

# Multi-Apartment Rent Division

Ariel D. Procaccia, Benjamin Schiffer, Shirley Zhang

Harvard University

## Abstract

Rent division is the well-studied problem of fairly assigning rooms and dividing rent among a set of roommates within a single apartment. A shortcoming of existing solutions is that renters are assumed to be considering apartments in isolation, whereas in reality, renters can choose among multiple apartments. In this paper, we generalize the rent division problem to the multi-apartment setting, where the goal is to both fairly choose an apartment among a set of alternatives and fairly assign rooms and rents within the chosen apartment. Our main contribution is a generalization of envy-freeness called *negotiated envy-freeness*. We show that a solution satisfying negotiated envy-freeness is guaranteed to exist and that it is possible to optimize over all negotiated envy-free solutions in polynomial time. We also define an even stronger fairness notion called *universal envy-freeness* and study its existence when values are drawn randomly.

## 1 Introduction

Rent division is a classic and intuitive problem within the space of fair division, in which a number of roommates are faced with the joint decision of how to assign rooms and split the total rent in a shared apartment. The problem is complicated by the fact that the rooms may vary widely and that players may have very different values for different rooms; for example, one room may have a better view, or one player may derive higher utility from having larger closets. Preferences over prices may also be complex, but it is commonly assumed that players’ utilities are *quasi-linear*: utility equals value minus price. The goal is to find an assignment of players and prices to rooms that takes such player preferences into account and satisfies a rigorous notion of fairness.

The gold standard for fairness in rent division is *envy-freeness*, which guarantees that each player has higher utility for their own room than for any other room, given the prices assigned to each room. In other words, in an envy-free allocation, no roommate would want to trade their room for any other room. Along with being easily justifiable (Procaccia 2019), a solution which satisfies envy-freeness is also guaranteed to exist for the rent division problem (Svensson 1983) under the assumption of quasi-linearity. Furthermore, such a

solution can be found in polynomial time (Aragones 1995), and it is possible to optimize linear objectives over envy-free solutions in polynomial time as well (Gal et al. 2017). Therefore, envy-freeness is both a compelling and computationally tractable fairness notion in the rent division setting.

The study of rent division is not driven merely by theoretical interest — it is a poster child for applications of fair division more broadly. In particular, the not-for-profit fair division website *Spliddit* (Goldman and Procaccia 2014) operated between 2014 and 2022 and offered “provably fair solutions to everyday problems.” Its rent division application implemented the algorithm of Gal et al. (2017); among the five applications on Spliddit, it was the most popular, with more than 30,000 instances solved (Peters, Procaccia, and Zhu 2022).

A shortcoming of the prevalent approach to fair rent division, however, is that it assumes that the players have already chosen an apartment, and the only question remaining is how to divide the rooms and rent. By contrast, groups looking for an apartment are often not considering each potential apartment separately. Rather, many groups can further optimize by choosing among a set of available apartments which fit their budget and location constraints. This setting gives rise to new sources of complexity as the players are required not only to assign rooms and prices in the chosen apartment, but also to collectively decide on which apartment to rent. Such a decision can be contentious. Imagine, for instance, three players whose apartment hunting priorities are a short commute, proximity to green spaces, and an exciting neighborhood, respectively. Suppose further that current apartment options include one apartment that is close to player 1’s workplace, one which neighbors the largest park in the city, and one in the center of Restaurant Row. The rooms in each apartment may be asymmetric as well. Given all of these factors, how should the players decide which apartment to rent and who gets which room?

Crucially, it does not suffice to merely compute an envy-free solution in each individual apartment. Consider the following example:

**Example 1.1.** There are two players and two apartments. Each apartment has total rent 300 and contains two symmetric rooms. The value of each player (rows) for each room

(columns) is shown below:

	$r_{11}$	$r_{12}$
1	200	200
2	100	100

	$r_{21}$	$r_{22}$
1	100	100
2	200	200

In this example, the only solution which is individually envy-free in each apartment assigns equal rent in each apartment. However, the players will then disagree on which apartment to rent, as player 1 will prefer apartment 1 while player 2 will prefer apartment 2. It also seems intuitively unfair to assign equal rent in each apartment, as this would result in the two players having unequal utilities, even though their utility functions are symmetric. Therefore, a new fairness notion is needed for the multi-apartment rent division problem.

### 1.1 Our Contributions

First and foremost, we present a formal model of the rent division problem in the multi-apartment setting, and show that new fairness notions are necessary in this model. Our main contribution is a generalization of envy-freeness called *negotiated envy-freeness*. We show that negotiated envy-freeness satisfies several desirable properties such as Pareto optimality and individual rationality, and reduces to envy-freeness in the single apartment setting. We provide an intuitive justification for negotiated envy-freeness based on negotiating rent between players who have different favorite apartments, and show how such negotiations allow players to reach a consensus apartment while maintaining fair rent burdens across players. We then show that a solution satisfying negotiated envy-freeness is guaranteed to exist and that it is possible to optimize linear objectives over all negotiated envy-free solutions in polynomial time, mirroring a similar result by Gal et al. (2017) for the single-apartment setting. Finally, we introduce strong negotiated envy-freeness, a variant of negotiated envy-freeness which imposes additional fairness constraints on negotiations.

We also explore *universal envy-freeness*, which is the most direct generalization of envy-freeness. Unlike negotiated envy-freeness, however, we show that a universal envy-free solution is not guaranteed to exist (and, in fact, there are many instances where such a solution does not exist). Therefore, we instead study the probability of such a solution existing when players’ utilities are drawn i.i.d. from an arbitrary distribution with a fixed number of players. For discrete distributions, we show that the probability that a universal envy-free solution exists approaches 1 as the number of apartments approaches infinity. For continuous distributions, on the other hand, we show that the probability that a universal envy-free solution exists does not converge to either 0 or 1. However, if we add apartments one by one, the probability that there exists a stopping point where a universal envy-free solution exists converges to 1, even for continuous distributions.

### 1.2 Related Work

As mentioned earlier, our work is most closely related to the paper of Gal et al. (2017). They study envy-free solutions for

(single-apartment) rent division under quasi-linear utilities, and—building on earlier work by Alkan, Demange, and Gale (1991)—single out the *maximin solution* (which maximizes the minimum utility subject to envy-freeness) as especially desirable. They also develop an algorithmic framework that allows for efficient computation of the maximin solution and a range of other objectives.

By contrast, Peters, Procaccia, and Zhu (2022) consider a different type of objective: they seek envy-free solutions that—rather than optimizing a welfare function—are robust to perturbations of the utilities. This objective is beyond the scope of our work.

There are several approaches to rent division that relax the assumption of quasi-linear utilities or provide (incomparable) alternatives. One line of work allows players to express a (hard or soft) budget constraint (Procaccia, Velez, and Yu 2018; Airiau et al. 2023; Velez 2022, 2023). Arunachaleswaran, Barman, and Rathi (2021) develop a fully-polynomial approximation scheme for envy-free rent division under the assumption that each player’s utility (as a function of price) is continuous, monotone decreasing, and piecewise-linear. Finally, classic work by Su (1999) uses Sperner’s Lemma to construct an algorithm for envy-free rent division under the “miserly tenants” assumption, which requires players to prefer a free room to any other room (even if its rent is \$1); while his approach is elegant, it has several shortcomings, including that preference elicitation requires repeated interaction with the players and that it is infeasible to optimize over envy-free solutions. Segal-Halevi (2022), however, shows that techniques developed for miserly tenants extend to the quasi-linear setting. All of these papers are orthogonal to ours, as we focus on the (widely used in practice) quasi-linear setting and instead extend the standard rent division problem to multiple apartments.

It is worth noting that even in the basic setting of a single apartment and quasi-linear utilities, envy-freeness is incompatible with strategyproofness. For this reason, work on incentives in fair rent division is relatively limited. A notable exception is the work of Velez (2018), who studies mechanisms whose equilibria give rise to envy-free solutions.

## 2 Model

A multi-apartment rent division instance is composed of a set of players  $[n] = \{1, \dots, n\}$  and a set of apartments  $[m] = \{1, \dots, m\}$ , where each apartment  $j$  consists of  $n$  rooms  $\{r_{j1}, \dots, r_{jn}\}$ . For each room  $r_{jk}$  in apartment  $j$ , player  $i$  has a non-negative value  $V_i(r_{jk})$ . Each apartment  $j$  also has a total rent  $R_j$ . Unless otherwise noted, we will assume that  $\sum_j \sum_k V_i(r_{jk}) = \sum_j R_j$  for all  $i$ . In other words, we do not assume that players have the same value for each apartment, but do assume that players have the same total value for all apartments under consideration. In our model, players are therefore able to express preferences over apartments, but their overall utility is still normalized as in the single apartment setting (Gal et al. 2017). The entire instance can be represented by a valuation matrix  $V \in \mathbb{M}_{n \times m \times n}(\mathbb{R}^+)$

and a rent vector  $R \in \mathbb{R}^n$ .

An apartment assignment  $A_j : [n] \rightarrow [n]$  is a mapping of players to rooms in apartment  $j$ , where  $A_j(i)$  is the room assigned to player  $i$  in apartment  $j$ . An assignment  $A = \{A_1, \dots, A_m\}$  is the vector of mappings for all apartments. An assignment for apartment  $j$  is *welfare-maximizing* if it maximizes  $\sum_{i=1}^n V_i(A_j(i)) - R_j$  over all possible assignments in  $j$ . Rent is allocated via the price matrix  $P \in \mathbb{R}^{m \times n}$ , and we denote the price of a specific room  $r$  as  $P(r)$ . We require that  $\sum_{k \in [n]} P(r_{jk}) = R_j$  for any valid price matrix, and will refer to the price vector for apartment  $j$  as  $P_j$ , where  $P_j$  is the  $j$ th row of  $P$ . For a specific  $A_j, P$ , the quasi-linear utility of player  $i$  is  $U_i(A_j, P) = V_i(A_j(i)) - P(A_j(i))$ . A *solution* for the multi-apartment rent division instance is a tuple  $(A, P, j^*)$ , where  $A$  and  $P$  contain all assignments and prices and  $j^*$  denotes the chosen apartment. We will often refer to a *partial solution*  $(A, P)$  where the apartment has not yet been chosen.

When there is only a single apartment ( $m = 1$ ), a solution  $(A_1, P)$  is *envy-free* (EF) if each player prefers her room over every other room. Formally, a solution is envy-free if for all  $i, i' \in [n]$ ,

$$V_i(A_1(i)) - P(A_1(i)) \geq V_i(A_1(i')) - P(A_1(i')). \quad (1)$$

In the single apartment setting, an envy-free solution always exists and can be found in polynomial time (Gal et al. 2017). The goal of this paper is to generalize the notion of envy-freeness in Equation (1) to the multi-apartment setting when  $m > 1$ . In the multi-apartment setting, one starting point is to simply enforce the single apartment definition of envy-freeness for each apartment separately, i.e. for all  $i, i' \in [n], j \in [m]$ ,

$$V_i(A_j(i)) - P(A_j(i)) \geq V_i(A_j(i')) - P(A_j(i')). \quad (2)$$

We will refer to partial solutions  $(A, P)$  that satisfy Equation (2) as *individually envy-free*. However, as we showed in the introduction, it may be impossible for the players to agree on the final apartment choice  $j^*$ , and therefore this does not lead to an obviously fair solution.

### 3 Universal Envy-Freeness

We first consider *universal envy-freeness*, a natural generalization of envy-freeness which captures the spirit of the original definition. Informally, a universal envy-free assignment guarantees that no player will want to switch her room in the chosen apartment with any room in any apartment. We formalize this definition below.

**Definition 3.1.** A solution  $(A, P, j^*)$  is **universal envy-free (UEF)** if for all  $i, i' \in [n]$  and  $j \in [m]$ ,

$$V_i(A_{j^*}(i)) - P(A_{j^*}(i)) \geq V_i(A_j(i')) - P(A_j(i')).$$

Unfortunately, a universal envy-free solution does not always exist, as can be seen in Example 1.1. In that example, the only way for both apartments to be individually envy-free is for rent to be assigned evenly in each apartment; i.e.

that the price of each room is 150. However, if rent is assigned in this way, then one of the players will be envious in the final apartment: if apartment 1 were chosen, then player 2 would rather have either room in apartment 2, while if apartment 2 were chosen, then player 1 would rather have either room in apartment 1. Therefore, no universal envy-free solution exists for Example 1.1. It turns out that this is not a knife's edge example, and that there are many instances which have no universal envy-free solution. For example, there are instances with no universal envy-free solution even when we restrict all players to have the same total value for each apartment ( $\sum_{k=1}^n V_i(r_{jk}) = R_j \forall i, j$ ). We note, however, that it is easy to check whether a universal envy-free solution exists via a simple linear program [Appendix A.4].

#### 3.1 Probabilistic Universal Envy-Freeness

In computational fair division, when researchers were faced with the unfortunate non-existence of envy-free allocations of indivisible goods, they have asked whether such solutions are *likely* to exist in random instances, at least in the large (Dickerson et al. 2014; Manurangsi and Suksompong 2017, 2020). In this section, we adopt the same approach in the context of universal envy-freeness.

Specifically, we assume that the valuation matrix  $V$  is drawn randomly, where each player's value for each room is drawn from a distribution  $\mathcal{D}$  supported on  $[0, 1]$ . In other words, for any player  $i$  and room  $r_{jk}$  in apartment  $j$ , the value of player  $i$  for room  $r_{jk}$  is  $V_i(r_{jk}) \stackrel{i.i.d.}{\sim} \mathcal{D}$ . Because each apartment is drawn symmetrically, we will also assume for this section only that the rent in every apartment is equal to  $R$ . Importantly, note that under this model we are no longer requiring player utilities to be normalized to add up to the total rent. However, as values are drawn i.i.d, every player will have the same total value for all of the rooms in expectation.

Define  $E_m$  as the event that there exists a universal envy-free solution with  $m$  apartments when values are drawn i.i.d from distribution  $\mathcal{D}$ . First, we consider the simpler case when  $\mathcal{D}$  is a discrete distribution with a finite number of values (see Appendix A for the proof).

**Theorem 3.2.** *If  $V_i(r_{jk}) \stackrel{i.i.d.}{\sim} \mathcal{D}$  for all  $i, j, k$  where  $\mathcal{D}$  is a discrete distribution that takes on  $\kappa < \infty$  distinct values, then for any constant  $n$ ,  $\Pr(E_m) \xrightarrow{m \rightarrow \infty} 1$ .*

The above result holds in the limit as  $m$  goes to infinity, which is unrealistic in the real-world setting of searching for an apartment. However, a group of roommates may consider dozens of apartments in their search, in which case limit laws such as this result can become useful approximations. We would also ideally like to generalize this result to continuous distributions. Somewhat surprisingly, the same result does not hold in the case of continuous distributions supported on  $[0, 1]$ . In fact for sufficiently large  $m$ , we can bound the probability that there exists a universal envy-free solution away from both 0 and 1 for a fixed value of  $n$ . Note that the lower bound is the same for any continuous distribution, while the upper bound is distribution-dependent.

**Theorem 3.3.** *Suppose that  $V_i(r_{jk}) \stackrel{i.i.d.}{\sim} \mathcal{D}$  for all  $i, j, k$  where  $\mathcal{D}$  is a continuous distribution supported on  $[0, 1]$ . Then there exists a  $p_0(n) > 0$  such that for any  $m$  and  $\mathcal{D}$ ,  $\Pr(E_m) \geq p_0(n)$ . Furthermore, for any distribution  $\mathcal{D}$ , there exists a  $p_1(n) < 1$  such that for any  $m \geq n + 1$ ,  $\Pr(E_m) \leq p_1(n)$ .*

*Proof sketch.* We provide a brief sketch of the proof of Theorem 3.3 and defer the formal proof to Appendix A. The structure of the proof is to construct two events  $\mathcal{F}$  and  $\mathcal{E}$  (both with probabilities independent of  $m$ ) such that under event  $\mathcal{F}$  a universal envy-free solution always exists and under event  $\mathcal{E}$  a universal envy-free solution never exists. The key idea behind both constructed events is to consider the highest welfare achievable by potentially non-bijective assignments that can assign multiple rooms within an apartment to the same player. Note that these are not necessarily valid regular assignments. We will denote the maximum total utility achievable by a potentially non-bijective assignment for a given apartment  $j$  as  $MUW(j)$ . Suppose the apartments are numbered in order of  $MUW$ , and therefore apartment 1 has the highest  $MUW$ . We define  $\mathcal{F}$  as the event when the potentially non-bijective assignment with the highest total utility in apartment 1 is actually bijective, and therefore is a valid regular assignment. Under event  $\mathcal{F}$ , a universal envy-free solution exists. We then show that  $\Pr(\mathcal{F})$  is always positive and independent of  $m$ . To show non-existence of a universal envy-free solution, we construct the event  $\mathcal{E}$  with the three following conditions on the first  $n + 1$  apartments based on the  $MUW$  ordering. The event  $\mathcal{E}$  occurs when apartment  $n + 1$  is the unique maximum welfare apartment, the potentially non-bijective assignment which achieves  $MUW(n + 1)$  is a bijective assignment, and for every  $j \in [n]$ , the potentially non-bijective assignment which achieves  $MUW(j)$  assigns every room to player  $j$ . Conditioned on event  $\mathcal{E}$ , we show that no universal envy-free solution exists. Once again,  $\Pr(\mathcal{E})$  is always positive and independent of  $m$ .  $\square$

Theorem 3.3 implies that simply starting with a very large number of apartments is not sufficient for guaranteeing existence of a universal envy-free solution. However, the construction of event  $\mathcal{F}$  in Theorem 3.3 also implies the following. Suppose there are  $n$  players trying to find a universal envy-free solution, starting with  $m_0$  apartments. If new apartments are added one at a time, and we check for universal envy-freeness before each new apartment is added, then with probability 1 this process will terminate with a universal envy-free solution in a finite number of apartments. The proof of this result leverages the fact that event  $\mathcal{F}$  relies only on the ordering of utilities within the apartment with the highest value of  $MUW$ . This result is formally outlined in Corollary 3.4 and proven in Appendix A.

**Corollary 3.4.** *Suppose there exists an infinite sequence of apartments with  $V_i(r_{jk})$  drawn i.i.d from a continuous distribution  $\mathcal{D}$ . Then for any constant  $m_0 > 0$ ,*

$$\Pr(\inf\{m \geq m_0 : E_m\} < \infty) = 1.$$

Conceptually, this corollary sends an encouraging message to apartment hunters: perseverance will likely lead to a fair outcome!

## 4 Consensus and Negotiated Envy-Freeness

As universal envy-freeness is often too strong of a requirement, we would like to find a weaker condition that is always feasible, but still acts as a natural extension of envy-freeness in the single apartment setting. In this section, we introduce a fairness condition for the multi-apartment rent division problem that satisfies multiple desirable properties and is always guaranteed to exist.

### 4.1 Definition and Motivation

First, we observe that a desirable condition in the multiple apartment setting is for all players to agree on the chosen apartment. In order for this to happen, every player must be at least as happy with their assignment in the chosen apartment as with their assignment in any other apartment. If this is the case, then we say that the chosen apartment is a *consensus* apartment.

**Definition 4.1.** A solution  $(A, P, j^*)$  satisfies **consensus** if every player weakly prefers their assignment in  $j^*$  to their assignment in any other apartment  $j \neq j^*$ . Formally, for every player  $i$ ,

$$V_i(A_{j^*}(i)) - P(A_{j^*}(i)) \geq \max_{j \in [m]} V_i(A_j(i)) - P(A_j(i)). \quad (3)$$

A partial solution  $(A, P)$  satisfies consensus if there exists an apartment  $j^*$  satisfying Equation (3). We will refer to such a  $j^*$  as a **consensus apartment** for  $(A, P)$ .

Given a partial solution that satisfies consensus, the set of consensus apartments is exactly the set of apartments which have the highest sum of player utilities. If there is more than one consensus apartment, every player has the same utility for their assigned room in each consensus apartment; the next lemma, whose proof is in Appendix B, formalizes this.

**Lemma 4.2.** *Let  $(A, P)$  be a partial solution satisfying consensus. Apartment  $j$  is a consensus apartment for  $(A, P)$  if and only if*

$$\sum_{i=1}^n V_i(A_j(i)) - R_j \geq \sum_{i=1}^n V_i(A_{j'}(i)) - R_{j'}$$

*for all  $j' \neq j$ . Furthermore, if  $j_1$  and  $j_2$  are both consensus apartments for  $(A, P)$ , then for all  $i$ ,  $U_i(A_{j_1}, P) = U_i(A_{j_2}, P)$ .*

Intuitively, consensus is a desirable property because a lack of consensus would imply that the players cannot decide on which apartment to rent. However, consensus alone is not sufficient. In the motivating example, for instance, a possible solution that satisfies consensus is to have player 1 pay 300 for room  $r_{11}$  and 200 for room  $r_{21}$ . However, this is clearly not a fair assignment, as player 1 has  $-100$  utility in both apartments while player 2 has 100 utility in both apartments, despite their valuations being perfectly symmetrical.

This example can be further extended to make player 1 arbitrarily unhappy in the consensus apartment. Therefore, we need a requirement stronger than just consensus to guarantee reasonable fairness for all players.

To prevent one player being significantly more unhappy in every apartment as in the previous example, we can instead try to reach consensus from a “fair” starting partial solution. One such natural starting point is a partial solution  $(A, P)$  which is individually envy-free. If this starting partial solution satisfies consensus, then the solution with a consensus apartment already satisfies universal envy-freeness. However, this starting partial solution may not satisfy consensus, as in Example 1.1. If no  $j^*$  exists such that  $(A, P, j^*)$  satisfies consensus, then we would like to adjust the rents  $P$  in a fair way such that the resulting solution  $(A, P', j^*)$  satisfies consensus. One way to do this is by negotiating a “fair” compromise by conducting fair negotiations that change the prices  $P$ . An example of when negotiating prices may be useful is when apartment  $j$  is player 1’s favorite apartment and apartment  $j'$  is player 2’s favorite apartment. Then a negotiation between these two players and apartments could help balance the prices and either convince player 1 to rent apartment  $j'$  or convince player 2 to rent apartment  $j$ , therefore bringing the partial solution closer to consensus.

We formalize this notion of negotiating as follows. Suppose we have a partial solution  $(A, P)$  and a negotiation tuple  $\tau = (\delta, i_1, i_2, j_1, j_2)$ , where  $\delta > 0$ ,  $i_1, i_2 \in [n]$ ,  $j_1, j_2 \in [m]$ . Then a negotiation will consist of player  $i_1$  increasing their rent in apartment  $j_1$  by  $\delta$  and decreasing their rent in apartment  $j_2$  by  $\delta$ , while player  $i_2$  decreases their rent in apartment  $j_1$  by  $\delta$  and increases their rent in apartment  $j_2$  by  $\delta$ . Importantly, this negotiation only affects the players’ rents and does not change the assignment  $A$ . Formally, the partial solution after making negotiation  $\tau$  is  $(A, P')$ , where  $P'(A_{j_1}(i_1)) = P(A_{j_1}(i_1)) + \delta$ ,  $P'(A_{j_2}(i_1)) = P(A_{j_2}(i_1)) - \delta$ ,  $P'(A_{j_1}(i_2)) = P(A_{j_1}(i_2)) - \delta$ ,  $P'(A_{j_2}(i_2)) = P(A_{j_2}(i_2)) + \delta$ , and for all other  $(i, j)$  pairs,  $P'(A_j(i)) = P(A_j(i))$ . Note that each negotiation involves a player increasing rent in one apartment by  $\delta$  and decreasing rent in another apartment by  $\delta$ , which enforces that the negotiations are fair. We define the partial solution  $(A, P)$  as *reachable by negotiation* from the partial solution  $(A, Q)$  if there exists a series of  $T$  negotiations  $\{(\delta^t, i_1^t, i_2^t, j_1^t, j_2^t)\}_{t=1}^T$  such that after making all  $T$  negotiations starting from  $(A, Q)$ , the resulting partial solution is  $(A, P)$ . Returning to our concept of fairness, we want to consider partial solutions  $(A, P)$  that are reachable by this form of fair negotiations from some individually envy-free starting partial solution  $(A, Q)$ . By construction of negotiations, the total utility of any player across all of their assigned rooms is the same in the final partial solution  $(A, P)$  as in the individually envy-free starting partial solution  $(A, Q)$ .

**Lemma 4.3.** *A partial solution  $(A, P)$  is reachable by negotiation from an individual envy-free starting partial solution  $(A, Q)$  if and only if there exists an individually envy-free solution  $(A, Q)$  such that for every player  $i$ ,*

$$\sum_{j=1}^m P(A_j(i)) = \sum_{j=1}^m Q(A_j(i)).$$

*Proof sketch.* The “only if” direction follows from the fact that each negotiation does not change any player’s total rent for all  $m$  of their assigned rooms. Therefore, if there exists a sequence of negotiations that reach  $(A, P)$  starting from  $(A, Q)$ , then every player has the same total rent for their  $m$  assigned rooms in  $(A, P)$  and  $(A, Q)$ .

The “if” direction requires a more technical construction, so we provide a brief proof sketch here and the formal proof to Lemma B.1 in Appendix B.2. Suppose we start with an assignment  $A$  and price matrices  $Q$  and  $P$  such that  $\sum_{j=1}^m P(A_j(i)) = \sum_{j=1}^m Q(A_j(i))$ . We want to construct a series of negotiations to transform  $Q$  into  $P$ . We do this for each apartment 1 through  $m$ , one at a time. First, we construct a series of negotiations involving only apartments 1 and  $m$  such that the resulting price matrix  $Q_1$  has the same prices as  $P$  for apartment 1. This is possible because the two price matrices have the same total rent within each apartment. We repeat this process for apartments  $j \in \{2, \dots, m-1\}$  by constructing negotiations between apartment  $j$  and apartment  $m$  such that the resulting price matrix  $Q_j$  has the same prices as  $P$  for apartments 1 through  $j$ . We then show that the final set of negotiations between apartment  $m-1$  and  $m$  results in a post-negotiation price matrix equal to the desired price matrix  $P$ . Therefore,  $(A, P)$  is reachable by negotiation from  $(A, Q)$ .  $\square$

We now formally present the notion of fairness that comes from starting at an individually envy-free solution and conducting a sequence of negotiations until reaching consensus.

**Definition 4.4.** A solution  $(A, P, j^*)$  satisfies **negotiated envy-freeness** if  $(A, P, j^*)$  satisfies consensus and there exists a price matrix  $Q$  such that  $(A, Q)$  is individually envy-free and for every player  $i$ ,

$$\sum_{j=1}^m P(A_j(i)) = \sum_{j=1}^m Q(A_j(i)).$$

By Lemma 4.3, a solution that satisfies negotiated envy-freeness also satisfies that  $(A, P)$  is reachable by negotiations from an individually envy-free partial solution  $(A, Q)$ .

Note that the final chosen apartment and prices in a solution satisfying negotiated envy-freeness do not have a meaningful fairness interpretation in isolation; that is, without the context of the original problem. Fundamentally, negotiated envy-freeness should be viewed as constructing a fair compromise when no single solution exists that every player prefers. The fairness of a compromise cannot (and should not) be evaluated without considering the full context, because all players must give something up in order to reach a compromise. For example, consider two friends A and B committed to renting an apartment together. To convince B to rent an apartment closer to A’s workplace, A might offer the larger bedroom to B while paying equal rent, even though A prefers the larger bedroom. To an outside observer, it seems unfair that A pays equal rent but gets the

smaller bedroom. However, in the context of the original decision, this was a natural compromise. Similarly, the fairness of the negotiated envy-free solution cannot be evaluated by the chosen apartment only, but also needs to account for the original set of apartments.

For the rest of this section, we will study solutions  $(A, P, j^*)$  that satisfy negotiated envy-freeness, and argue that this is a good generalization of single apartment envy-freeness.

## 4.2 Properties of Negotiated Envy-Freeness

We will first show that a solution which satisfies negotiated envy-freeness also satisfies several desirable properties. In particular, such a solution satisfies Pareto optimality and individual rationality, and reduces to an envy-free solution in the single-apartment setting. Proofs of the below can be found in Appendix B.

**Property 4.5** (Pareto optimality). *A solution  $(A, P, j^*)$  which satisfies negotiated envy-freeness also satisfies Pareto optimality, in that there exists no other solution  $(A', P', j')$  such that  $U_i(A'_j(i), P') \geq U_i(A_{j^*}(i), P) \quad \forall i$  and the inequality is strict for at least one player.*

**Property 4.6** (Individual rationality). *A solution  $(A, P, j^*)$  which satisfies negotiated envy-freeness also satisfies individual rationality, in that all players have non-negative utility for their assigned rooms in the chosen apartment  $j^*$ .*

**Property 4.7** (Reduces to single-apartment setting). *In the single-apartment setting (when  $m = 1$ ), a solution  $(A, P, j^*)$  satisfies negotiated envy-freeness if and only if  $(A, P, j^*)$  is envy-free within the single apartment.*

## 4.3 Existence of Negotiated Envy-Freeness

We have shown that solutions which satisfy negotiated envy-freeness have both an intuitive explanation based on fair negotiating and desirable properties including Pareto optimality and individual rationality. Crucially, and in contrast to universal envy-freeness, it is also always possible to find a solution which satisfies negotiated envy-freeness.

**Theorem 4.8.** *There exists a solution which satisfies negotiated envy-freeness for every multi-apartment rent division instance.*

*Proof sketch.* We provide a brief sketch of the proof of Theorem 4.8 and defer the formal proof to Appendix B. To prove this result, we will construct a solution  $(A, P^*, j^*)$  that satisfies negotiated envy-freeness for any instance of the problem. To do this, we first start with a partial solution  $(A, Q)$  that is individually envy-free within each apartment. Such a solution is guaranteed to exist. We then consider the solution  $(A, P)$ , where  $P$  is the price matrix that, for each apartment, induces equal utilities for all players in that apartment.  $(A, P)$  satisfies consensus, as the apartment which gives the highest utility for any player will be a consensus apartment. While  $(A, P)$  may not satisfy negotiated envy-freeness, we can redistribute the prices evenly in  $(A, P)$  through negotiating to get a new solution  $(A, P^*)$  that satisfies the following two properties. First, the utility of a given player for

each apartment in  $(A, P^*)$  will be the same as in  $(A, P)$ , except that the utility may be scaled up or down by the same additive factor for all apartments. Second, we have  $\sum_{j=1}^m P^*(A_j(i)) = \sum_{j=1}^m Q(A_j(i))$ . By this construction, we can conclude by showing that  $(A, P^*)$  is a valid price assignment and satisfies negotiated envy-freeness and consensus.  $\square$

## 4.4 Polynomial-Time Optimization

We have shown that there always exists a solution which satisfies negotiated envy-freeness, and the proof provides a constructive way to find such a solution in polynomial time. However, there may be many solutions which satisfy negotiated envy-freeness. In the single apartment setting, there exists a polynomial-time algorithm that optimizes a linear objective function over all envy-free solutions in polynomial time (Gal et al. 2017). This raises the question of whether it is also possible to find a solution that optimizes a linear objective function in polynomial time over all negotiated envy-free solutions. As in the single-apartment setting, objective functions of special interest include maximin, which maximizes the utility of the least happy player, and equitability, which minimizes the disparity between players' utilities. Our main theorem of this section, Theorem 4.9, generalizes Theorem 3.1 of Gal et al. (2017) to the multi-apartment setting with negotiated envy-freeness.

**Theorem 4.9.** *Let  $f_1, \dots, f_t : \mathbb{R}^{n \times m} \rightarrow \mathbb{R}$  be linear functions, where  $t$  is polynomial in  $n$  and  $m$ . Given a multi-apartment rent division instance, a solution  $(A, P, j^*)$  that maximizes the minimum of  $f_q(U_1(A_{j^*}, P), \dots, U_n(A_{j^*}, P))$  over all  $q \in [t]$  subject to negotiated envy-freeness can be computed in time polynomial in both  $n$  and  $m$ .*

Under this formalization, the maximin objective function can be represented by the linear functions  $f_i(U_1(A_{j^*}, P), \dots, U_n(A_{j^*}, P)) = U_i(A_{j^*}, P)$  for all  $i \in [n]$ , and equitability can be represented by  $f_{i,i'}(U_1(A_{j^*}, P), \dots, U_n(A_{j^*}, P)) = U_i(A_{j^*}, P) - U_{i'}(A_{j^*}, P)$  for all  $i, i' \in [n]$ .

While we defer the proof of Theorem 4.9 to Appendix B.8, we will state the key lemma used in the proof (Lemma 4.10) and the algorithm that achieves the result.

Informally, Lemma 4.10 states that the set of solutions which satisfy negotiated envy-freeness are equivalent for all welfare-maximizing assignments, in the sense that the same set of player utilities can always be found for any choice of welfare-maximizing assignment. Note that Lemma 4.10 has a similar flavor to the 2nd Welfare Theorem (Gal et al. 2017; Mas-Colell, Whinston, and Green 1995).

**Lemma 4.10.** *Let  $A, A'$  be two assignments that maximize welfare in every apartment, and let  $P$  be a price matrix such that  $(A, P, j^*)$  satisfies negotiated envy-freeness. Then there exists a price matrix  $P'$  such that  $(A', P', j^*)$  satisfies negotiated envy-freeness and  $U_i(A_j, P) = U_i(A'_j, P')$  for all  $i, j$ .*

---

Algorithm 1: Optimizing an objective function subject to Negotiated Envy-Freeness

---

```

for  $j \leftarrow 1$  to  $m$  do
     $A_j \leftarrow$  welfare-maximizing assignment in  $j$  via max-
    weight bipartite matching
end for
 $A \leftarrow \{A_1, \dots, A_m\}$ 
Choose  $j_A^*$  such that  $\sum_{i=1}^n V_i(A_{j_A^*}(i)) - R_{j_A^*} \geq$ 
 $\max_j \sum_{i=1}^n V_i(A_j(i)) - R_j$ 
 $P^*, Q^* \leftarrow \arg \max_{P, Q} Z$ 

s.t.  $Z \leq f_q(U_1(A_{j'}, P), \dots, U_n(A_{j'}, P))$ 
 $V_i(A_{j'}(i)) - P(A_{j'}(i)) \geq V_i(A_j(i)) - P(A_j(i))$ 
 $V_i(A_j(i)) - Q(A_j(i)) \geq V_i(A_{j'}(i)) - Q(A_{j'}(i))$ 
 $\sum_j P(A_j(i)) = \sum_j Q(A_j(i))$ 
 $\sum_i P(A_j(i)) = R_j$ 

return  $(A, P^*, j^*)$ 

```

---

The proof of Lemma 4.10 can be found in Appendix B.7. The algorithm used to prove Theorem 4.9 is the following.

#### 4.5 Strong Negotiated Envy-freeness

We have so far defined two notions of fairness for multi-apartment rent division: universal envy-freeness and negotiated envy-freeness. Universal envy-freeness is simpler to interpret, but may not always exist. On the other hand, negotiated envy-freeness always exists, but there may be envy in the chosen apartment. In this section, we briefly describe an extension of negotiated envy-freeness (strong negotiated envy-freeness) that provides an interpretation for envy within the chosen apartment. Due to space constraints we defer the formal definition and proofs to Appendix G.

The definition of strong negotiated envy-freeness is similar to negotiated envy-freeness in that the solution  $(A, P, j^*)$  must satisfy consensus and there must be an individually envy-free solution  $(A, Q)$  that is reachable by negotiations. However, strong negotiated envy-freeness has an additional requirement that bounds the price differences between  $P$  and  $Q$  in the consensus apartment  $j^*$ . Due to this additional requirement, any envy in the final chosen apartment can be interpreted as being “necessary” to reach consensus. We also show that a solution satisfying strong negotiated envy-freeness always exists and that optimizing an objective subject to strong negotiated envy-freeness can be done in polynomial time (as in Theorem 4.9).

## 5 Discussion

### 5.1 Limitations

Our model suffers the same limitations as that of Gal et al. (2017), including that envy-freeness is not strategy-proof

and that the quasi-linear utility model is a strong simplifying assumption. We also acknowledge that negotiated envy-freeness, despite having a reasonable justification as reachable by negotiations, is less easily explainable as a fairness notion than envy-freeness, and that envy within the consensus apartment may still lead to discontent. From an application standpoint, it may be useful to let users find a negotiated envy-free solution themselves, by presenting users with individually envy-free solutions in each apartment and enabling the group to negotiate amongst themselves in the way dictated in the paper.

### 5.2 Extensions

We have required throughout this work that solutions to the multi-apartment rent division problem are of the form  $(A, P, j^*)$ , where  $j^*$  is the chosen apartment. A natural generalization would be a distributional solution of the form  $(A, P, D)$ , where  $D$  is a distribution over all  $m$  apartments. We would then want to find a distributional solution  $(A, P, D)$  such that in expectation, every player prefers their assignment to any other player’s assignment. This notion of *distributional envy-freeness* is a weaker notion of fairness than universal envy-freeness, and more instances of the problem have a distributional envy-free solution than a universal envy-free solution (including Example 1.1). See Appendix C for more details on distributional envy-freeness.

In our work, we assumed that each group of  $n$  players is determined to room together in an apartment with  $n$  bedrooms. One extension is to have various sizes of apartments and allow the group of players to be split into multiple smaller apartments. This setting can be modeled as a cooperative game with transferable utility, and a natural question is then whether the core is always non-empty. In Appendix E, we show that even when there are infinite copies of each apartment, there still exist instances where the core is empty.

Another interesting question is how the solution changes as additional apartments are added (e.g. appear on the market). An additional apartment could change the set of negotiated envy-free solutions, which could in turn change the chosen apartment when optimizing objective functions such as maximin or equitability. We show in Appendix D that the maximin solution under negotiated envy-freeness does not satisfy apartment monotonicity, in that an additional apartment could either raise or lower the achievable maximin value over all negotiated envy-free solutions.

In Section 3, we studied the existence of universal envy-free solutions when the values are drawn i.i.d. at random. In Appendix F, we study the probability of the existence of a universal envy-free solution when there are two players with binary valuations. As the correlation between player values decreases, the probability of the existence of a universal envy-free solution generally increases. Interestingly, however, the increase is not monotonic. Therefore, sometimes the highest likelihood that there does not exist a universal envy-free solution occurs with correlation strictly between  $-1$  and  $1$ , implying that correlation between players is not strictly better for finding a universal envy-free solution.

## Acknowledgements

Procaccia gratefully acknowledges research support by the National Science Foundation under grants IIS-2147187, IIS-2229881 and CCF-2007080; and by the Office of Naval Research under grant N00014-20-1-2488. Schiffer was supported by an NSF Graduate Research Fellowship. Zhang was supported by an NSF Graduate Research Fellowship.

## References

- Airiau, S.; Gilbert, H.; Grandi, U.; Lang, J.; and Wilczynski, A. 2023. Fair Rent Division on a Budget Revisited. In *Proceedings of the 26th European Conference on Artificial Intelligence (ECAI)*, 52–59.
- 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.
- Arunachaleswaran, E. R.; Barman, S.; and Rathi, N. 2021. Polynomial-Time Approximation Schemes for Fair Rent Division. *Mathematics of Operations Research*, 47(3): 1970–1998.
- Dickerson, J. P.; Goldman, J.; Karp, J.; Procaccia, A. D.; and Sandholm, T. 2014. The Computational Rise and Fall of Fairness. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, 1405–1411.
- 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.
- Goldman, J.; and Procaccia, A. D. 2014. Spliddit: Unleashing Fair Division Algorithms. *SIGecom Exchanges*, 13(2): 41–46.
- Manurangsi, P.; and Suksompong, W. 2017. Asymptotic Existence of Fair Divisions for Groups. *Mathematical Social Sciences*, 89: 100–108.
- Manurangsi, P.; and Suksompong, W. 2020. When Do Envy-Free Allocations Exist? *SIAM Journal on Discrete Mathematics*, 34(3): 1505–1521.
- Mas-Colell, A.; Whinston, M. D.; and Green, J. R. 1995. *Microeconomic Theory*, volume 1. Oxford University Press.
- Peters, D.; Procaccia, A. D.; and Zhu, D. 2022. Robust Rent Division. In *Proceedings of the 36th Annual Conference on Neural Information Processing Systems (NeurIPS)*.
- Procaccia, A. D. 2019. Axioms Should Explain Solutions. In Laslier, J.-F.; Moulin, H.; Sanver, R.; and Zwicker, W. S., eds., *The Future of Economic Design*. Springer.
- Procaccia, A. D.; Velez, R. A.; and Yu, D. 2018. Rent Division on a Budget. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence (AAAI)*, 1177–1184.
- Segal-Halevi, H. 2022. Generalized Rental Harmony. *American Mathematical Monthly*, 129(5): 403–414.
- 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*, 6: 1–25.
- Velez, R. A. 2022. A Polynomial Algorithm for Maxmin and Minmax Envy-Free Rent Division on a Soft Budget. *Social Choice and Welfare*, 59: 93–118.
- Velez, R. A. 2023. Equitable Rent Division on a Soft Budget. *Games and Economic Behavior*, 139: 1–14.