mathsf{argmax}_{(\sigma,p)}\ \mathsf{slack}(\sigma, p).$$

This rule always selects an envy-free allocation: since envy-free allocations exist for every v, there exists an allocation with non-negative slack, and hence the maxislack solution also has non-negative slack. A maxislack solution can be found in polynomial time by computing an optimal assignment σ and then solving the following LP:

$$\begin{aligned}\max L \text{ s.t. } &(v_{i\sigma(i)} - p_{\sigma(i)}) - (v_{i\sigma(j)} - p_{\sigma(j)}) \geqslant L\ \forall i \neq j\\ &\textstyle\sum_{r\in R} p_r = 1\end{aligned}$$

However, there are a few drawbacks to the maxislack rule. First, the rule is not essentially single-valued: there may be several maxislack allocations which induce different utilities. This is unlike the maximin rule which is used on Spliddit, where all maximin allocations induce the same vector of utilities (Alkan, Demange, and Gale 1991). Second, there may be maxislack solutions that do not maximize robustness for all agents. To see this, suppose that there are two agents i_1 and i_2 who agree on the valuation for every room. Then in any envy-free allocation, the utility they assign to the two bundles allocated to them must be equal. Hence the maximum slack attainable is 0, and thus *every* envy-free allocation is a maxislack allocation. However, there may be allocations for which the slack between other pairs of agents is larger than 0, and if we choose such an allocation, it is more robustly fair.

In this spirit, to obtain robustness for a larger collection of agents (or of agent pairs), we can refine the maxislack solution using a leximin strategy. We call the resulting solution the *lexislack rule*. Define $\Delta(\sigma, p) = (\Delta_{ir}(\sigma, p))_{i\in N, r\in R}$. The lexislack rule selects an allocation (σ, p) that maximizes the smallest of the n^2 values in $\Delta(\sigma, p)$, and subject to that maximizes the second-smallest of these values, and so on.

In contrast to the maxislack rule, the lexislack rule is essentially single-valued.

Theorem 3. *The lexislack rule is essentially single-valued.*

Proof. For now, fix an optimal room assignment σ. We show that there is a unique price vector p such that (σ, p) is a lexislack solution. Because the leximin relation over vectors is strictly convex, there is a unique vector $\Delta = \Delta(\sigma, p)$ maximizing the leximin objective, since if there were two different ones, a convex combination of the two would be strictly better. But Δ uniquely specifies a price vector: Δ specifies the differences $p_r - p_{r'}$ between any pair of prices, and with $\sum_r p_r = 1$ this gives a unique price vector.

Next, we show that if σ_1 and σ_2 are optimal room assignments, and (σ_1, p) is an envy-free allocation, then $\Delta(\sigma_1, p) = \Delta(\sigma_2, p)$. To see this, note that by Theorem 1(c), (σ_2, p) is an allocation where every agent obtains the same utility as under (σ_1, p). Let $i \in N$. If $\sigma_1(i) = \sigma_2(i)$, then clearly the values Δ_{ir} are the same in both allocations. If $r_1 = \sigma_1(i) \neq \sigma_2(i) = r_2$, then equal utility under both allocations implies $v_{ir_1} - p_{r_1} = v_{ir_2} - p_{r_2}$, and hence $\Delta_{ir_2}(\sigma_1, p) = 0$ and $\Delta_{ir_1}(\sigma_2, p) = 0$. By definition, also $\Delta_{ir_1}(\sigma_1, p) = \Delta_{ir_2}(\sigma_2, p) = 0$, so the values of Δ_{ir_1} and Δ_{ir_2} agree on both allocations. For $r \in R\setminus\{r_1, r_2\}$, we have that the value of Δ_{ir} agrees on both allocations by the equal utility property. Hence $\Delta(\sigma_1, p) = \Delta(\sigma_2, p)$.

Thus, any vector $\Delta \geqslant 0$ achievable on one optimal room assignment can be achieved on any other optimal room assignment, with the same utility vector. This holds in particular for the lexislack vector Δ. We have seen that for any fixed room assignment, there is a unique lexislack utility vector. Hence the lexislack utility vector is unique. □

In addition, this rule remains easy to compute.

Theorem 4. *A lexislack allocation can be found in polynomial time by solving $O(n^4)$ linear programs.*

Proof sketch. This can be done using standard techniques (see Kurokawa, Procaccia, and Shah 2018, Section 5). We give an overview of the algorithm. Start by computing an optimal σ. We will decide on the best value of Δ_{ir} one-by-one. Let $F \leftarrow \emptyset$ be the set of (i, r) pairs for which we have fixed their value. Use linear programming to find a price vector such that (σ, p) maximizes the smallest of the non-fixed values Δ_{ir}, subject to keeping the other Δ_{ir} at their fixed value. Say the optimum is L. Now we need to find a pair $(i, r) \notin F$ such that necessarily $\Delta_{ir} = L$ in any lexislack allocation. This can again be done by linear programs that check whether it is possible that $\Delta_{ir} > L$. One can show that at least one such pair $(i, r) \notin F$ must exist; we then add it to F and fix its value to L, and repeat. □

4 Maximizing Probability of Envy-Freeness

In the previous section, we defined robustness using a measure of closeness based on the ℓ_1-distance. We now look at a more flexible model where the valuations are drawn from a distribution. Given a distribution $\mathcal{D}$ over valuations v, our goal will be to find an allocation (σ, p) that maximizes the probability of being envy-free, i.e., one that maximizes

$$\mathsf{EFrate}_{\mathcal{D}}(\sigma, p) = \Pr_{v \sim \mathcal{D}}[(\sigma, p) \text{ is envy free under } v].$$

In the simplest case, the distribution $\mathcal{D}$ could be obtained by asking users for valuations, and then adding noise (e.g. Gaussian or uniform) around those valuations.

We will only assume to have access to $\mathcal{D}$ by being able to sample from it. This means our results will work for a wide variety of distributions. For example, the user interface could allow agents to specify how much uncertainty they face about the valuation of each room. Our theory can also accommodate more complicated distributions, where valuations of different agents or of different rooms may be correlated.

A natural way to find an allocation with high probability of envy-freeness is to obtain a sample S of m valuation profiles sampled from $\mathcal{D}$, and to compute an allocation that is envy-free on the most profiles in S, i.e., one that maximizes

$$\mathsf{EFrate}_S(\sigma, p) = \tfrac{1}{m} \cdot |\{v \in S : (\sigma, p) \text{ is envy free under } v\}|.$$

If the number m of samples is sufficiently high, we may hope that the best allocation on the sample S is also approximately the best on the distribution $\mathcal{D}$. In this section, we will give a bound for the sample size m to be sufficient to ensure this property, and then we will discuss the computational problem of finding the best allocation for a given sample.

4.1 Sample Complexity

In this section, we will prove the following result.

Theorem 5. *Let $\varepsilon, \delta > 0$. Then there is a value $m \in \mathbb{N}$ with*

$$m = O\left(\frac{n^2 \log n + \log(1/\delta)}{\varepsilon^2}\right)$$

such that for every probability distribution $\mathcal{D}$ over valuation profiles, if S is a collection of at least m samples drawn i.i.d. from $\mathcal{D}$, and (σ^*, p^*) is the allocation that maximizes EFrate_S, then with probability at least $1 - \delta$,

$$\mathsf{EFrate}_{\mathcal{D}}(\sigma^*, p^*) \geqslant \max_{(\sigma,p)} \mathsf{EFrate}_{\mathcal{D}}(\sigma, p) - \varepsilon.$$

We prove this theorem by adapting standard tools from learning theory. Let X be any set, with an unknown ground truth labeling $\tau : X \to \{0, 1\}$. A *hypothesis* is a function $h : X \to \{0, 1\}$. Given a *sample* $S = (x_1, \dots, x_m)$ of m elements of X (not necessarily distinct), write $\mathsf{err}_S(h) = \frac{1}{m}|\{x_i : h(x_i) \neq \tau(x_i)\}|$ for the fraction of samples that h labeled incorrectly. For a probability distribution $\mathcal{D}$ over X, write $\mathsf{err}_{\mathcal{D}}(h) = \Pr_{x \sim D}[h(x) \neq \tau(x)]$ for the probability that h incorrectly labels a point sampled from $\mathcal{D}$.

A *hypothesis class* $\mathcal{H}$ is a set of hypotheses. Given a random sample S drawn i.i.d. from $\mathcal{D}$, and knowledge of the true labeling τ of those samples, our goal is to find a hypothesis $h \in H$ that approximately minimizes $\mathsf{err}_{\mathcal{D}}(h)$, with high probability. Note that the ground truth τ, interpreted as a hypothesis, need not be a member of $\mathcal{H}$. In learning theory, this setup corresponds to "agnostic PAC learning", where the realizability assumption need not hold.

We say that a set $C \subseteq X$ is *shattered* by $\mathcal{H}$ if for all $S \subseteq C$, there exists $h \in \mathcal{H}$ with $h(x) = 1$ if $x \in S$ and $h(x) = 0$ if $x \in C \setminus S$. In other words, if we restrict the hypotheses in $\mathcal{H}$ to the set C, then all possible labelings of C are part of $\mathcal{H}$. The *VC dimension* $\mathsf{VCdim}(\mathcal{H})$ of $\mathcal{H}$ is the cardinality of the largest subset of X that is shattered by $\mathcal{H}$. We are interested in VC dimension due to the following standard result, adapted from Shalev-Shwartz and Ben-David (2014, Theorem 6.8), which says that PAC learning is possible on hypothesis classes of finite VC dimension.

Theorem 6. *Let $\varepsilon, \delta > 0$. Let $\mathcal{H}$ be a hypothesis class with $\mathsf{VCdim}(\mathcal{H}) = d$. Then there exists a value $m \in \mathbb{N}$ with*

$$m = O\left(\frac{d + \log(1/\delta)}{\varepsilon^2}\right)$$

such that for every probability distribution $\mathcal{D}$ over X, if S is a collection of at least m samples drawn i.i.d. from $\mathcal{D}$, and $h^* \in H$ is the hypothesis that minimizes err_S, then with probability at least $1 - \delta$,

$$\mathsf{err}_{\mathcal{D}}(h^*) \leqslant \min_{h \in \mathcal{H}} \mathsf{err}_{\mathcal{D}}(h) + \varepsilon.$$

For our application, we let X be the set of all valuation profiles v. The "correct" labeling is $\tau(v) = 1$ for all v. We identify allocations with hypotheses: For an allocation (σ, p), we define the hypothesis $h_{(\sigma,p)}$ so that for each v,

$$h_{(\sigma,p)}(v) = \begin{cases} 1 & \text{if } (\sigma, p) \text{ is envy-free under } v, \\ 0 & \text{otherwise.} \end{cases}$$

Note that these definitions mean that for all S and $\mathcal{D}$, we have

$$\mathsf{EFrate}_S(\sigma, p) = 1 - \mathsf{err}_S(h_{(\sigma,p)}), \text{ and}$$
$$\mathsf{EFrate}_{\mathcal{D}}(\sigma, p) = 1 - \mathsf{err}_{\mathcal{D}}(h_{(\sigma,p)}).$$

We study the hypothesis class $\mathcal{H}$ of all such hypotheses:

$$\mathcal{H} = \{h_{(\sigma,p)} : \text{allocations } (\sigma, p)\}.$$

To bound its VC dimension, the following result is useful:

Lemma 7 (Shalev-Shwartz and Ben-David 2014, Exercise 6.11)**.** *Let $\mathcal{H}_1, \dots, \mathcal{H}_t$ be hypothesis classes over X, with $\mathsf{VCdim}(\mathcal{H}_i) \leqslant d$ for each $i = 1, \dots, t$. Then*

$$\mathsf{VCdim}(\mathcal{H}_1 \cup \dots \cup \mathcal{H}_t) \leqslant 4d \log(2d) + 2 \log(t).$$

We can now bound the VC dimension of $\mathcal{H}$.

Lemma 8. $\mathsf{VCdim}(\mathcal{H}) = O(n^2 \log n)$.

Proof. For each room assignment σ, define the hypothesis class $\mathcal{H}_\sigma = \{h_{(\sigma,p)} : p \in \mathbb{R}^n\}$ corresponding to allocations whose room assignment is σ. Then $\mathcal{H} = \bigcup_\sigma \mathcal{H}_\sigma$ where the union ranges over all room assignments. We will show that $\mathsf{VCdim}(\mathcal{H}_\sigma) \leqslant n^2$ for each σ. Since there are $n!$ different room assignments and $\log n! = O(n \log n)$, it follows from Lemma 7 that $\mathsf{VCdim}(\mathcal{H}) = O(n^2 \log n)$, as required.

Let σ be a room assignment. Without loss of generality assume that $\sigma(i) = i$. Let $d \geqslant n^2 + 1$. Consider a collection of d distinct valuation profiles $v^{(1)}, \dots, v^{(d)}$. We show that this collection cannot be shattered by $\mathcal{H}_\sigma$.

For $i, j \in N$, say $v^{(k)}$ is *uniquely restricting for* (i, j) if

$$v_{ij}^{(k)} - v_{ii}^{(k)} > v_{ij}^{(\ell)} - v_{ii}^{(\ell)} \quad \text{for all } \ell \neq k.$$

Thus, such a profile uniquely maximizes the amount by which agent i prefers j's room to her own room, ignoring prices. Clearly, for any pair $i, j \in N$, at most one profile can be uniquely restricting for it. Since there are n^2 many pairs (i, j) and $d > n^2$, there is at least one profile which is not uniquely restricting for any pair, say $v^{(1)}$.

We now ask if there is an allocation (σ, p) that is envy-free under $v^{(2)}, \dots, v^{(d)}$, but not envy-free under $v^{(1)}$. We show that the answer is no, so $\mathcal{H}_\sigma$ fails to shatter this collection.

Assume for a contradiction that (σ, p) is such an allocation. Since it is not envy-free under $v^{(1)}$, there is a pair $i, j \in N$ with $v_{ij}^{(1)} - p_j > v_{ii}^{(1)} - p_i$ or equivalently

$$v_{ij}^{(1)} - v_{ii}^{(1)} > p_j - p_i. \tag{2}$$

As $v^{(1)}$ is not uniquely restricting for (i, j), for some $\ell \neq 1$,

$$v_{ij}^{(\ell)} - v_{ii}^{(\ell)} \geqslant v_{ij}^{(1)} - v_{ii}^{(1)}. \tag{3}$$

Combining (2) and (3), it follows that $v_{ij}^{(\ell)} - v_{ii}^{(\ell)} > p_j - p_i$. Thus, (σ, p) is not envy-free under $v^{(\ell)}$, a contradiction. $\square$

Our main result in this section, Theorem 5, now follows immediately from Theorem 6.

4.2 Computational Complexity

To make use of Theorem 5, we need an algorithm that, given a collection $S = (v^{(1)}, \dots, v^{(m)})$ of valuation profiles sampled from $\mathcal{D}$, finds an allocation that maximizes $\mathsf{EFrate}_S(\sigma, p)$. This problem can be encoded as an integer linear program via standard encoding techniques, using binary variables x_{ir} encoding that agent i receives room r, continuous variables p_r encoding the prices, and a binary variable y_ℓ for each sample $\ell \in [m]$, indicating whether the produced allocation will be envy-free under $v^{(\ell)}$. The full encoding appears in the appendix.

Instead of an ILP approach, can we hope for a polynomial time algorithm finding the best allocation? Let us formulate our optimization problem as a decision problem as follows.

> EF-RATE MAXIMIZATION
> **Input:** Set N of agents, set R of rooms, a list of m valuation profiles $v^{(1)}, \dots, v^{(m)}$, number B.
> **Question**: Does there exist an allocation that is envy-free for at least B of the m valuation profiles?

Unfortunately, this problem is computationally hard.

Theorem 9. EF-RATE MAXIMIZATION *is NP-complete, even for binary valuation profiles ($v_{ir} \in \{0, 1\}$).*

Proof. Membership in NP is clear. We give a reduction from CLIQUE. Let $G = (V, E)$ be a graph with n vertices and m edges and let k be the target clique size.

We make each vertex an agent, $N = V$. The set of rooms is $R = \{r_1, \dots, r_k, d_1, \dots, d_{n-k}\}$ consisting of k *slot rooms* and of $n - k$ *dummy rooms*. Writing $E = \{e_1, \dots, e_m\}$, we construct m valuation profiles, one per edge. For $\ell \in [m]$, write $e_\ell = \{u, v\}$; the valuation profile $v^{(\ell)}$ is defined by

$$v_{i,r}^{(\ell)} = \begin{cases} 1 & \text{if } i \in \{u, v\} \text{ and } r \in \{r_1, \dots, r_k\}, \\ 0 & \text{otherwise.} \end{cases}$$

Thus, in the ℓth valuation profile, the two agents corresponding to the endpoints of the ℓth edge want to be in a slot room. All other agents do not care. Finally, set $B = \binom{k}{2}$.

We prove that G has a k-clique iff there is an allocation that is envy-free in at least B of the valuation profiles.

($\Leftarrow$): Suppose (σ, p) is envy free for B profiles. Let $C \subseteq V$ be the set of k agents/vertices that are assigned to slot rooms under σ; write $C = \{i_1, \dots, i_k\}$. Let $\ell_1, \dots, \ell_B$ be the collection of indices corresponding to valuation profiles under which the allocation is envy free. We claim that C is a clique, by showing that $e_{\ell_t} \subseteq C$ for each $t \in [B]$. This suffices sine a set of k vertices with $\binom{k}{2}$ edges is a clique.

Let $t \in [B]$. Since (σ, p) is envy-free under $v^{(\ell_t)}$, by Theorem 1(b), σ is optimal under $v^{(\ell_t)}$. This implies that σ has welfare 2, which happens only if both endpoints of edge e_{ℓ_t} get a slot room. So by definition of C, $e_{\ell_t} \subseteq C$, as desired.

($\Rightarrow$): Suppose there is a clique $C \subseteq V$ of size k in G; write $C = \{i_1, \dots, i_k\}$. Make a room assignment σ in which we assign agent i_s to slot room r_s, for each $s \in [k]$. The remaining agents can be assigned arbitrarily to dummy rooms. We set $p_r = \frac{1}{n}$ for each $r \in R$.

Write $e_{\ell_1}, \ldots, e_{\ell_B}$ for the set of edges within C; there are exactly B of them since C is a clique. Let $t \in [B]$, and write $e_{\ell_t} = \{i_a, i_b\}$. We show that (σ, p) is envy-free under $v^{(\ell_t)}$. All agents except i_a and i_b are indifferent between all rooms, and since all rents are the same, they are not envious. Agents i_a and i_b both receive a room that they most prefer, and since all rents are the same, they are not envious. □

There are two sources of computational difficulty for solving EF-RATE MAXIMIZATION: we have to decide on one of the $n!$ possible room assignments, and we have to decide on which subset of valuation profiles we are aiming to be envy-free on. But in practice, there is a way to avoid the first source of hardness. Suppose the m valuation profiles are sampled from a *continuous* distribution $\mathcal{D}$. Then with probability 1, for each sampled profile $v^{(\ell)}$ there is a *unique* optimal room assignment $\sigma^{(\ell)}$. Any solution to the EF-rate maximization problem must use a room assignment that is optimal for at least one of the given valuation profiles. Thus, at most m different room assignments are candidates, and we can find an optimal solution using m calls to the following problem:

> EF-RATE MAXIMIZATION WITH FIXED ASSIGNMENT
> **Input:** A list of m valuation profiles $v^{(1)}, \ldots, v^{(m)}$, number B, room assignment σ.
> **Question**: Is there a price vector p such that (σ, p) is envy-free for at least B of the m valuation profiles?

Unfortunately, this version of the problem is also hard, and so this trick for continuous distributions does not help. We prove this by reduction from the feedback arc set problem.

Theorem 10. EF-RATE MAXIMIZATION WITH FIXED ASSIGNMENT *is NP-complete.*

Proof. Membership in NP is clear. We give a reduction from FEEDBACK ARC SET, which can be stated as follows.

> **Input:** Digraph $D = (V, E)$, number B.
> **Question:** Is there an ordering $(x_1, \ldots, x_n)$ of V such that at least B arcs from E point from left to right? (An arc $(x_k \to x_s) \in E$ points from left to right if $k < s$.)

Consider an instance of this problem: Let $D = (V, E)$ be a digraph and let B be a number. We construct a rent division instance where V is both the set of agents and of rooms. Let $\sigma(x) = x$ be the identity room assignment.

Let $\varepsilon = 2/(n(n+1))$, chosen so that $\varepsilon + 2\varepsilon + \cdots + n\varepsilon = 1$.

Label the arc set $E = \{a_1, \ldots, a_m\}$. We define one valuation profile for each arc $a_\ell = (x \to y)$, with

$$v_{xy}^{(\ell)} = 1 + \varepsilon, \qquad v_{zz}^{(\ell)} = 1 \text{ for all } z \in V,$$

and valuation 0 for all unspecified combinations.

We now prove that there is an ordering $(x_1, \ldots, x_n)$ of V with at least B arcs pointing from left to right if and only if there exists a price vector p that makes the identity room assignment σ envy-free in at least B of the valuation profiles.

($\Rightarrow$): Let $(x_1, \ldots, x_n)$ of V be an ordering such that (wlog) the arcs $a_1, \ldots, a_B$ point from left to right.

Consider the price vector $p = (\varepsilon, 2\varepsilon, \ldots, n\varepsilon)$, so room x_i has rent $i \cdot \varepsilon$. This is a valid price vector because it sums up to 1 by choice of ε. We claim that this price vector is envy-free for valuation profiles $v^{(1)}, \ldots, v^{(B)}$. Let $\ell \in [B]$. Since a_ℓ points from left to right, we have $a_\ell = (x_k \to x_s)$ for some $k < s$. First, note that any agent $x_i \neq x_k$ does not envy another agent because x_i values her assigned room x_i at utility 1 higher than other rooms, which avoids envy because room prices differ by less than 1. By the same argument, agent x_k never envies any agent except perhaps x_s. Finally, we check that agent x_k does not envy agent x_s: Note that the rent of room x_s is $s \cdot \varepsilon$, which is at least ε higher than the rent $k \cdot \varepsilon$ of room x_k. Since x_k values room x_s only ε more than her assigned room x_k, she does not envy agent x_s. Thus p is envy free for valuation profile $v^{(\ell)}$, as required.

($\Leftarrow$): Suppose there is a price vector p that is envy free for valuation profiles $v^{(1)}, \ldots, v^{(B)}$ (relabeled for convenience). Label the vertices $x_1, \ldots, x_n$ in order of increasing price, i.e., such that $p_{x_1} \leqslant p_{x_2} \leqslant \cdots \leqslant p_{x_n}$ with ties broken arbitrarily. Let $\ell \in [B]$ and consider arc $a_\ell = (x_k \to x_s)$. We show that a_ℓ points from left to right, i.e., $k < s$. As p is envy-free for agent x_k under $v^{(\ell)}$, we have

$$v_{x_k,x_s}^{(\ell)} - p_{x_s} \leqslant v_{x_k,x_k}^{(\ell)} - p_{x_k} \iff 1 + \varepsilon - p_s \leqslant 1 - p_k$$
$$\iff p_{x_s} \geqslant p_{x_k} + \varepsilon.$$

In particular $p_{x_k} < p_{x_s}$. By our choice of ordering, it follows that $k < s$, as required. □

5 Minimizing Expected Envy

In Section 4, we defined robust envy-freeness as allocations that have a high probability of being envy-free when valuations come from a given distribution $\mathcal{D}$. In this section, we consider a different objective function that is more fine-grained. In measuring the probability of envy-freeness, we implicitly treat all failures of envy-freeness equally. We will now minimize *expected envy*, which treats cases where one agent envies another by a lot as more severe.

Given a valuation profile v and an allocation (σ, p), we define the allocation's *(maximum) envy*, $\mathsf{envy}_v(\sigma, p)$, to be

$$\max\left\{0, \max_{i,j \in N} \left[(v_{i\sigma(j)} - p_{\sigma(j)}) - (v_{i\sigma(i)} - p_{\sigma(i)})\right]\right\}.$$

This quantity, which is related to slack as considered in Section 3, measures the biggest amount by which one agent prefers another's bundle. In principle one could allow negative values of $\mathsf{envy}_v(\sigma, p)$ for allocations that have positive slack, but we chose to force these values to be non-negative, since our focus is on avoiding envy. Note that an allocation is envy-free if and only if $\mathsf{envy}_v(\sigma, p) = 0$.

Our goal in this section is to find an allocation minimizing the expected envy with respect to $\mathcal{D}$, defined as

$$\mathsf{envy}_{\mathcal{D}}(\sigma, p) = \mathbb{E}_{v \sim \mathcal{D}}[\mathsf{envy}_v(\sigma, p)].$$

Our approach will be similar to before: we obtain a sufficiently large sample S of m profiles from $\mathcal{D}$ and select the allocation that does best on the sample, i.e. it minimizes

$$\mathsf{envy}_S(\sigma, p) = \tfrac{1}{m} \textstyle\sum_{v \in S} \mathsf{envy}_v(\sigma, p).$$

Unlike before, to bound the sample complexity, we use a direct approach that does not work for EF rate (as there is no analog for Lemma 13 below). We believe that an approach based on extensions of VC dimension to real-valued functions (such as pseudo-dimension) could give similar bounds.

5.1 Sample Complexity

In this section, we assume that valuations v are normalized: let $v_{ir} \geqslant 0$ for all $i \in N$ and $r \in R$, and $\sum_{r\in R} v_{ir} = 1$ for all $i \in N$. We are going to prove the following result:

Theorem 11. *Let $\varepsilon, \delta > 0$, and let $\mathcal{D}$ be a distribution. If we draw $m = O(\frac{n}{\varepsilon^2} \log \frac{n}{\varepsilon\delta})$ samples i.i.d. from $\mathcal{D}$ and let (σ^*, p^*) be the allocation minimizing envy_S , then with probability at least $1 - \delta$, we have*

$$\mathsf{envy}_{\mathcal{D}}(\sigma^*, p^*) < \min_{(\sigma,p)} \mathsf{envy}_{\mathcal{D}}(\sigma, p) + \varepsilon.$$

Thus, if we draw sufficiently many samples, then with high probability the allocation minimizing expected envy on the sample will, up to ε, be minimizing with respect to $\mathcal{D}$. We will prove this result by discretizing the space of allocations, and using a concentration inequality to show that w.h.p. the expected envy with respect to $\mathcal{D}$ is close to the expected envy with respect to the sample S. We start by proving a few technical lemmas. First, define Λ to be the set all allocations (σ, p) with $-2 \leqslant p_r \leqslant 2$ for all $r \in R$. We call such allocations *reasonable*. Our first lemma shows that we may restrict attention to reasonable allocations only: in particular, if an allocation minimizes $\mathsf{envy}_{\mathcal{D}}$ then it must be reasonable.

Lemma 12. *Let σ be a room assignment and v a profile.*

(a) If p is a price vector with $|p_r - p'_r| > 2$ for some $r, r' \in R$, then $\mathsf{envy}_v(\sigma, p) > 1$.
(b) If $p = (\frac{1}{n}, \ldots, \frac{1}{n})$, then $\mathsf{envy}_v(\sigma, p) \leqslant 1$.
(c) If (σ, p) is reasonable, then $\mathsf{envy}_v(\sigma, p) \leqslant 5$.

Proof. Note that since valuations are assumed to sum to 1, we have $v_{ir} - v_{ir'} \leqslant 1$ for all $r, r' \in R$.

(a) Say $p_r > p_{r'} + 2$ and $\sigma(i) = r$. Since $v_{ir} \leqslant v_{ir'} + 1$ as just noted, we have $v_{ir} - p_r < v_{ir'} - p_{r'} - 1$.

(b) If prices for all rooms are equal, then $\mathsf{envy}_v(\sigma, p) = \max\{0, \max_{i,j}(v_{i\sigma(j)} - v_{i\sigma(i)})\} \leqslant \max\{0, 1\} = 1$.

(c) $|v_{i\sigma(j)} - v_{i\sigma(i)}| + |p_{\sigma(j)} - p_{\sigma(i)}| \leqslant 1 + 4 = 5$. ☐

From now on, we assume all allocations to be reasonable.

Our second lemma says that if two allocations have similar price vectors, then they have similar expected envy.

Lemma 13. *Let $p, p' \in \mathbb{R}^n$ be price vectors with $|p_r - p'_r| \leqslant t$ for all $r \in R$. Then for any sample S and distribution $\mathcal{D}$,*

$$|\mathsf{envy}_{\mathcal{D}}(\sigma, p) - \mathsf{envy}_{\mathcal{D}}(\sigma, p')| \leqslant 2t,$$
$$|\mathsf{envy}_S(\sigma, p) - \mathsf{envy}_S(\sigma, p')| \leqslant 2t.$$

Proof. First, we claim that for any valuation profile v,

$$|\mathsf{envy}_v(\sigma, p) - \mathsf{envy}_v(\sigma, p')| \leqslant 2t.$$

This holds since the value of $(v_{i\sigma(j)} - p_{\sigma(j)}) - (v_{i\sigma(i)} - p_{\sigma(i)})$ changes by at most $\pm 2t$ if we move from p to p', and thus the same holds after taking the maximum.

Now let $\mathcal{D}$ be a distribution. By linearity of expectation, and since $|\mathbb{E}[X]| \leqslant \mathbb{E}[|X|]$ by Jensen's inequality,

$$\begin{aligned} &|\mathsf{envy}_{\mathcal{D}}(\sigma, p) - \mathsf{envy}_{\mathcal{D}}(\sigma, p')| \\ &= |\mathbb{E}_{v\sim\mathcal{D}}[\mathsf{envy}_v(\sigma, p) - \mathsf{envy}_v(\sigma, p')]| \\ &\leqslant \mathbb{E}_{v\sim\mathcal{D}}\left[|\mathsf{envy}_v(\sigma, p) - \mathsf{envy}_v(\sigma, p')|\right] \leqslant 2t, \end{aligned}$$

where the last inequality follows by our claim. This proves the first statement. The second statement follows from the first by taking $\mathcal{D}$ to be the uniform distribution over S. ☐

We also need a standard concentration inequality.

Lemma 14 (Hoeffding's inequality)**.** *Let $X_1, \ldots, X_m$ be i.i.d. random variables with $0 \leqslant X_k \leqslant c$ and $\mathbb{E}[X_k] = \mu$ for all $k \in [m]$. Then for all $\varepsilon > 0$,*

$$\Pr\left[|\mu - \tfrac{1}{m}\textstyle\sum_{k=1}^m X_i| \geqslant \varepsilon\right] \leqslant 2\exp(-2m\varepsilon^2/c^2).$$

We are now ready to prove our main result of this section.

Proof of Theorem 11. Let S be a random sample from $\mathcal{D}$ of size m, where

$$m = \tfrac{200}{\varepsilon^2} \ln\left[\left(\tfrac{60}{\varepsilon}\right)^n \tfrac{2n!}{\delta}\right] = O\left(\tfrac{n}{\varepsilon^2} \ln\left(\tfrac{n}{\varepsilon\delta}\right)\right).$$

Let $t = 1/\lceil 12/\varepsilon\rceil$. Let $\Lambda^t \subseteq \Lambda$ be the set of all allocations (σ, p) where p_r is an integer multiple of t. We call these *discretized allocations*. Note that for any $(\sigma, p) \in \Lambda$, there is a discretized allocation $(\sigma, p') \in \Lambda^t$ with $|p_r - p'_r| \leqslant t$ for all r (call such allocations *t-close*), which can be found by rounding the values p_r up or down to ensure that $\sum_r p'_r = 1$.

Now write:

- $\mathrm{OPT}_{\mathcal{D}}$ for an allocation (σ, p) minimizing $\mathsf{envy}_{\mathcal{D}}$,
- OPT_S for an allocation minimizing envy_S (which depends on the random choice of S),
- $\overline{\mathrm{OPT}}_{\mathcal{D}} \in \Lambda^t$ for a discretized allocation t-close to $\mathrm{OPT}_{\mathcal{D}}$,
- $\overline{\mathrm{OPT}}_S \in \Lambda^t$ for a discretized allocation t-close to OPT_S.

Let $(\sigma, p) \in \Lambda$. For $k \in [m]$, let X_k be the random variable taking the value $\mathsf{envy}_v(\sigma, p)$, where v is the kth sample in S. By reasonableness and Lemma 12, $0 \leqslant X_k \leqslant 5$. Then Hoeffding's inequality implies that

$$\Pr\left[|\mathsf{envy}_S(\sigma, p) - \mathsf{envy}_{\mathcal{D}}(\sigma, p)| \geqslant \tfrac{\varepsilon}{4}\right] \leqslant 2\exp(-\tfrac{2}{25}m(\tfrac{\varepsilon}{4})^2).$$

Let E be the event that $|\mathsf{envy}_S(\sigma, p) - \mathsf{envy}_{\mathcal{D}}(\sigma, p)| < \varepsilon/4$ holds for all discretized allocations $(\sigma, p) \in \Lambda^t$ simultaneously. By Hoeffding's inequality as above and a union bound over all $|\Lambda^t| \leqslant (\frac{4}{t})^n n!$ discretized allocations, we get

$$\Pr[E] \geqslant 1 - (\tfrac{4}{t})^n n!\, 2\exp(-\tfrac{2}{25}m(\tfrac{\varepsilon}{4})^2) \geqslant 1 - \delta.$$

where the second inequality holds by choice of m.

Suppose that the event E attains. In this case we have:

$$\begin{aligned} &\mathsf{envy}_{\mathcal{D}}(\mathrm{OPT}_S) - \mathsf{envy}_{\mathcal{D}}(\mathrm{OPT}_{\mathcal{D}}) \\ &= (\mathsf{envy}_{\mathcal{D}}(\mathrm{OPT}_S) - \mathsf{envy}_{\mathcal{D}}(\overline{\mathrm{OPT}}_S)) && \text{(Lemma 13)} \\ &\quad + (\mathsf{envy}_{\mathcal{D}}(\overline{\mathrm{OPT}}_S) - \mathsf{envy}_S(\overline{\mathrm{OPT}}_S)) && \text{(E attains)} \\ &\quad + (\mathsf{envy}_S(\overline{\mathrm{OPT}}_S) - \mathsf{envy}_S(\mathrm{OPT}_S)) && \text{(Lemma 13)} \\ &\quad + (\mathsf{envy}_S(\mathrm{OPT}_S) - \mathsf{envy}_S(\overline{\mathrm{OPT}}_{\mathcal{D}})) && \text{(optimality)} \\ &\quad + (\mathsf{envy}_S(\overline{\mathrm{OPT}}_{\mathcal{D}}) - \mathsf{envy}_{\mathcal{D}}(\overline{\mathrm{OPT}}_{\mathcal{D}})) && \text{(E attains)} \\ &\quad + (\mathsf{envy}_{\mathcal{D}}(\overline{\mathrm{OPT}}_{\mathcal{D}}) - \mathsf{envy}_{\mathcal{D}}(\mathrm{OPT}_{\mathcal{D}})) && \text{(Lemma 13)} \\ &< 2t + \varepsilon/4 + 2t + 0 + \varepsilon/4 + 2t \\ &= 6/\lceil 12/\varepsilon\rceil + \varepsilon/2 \leqslant \varepsilon. \end{aligned}$$

The references on the right indicate what we have used to bound the respective term to obtain the strict inequality. "Optimality" refers to the fact that OPT_S minimizes envy_S. Because event E implies the above inequality, it follows that

$$\Pr[\mathsf{envy}_{\mathcal{D}}(\text{OPT}_S) - \mathsf{envy}_{\mathcal{D}}(\text{OPT}_{\mathcal{D}}) < \varepsilon] \geqslant \Pr[E] \geqslant 1 - \delta.$$

This proves the result. □

5.2 Computational Complexity

Again, our sample complexity result needs an algorithm that finds the best allocation for a given sample S. Like for EFrate, we can solve this problem using integer linear programming. The formulation appears in the appendix. For the formal complexity analysis, consider the following decision problem:

EXPECTED ENVY MINIMIZATION
Input: List $S = (v^{(1)}, \dots, v^{(m)})$, number B.
Question: Is there (σ, p) with $\mathsf{envy}_S(\sigma, p) \leqslant B$?

This problem is again NP-complete, using a similar reduction as for EF-RATE MAXIMIZATION.

Theorem 15. EXPECTED ENVY MINIMIZATION *is NP-complete, even for binary valuation profiles.*

Proof. Membership in NP is clear. Reduction from CLIQUE.

Let $G = (V, E)$ be a graph with n vertices and m edges and target clique size k. Set the target envy amount to be $B = m - \binom{k}{2}$. Let M be a large integer, $M > (B+1)^2$. Write $\varepsilon = (B+1)/M$. Our choices of these numbers imply the following estimates which we will need later:

- $M\varepsilon > B$, since $M\varepsilon = B + 1 > B$.
- $\varepsilon(B+1) < 1$, since $\varepsilon(B+1) = (B+1)^2/M < 1$.

The set of agents is V. The set of rooms is $R = \{o_1, \dots, o_k, d_1, \dots, d_{n-k}\}$ consisting of k *slot rooms* and $n - k$ dummy rooms. Write $E = \{e_1, \dots, e_m\}$. We construct a sample S of $m + M$ valuation profiles. For $j \in [m]$, write $e_j = \{u, v\}$; then valuation profile $v^{(j)}$ is defined by

$$v_{i,r}^{(j)} = \begin{cases} 1 & \text{if } i \in \{u, v\} \text{ and } r \in \{o_1, \dots, o_k\}, \\ 0 & \text{otherwise.} \end{cases}$$

For $j = m+1, \dots, m+M$, let $v^{(j)}$ be a *uniform profile*:

$$v_{i,r}^{(j)} = 0 \quad \text{for all } i \in V \text{ and } r \in R.$$

($\Rightarrow$): Suppose there is a clique $C \subseteq V$ of size k in G; write $C = \{i_1, \dots, i_k\}$. We construct an allocation (σ, p) that will be envy free for B profiles. In the room assignment, we will assign agent i_r to slot room o_r, for $r \in [k]$. The remaining agents can be assigned arbitrarily to dummy rooms. We'll say that each room costs the same rent so $p_r = \frac{1}{n}$ for all $r \in R$.

- For each of the M uniform profiles, $\mathsf{envy}_v(\sigma, p) = 0$.
- For a profile v corresponding to an edge $e_j = \{i_a, i_b\}$ with $i_a, i_b \in C$ (i.e. contained in the clique), $\mathsf{envy}_v = 0$.
- For one of the $m - \binom{k}{2}$ profiles v corresponding to edges not contained in a clique, we have $\mathsf{envy}_v = 1$.

Summing these up, we have $\mathsf{envy}_S(\sigma, p) = m - \binom{k}{2} = B$.

($\Leftarrow$): Suppose there is an allocation (σ, p) with $\mathsf{envy}_S(\sigma, p) \leqslant B$. Let $C \subseteq V$ be the set of k agents/vertices assigned to slot rooms under σ; write $C = \{i_1, \dots, i_k\}$.

First we show that the rents $p = (p_1, \dots, p_n)$ of the rooms are close to uniform, in the sense that $|p_r - p_{r'}| \leqslant \varepsilon$ for all $r, r' \in R$. Assume for a contradiction that there are $r, r' \in R$ with $p_r > p_{r'} + \varepsilon$. Then in each uniform profile, the agent assigned to room r envies the agent assigned to room r' by at least ε, and hence the max envy in a uniform profile is at least ε. Since we have introduced M uniform profiles, it follows that $\mathsf{envy}_S(\sigma, p) \geqslant M\varepsilon > B$, a contradiction.

Now we show that C is a clique. Suppose not. Then there are at least $m - \binom{k}{2} + 1 = B + 1$ edges that are not completely contained in C. For each profile corresponding to such an edge, the agent corresponding to the endpoint not in C envies other agents who are assigned a slot room by at least $1 - \varepsilon$. Hence $\mathsf{envy}_S(\sigma, p) \geqslant (B+1)(1-\varepsilon) = B + 1 - \varepsilon(B+1) > B$, because $\varepsilon(B+1) < 1$. This is a contradiction. □

Interestingly, this problem becomes easy once we fix a room assignment σ, because the best price vector can then be computed by linear programming. In particular, this means that the problem can be solved in time $n! \cdot \text{poly}(n, m)$, and thus is fixed-parameter tractable with respect to the number of agents n. This is good news: instances will often have a small number of agents, but we will want to consider as large a sample as feasible to ensure low maximum envy.

6 Experiments

We evaluated our rules on user data taken from Spliddit. We studied distributions obtained by adding noise to valuations. We started by selecting 500 instances v at random, to speed up computations. The same selection is used for each experiment. For each instance, we normalize the rent to 1, and normalize valuations to sum to 1. We considered three noise models, each parameterized by a choice of noise level $\varepsilon \in \{0, 0.01, \dots, 0.09\}$.

$$v_{ir}^{(\ell)} \sim v_{ir} \cdot (1 + \text{Uniform}[-\varepsilon, +\varepsilon]) \qquad \text{(Uniform)}$$

$$v_{ir}^{(\ell)} \sim v_{ir} \cdot (1 + \text{N}[0, \varepsilon]) \qquad \text{(Normal)}$$

$$v_{ir}^{(\ell)} \sim v_{ir} \cdot (1 + r \cdot \text{N}[0, \varepsilon]) \qquad \text{(Biased Normal)}$$

In each of these noise models, valuations are increased or decreased by a random fraction. Here, $N[\mu, \sigma]$ is a normal distribution with mean μ and standard deviation σ. For the biased normal noise model, we put rooms in an arbitrary fixed order and label them with integers $0, 1, \dots, n-1$. Rooms with a higher index have more noise.

For each noise model and choice of ε, we produced a sample S of $m = 100$. We then computed allocations maximizing EFrate_S and minimizing envy_S. We also computed the maximin and lexislack rules based on the input profile v. For each of the four allocations, we calculated their value of EFrate_S, and of envy_S. The results are shown in Table 2, where for each choice of ε, we average over all 500 instances. By definition, on each of the two metrics, the rule optimizing

Probability of Envy-Freeness (EFrate_S)	Expected Envy (envy_S)

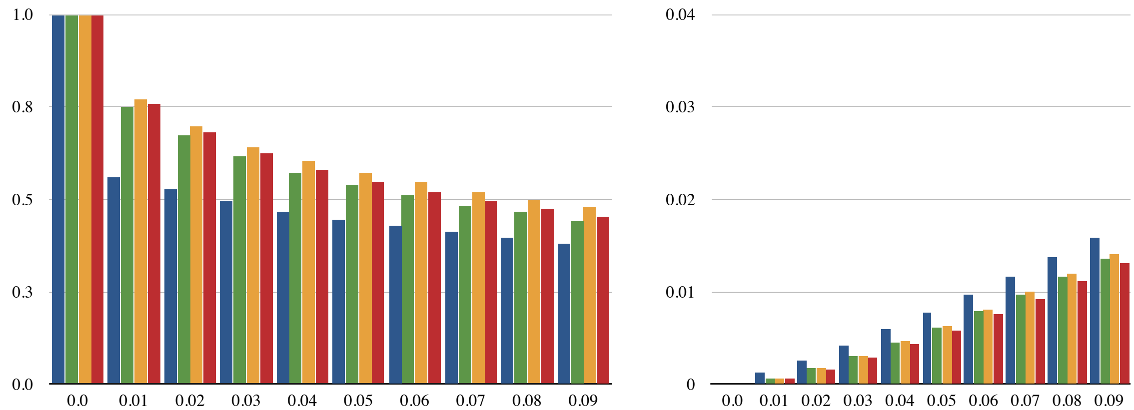

(a) Multiplicative Uniform Noise

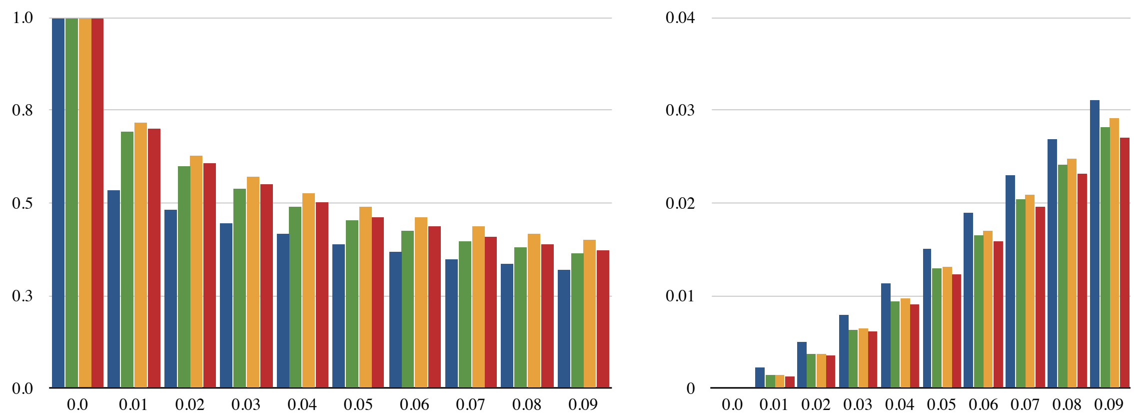

(b) Multiplicative Normal Noise

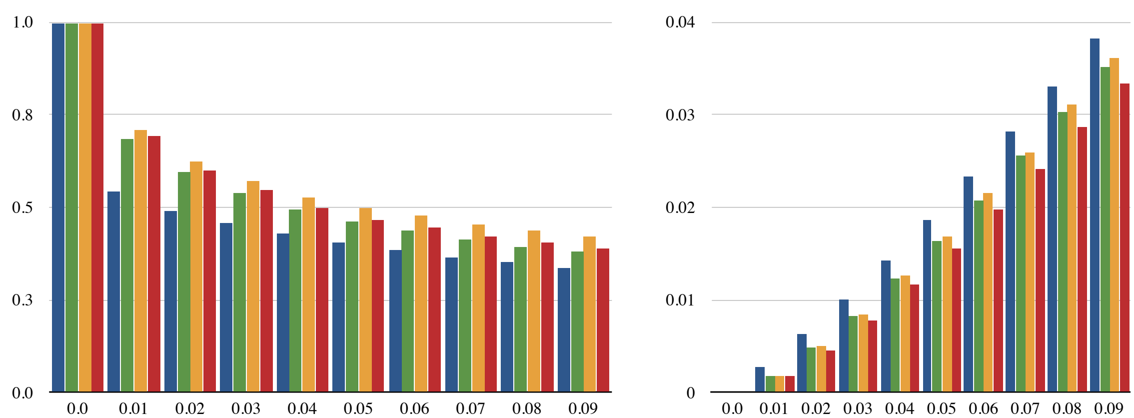

(c) Biased Multiplicative Normal Noise

Table 2: Results of experiments for three noise models.

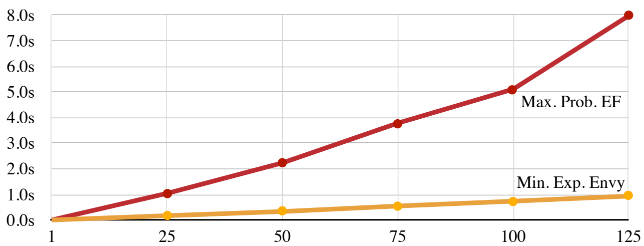

Figure 1: Computation time depending on sample size

it does best, but all three rules aiming for robustness do similarly well. Spliddit's maximin rule does significantly worse on this metric. Before running the experiments, we expected that the lexislack rule would do worse for the biased noise model, but this does not appear to be the case.

Figure 1 shows average computation time to compute allocations optimizing EFrate_S and envy_S, using Gurobi 9.1.2 on four threads of an AMD Ryzen 2990WX (128 GB RAM) with the ILP formulations from the appendix. The results were obtained for a random selection of 300 Spliddit instances with $n = 4$, with the Uniform noise model for $\varepsilon = 0.05$, and sample sizes m varying from 1 to 125. Minimizing envy is much faster, presumably due to fewer integral variables.

7 Future Directions

Our approach should be applicable to many settings beyond rent division, such as homogeneous divisible goods, cake cutting, or even indivisible goods. For example, the lexislack rule can be adapted to these settings, and similar results as in our distribution-based approach might be achievable.

We have shown that the lexislack rule shares some key properties with the maximin rule, such as essential single-valuedness and polynomial-time computability. It would be interesting to axiomatically contrast the two solutions, to see if one might be better behaved than the other.

For our distribution-based approach, we assumed that we have access to $\mathcal{D}$ only via sampling. Often we may know $\mathcal{D}$ more explicitly, for example if we are just adding noise to reported valuations. For such well-behaved $\mathcal{D}$, can we design direct algorithms for finding optimal allocations with respect to our two objectives, without going through samples?

References

Alkan, A.; Demange, G.; and Gale, D. 1991. Fair allocation of indivisible goods and criteria of justice. *Econometrica*, 59(4): 1023–1039.

Aragones, E. 1995. A derivation of the money Rawlsian solution. *Social Choice and Welfare*, 12: 267–276.

Bei, X.; Li, Z.; Liu, J.; Liu, S.; and Lu, X. 2021. Fair division of mixed divisible and indivisible goods. *Artificial Intelligence*, 293: Article 103436.

Bredereck, R.; Faliszewski, P.; Kaczmarczyk, A.; Niedermeier, R.; Skowron, P.; and Talmon, N. 2021. Robustness among multiwinner voting rules. *Artificial Intelligence*, 290: Article 103403.

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

Caragiannis, I.; Kurokawa, D.; Moulin, H.; Procaccia, A. D.; Shah, N.; and Wang, J. 2019. The unreasonable fairness of maximum Nash welfare. *ACM Transactions on Economics and Computation (TEAC)*, 7(3): 1–32.

Chen, J.; Skowron, P.; and Sorge, M. 2019. Matchings under preferences: Strength of stability and trade-offs. In *Proceedings of the 2019 ACM Conference on Economics and Computation (EC)*, 41–59.

Critch, A. 2015. Robust rental harmony. *https://acritch.com/rent/*. Archived at https://perma.cc/Q9MD-HMAY.

Flanigan, B.; Gölz, P.; Gupta, A.; Hennig, B.; and Procaccia, A. D. 2021. Fair algorithms for selecting citizens' assemblies. *Nature*, 596: 548–552.

Gal, Y.; Mash, M.; Procaccia, A. D.; and Zick, Y. 2017. Which is the fairest (rent division) of them all? *Journal of the ACM*, 64(6): Article 39.

Gawron, G.; and Faliszewski, P. 2019. Robustness of approval-based multiwinner voting rules. In *Proceedings of the 6th International Conference on Algorithmic Decision Theory (ADT)*, 17–31.

Goldman, J.; and Procaccia, A. D. 2014. Spliddit: Unleashing fair division algorithms. *SIGecom Exchanges*, 13(2): 41–46.

Kurokawa, D.; Procaccia, A. D.; and Shah, N. 2018. Leximin allocations in the real world. *ACM Transactions on Economics and Computation*, 6(3–4): Article 11.

Mai, T.; and Vazirani, V. V. 2018. Finding stable matchings that are robust to errors in the input. In *Proceedings of the 26th Annual European Symposium on Algorithms (ESA)*, Article No. 60.

Menon, V.; and Larson, K. 2020. Algorithmic Stability in Fair Allocation of Indivisible Goods Among Two Agents. *arXiv:2007.15203*.

Misra, N.; and Sonar, C. 2019. Robustness radius for Chamberlin–Courant on restricted domains. In *Proceedings of the 45th International Conference on Current Trends in Theory and Practice of Informatics (SOFSEM)*, 341–353.

Moulin, H. 2019. Fair division in the internet age. *Annual Review of Economics*, 11: 407–441.

Procaccia, A. D.; Velez, R. A.; and Yu, D. 2018. Fair rent division on a budget. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, 1177–1184.

Shalev-Shwartz, S.; and Ben-David, S. 2014. *Understanding Machine Learning: From Theory to Algorithms*. Cambridge University Press.

Shiryaev, D.; Yu, L.; and Elkind, E. 2013. On elections with robust winners. In *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, 415–422.

Su, F. E. 1999. Rental harmony: Sperner's lemma in fair division. *American Mathematical Monthly*, 106(10): 930–942.

Svensson, L.-G. 1983. Large indivisibles: An analysis with respect to price equilibrium and fairness. *Econometrica*, 51(4): 939–954.

Velez, R. A. 2018. Equitable rent division. *ACM Transactions on Economics and Computation (TEAC)*, 6(2): Article 9.

Wolsey, L. A. 1998. *Integer Programming*. John Wiley & Sons.

A ILP Formulations

A.1 Envy-Free Rate

Write $v_{\max} = \max_{i,\ell,r} v_{ir}^{(\ell)}$ (if we use normalized valuations, this value is at most 1). Note that an allocation (σ, p) where $p_r < -v_{\max}$ for some $r \in R$ cannot be envy-free for any of the valuation profiles: If r' is a room with $p_{r'} > 0$, then the person receiving r' values r' at most $v_{\max}$ more than room r, and hence by the price difference will envy the agent receiving room r. So in solving our maximization problem, we can restrict attention to price vectors with $p_r \geqslant -v_{\max}$ for all $r \in R$. Similarly we can assume $p_r \leqslant v_{\max}$.

Using such price vectors, note that the envy between any pair of players is at most $3v_{\max}$. Write $M = 3v_{\max}$.

$$
\begin{aligned}
\max\ & \textstyle\sum_{\ell \in [m]} y_\ell & \\
\text{s.t. } & \textstyle\sum_{r \in R} x_{ir} = 1 & \text{for all } i \in N \\
& \textstyle\sum_{i \in N} x_{ir} = 1 & \text{for all } r \in R \\
& v_{ir}^{(\ell)} - p_r \geqslant v_{ir'}^{(\ell)} - p_{r'} - M(1 - y_\ell) - M(1 - x_{ir}) & \text{for all } i \in N, r, r' \in R, \ell \in [m] \\
& \textstyle\sum_{r \in R} p_r = 1 & \\
& -v_{\max} \leqslant p_r \leqslant v_{\max} & \text{for all } r \in R \\
& x_{ir} \in \{0, 1\} & \text{for all } i \in N, r \in R \\
& y_\ell \in \{0, 1\} & \text{for all } \ell \in [m]
\end{aligned}
$$

A.2 Minimize Expected Envy

As in the text, assume that valuations are normalized, and hence restrict attention to reasonable allocations with $-2 \leqslant p_r \leqslant 2$ for all r. Then envy is at most 5. Let $M = 5$.

$$
\begin{aligned}
\min\ & \textstyle\sum_{\ell \in [m]} B_\ell & \\
\text{s.t. } & \textstyle\sum_{r \in R} x_{ir} = 1 & \text{for all } i \in N \\
& \textstyle\sum_{i \in N} x_{ir} = 1 & \text{for all } r \in R \\
& (v_{ir'}^{(\ell)} - p_{r'}) - (v_{ir}^{(\ell)} - p_r) \leqslant B_\ell + M(1 - x_{ir}) & \text{for all } i \in N, r, r' \in R, \ell \in [m] \\
& \textstyle\sum_{r \in R} p_r = 1 & \\
& -2 \leqslant p_r \leqslant 2 & \text{for all } r \in R \\
& x_{ir} \in \{0, 1\} & \text{for all } i \in N, r \in R \\
& B_\ell \geqslant 0 & \text{for all } \ell \in [m]
\end{aligned}
$$