DOI:10.1145/3166068

# Which Is the Fairest (Rent Division) of Them All?

By Kobi Gal, Ariel D. Procaccia, Moshe Mash, and Yair Zick

**"Mirror mirror on the wall, who is the fairest of them all?"**
***The Evil Queen***


## Abstract

**What is a *fair* way to assign rooms to several housemates, and divide the rent between them? This is not just a theoretical question: many people have used the *Spliddit* website to obtain *envy-free* solutions to rent division instances. But envy freeness, in and of itself, is insufficient to guarantee outcomes that people view as intuitive and acceptable. We therefore focus on solutions that optimize a criterion of social justice, subject to the envy freeness constraint, in order to pinpoint the "fairest" solutions. We develop a general algorithmic framework that enables the computation of such solutions in polynomial time. We then study the relations between natural optimization objectives, and identify the *maximin* solution, which maximizes the minimum utility subject to envy freeness, as the most attractive. We demonstrate, using experiments on real data from Spliddit, that the maximin solution gives rise to significant gains in terms of our optimization objectives. Finally, a user study with Spliddit users as subjects demonstrates that people find the maximin solution to be significantly fairer than arbitrary envy-free solutions; this user study is unprecedented in that it asks people about their real-world rent division instances. Based on these results, the maximin solution has been deployed on Spliddit since April 2015.**


## 1. INTRODUCTION

Many a reader may have personally experienced the *rent division problem*: several housemates move in together, and need to decide who gets which room, and at what price. The problem becomes interesting—and, more often than not, a source of frustration—when the rooms differ in quality. The challenge is then to achieve "rental harmony"[18] by assigning the rooms and dividing the rent *fairly*.

In more detail, suppose each player $i$ has value $v_{ij}$ for room $j$, such that each player's values for the rooms sum up to the total rent (see Figure 1a). The (quasilinear) utility of player $i$ for getting room $j$ at price $p_j$ is $v_{ij} - p_j$ (see Figure 1b). A solution (i.e. an assignment of the rooms and division of the rent) is *envy free*[8] if the utility of each player for getting his room at its price is at least as high as getting any other room at the price of that room (see Figure 1c). More generally, one can think of this problem as allocating indivisible goods and splitting a sum of money—but we adopt the rent division terminology, which grounds the problem and justifies our assumptions.

Envy freeness is undoubtedly a compelling fairness notion. But what makes it truly powerful in the context of rent division is that an envy-free solution to a rent division problem always exists.[19] Even better, such a solution can be computed in polynomial time.[3]

However, envy freeness in and of itself is insufficient to guarantee satisfactory solutions. For example, consider an apartment with three rooms and total rent of \$3000. Each player $i$ has value \$3000 for room $i$, and value \$0 for the two other rooms. Furthermore, consider the solution that assigns room 1 to player 1 at \$3000, and, for $i \in \{2, 3\}$, gives room $i$ to player $i$ for free. This solution is envy free: players 2 and 3 are obviously overjoyed, while player 1 is indifferent between the three rooms. However, from an interpersonal perspective, this solution is not fair at all, as the distribution of prices between players is unequal. An intuitive alternative solution here would be to keep the same assignment of rooms, but equally split the rent between the different rooms—\$1000 per room—thereby equalizing the utilities of the players.

The challenge, therefore, is to choose among many possible envy-free solutions. And, arguably, the most natural way to do this is to optimize a function of the utilities that meets desirable social criteria, subject to the envy freeness constraint.[2] In particular, if we were to maximize the minimum utility of any player subject to envy freeness, or if we were to minimize the maximum difference in utilities subject to envy freeness, we would obtain the aforementioned solution in the example. This focus on optimization in rent division motivates us to

> *... design polynomial-time algorithms for optimization under the envy freeness constraint; understand the relationship between natural optimization objectives; and measure the benefits of optimization in rent division.*

### 1.1. Real-world connections and implications: The Spliddit service

The above challenges are especially pertinent when put in the context of Spliddit (www.spliddit.org), a not-for-profit fair division website.[9] Spliddit offers "provably fair solutions" for the division of credit, indivisible goods, chores, fare—and, of course, rent. Since its launch in November 2014, Spliddit has attracted more than 100,000 users, who, in particular, have created 27,344 rent division instances (as of July 6, 2017).

Until April 2015, Spliddit's rent division application relied on the algorithm of Abdulkadiroğlu et al.,[1] which elicits the values of the players for the rooms, and computes an envy-free solution assuming quasi-linear utilities. While many users were satisfied with the results (based on their

The original version of this paper was published in the *Proceedings of the 17th ACM Conference on Economics and Computation* (EC '16).

**Figure 1. Envy-free solutions, illustrated.**

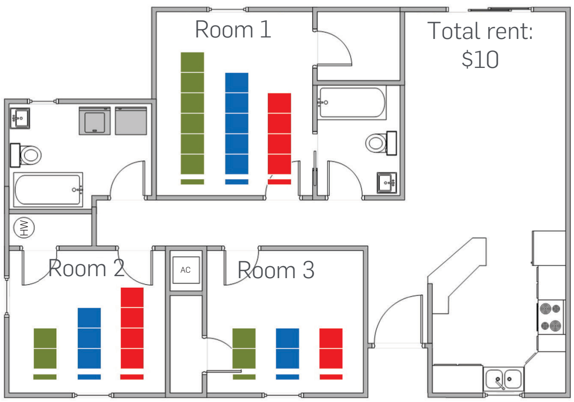


(a) Values, shown by colored boxes. For example, the values of the green, blue, and red players for Room 1 are 6, 5, and 4, respectively. Note that the values of each player for the three rooms add up to the total rent of 10.

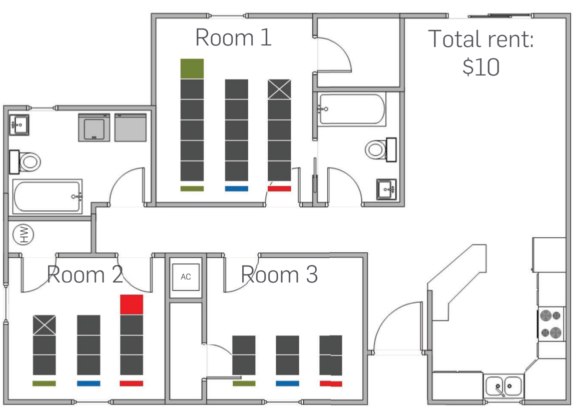


(b) Prices, shown by gray boxes, and utilities. For example, the price of Room 1 is 5, and the utility of the green player for Room 1 at its price is 6 – 5 =1.

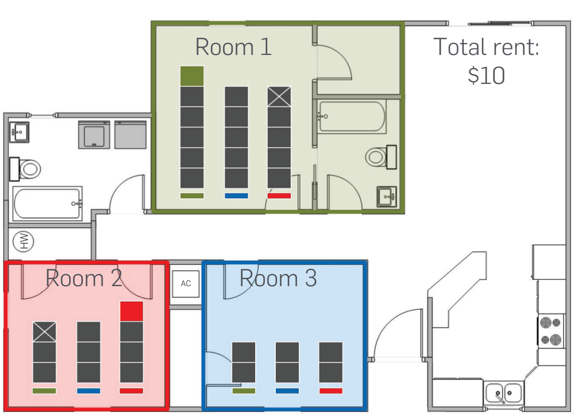


(c) An envy-free solution. For example, the green player is not envious because he has utility 1 for his room (Room 1), utility – 1 for Room 2 (indicated by an "X" in the top box), and utility 0 for Room 3.

reported evaluations), the algorithm does provide nonintuitive solutions in some cases. This prompted an investigation of alternative approaches, and ultimately led to the deployment of a new algorithm in April 2015, based entirely on the results presented in (the original version of) this paper.

It is important to point out that Spliddit not only motivates our research questions, but also helps answer them. Indeed, while Spliddit's primary goals are making fair division methods accessible to people, and outreach, a secondary goal is the collection of an unprecedented dataset for fair division research.[9] This real-world dataset is exciting because, as noted by Herreiner and Puppe,[12] fair division is hard to study in the lab: researchers can tell subjects what their valuations are for different goods, but these values are not ecologically realistic, in that they do not represent subjects' actual preferences. To quote Herreiner and Puppe,[12] "the goods in the lab are not really distributed among participants, but serve as temporary substitutes for money." In contrast, Spliddit instances are ecologically valid, as they are posed by real people facing real division problems. Thus the Spliddit data enables studies at a realistic level and scale that was not possible before. Even better, we can ask Spliddit users to evaluate different solutions based on the actual instances they participated in. This is exactly what we do in this paper.

### 1.2. Our results

We start, in Section 3, by constructing a general yet simple algorithmic framework for optimization under the envy freeness constraint. Specifically, our algorithm maximizes the minimum of linear functions of the utilities, subject to envy freeness, in polynomial time. We do this by using the Second Welfare Theorem to argue that we can employ any welfare-maximizing assignment of players to rooms, and then solve a linear program to compute the optimal envy free prices.[a]

Our main goal in Section 4 is to understand the relation between two solution concepts: the *maximin* solution,[2] which maximizes the minimum utility of any player subject to envy freeness; and the *equitable* solution, which minimizes *disparity*—the maximum difference in utilities—subject to envy freeness. (Our algorithm can compute either solution in polynomial time.) Our most significant result in this section is proving that the maximin solution is also equitable, but not every equitable solution is maximin.

Based on these results, we have implemented the polynomial-time algorithm of Section 3, with the maximin objective function.[b] As noted above, it has been deployed on Spliddit since April 2015.

[a] It is interesting to note that, even though the instances on Spliddit are small, computational tractability does play a key role, as there are many instances and computation incurs a cost (Spliddit uses Amazon Web Services to run all its algorithms).

[b] To be completely precise, the algorithm deployed on Spliddit first tries to maximize the minimum utility, subject to envy freeness as well as an additional constraint: prices must be non-negative. If an envy-free solution with non-negative prices does not exist [4], it removes the non-negative price constraint (in which case a solution always exists). Most of our results go through even when prices are assumed to be non-negative. In any case, real-world instances where negative prices actually help are extremely rare, so throughout the paper prices are unconstrained.

The remainder of the paper focuses on demonstrating that the foregoing approach is indeed effective. Here our contribution is twofold. First, we show that real-world instances give rise to significant differences, according to both the maximin and equitability objectives, between the maximin solution (which optimizes both objectives simultaneously) and an arbitrary envy-free solution (which does not attempt to optimize either objective).

Second, we report results from a user study. We contacted Spliddit users, and asked them to compare two solutions: the maximin solution, and an arbitrary envy-free solution. Crucially, the two solutions were computed on each user's actual Spliddit instance (the values of other tenants were perturbed to preserve privacy). Subjects were asked to subjectively rate the solutions in terms of fairness to themselves, and fairness to others. The results show a significant advantage for the maximin solution in both questions, thereby demonstrating the added value of optimization and supporting the decision to use the maximin solution on Spliddit.

### 1.3. Related work

The idea of refining envy-free solutions has been explored in several papers,[2, 20, 21, 22] typically from an axiomatic viewpoint. We focus on the work of Alkan et al.,[2] who study the more general problem of allocating goods and dividing money. They start by proving the existence of envy-free solutions in this setting, but, like us, they ultimately employ criteria of justice in order to find the "best" envy-free solutions. They are especially interested in the maximin solution, which they call the *value-Rawlsian* solution; and the solution that maximizes the minimum amount of money allocated to any player, subject to envy freeness, which they call the *money-Rawlsian* solution. They show that the maximin solution is unique, as are a number of less attractive solutions (minimize the maximum utility, maximize the utility of one particular player). Finally, they show that these criteria imply solutions with a monotonicity property: if the amount of money is increased, the utility of all players is strictly higher (this property is moot in our setting). Alkan et al.[2] do not provide algorithmic results.

Aragones[3] designs a polynomial-time algorithm for computing the money-Rawlsian solution of Alkan et al.[2] Her combinatorial algorithm does not extend to other criteria. In contrast, our LP-based framework is significantly more general, and, in particular, allows us to compute the maximin solution (which we view as the most attractive) in polynomial time. Our algorithmic approach is also much simpler. It is worth noting that Klijn[14] gives a different polynomial-time algorithm for computing envy-free solutions, without guaranteeing any additional properties (other than being extreme points of a certain polytope).

There are (at least) three *marketlike* mechanisms for computing solutions for the rent division problem assuming quasi-linear utilities, by Brams and Kilgour,[4] Haake et al.,[11] and Abdulkadiroğlu et al.[1] All three do not consider optimization criteria; in the case of the mechanism of Brams and Kilgour,[4] the solution may not be envy free. As mentioned above, the mechanism of Abdulkadiroğlu et al.[1] was deployed on Spliddit until April 2015.

One fundamentally different approach to rent division that we would like to discuss in more detail is that of Su.[18] He does not assume quasi-linear utilities; rather, his main assumption is that a player would always prefer getting a free room to getting another room at a positive price (the so-called *miserly tenants* assumption). Under this assumption, Su[18] designs an algorithm that converges to an (approximately) envy-free solution, by iteratively querying players about their favorite room at given prices. While eschewing the quasi-linear utilities assumption is compelling, a (crucial, in our view) disadvantage of this approach is that preference elicitation is very cumbersome. Interestingly, Su's method was implemented by the New York Times.[c]

Relatively few papers explore fair allocations among people in lab settings, and there is inconclusive evidence about the types of solution criteria that are favored by people. Dupuis-Roy and Gosselin[7] report that fair division algorithms were rated less desirable than imperfect allocations that did not employ any fairness criterion, while Schneider and Krämer[17] find that subjects preferred envy-free solutions to a divide-and-choose method that does not guarantee envy freeness. Herreiner and Puppe[12, 13] find that envy freeness was a dominant factor in the allocations favored by subjects, but that it was a secondary criterion to Pareto optimality or inequality minimizing allocations. Kohler[15] proposes an equilibrium strategy for repeated negotiation that incorporates fairness and envy concerns. In all of these papers, the studies were conducted in a controlled lab setting in which subjects' valuations over goods were imposed on the subjects, or the goods to be allocated were chosen by the experimenters themselves.

## 2. THE MODEL

We are interested in rent division problems involving a set of players $[n] = \{1, \ldots, n\}$, and a set of rooms $[n]$. Each player $i$ has a non-negative value $v_{ij} \in \mathbb{R}^+$ for each room $j$. We assume without loss of generality that the total rent is 1, and also assume (with loss of generality) that for all $i \in [n]$, $\sum_{j=1}^{n} v_{ij} = 1$. We can therefore represent an instance of the rent division problem as a right stochastic (rows sum to 1) matrix $V \in \mathbb{M}_{n\times n}(\mathbb{R}^+)$.

An assignment of the rooms is a permutation $\sigma: [n] \to [n]$, where $\sigma(i)$ is the room assigned to player $i$. The division of rent is represented through a vector of (possibly negative) prices $\mathbf{p} \in \mathbb{R}^n$ such that $\sum_{i=1}^{n} p_i = 1$; $p_j$ is the price of room $j$.

Given a solution $(\sigma, \mathbf{p})$ for a rent division problem $V$, the quasi-linear *utility* of player $i$ is denoted $u_i(\sigma, \mathbf{p}) = v_{i\sigma(i)} - p_{\sigma(i)}$. A solution is *Envy Free (EF)* if the utility of each player for her room is at least as high as any other room. Formally, $(\sigma, \mathbf{p})$ is EF if and only if

$$\forall i, j \in [n],\ \ v_{i\sigma(i)} - p_{\sigma(i)} \geq v_{ij} - p_j. \tag{1}$$

## 3. COMPUTATION OF OPTIMAL ENVY-FREE SOLUTIONS

As noted above, it is possible to compute an envy-free solution to a given rent division problem in polynomial time.[3] We are interested in choosing among envy-free allocations by optimizing an objective function, subject to the envy freeness constraint. Our goal in this section is to show that this

[c] http://goo.gl/Xp3omV. This article also discusses the then under-construction Spliddit.

can be done in polynomial time, when the objective function is the minimum of linear functions of the utilities.

THEOREM 3.1. *Let* $f_1, \ldots, f_t\colon \mathbb{R}^n \to \mathbb{R}$ *be linear functions, where t is polynomial in n. Given a rent division instance V, a solution* $(\sigma, \mathbf{p})$ *that maximizes the minimum of* $f_q(u_1(\sigma, \mathbf{p}), \ldots, u_n(\sigma, \mathbf{p}))$ *over all* $q \in [t]$ *subject to envy freeness can be computed in polynomial time.*

Natural examples of objective functions of the form specified in the theorem are maximizing the minimum utility, and minimizing the maximum difference in utilities; we discuss these objectives in detail in Section 4. The former objective can be directly captured by setting $t = n$, and $f_i(u_1(\sigma, \mathbf{p}), \ldots, u_n(\sigma, \mathbf{p})) = u_i(\sigma, \mathbf{p})$ for all $i \in [n]$. The latter criterion is also captured by setting $t = n^2$ and,

$$f_{ij}(u_1(\sigma,\mathbf{p}),\ldots,u_n(\sigma,\mathbf{p})) = u_i(\sigma,\mathbf{p}) - u_j(\sigma,\mathbf{p}).$$

Indeed,

$$\min_{i,j\in[n]} f_{ij}(u_1(\sigma,\mathbf{p}),\ldots,u_n(\sigma,\mathbf{p})) = -\max_{i,j\in[n]} \{u_i(\sigma,\mathbf{p}) - u_j(\sigma,\mathbf{p})\},$$

so maximizing the minimum of these linear functions is equivalent to minimizing the maximum difference in utilities.

Our polynomial-time algorithm relies on a connection between envy-free rent division and the concept of *Walrasian equilibrium*. To understand this connection, imagine a more general setting where a set of buyers $[n]$ are interested in purchasing bundles of goods $G$; here, each buyer $i$ has a valuation function $v_i\colon 2^G \to \mathbb{R}$, assigning a value $v_i(S)$ to every bundle of goods. A Walrasian equilibrium is an allocation $\mathbf{A} = (A_1, \ldots, A_n)$ of the goods to buyers (where $A_i \subseteq G$ is the bundle given to buyer $i$), coupled with a price vector $\mathbf{p}$ that assigns a price to each good, such that each player receives the best bundle of goods that she can buy for the price $\mathbf{p}$; formally:

$$\forall i\in[n], S\subseteq G,\ v_i(A_i) - p(A_i) \geq v_i(S) - p(S). \qquad (2)$$

We say that an allocation $\mathbf{A}$ is *welfare-maximizing* if it maximizes $\sum_{i=1}^n v_i(A_i)$. The following properties of Walrasian equilibria are well known; see, for example, the book of Mas-Colell et al.[16] (Chapter 16).

THEOREM 3.2 (1ST WELFARE THEOREM). *If* $(\mathbf{A}, \mathbf{p})$ *is a Walrasian equilibrium, then* $\mathbf{A}$ *is a welfare-maximizing allocation.*

THEOREM 3.3 (2ND WELFARE THEOREM). *If* $(\mathbf{A}, \mathbf{p})$ *is a Walrasian equilibrium, and* $\mathbf{A}'$ *is a welfare-maximizing allocation, then* $(\mathbf{A}', \mathbf{p})$ *is a Walrasian equilibrium as well. Furthermore,* $v_i(A_i) - p(A_i) = v_i(A_i') - p(A_i')$ *for all* $i \in [n]$.

Now, an EF solution in the rent division setting is a Walrasian equilibrium in the setting where the goods are the rooms, and the valuation function of each player for a subset $S \subseteq [n]$ of rooms is given by $v_i(S) = \max_{j\in S} v_{ij}$ (these are *unit demand* valuations)—it is easily seen that Equation (1) coincides with Equation (2) in this case. This means that we can apply the welfare theorems to EF allocations. For example, we can immediately deduce a simple result of Svensson[19]: any EF solution $(\sigma, \mathbf{p})$ is Pareto efficient, in the sense that there is no other solution $(\sigma', \mathbf{p}')$ such that $u_i(\sigma', \mathbf{p}') \geq u_i(\sigma, \mathbf{p})$ for all $i \in [n]$, with strict inequality for at least one $i \in [n]$. To see this, note that $\sigma$ is welfare-maximizing by Theorem 3.2, and the sum of prices is 1 under both $\mathbf{p}$ and $\mathbf{p}'$.

We are now ready to present our polynomial-time algorithm for maximizing the minimum of linear functions $f_1, \ldots, f_t$ of the utilities, subject to EF; it is given as Algorithm 1.

**ALGORITHM 1:**

1. Let $\sigma \in \arg\max_\pi \left\{\sum_{i=1}^n v_{i\pi(i)}\right\}$ be a welfare-maximizing assignment
2. Compute a price vector $\mathbf{p}$ by solving the linear program

$$\begin{aligned}
&\max R \\
&\text{s.t.: } R \leq f_q\left(v_{1\sigma(1)} - p_{\sigma(1)}, \ldots, v_{n\sigma(n)} - p_{\sigma(n)}\right) && \forall q\in[t] \\
&\qquad v_{i\sigma(i)} - p_{\sigma(i)} \geq v_{ij} - p_j && \forall i,j\in[n] \\
&\qquad \sum_{j=1}^n p_j = 1
\end{aligned}$$

The algorithm starts by computing a welfare-maximizing assignment $\sigma$ of players to rooms; this can be done in polynomial time, as this task reduces to the maximum weight bipartite matching problem, with players on one side of the graph, rooms on the other, and a weight $v_{ij}$ on each edge $(i, j)$. It then solves (in polynomial time) a linear program, with variables $p_1, \ldots, p_n$, which computes optimal envy-free prices *with respect to* $\sigma$. The first constraint sets (in an optimal solution) the objective $R$ to the minimum of the linear functions $f_q(\cdot)$. Envy-freeness is enforced by the second constraint, and the third constraint guarantees that the prices sum to 1.

However, it may not be immediately clear why starting from an arbitrary welfare-maximizing assignment allows us to compute the optimal solution subject to envy freeness. In a nutshell, the reason is the second Welfare Theorem: If $(\sigma', \mathbf{p})$ is an optimal EF solution, and $\sigma$ is an arbitrary welfare-maximizing assignment, then $(\sigma, \mathbf{p})$ is EF (so $\mathbf{p}$ is a feasible solution to the linear program) and induces the same utilities as $(\sigma', \mathbf{p})$, that is, it achieves the same objective function value.

## 4. RELATIONS BETWEEN THE FAIREST SOLUTIONS

Algorithm 1 allows us to maximize the minimum of linear functions of the utilities, subject to EF, in polynomial time. With the potential computational barrier out of the way, we would like to understand which optimization objective to use. Specifically, we focus on two natural optimization objectives, and evaluate their properties.

We refer to the first objective as *equitability*. Let $EF(V)$ be the set of all EF solutions for $V$. Given a solution $(\sigma, \mathbf{p}) \in EF(V)$, we define $D(\sigma, \mathbf{p})$ as the difference between the utilities of the happiest player and the worst off player under the solution $(\sigma, \mathbf{p})$, that is,

$$D(\sigma,\mathbf{p}) = \max_{i,j\in N} \{u_i(\sigma,\mathbf{p}) - u_j(\sigma,\mathbf{p})\}.$$

In more general terms, the function $D$ measures the social *disparity* under the solution $(\sigma, \mathbf{p})$; we would like to minimize

this quantity. An outcome $(\sigma^\star, \mathbf{p}^\star)$ is called *equitable* if it minimizes $D$ over $EF(V)$, that is,

$$(\sigma^\star, \mathbf{p}^\star) \in \arg\min\{D(\sigma, \mathbf{p}) \mid (\sigma, \mathbf{p}) \in EF(V)\}.$$

Herreiner and Puppe[12] demonstrate via experiments with human subjects that equitability is of great importance in determining whether an allocation is perceived to be fair by people.

Alternatively, instead of minimizing social disparity, one might be interested in maximizing the utility of the worst off player. More formally, given an EF solution $(\sigma, \mathbf{p})$, we let $U(\sigma, \mathbf{p}) = \min_{i\in N} u_i(\sigma, \mathbf{p})$; if

$$(\sigma^\star, \mathbf{p}^\star) \in \arg\max\{U(\sigma, \mathbf{p}) \mid (\sigma, \mathbf{p}) \in EF(V)\} \tag{3}$$

then we say that $(\sigma^\star, \mathbf{p}^\star)$ is a *maximin* solution.

Alkan et al.[2] argue that the maximin solution—which they call the *value-Rawlsian solution*—is compelling on philosophical grounds. Mathematically, they demonstrate that the maximin solution is associated with a unique vector of utilities, making this solution even more appealing.

The fact that equitable and maximin allocations are constrained to be EF again allows us to employ the Second Welfare Theorem (Theorem 3.3) to great effect. Indeed, if $(\sigma^\star, \mathbf{p}^\star)$ is equitable (resp., maximin), and $\sigma'$ is a welfare-maximizing assignment, then $(\sigma', \mathbf{p}^\star)$ is equitable (resp., maximin). Therefore, hereinafter we assume without loss of generality that the identity assignment $\sigma(i) = i$ is welfare maximizing, and simply use $D(\mathbf{p})$ or $U(\mathbf{p})$ to refer to these measures under the identity assignment. In particular, we can talk about equitable or maximin vectors of prices with respect to the identity assignment.

At first glance, the equitability and maximin criteria seem equally appealing. Which one leads to fairer solutions? The next theorem shows that we do not have to choose—the maximin solution is equitable.

THEOREM 4.1. *If $\mathbf{p}^\star$ is a maximin vector of prices, then it is also equitable.*

By contrast, an equitable solution may not be maximin, as the following example shows.

EXAMPLE 4.2. This example is particularly appealing, as it is a real-world instance submitted by Spliddit users.

$$\begin{pmatrix} 2227 & 708 & 0 \\ 258 & 1378 & 1299 \\ 1000 & 1000 & 935 \end{pmatrix}$$

Note that the total rent is \$2935. The optimal room assignment gives room $i$ to player $i$; the maximin rent division is $\mathbf{p}^\star = (1813\frac{1}{3}, 600\frac{1}{3}, 521\frac{1}{3})$, with a utility vector of $u_1(id, \mathbf{p}^\star) = 413\frac{2}{3}$, $u_2(id, \mathbf{p}^\star) = 777\frac{2}{3}$, $u_3(id, \mathbf{p}^\star) = 413\frac{2}{3}$. We have $D(\mathbf{p}^\star) = 777\frac{2}{3} - 413\frac{2}{3} = 364$, and by Theorem 4.1 any solution that has the same disparity is equitable. However, the price vector $\mathbf{p}' = (1570\frac{2}{3}, 721\frac{2}{3}, 642\frac{2}{3})$ is an EF rent division resulting in $u_1(id, \mathbf{p}') = 656\frac{1}{3}$, $u_2(id, \mathbf{p}') = 656\frac{1}{3}$, $u_3(id, \mathbf{p}') = 292\frac{1}{3}$, and $D(\mathbf{p}') = 656\frac{1}{3} - 292\frac{1}{3} = 364$ as well, that is, it is equitable, but the minimum utility is (much) smaller than that under $\mathbf{p}^\star$.

Let us now discuss a third optimization objective, the *money-Rawlsian solution*, which is mentioned by Alkan et al.,[2] and implemented in polynomial time by Aragones.[3] The latter author describes the following procedure for finding EF solutions. Begin by finding a welfare-maximizing assignment of rooms (again, assume without loss of generality that room $i$ goes to player $i$); next, find a vector $\mathbf{q}^\star \in \mathbb{R}_+^n$ of non-negative values such that $v_{ii} + q_i^\star \geq v_{ij} + q_j^\star$ and $Q^\star = \sum_{i=1}^n q_i^\star$ is minimized. That is, each player $i$ pays a value of $-q_i^\star$. Next, increase the prices of all players by a quantity $\alpha$ such that $n\alpha - Q^\star = 1$, that is the vector $(\alpha, \ldots, \alpha) - \mathbf{q}^\star$ is a valid price vector.

While the money-Rawlsian solution is interesting, it may be "maximally unfair" in terms of disparity, as the following example shows.

EXAMPLE 4.3. We analyze the following rent division instance:

$$V = \begin{pmatrix} 1 & 0 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}$$

The welfare-maximizing assignment allocates room $i$ to player $i$, and $\mathbf{q}^\star = (0, \ldots, 0)$. A uniform increase in rent will ensue, resulting in the price vector (1/2, 1/2) and the utility vector (1/2, 0). Crucially, the money-Rawlsian price vector *maximizes* disparity among all EF solutions. Note that the maximin price vector is (3/4, 1/4), which, of course, minimizes disparity.

To conclude, so far we know that the maximin solution, the equitable solution, and the money-Rawlsian solution can be computed in polynomial time. Moreover, Theorem 4.1 shows that the maximin solution, which by definition maximizes the minimum utility, also minimizes disparity (among all EF solutions)—so it is a refinement of the equitable solution. In stark contrast, the money-Rawlsian solution may *maximize* disparity (among all EF solutions). We therefore view the maximin solution as the clear choice, and focus on analyzing its effectiveness hereinafter.

## 5. ON THE IMPORTANCE OF BEING EQUITABLE

Our goal in this section is to understand how much better the maximin solution is, in terms of the maximin and disparity objectives, compared to suboptimal solutions *on average*. The original version of the paper includes a theoretical analysis in a formal probabilistic model. Here we focus on empirical results, which demonstrate the practical benefit of the maximin solution with respect to real-world instances that were submitted by Spliddit users.

In our experiments, we compare the maximin solution to an *arbitrary* EF solution, which is obtained by solving a feasibility linear program without an optimization objective. (We note that similar empirical results are obtained when comparing the maximin solution to the algorithm of Abdulkadiroğlu et al.[1]) The comparison is in terms of both of our main objectives, $D$ and $U$ (which are simultaneously optimized by the maximin solution). We expected that $D$ would be *significantly lower*, and $U$ *significantly higher*, in the maximin solution compared to an arbitrary EF solution.

We focus our analysis on 1,358 rent division instances involving 3,682 players, which were submitted on Spliddit

between January 2015 and December 2015. The number of instances for each number of players 2, 3, 4, 5, 6, 7, 8, 9 is 698, 445, 160, 35, 9, 8, 1, 2, respectively. We only use instances that include 2, 3 or 4 players, for which we have at least 160 instances in the database and for which obtaining statistical significance was possible. Importantly, note that this is a small subset of the roughly 13,000 instances created by Spliddit users by the time the experiment was run in December 2015; this is because we selected instances very conservatively, to ensure the ecological validity of our analysis. For example, Spliddit allows a "live demo" mode of interaction, and we excluded instances created that way.

To illustrate users' values for rooms in the Spliddit dataset, we present Figure 2, which visualizes the distribution for two-player instances. The x axis shows the value of player 1 for room 1, and the y axis shows the value of player 2 for room 1. The total rent is normalized to \$1, so each player's value for room 2 is simply the complement of the displayed value; that is, the point $(x, y)$ corresponds to an instance where the values of player 1 are $(x, 1 - x)$, and those of player 2 are $(y, 1 - y)$. The diagonal from points $(0, 0)$ to $(1, 1)$ represents the points in which players completely agree on the rooms' values. We color each instance according to its distance from this line, using shades of red for shorter distances, and shades of blue for longer distances.

The figure reveals several interesting phenomena. First, there is a significant cluster of instances which is centered on or close to the (0.5, 0.5) mark, implying that both players are indifferent between the two rooms. Second, we see a "cross" centered at the (0.5, 0.5) point, in which one of the players is indifferent, while the other player prefers one of the two rooms. Third, there are some instances in which one or both of the players are obstinate (i.e., $x \in \{0, 1\}$ or $y \in \{0, 1\}$), that is, they desire a specific room at any cost.

Let us now turn to the comparison we promised above. Given a rent division instance $V$, let $\mathbf{p}^\star$ denote the price vector associated with the maximin solution, and $\mathbf{p}^{EF}$ denote the price vector associated with an arbitrary EF solution, as discussed earlier. As before, we let $D(\mathbf{p})$ and $U(\mathbf{p})$ denote the social disparity and utility of the worst-off player under price vector $\mathbf{p}$ (assuming a welfare-maximizing assignment of players to rooms). The improvement in social disparity $D$ from using the maximin price vector over the EF vector is defined as $D(\mathbf{p}^{EF})-D(\mathbf{p}^\star)$, and the improvement in the utility of the worst-off player $U$ from using the maximin price vector over the EF vector is defined as $U(\mathbf{p}^\star) - U(\mathbf{p}^{EF})$.

Figure 3 shows the percentage of improvement *out of the total rent* in $D$ and $U$. As shown by the figure, for $n = 2, 3, 4$, the disparity associated with the maximin solution is significantly lower than that of the EF solution (9% of the total rent on average), and the utility of the worst-off player associated with the maximin solution is significantly higher than that of the EF solution (4% of the total rent on average). This trend is exhibited with respect to each value of $n$.

We note the following points. First, the degree of improvement in both $D$ and $U$ becomes smaller as the number of players grows. However, even in cases where the improvement is relatively small, it still makes a *qualitative* difference, for example, when the maximin solution achieves zero disparity, and the arbitrary EF solution achieves strictly positive disparity (we discuss this fact in the next section). In addition, as noted above, the vast majority of Spliddit instances include two or three players, for which the improvement in $D$ and $U$ is higher than four players. Lastly, although this is not shown in the figure, an improvement in both $D$ and $U$ occurs in over 90% of the instances, for $n \in \{2, 3, 4\}$.

## 6. USER STUDY

In the previous sections, we established, both theoretically and empirically, the benefits of the maximin approach to computing envy-free solutions for rent division problems. The question addressed by this section is, are people willing to accept such solutions in practice? To answer this question, we conducted the following user study.

### 6.1. Study design

People who used the Spliddit service during the year 2015 were invited (via email) to participate in a short study to evaluate the new allocation method. We targeted users who participated in rent division instances on Spliddit that included 2, 3 or 4 players. In order to use Spliddit one need not supply an email

**Figure 2. The distribution of values for two player Spliddit instances (normalized to a total rent of \$1).**

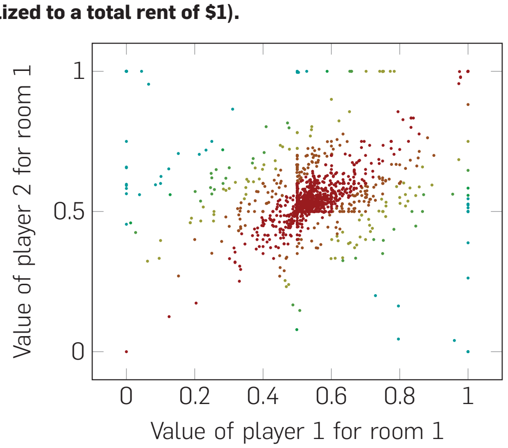


**Figure 3. Average percentage of improvement (out of the total rent) in social disparity $D$ and utility of the worst-off player $U$ when using the price vector associated with the maximin solution, compared to an arbitrary EF solution, on Spliddit instances.**

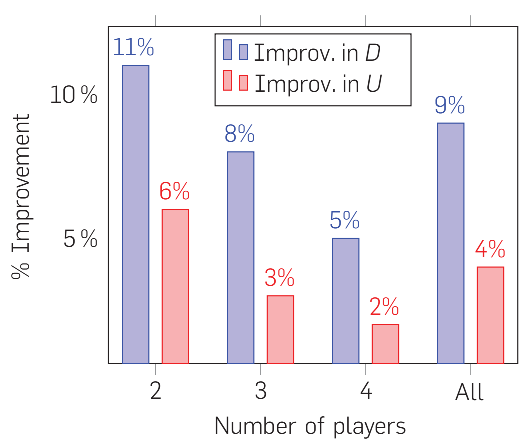

address; users can opt to send out URLs to other users, which is what the vast majority of users choose to do. We only contacted users who supplied their email address—a relatively small subset of the users who were involved in rent division instances.

All participants were given a \$10 compensation that did not depend on their responses. In total, the invitation email was sent to 344 Spliddit users, of which 46 users (13%) chose to participate. The study was approved by the Institutional Review Board (IRB) of Carnegie Mellon University.

The study followed a within-subject design, by which each of the subjects was shown, in random order, an arbitrary EF solution (as discussed in Section 5) and the maximin solution, applied to their *original problem instance*.

Importantly, we wished to preserve the privacy of players regarding their evaluations over the different rooms. Therefore, each player who participated in the study was shown a slightly modified version of their own rent division problem. Information that was already known to each subject was identical to the original Spliddit instance, including the total rent, the number of rooms, their names, the subject's *own* values for the different rooms, and the allocation of the rooms to the players. Information that was perturbed to preserve the privacy of the other players included their names, which were changed to "Alice", "Bob" or "Claire", depending on whether there were 2, 3, or 4 players; and the other players' valuations, which were randomly increased or decreased by a value of up to 15% under the constraint that the total rent is unchanged, and that player valuations are still valid (non-negative and sum to the total rent).

The subjects were shown the two solutions—maximin and arbitrary EF—for the instance presented to them. Both solutions include the same room allocation, but possibly differ in the prices paid by the players. The two solution outcomes were shown in sequence, and in random order.

The subjects were asked to rate two different aspects of each of the two solutions on a scale from 1 to 5, with 1 being least satisfied and 5 being most satisfied. The two aspects are the subject's individual allocation, and the allocations of the other players. The two questions were phrased as follows (using an example with $n = 3$):

1. **Individual:** This question relates to your own allocation. In other words, we would like you to pay attention *only* to your own benefit. How happy are you with getting the room called $\langle$*RoomName*$\rangle$ for \$$\langle$*price*$\rangle$?
2. **Others:** This question relates to the allocation for *everyone else*. How fair do you rate the allocation for Bob and Claire?

In both questions, players were able to write an argument or justification for their rating. To cancel order effects, the two questions were presented in random order.

**Figure 4. Results of the user study.**

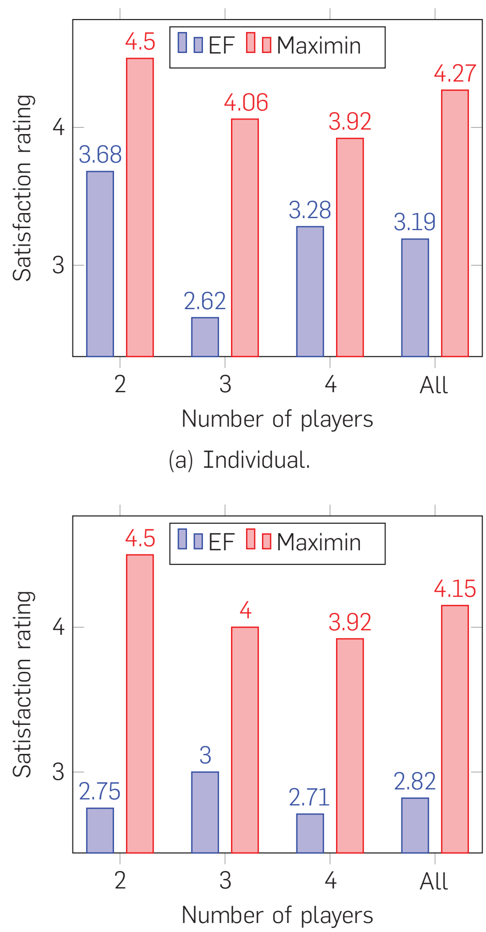


(a) Individual.

(b) Others.

### 6.2. Results

We hypothesized that players would rate their own allocation under the maximin solution significantly higher than under the EF solution, and similarly for the allocation of the other participants. Figure 4 shows the results of the user study. For each number of players (2, 3, 4) we show the average satisfaction level reported for the arbitrary EF solution and maximin solution when relating to each player's individual outcome (left chart), and others' outcomes (right chart). In all cases, the maximin solution is rated significantly higher than the envy-free solution for both questions, passing a Wilcoxon signed-rank test with $p < 0.04$.

Why did players overwhelmingly prefer the prices from the maximin solution over the arbitrary EF solution? Given the high importance attributed to social disparity when reasoning about fair division,[12] we hypothesized that the price vectors of the maximin solution exhibited significantly lower disparity than the price vectors of the EF solution. This was supported by many of the textual comments relating to social disparity. Figure 5 shows the cumulative distribution of disparity across all instances that were included in the user study. The x axis indicates the disparity as percentage of the total rent. As shown by the figure, the disparity associated with the maximin solution is indeed significantly lower. In fact, in many instances the disparity is zero under the maximin solution (this is guaranteed to be true when $n = 2$, as we show in the original version of the paper). We believe that this large difference in disparity played a key role in subjects' preference for the maximin solution, trumping the relatively small improvement in utilities.

**Figure 5. Cumulative distribution over the social disparity across all instances that were included in the user study. The *x*-axis indicates the percentage of social disparity out of the total rent price.**

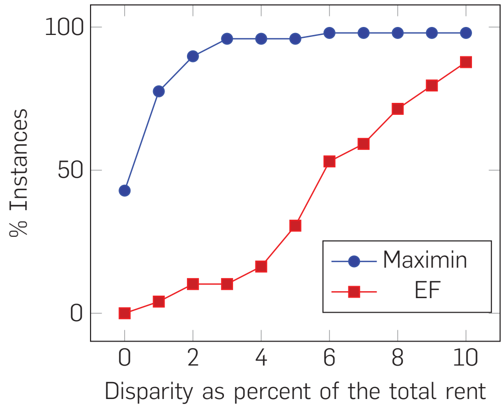


## 7. DISCUSSION

There are two practical questions that inevitably come up when we present our work on rent division, and its deployed application.

The first question is whether participants can achieve a better outcome by misreporting their values. Indeed, they can, and the reason we do not address such game-theoretic concerns is twofold. First, envy freeness is inherently incompatible with strategyproofness (immunity to strategic manipulation). This follows from the classic result of Green and Laffont[10] and the fact that envy freeness implies Pareto efficiency in our setting. More importantly, we believe that, in rent division, strategic behavior does not play a significant role in practice. In particular, most Spliddit users do not know how the algorithm works, as we do not attempt to explain the algorithm itself, only its fairness guarantees. While users can experiment with Spliddit's demo mode to determine the impact of various reported values on the outcome, doing this effectively would require an accurate estimate of the values submitted by others, and seems quite unwieldy in general. That said, being able to give some game-theoretic guarantees would be desirable, of course.

The second question is whether the quasi-linear utility model truly captures people's preferences. For example, one participant might believe that it is unfair that he is paying more for a room he values highly, when his housemate values the two rooms equally (this happens under the maximin solution in Example 4.3); or some participants may have budget constraints—they simply cannot pay more than a certain price. Clearly, these are valid concerns. However, there is a tradeoff between expressiveness and ease of elicitation. We believe that quasi-linear utilities hit a sweet spot between the two, in the sense that they are reasonably expressive, yet very easy to elicit (each user simply reports a value for each room). Nevertheless, some of us are studying rent division algorithms that support richer utility functions.

Taking a broader viewpoint, we believe that computational fair division is a prime example of how the interaction between computer science and economics can lead to novel applications. We find it particularly exciting that fundamental theoretical questions in this field have direct real-world implications, both on Spliddit,[6] and beyond. (Ref. Budish et al.[5]) The current paper (or the original version thereof) takes the computational fair division agenda a step further, by tying together theory, experiments on real data, a carefully designed user study, and a deployed application.

### Acknowledgments

This work was supported by EU FP7 FET project, grant agreement n.600854; by the National Science Foundation under grants IIS-1350598, CCF-1215883, and CCF-1525932; by the Office of Naval Research under grants N00014-16-1-3075 and N00014-17-1-2428; and by a Sloan Research Fellowship.

**References**

1. Abdulkadiroğlu, A., Sönmez, T., Ünver, M.U. Room assignment-rent division: A market approach. *Soc. Choice Welf. 22*, 3 (2004), 515–538.
2. Alkan, A., Demange, G., Gale, D. Fair allocation of indivisible goods and criteria of justice. *Econometrica 59*, 4 (1991), 1023–1039.
3. Aragones, E. A derivation of the money Rawlsian solution. *Soc. Choice Welf. 12*, (1995), 267–276.
4. Brams, S.J., Kilgour, D.M. Competitive fair division. *J. Polit. Econ. 109*, (2001), 418–443.
5. Budish, E., Cachon, G.P., Kessler, J., Othman, A. Course match: A large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation. *Oper. Res. 65*, 2 (2017), 314–336.
6. Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A.D., Shah, N., Wang, J. The unreasonable fairness of maximum Nash product. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)* (2016), 305–322.
7. Dupuis-Roy, N., Gosselin, F. The simpler, the better: A new challenge for fair-division theory. In *Proceedings of the 33rd Annual Meeting of the Cognitive Science Society (CogSci)* (2011), 3229–3234.
8. Foley, D. Resource allocation and the public sector. *Yale Econ. Essays* 7, (1967), 45–98.
9. Goldman, J., Procaccia, A.D. Spliddit: Unleashing fair division algorithms. *SIGecom Exchanges 13*, 2 (2014), 41–46.
10. Green, J.R. Laffont, J.-J. *Incentives in Public Decision Making.* North Holland, Amsterdam, The Netherlands, 1979.
11. Haake, C.-J., Raith, M.G., Su, F.E. Bidding for envy-freeness: A procedural approach to *n*-player fair-division problems. *Soc. Choice Welf. 19*, (2002), 723–749.
12. Herreiner, D.K., Puppe, C.D. Envy freeness in experimental fair division problems. *Theo. Decis. 67*, 1 (2009), 65–100.
13. Herreiner, D.K., Puppe, C.D. Inequality aversion and efficiency with ordinal and cardinal social preferences—An experimental study. *J. Econ. Behav Organ. 76*, 2 (2010), 238–253.
14. Klijn, F. An algorithm for envy-free allocations in an economy with indivisible objects and money. *Soc. Choice Welf. 17*, (2000), 201–215.
15. Kohler, S. Envy can promote more equal division in alternating-offer bargaining. *J. Neurosci. Psychol. Econ. 1*, 6 (2013), 31–41.
16. Mas-Colell, A., Whinston, M.D., Green, J.R. *Microeconomic Theory*. Oxford University Press, Oxford, U.K., 1995.
17. Schneider, G., Krämer, U. The limitations of fair division: An experimental evaluation of three procedures. *J. Confl. Resolut. 48*, 4 (2004), 506–524.
18. Su, F.E. Rental harmony: Sperner's lemma in fair division. *Am. Math. Mon. 106*, 10 (1999), 930–942.
19. Svensson, L.-G. Large indivisibles: An analysis with respect to price equilibrium and fairness. *Econometrica 51*, 4 (1983), 939–954.
20. Tadenuma, K., Thomson, W. No-envy and consistency in economies with indivisible goods. *Econometrica 59*, 6 (1991), 1755–1767.
21. Tadenuma, K., Thomson, W. Refinements of the no-envy solution in economies with indivisible goods. *Theor. Decis. 39*, 2 (1995), 189–206.
22. Velez, R. Sharing an increase of the rent fairly. *Soc. Choice Welf. 48*, 1 (2017), 59–80.

**Kobi Gal and Moshe Mash** ([kobig, mashm]@bgu.ac.il), Dept. of Software and Information Systems Engineering, Ben-Gurion University, Beer-Sheva, Israel.

**Ariel D. Procaccia** (arielpro@cs.cmu.edu), Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, USA.

**Yair Zick** (zick@comp.nus.edu.sg), School of Computing, National University of Singapore. The work was done while Zick was at Carnegie Mellon University, Pittsburgh, PA, USA.