

Fair Rent Division on a Budget

Ariel D. Procaccia

Computer Science Department
Carnegie Mellon University

Rodrigo A. Velez

Department of Economics
Texas A&M University

Dingli Yu

Institute for Interdisciplinary
Information Sciences
Tsinghua University

Abstract

The standard approach to fair rent division assumes that agents have quasi-linear utilities, and seeks allocations that are envy free; it underlies an algorithm that is widely used in practice. However, this approach does not take budget constraints into account, and, therefore, may assign agents to rooms they cannot afford. By contrast, we design a polynomial-time algorithm that takes budget constraints as part of its input; it determines whether there exist envy-free allocations that satisfy the budget constraints, and, if so, computes one that optimizes an additional criterion of justice. In particular, this gives a polynomial-time implementation of the *budget-constrained maximin solution*, where the maximization objective is the minimum utility of any agent. We show that, like its non-budget-constrained counterpart, this solution is unique in terms of utilities (when it exists), and satisfies additional desirable properties.

1 Introduction

A recent story in *New Scientist*¹ nicely introduces a familiar dilemma:

Every renter knows that not all rooms are created equally. One bedroom might have an en-suite bathroom and a view of the seaside, whereas another is half the size and overlooks a dump. It's only fair that the person who gets the better room pays a larger share of the rent. The question [...] is: how much?

The ubiquity and simplicity of the rent division problem have made it a poster child for *computational fair division*. Indeed, of the five applications on the not-for-profit fair division website Spliddit.org (Goldman and Procaccia 2014), the rent calculator is the most popular — almost 30,000 rent division instances have been created by users since the website's launch in November 2014.

The standard rent division setup involves n agents, who share an apartment with n rooms. Every agent has a value for each room. An allocation matches the agents with rooms, and assigns a price to each room, such that the prices add up to the total rent. We make the standard assumption that the utility functions of the agents are *quasi-linear*: the utility of

an agent for getting a room he values at x for the price of y is $x - y$. The goal is to find an allocation that is *envy free*: the utility of each agent for his room at its price should be at least as high as his utility for any other room, at the price of that room.

It is well known that an envy-free allocation always exists (Svensson 1983), and one can be computed in polynomial time (Aragones 1995). In fact, there is an embarrassment of riches, as there are typically many envy-free allocations, which necessitates selecting between them. A natural choice is the *maximin solution* of Alkan et al. (1991), which selects the allocation that maximizes the minimum utility of any agent, subject to the envy-freeness constraint. A linear-programming-based algorithm of Gal et al. (2017) can compute the maximin solution in polynomial time; this algorithm has been deployed on Spliddit since May 2015.

One advantage of the foregoing approach (especially the assumption of quasi-linear utilities) is that preference elicitation is straightforward: participants only need to report their value for each room. In many ways, it does hit a sweet spot between expressiveness and ease of elicitation. But it certainly does not address the needs of some users, and its most prominent shortcoming is ignoring *budget constraints*. This is reflected in feedback from Spliddit users — the following representative example was submitted in December 2016:

I've tried to use your fair calculator, but it did not work in my case. In our situation, I am the guy with the most tight [sic] budget. Unfortunately, your system does not take into account the 'maximum' budget restrictions. I was assigned an option that was too expensive for me, so it did not help. Please advise if there is a way to use the system, taking that kind of limitation into account.

Unfortunately, when agents are budget-constrained, an envy-free allocation is no longer guaranteed to exist. For example, suppose there are two agents; a total rent of \$1000; two rooms where the first is valued at \$800 by both agents, and the second is valued at \$200 by both agents; and budget constraints of \$600 for both agents. To avoid envy, the desirable room must be priced at \$800, but this is more than either agent can pay.

Our goal in this paper, therefore, is to gain a rigorous understanding of the budget-constrained rent division problem.

¹<https://www.newscientist.com/article/2137642-maths-website-stops-you-being-ripped-off-by-your-flatmates>

In particular, we would like to be able to say (algorithmically) when the set of envy-free allocations is nonempty under budget constraints, and, when it is, select among these allocations in a principled way.

1.1 Our Results

In Section 3, we construct a polynomial-time (LP-based) algorithm that computes an optimal allocation with respect to a given (linear) criterion of justice, subject to the envy-freeness constraints and the given budget constraints, when a feasible allocation exists. This algorithm extends that of Gal et al. (2017) to the budget-constrained setting, but, for reasons that we explain in detail, its design and analysis are significantly more involved.

The obvious next question is: What should we use as the optimization objective? The foregoing algorithm can be used to compute the *budget-constrained maximin solution* – which maximizes the minimum utility of any agent subject to envy-freeness and the budget constraints – when a feasible solution exists. It can also compute the budget-constrained version of the *equitable solution* (Gal et al. 2017), where the objective is to minimize the maximum difference in utilities between any two agents. We show that the properties that make the maximin solution attractive in the non-budget-constrained setting extend to the budget constrained setting: it is unique in terms of utilities, and it dominates the equitable solution, in the sense that any budget-constrained maximin solution is a budget-constrained equitable solution, but not vice versa. Finally, we extend our algorithms in order to compute, in polynomial time, an optimal allocation, subject to the envy-freeness constraint and the minimization of the maximum budget violation among agents — this recommendation exists for all possible instances.

1.2 Significance in AI

Work on computational fair division has been featured prominently in recent AI conferences (Bliem, Bredereck, and Niedermeier 2016; Kurokawa, Procaccia, and Wang 2016; Brânzei et al. 2016; Aziz et al. 2017; Bei, Qiao, and Zhang 2017; Bouveret et al. 2017; Menon and Larson 2017; Aleksandrov and Walsh 2017). The surge in excitement is driven in part by a newfound understanding that, beyond potential applications to multiagent systems, fair division algorithms can play a role in helping groups of people make decisions on a large scale. Our paper contributes to this line of work in that it is driven by real-world needs, and provides practicable solutions that are likely to be deployed.

2 The Model

We study the allocation of a set of n objects A of indivisible goods, that we refer to as *rooms*, among a set of n agents $N \equiv \{1, \dots, n\}$. Generic rooms are denoted by letters such as a, b . Each agent is to receive a room and pay an amount of money for it. Agents have quasi-linear preferences on their allotments, i.e., agent i 's utility from receiving room a and paying p_a for it is $v_{ia} - p_a$. The profile of agents values is $v \equiv (v_{ia})_{i \in N, a \in A}$. We assume that individual payments

should add up to a fixed amount that we refer to as the *house rent* r . An *economy* is then a tuple $e \equiv (N, A, v, r)$.

An *allocation* for $e \equiv (N, A, v, r)$ is a pair (σ, p) , where the *assignment* $\sigma : N \rightarrow A$ is a bijection, and the vector of *prices* $p \equiv (p_a)_{a \in A}$ is such that $\sum_{a \in A} p_a = r$. That is, under the allocation (σ, p) , agent i is assigned the room $\sigma(i)$ for the price $p_{\sigma(i)}$.

An allocation is *envy free* if no agent prefers the consumption of any other agent at the allocation, i.e., for all $i, j \in N$,

$$v_{i\sigma(i)} - p_{\sigma(i)} \geq v_{i\sigma(j)} - p_{\sigma(j)}.$$

The set of envy-free allocations for the economy e is $F(e)$.

We assume that agents may be budget constrained. For each i , there is $b_i \in \mathbb{R} \cup \{+\infty\}$ which is the maximum amount that the agent can pay for any room. The vector of budget constraints is $b \equiv (b_i)_{i \in N}$. An allocation (σ, p) satisfies budget constraints b if for each $i \in N$, $p_{\sigma(i)} \leq b_i$. For a vector of budget constraints b , the set of envy-free budget-constrained allocations is $F_b(e)$.

It is well known that for each e , $F(e)$ is a non-empty compact set. Moreover, one can calculate directly from an agent's utility function an upper bound for his consumptions in each allocation in $F(e)$. Thus, we can assume without loss of generality that $b \in \mathbb{R}^n$.

3 A Polynomial-Time Algorithm

For a given economy, $e \equiv (N, A, v, r)$ and a vector of budget constraints b , we are primarily interested in determining whether $F_b(e) \neq \emptyset$. If so, we are interested in selecting an element of $F_b(e)$ that optimizes an additional criterion of justice. The rest of the section is devoted to understanding and proving the following theorem, which is our main result:

Theorem 1. *Let $e \equiv (N, A, v, r)$, $b \equiv (b_i)_{i \in N}$, and, for $s \in [t]$, where t is polynomial in n , let $f_s : \mathbb{R}^n \rightarrow \mathbb{R}$ be a linear function. Then there is a polynomial-time algorithm that determines whether $F_b(e) \neq \emptyset$, and, if so, outputs an allocation $(\sigma, p) \in F_b(e)$ that maximizes*

$$\min_{s \in [t]} f_s(v_{1\sigma(1)} - p_{\sigma(1)}, \dots, v_{n\sigma(n)} - p_{\sigma(n)})$$

over all allocations in $F_b(e)$.

Gal et al. (2017) pay special attention to two particular optimization objectives: maximizing the minimum utility of any agent, which is obtained by simply letting

$$f_i(v_{1\sigma(1)} - p_{\sigma(1)}, \dots, v_{n\sigma(n)} - p_{\sigma(n)}) = v_{i\sigma(i)} - p_{\sigma(i)}$$

for all $i \in N$; and minimizing the *social disparity* — defined as the maximum difference in utilities between any two agents — which can also be easily captured using the same formalism. We discuss the relation between these two objectives in Section 4; for now we focus on the algorithmic question.

Our starting point is a simple algorithm of Gal et al. (2017), which achieves the equivalent result to Theorem 1 without budget constraints, i.e., it finds the optimal allocation in $F(e)$ (which is always nonempty, under their

assumptions on valuations). However, we provide the requisite background from a different viewpoint, in order to be consistent with the subsequent presentation of our own algorithm.

For $e \equiv (N, A, v, r)$, an assignment σ is *efficient* if it maximizes the sum of agents values, i.e., $\sum_{i \in N} v_{i\sigma(i)}$, among all assignments. Basic LP theory shows that the dual variables in the problem determining an efficient assignment provide an envy-free allocation (Alkan, Demange, and Gale 1991, Duality Theorem). This implies that each efficient assignment can be completed into an envy-free allocation. Conversely, an assignment in an envy-free allocation is always efficient (Svensson 1983). Two additional facts solve the puzzle. For any two allocations $\{(\sigma, p), (\mu, q)\} \subseteq F(e)$, both (μ, p) and (σ, q) are also in $F(e)$ (Svensson 2009).² Moreover, each agent is indifferent between (σ, p) and (μ, p) (this follows from Lemma 1 below). Thus, as observed by Gal et al. (2017), it is sufficient to calculate an efficient assignment σ (which can be done in polynomial time by reducing the problem to maximum weight matching in a bipartite graph), and then solving the following LP, which uses the same notations as Theorem 1:

$$\begin{aligned} & \max_{x, p_{\sigma(1)}, \dots, p_{\sigma(n)}} x \\ & \text{s.t.} \quad \sum_{i \in N} p_{\sigma(i)} = r \\ & \quad \forall s \in [t], \quad x \leq f_s(v_{1\sigma(1)} - p_{\sigma(1)}, \dots, v_{n\sigma(n)} - p_{\sigma(n)}) \\ & \quad \forall i, j \in N, \quad v_{i\sigma(i)} - p_{\sigma(i)} \geq v_{i\sigma(j)} - p_{\sigma(j)} \end{aligned}$$

To be clear, the second constraint guarantees that x is set to the minimum of the given linear functions, and the third ensures that the allocation (σ, p) is envy free.

It is tempting think that one can adapt the algorithm of Gal et al. (2017) to our problem by simply adding the budget constraints to the above LP:

$$\begin{aligned} & \max_{x, p_{\sigma(1)}, \dots, p_{\sigma(n)}} x \\ & \text{s.t.} \quad \sum_{i \in N} p_{\sigma(i)} = r \\ & \quad \forall s \in [t], \quad x \leq f_s(v_{1\sigma(1)} - p_{\sigma(1)}, \dots, v_{n\sigma(n)} - p_{\sigma(n)}) \\ & \quad \forall i, j \in N, \quad v_{i\sigma(i)} - p_{\sigma(i)} \geq v_{i\sigma(j)} - p_{\sigma(j)} \\ & \quad \forall i \in N, \quad p_{\sigma(i)} \leq b_i \end{aligned} \tag{1}$$

However, the basic premise behind their algorithm — that one can start from any efficient assignment — is false in our setting, as shown by the following example.

Example 1. Let $N = \{1, 2\}$, $A = \{a, b\}$, and $r = 1$. For $i = 1, 2$, $v_{ia=1}$ and $v_{ib=0}$. Moreover, $b_1 = 1$ and $b_2 = 0$. The allocation that assigns room a to agent 1 for $p_a = 1$, and room b to agent 2 for $p_b = 0$, is envy free and satisfies the budget constraints. However, if we start from assigning room a to agent 2, and room b to agent 1 — which is also efficient — then there is no vector of prices that leads to an

²We provide a simple proof of this fact embedded in our proof of Lemma 2.

envy-free allocation satisfying the budget constraints: Due to the budget constraint of agent 2, we must have $p_{\sigma(2)} = 0$, even though he is now assigned the valuable room, so agent 1 must be envious.

Example 1 suggests that we need to develop a new algorithmic approach in order to prove Theorem 1. We start with several lemmas. The first (Lemma 1) is well known. We present a proof of Lemma 3 for completeness. The rest are new.

Lemma 1 (Alkan, Demange, and Gale 1991). *Let $\{(\sigma, p), (\mu, q)\} \subseteq F(e)$. Then σ and μ are both bijections between the following sets:*

$$\begin{aligned} & \{i \in N : v_{i\sigma(i)} - p_{\sigma(i)} > v_{i\mu(i)} - q_{\mu(i)}\} \text{ and} \\ & \{a \in A : p_a < q_a\}, \\ & \{i \in N : v_{i\sigma(i)} - p_{\sigma(i)} = v_{i\mu(i)} - q_{\mu(i)}\} \text{ and} \\ & \{a \in A : p_a = q_a\}, \\ & \{i \in N : v_{i\sigma(i)} - p_{\sigma(i)} < v_{i\mu(i)} - q_{\mu(i)}\} \text{ and} \\ & \{a \in A : p_a > q_a\}. \end{aligned}$$

Before proceeding, we introduce a concept that has played a role in several previous papers. For each $(\sigma, p) \in F(e)$ define the *weak envy graph* by $\Gamma(\sigma, p) \equiv (N, E)$, where $(i, j) \in E$ if and only if $v_{i\sigma(i)} - p_{\sigma(i)} = v_{i\sigma(j)} - p_{\sigma(j)}$. That is, it is the graph where the vertices are agents, and there is an edge between i and j if i weakly envies j under (σ, p) . Note that i cannot strictly envy j because the allocation is envy free.

Lemma 2. *Let $\{(\sigma, p), (\mu, q)\} \subseteq F(e)$. for all $i \in N$, $j = \sigma^{-1}(\mu(i))$, i and j are strongly connected in both $\Gamma(\sigma, p)$ and $\Gamma(\mu, p)$.*

Proof. Since $\{(\sigma, p), (\mu, q)\} \subseteq F(e)$, $(\mu, p) \in F(e)$.³ By Lemma 1 applied to (σ, p) and (μ, p) , for each $i \in N$,

$$v_{i\sigma(i)} - p_{\sigma(i)} = v_{i\mu(i)} - p_{\mu(i)},$$

i.e., $(i, \sigma^{-1}(\mu(i)))$ is an edge of $\Gamma(\sigma, p)$. Let $\nu = \sigma^{-1} \circ \mu$, then $i \rightarrow \nu(i) = j \rightarrow \nu^2(i) \rightarrow \dots \rightarrow i$ is a cycle in $\Gamma(\sigma, p)$, because ν is a permutation on N and can be decomposed into cycles. Similarly, we have $(\nu(i), i)$ is an edge of $\Gamma(\mu, p)$ for all $i \in N$, and hence $i, \nu(i) = j, \nu^2(i), \nu^3(i), \dots$ are also strongly connected in $\Gamma(\mu, p)$. \square

Lemma 3 (Andersson, Ehlers, and Svensson 2014). *The set of strongly connected components of $\Gamma(\sigma, p)$ is invariant for each allocation in $F(e)$.*

Proof. Suppose that $\{(\sigma, p), (\mu, q)\} \subseteq F(e)$. Then, $(\sigma, q) \in F(e)$ (Svensson 2009). Suppose without loss of generality that $1, \dots, K$ are the vertices of cycle $\{(1, 2), \dots, (K, 1)\} \subseteq \Gamma(\sigma, p)$. Since $(\sigma, q) \in F(e)$, for each $k = 1, \dots, K$, $v_{k\sigma(k)} -$

³This was proved by Svensson (2009), and also follows from the Second Welfare Theorem, as noted by Gal et al. (2017). The following is a simple proof: Suppose that $\{(\sigma, p), (\mu, q)\} \subseteq F(e)$. Since $(\mu, q) \in F(e)$, for each $i \in N$, $v_{i\sigma(i)} - q_{\sigma(i)} \leq v_{i\mu(i)} - q_{\mu(i)}$. Adding up these n inequalities we find expressions that are equal because both σ and μ are efficient. Thus, for each $i \in N$, $v_{i\sigma(i)} - q_{\sigma(i)} = v_{i\mu(i)} - q_{\mu(i)}$, and it follows that $(\sigma, q) \in F(e)$.

$q_{\sigma(k)} \geq v_{k\sigma(k+1)} - q_{\sigma(k+1)}$ where $K + 1 = 1$. Adding up these K inequalities we find expressions that are equal because both σ and the assignment obtained by reshuffling along the cycle are efficient. Thus, for each for each $k = 1, \dots, K$, $v_{k\sigma(k)} - q_{\sigma(k)} = v_{k\sigma(k+1)} - q_{\sigma(k+1)}$. It follows that the set of strongly connected components of $\Gamma(\sigma, p)$ is contained in that of $\Gamma(\sigma, q)$.

Let (i, j) be an edge of $\Gamma(\sigma, q)$. Let $k \in N$ be such that $\mu(k) = \sigma(j)$, and, therefore,

$$v_{i\sigma(i)} + q_{\sigma(i)} = v_{i\mu(k)} + q_{\mu(k)}.$$

By Lemma 1, and using $\{(\mu, q), (\sigma, q)\} \subseteq F(e)$,

$$v_{i\mu(i)} + q_{\mu(i)} = v_{i\sigma(i)} + q_{\sigma(i)}.$$

Combining the two equalities, we get

$$v_{i\mu(i)} + q_{\mu(i)} = v_{i\mu(k)} + q_{\mu(k)},$$

that is, (i, k) is an edge of $\Gamma(\mu, q)$. Now, by Lemma 2, since $j = \sigma^{-1}(\mu(k))$, j and k are strongly connected in $\Gamma(\mu, q)$. We conclude that there is a path from i to j in $\Gamma(\mu, q)$. Since this is true for every edge (i, j) of $\Gamma(\sigma, q)$, the set of strongly connected components of $\Gamma(\sigma, q)$ is contained in that of $\Gamma(\mu, q)$. \square

Lemma 3 allows us to define the set of strongly connected components of the envy-free graph associated with a triple (N, A, v) as $\Gamma(\sigma, p)$ for some $r \in \mathbb{R}$ and $(\sigma, p) \in F(N, A, v, r)$. Let $\mathcal{C}(N, A, v)$ be the partition of N into strongly connected components of the envy-free graph in the set of economies $\{(N, A, v, r) : r \in \mathbb{R}\}$.

For technical reasons which will become clear momentarily, we are especially interested in the class of economies where all agents are strongly connected. The following lemma states some useful properties of envy-free allocations in these economies.

Lemma 4. *Let N, A, v , and suppose that $\mathcal{C}(N, A, v) = \{N\}$. Then for all $r \in \mathbb{R}$:*

1. *If $\{(\sigma, p), (\mu, q)\} \subseteq F(N, A, v, r)$, then $p = q$.*
2. *Each agent is indifferent among all allocations in $F(N, A, v, r)$.*
3. *Let $(\sigma, p) \in F(N, A, v, r)$. For any $s \in \mathbb{R}$, $(\mu, q) \in F(N, A, v, s)$, and $a \in A$, $q_a = p_a + (s - r)/n$.*
4. *The rent of each room in any element of $F(N, A, v, r)$ is an increasing function of r .*

Proof. We first prove Statement 1. Suppose (i, j) is an edge of $\Gamma(\sigma, p)$ and $p_{\sigma(j)} \geq q_{\sigma(j)}$. Then

$$v_{i\sigma(i)} - p_{\sigma(i)} = v_{i\sigma(j)} - p_{\sigma(j)} \leq v_{i\sigma(j)} - q_{\sigma(j)} \leq v_{i\sigma(i)} - q_{\sigma(i)},$$

where the first transition follows from the former assumption, the second from the latter, and the third from $(\sigma, q) \in F(N, A, v, r)$ (Svensson 2009). Thus, $p_{\sigma(i)} \geq q_{\sigma(i)}$. Thus, if there is a path from i to j in $\Gamma(\sigma, p)$ and $p_{\sigma(j)} \geq q_{\sigma(j)}$, then $p_{\sigma(i)} \geq q_{\sigma(i)}$. Since $\sum_{a \in A} p_a = \sum_{a \in A} q_a$, there is $a \in A$ such that $p_a \geq q_a$. Thus, using the assumption that $\mathcal{C}(N, A, v) = \{N\}$, for each $a \in A$, $p_a \geq q_a$. We conclude that $p = q$.

Statement 2 follows directly from 1 and Lemma 1. Statement 3 follows from Statement 1 and the observation that, since preferences are quasi-linear, $(\mu, q) \in F(N, A, v, s)$, where q is defined by $q_a = p_a + (s - r)/n$ for all $a \in A$. Statement 4 follows directly from Statement 3. \square

Observe that the envy-free graph could have been defined by identifying vertices with rooms, not agents as we did. Doing so would also produce a partition of the set of rooms into strongly connected components that is common for all envy-free allocations in an economy. There is a one-to-one correspondence between the two approaches.

Lemma 5. *Let $e \equiv (N, A, v, r)$ and $C \in \mathcal{C}(N, A, v)$. Then, for any two efficient assignments for e , σ and μ , $\sigma(C) = \mu(C)$.*

Proof. Let $\{(\sigma, p), (\mu, q)\} \subseteq F(e)$. Let $i \in N$ and let $C \in \mathcal{C}(N, A, v)$ be such that $i \in C$. Let j be such that $\mu(j) = \sigma(i)$. Recall that $(\sigma, q) \in F(e)$ (Svensson 2009). By Lemma 2, since $i = \sigma^{-1}(\mu(j))$, i, j must be strongly connected in $\Gamma(\sigma, q)$. By Lemma 3, $j \in C$ as well. We conclude that $\sigma(C) \subseteq \mu(C)$. \square

The previous lemma implies that even though there can be many efficient assignments in an economy, there is only one efficient assignment of rooms among the strongly connected components of the envy-free graph of the economy. Given N, A, v , and $C \in \mathcal{C}(N, A, v)$, let $A(C)$ be the set of rooms received by agents in C at some envy-free allocation for e . Furthermore, let v_C and b_C be the restrictions of the two vectors to C , and let $\delta_C \in \mathbb{R}$ the maximum $s \in \mathbb{R}$ for which $F_{b_C}(C, A(C), v_C, s) \neq \emptyset$. Note that δ_C exists because utility functions are continuous.

Lemma 6. *Let $e \equiv (N, A, v, r)$ and a vector of budget constraints b . For each $C \in \mathcal{C}(N, A, v)$, let $b_C \equiv (b_i)_{i \in C}$, $v_C \equiv (v_{ia})_{i \in C, a \in A(C)}$, and suppose that $(\mu_C, p_{A(C)}) \in F_{b_C}(C, A(C), v_C, \delta_C)$. Let μ be the allocation that for each $i \in N$, $\mu(i) = \mu_C(i)$ where $C \in \mathcal{C}(N, A, v)$ and $i \in C$. Then, $F_b(e) \neq \emptyset$ if and only if LP (1) for assignment μ is feasible.*

Proof. Clearly $F_b(e) \neq \emptyset$ if LP (1) is feasible. In the other direction, suppose that $(\sigma, q) \in F_b(e)$. We claim that $(\mu, q) \in F_b(e)$. We know that $(\mu, q) \in F(e)$ (Svensson 2009). To show that the budget constraints are satisfied, let $C \in \mathcal{C}(N, A, v)$. Then

$$s \equiv \sum_{i \in C} q_{\mu(i)} = \sum_{i \in C} q_{\sigma(i)} \leq \delta_C,$$

where the equality holds due to Lemma 5, and the inequality holds because otherwise $(\sigma_C, q_{A(C)}) \in F_{b_C}(C, v_C, s)$ with $s > \delta_C$. By Statement 4 in Lemma 4, for each $i \in C$, $q_{\mu(i)} \leq p_{\mu(i)} \leq b_i$. \square

Thus, we can restrict the solution of our problem to economies that are fully connected. Our first algorithm, given as Algorithm 1, applies to such economies. The algorithm uses a budget-aware type of weak envy graph, which we denote by $\Gamma_b(\sigma, p) = (N, E)$, where

$$E = \{(i, j) : v_{i\sigma(i)} - p_{\sigma(i)} = v_{i\sigma(j)} - p_{\sigma(j)} \text{ and } p_{\sigma(j)} < b_i\}.$$

```

1: compute  $(\sigma, p) \in F(N, A, v, r)$  for some  $r \in \mathbb{R}$ 
2: let  $\Delta \in \mathbb{R}$  such that
   
$$(\sigma, (p_a - \Delta)_{a \in A}) \in F_b(N, C, v, r - n\Delta)$$

   and there is  $i \in N$  such that  $p_{\sigma(i)} = b_i$ 
3:  $p \leftarrow (p_{\sigma(i)} - \Delta)_{i \in N}$ 
4:  $r \leftarrow r - n\Delta$ 
5: while  $\nexists i \in N$  s.t.  $p_{\sigma(i)} = b_i$  and  $i$  is not a vertex of a
   cycle of  $\Gamma_b(\sigma, p)$  do
6:   if  $\nexists i \in N$  s.t.  $p_{\sigma(i)} = b_i$  then {Case 1}
7:      $\Delta \leftarrow \min_{i \in N} (b_i - p_{\sigma(i)})$ 
8:      $p \leftarrow (p_a + \Delta)_{a \in A}$ 
9:      $r \leftarrow r + n\Delta$ 
10:  else {Case 2}
11:    find  $i \in N$  s.t.  $p_{\sigma(i)} = b_i$  and  $i$  is a vertex of a
    cycle  $C$  of  $\Gamma_b(\sigma, p)$ 
12:     $\sigma \leftarrow$  reshuffle of  $\sigma$  along  $C$ 
13:  end if
14: end while
15: return  $r$  and  $(\sigma, p)$ 

```

Algorithm 1: Maximum-rent envy-free allocation in a fully connected economy.

Lemma 7. For each N, A, v such that $\mathcal{C}(N, A, v) = \{N\}$, and $b \equiv (b_i)_{i \in N}$, Algorithm 1 terminates in polynomial time. If r and (σ, p) are the output of Algorithm 1, $(\sigma, p) \in F_b(N, A, v, r)$, and $\delta_N = r$, i.e., for each $s > r$, $F_b(N, A, v, s) = \emptyset$.

Proof. We first prove that Algorithm 1 terminates in polynomial time. The initial computation of (σ, p) in Line 1 can be done, e.g., using the algorithm of Gal et al. (2017). Case 2 (Line 10) is repeated at most n times before the algorithm either terminates or Case 1 (Line 6) holds. Case 1 holds at most n^2 times, because here an agent meets his budget constraint for a certain room, and the prices of rooms strictly increase each time Case 1 holds.

We now turn to the claims about the output of the algorithm; suppose r and (σ, p) are that output. By construction, $(\sigma, p) \in F_b(N, A, v, r)$, as this holds before the while loop, and every subsequent operation maintains envy-freeness and budget feasibility.

Next, we claim that for each $s > r$, $F_b(N, A, v, s) = \emptyset$. Suppose for the sake of contradiction that there is $s > r$ and $(\mu, q) \in F_b(N, A, v, s)$. Let $\Delta \equiv (s - r)/n$. By Statement 3 of Lemma 4, for each $a \in A$, $q_a = p_a + \Delta$. Recall that $(\mu, p) \in F(N, A, v, r)$ (Svensson 2009). Similarly to the proof of Lemma 2, by Lemma 1, we have that for all $i \in N$,

$$v_{i\sigma(i)} - p_{\sigma(i)} = v_{i\mu(i)} - p_{\mu(i)}.$$

Along with

$$p_{\mu(i)} = q_{\mu(i)} - \Delta < b_i,$$

it follows that $(i, \sigma^{-1}(\mu(i)))$ must be an edge of $\Gamma_b(\sigma, p)$ for any i . Let $\nu = \sigma^{-1} \cdot \mu$, then for all $i \in N$, $i \rightarrow \nu(i) \rightarrow \nu^2(i) \rightarrow \dots \rightarrow i$ is a cycle of $\Gamma_b(\sigma, p)$. Also notice that the while loop only terminates when there is $i \in N$ such

that $p_{\sigma(i)} = b_i$. Therefore, agent i satisfies the conditions of Case 2, and the algorithm should not have terminated. \square

We are finally ready to present our main algorithm, given as Algorithm 2, and its analysis, which completes the proof of Theorem 1.

```

1: compute  $(\sigma, p) \in F(N, A, v, r)$ 
2:  $\mathcal{C} \leftarrow$  strongly connected components of  $\Gamma(\sigma, p)$ 
3: for each  $C \in \mathcal{C}$  do
4:    $(\mu_C, p_C) \leftarrow$  output of Algorithm 1 on  $C, A(C), v_C,$ 
   and  $b_C$ 
5: end for
6: let  $\mu : N \rightarrow A$  s.t. for all  $i \in N$ ,  $\mu(i) = \mu_C(i)$  for
    $C \in \mathcal{C}$  s.t.  $i \in C$ 
7: if LP (1) for  $\mu$  is feasible then
8:    $p \leftarrow$  solution of LP (1) for  $\mu$ 
9:   return  $(\mu, p)$ 
10: else
11:   return “no solution”
12: end if

```

Algorithm 2: Optimal envy-free allocation subject to budget constraints.

Proof of Theorem 1. We first claim that if the algorithm outputs “no solution” then $F_b(e) = \emptyset$. Indeed, by Lemma 7, the rent r output by Algorithm 1 on component C is exactly the δ_C in the statement of Lemma 6, hence the claim follows from Lemma 6.

Next, suppose that the algorithm outputs (μ, p) . Since this allocation solves LP (1), $(\mu, p) \in F_b(e)$. Let (η, q) be a maximizer of the minimum of f_1, \dots, f_t applied to the vector of utilities among all allocations in $F_b(e)$. Then as a byproduct of Lemma 6, we have $(\mu, q) \in F_b(e)$, which means that (μ, q) is an optimal solution to LP (1). Moreover, by Lemma 1, all agents are indifferent between (μ, q) and (η, q) , so the objective value under (μ, p) must be at least the objective value under (η, q) . \square

4 Fairest Budget-Constrained Allocations

The algorithm of Section 3 provides a practical method for selecting an allocation in $F_b(e)$ maximizing an objective function, when $F_b(e) \neq \emptyset$. Gal et al. (2017) have identified the *maximin solution*, which maximizes the minimum utility of any agent subject to envy-freeness, as an especially desirable solution. Our goal in this section is to show that this solution has similar properties even under budget constraints.

Formally, given an allocation $(\sigma, p) \in F(N, A, v, r)$, let

$$U(\sigma, p) = \min_{i \in N} v_{i\sigma(i)} - p_{\sigma(i)}.$$

The *budget-constrained maximin solution* returns an allocation in $\arg\max_{(\sigma, p) \in F_b(e)} U(\sigma, p)$ when $F_b(e) \neq \emptyset$.

The following theorem asserts that agent utilities under the budget-constrained maximin solution are unique, thus eliminating the potential need to select among multiple solutions; it extends a result of Alkan et al. (1991) to the budgeted setting.

Theorem 2. *Let*

$$\{(\sigma, p), (\mu, q)\} \subseteq \operatorname{argmax}_{(\sigma', p') \in F_b(e)} U(\sigma', p').$$

Then for all $i \in N$, $v_{i\sigma(i)} - p_{\sigma(i)} = v_{i\mu(i)} - q_{\mu(i)}$.

Proof. We first claim that it is sufficient to show that for every $\{(\sigma, p), (\sigma, q)\} \subseteq \operatorname{argmax}_{(\sigma', p') \in F_b(e)} U(\sigma', p')$, $p = q$, i.e., we can assume the two assignments are identical. Indeed, let (σ, p) and (μ, q) be budget-constrained maximin solutions, and let (η, o) be the output of Algorithm 2. Then, as a byproduct of Lemma 6, $\{(\eta, p), (\eta, q)\} \subseteq F_b(e)$. By Lemma 1, all agents are indifferent between (σ, p) and (η, p) , and between (μ, q) and (η, q) , that is,

$$\{(\eta, p), (\eta, q)\} \subseteq \operatorname{argmax}_{(\sigma', p') \in F_b(e)} U(\sigma', p').$$

Now, suppose for contradiction that there exist

$$\{(\sigma, p), (\sigma, q)\} \subseteq \operatorname{argmax}_{(\sigma', p') \in F_b(e)} U(\sigma', p')$$

such that $p \neq q$. Consider the price vector o such that $\forall a \in A$, $o_a = \max\{p_a, q_a\}$. Although $\sum_{i \in N} o_{\sigma(i)} > r$, o clearly satisfies the budget constraints, i.e., for all $i \in N$, $o_{\sigma(i)} \leq b_i$. Moreover, (σ, o) is envy free, because for all $i, j \in N$,

$$\begin{aligned} v_{i\sigma(i)} - o_{\sigma(i)} &= \min\{v_{i\sigma(i)} - p_{\sigma(i)}, v_{i\sigma(i)} - q_{\sigma(i)}\} \\ &\geq \min\{v_{i\sigma(j)} - p_{\sigma(j)}, v_{i\sigma(j)} - q_{\sigma(j)}\} \\ &= v_{i\sigma(j)} - o_{\sigma(j)}. \end{aligned}$$

Let $s = \sum_{i \in N} o_{\sigma(i)}$, and define a price vector o' by $o'_a = o_a - (s - r)/n$ for all $a \in A$. Then

$$\sum_{i \in N} o'_{\sigma(i)} = \sum_{i \in N} (o_{\sigma(i)} - (s - r)/n) = s - s + r = r.$$

In addition, $(s - r)/n > 0$, so $(\sigma, o') \in F_b(N, A, v, r)$. Finally, since for all $i \in N$,

$v_{i\sigma(i)} - o'_{\sigma(i)} > v_{i\sigma(i)} - o_{\sigma(i)} = v_{i\sigma(i)} - \max\{p_{\sigma(i)}, q_{\sigma(i)}\}$, it holds that

$$\begin{aligned} U(\sigma, o') &= \min_{i \in N} (v_{i\sigma(i)} - o'_{\sigma(i)}) \\ &> \min_{i \in N} (v_{i\sigma(i)} - \max\{p_{\sigma(i)}, q_{\sigma(i)}\}) \\ &= \min\{U(\sigma, p), U(\sigma, q)\}. \end{aligned}$$

We conclude that that (σ, p) and (σ, q) are not budget-constrained maximin solutions — a contradiction to our assumption. \square

While the budget-constrained maximin solution is conceptually appealing, previous work in behavioral economics (Herreiner and Puppe 2009) has shown that *social disparity* — the maximum difference in utilities between any two agents — plays a major role in whether allocations are perceived as fair. Formally, define

$$D(\sigma, p) = \max_{i \in N} (v_{i\sigma(i)} - p_{\sigma(i)}) - \min_{i \in N} (v_{i\sigma(i)} - p_{\sigma(i)}).$$

The *budget-constrained equitable solution* returns an allocation in $\operatorname{argmin}_{(\sigma, p) \in F_b(e)} D(\sigma, p)$ when $F_b(e) \neq \emptyset$. Surprisingly, Gal et al. (2017) show that, without budget constraints, any maximin solution is also an equitable solution. Here we generalize their result to the budget-constrained setting.

Theorem 3. *If (σ, p) is a budget-constrained maximin solution, then it is a budget-constrained equitable solution.*

Proof. Let $G = (N \cup \{0\}, E)$ be a directed graph such that there is an edge from $i \in N$ to 0 if and only if $v_{i\sigma(i)} - p_{\sigma(i)} = U(\sigma, p)$ or $p_{\sigma(i)} = b_i$; and there is an edge from $i \in N$ to $j \in N$ if and only if $v_{i\sigma(i)} - p_{\sigma(i)} = v_{i\sigma(j)} - p_j$.

We first claim that for all $i \in N$, there is a path from i to 0. Indeed, assume for contradiction that this is not the case, and let S be the set of all $i \in N$ that have paths to 0 in G . Since $S \neq \emptyset$ and $N - S \neq \emptyset$, we can define another price vector p' such that

$$p'_{\sigma(i)} = \begin{cases} p_{\sigma(i)} - \varepsilon & i \in S \\ p_{\sigma(i)} + \frac{|S|}{n-|S|} \varepsilon & i \notin S \end{cases}$$

where $\varepsilon > 0$. Clearly $\sum_{i \in N} p'_{\sigma(i)} = r$. Also, we can set ε small enough to satisfy the following properties:

- $\forall i \notin S, v_{i\sigma(i)} - p'_{\sigma(i)} > U(\sigma, p)$;
- $\forall i \notin S, p'_{\sigma(i)} < b_i$;
- $\forall i \notin S, \forall j \in S, v_{i\sigma(i)} - p'_{\sigma(i)} > v_{i\sigma(j)} - p'_{\sigma(j)}$. (This is because no $i \notin S$ weakly envies $j \in S$ in (σ, p) , otherwise i would have a path to 0.)

Therefore, $(\sigma, p') \in F_b(N, A, v, r)$. Moreover, since for all $i \in S$,

$$v_{i\sigma(i)} - p'_{\sigma(i)} > v_{i\sigma(i)} - p_{\sigma(i)} \geq U(\sigma, p),$$

$U(\sigma, p') > U(\sigma, p)$, which contradicts the assumption that (σ, p) is a budget-constrained maximin solution.

Next, let (σ', p') be a budget-constrained equitable solution. Using similar arguments to those at the beginning of the proof of Theorem 2, we can assume without loss of generality that $\sigma' = \sigma$, that is, it is sufficient to show that $D(\sigma, p) \leq D(\sigma, p')$.

Let $\varepsilon = U(\sigma, p) - U(\sigma, p')$. Since (σ, p) is a budget-constrained maximin solution, $\varepsilon \geq 0$. We claim that for all $a \in A$, $p'_a \leq p_a + \varepsilon$. Indeed, consider a path $i_1 \rightarrow i_2 \rightarrow \dots \rightarrow i_\ell \rightarrow 0$ in G . Since such a path exists for all $i_1 \in N$, it is sufficient to prove that $p'_{\sigma(i_1)} \leq p_{\sigma(i_1)} + \varepsilon$.

We prove the claim by induction on the length of the path ℓ . First, when $\ell = 1$, there are two cases for the edge from i_1 to 0: if $p_{\sigma(i_1)} = b_{i_1}$, then $p'_{\sigma(i_1)} \leq b_{i_1} \leq p_{\sigma(i_1)} + \varepsilon$; and if $v_{i_1\sigma(i_1)} - p_{\sigma(i_1)} = U(\sigma, p)$,

$$\begin{aligned} p'_{\sigma(i_1)} &\leq v_{i_1\sigma(i_1)} - U(\sigma, p') \\ &= v_{i_1\sigma(i_1)} - U(\sigma, p) + \varepsilon \\ &= p_{\sigma(i_1)} + \varepsilon. \end{aligned}$$

When $\ell > 1$, the induction assumptions gives us $p'_{\sigma(i_2)} \leq p_{\sigma(i_2)} + \varepsilon$. Then

$$\begin{aligned} p'_{\sigma(i_1)} &\leq v_{i_1\sigma(i_1)} - v_{i_1\sigma(i_2)} + p'_{\sigma(i_2)} \\ &\leq v_{i_1\sigma(i_1)} - v_{i_1\sigma(i_2)} + p_{\sigma(i_2)} + \varepsilon \\ &= p_{\sigma(i_1)} + \varepsilon, \end{aligned}$$

where the first transition holds because (σ, p') is envy free, and the last transition holds because (i_1, i_2) is an edge of G .

We can now conclude that

$$\begin{aligned} D(\sigma, p') &= \max_{i \in N} \{v_{i\sigma(i)} - p'_{\sigma(i)}\} - U(\sigma, p') \\ &\geq \max_{i \in N} \{v_{i\sigma(i)} - p_{\sigma(i)} - \varepsilon\} - U(\sigma, p) + \varepsilon \\ &= D(\sigma, p). \end{aligned}$$

□

5 Extensions

In order to apply our results in practice, one would ask participants for both values and budget constraints, instead of just values. Now, there are two options. The first is to compute the budget-constrained maximin solution (Algorithm 2), and if $F_b(e)$ is empty, simply compute the (non-budget-constrained) maximin solution, and let users know that it is impossible to satisfy the budget constraints. The second option is to inform the choice of an envy-free allocation with the budget constraints, even when these cannot be satisfied. This suggests the following general approach.

For an allocation (σ, p) for e , define its *budget violation penalty* as $\max_{i \in N} \{0, p_{\sigma(i)} - b_i\}$. We first find the allocations that minimize the budget violation penalty among all envy-free allocations for e . We then maximize a given criterion of justice among these allocations. Note that selecting an element in $F_b(e)$ that optimizes an objective function whenever $F_b(e) \neq \emptyset$ can be seen as an instance of this two-step approach. And it turns out that the more general approach is still implementable in polynomial time.

Theorem 4. *Let $e \equiv (N, A, v, r)$, $b \equiv (b_i)_{i \in N}$, and, for $s \in [t]$, where t is polynomial in n , let $f_s : \mathbb{R}^n \rightarrow \mathbb{R}$ be a linear function. Then there is a polynomial-time algorithm that outputs an allocation $(\sigma, p) \in F(e)$ that maximizes*

$$\min_{s \in [t]} f_s(v_{1\sigma(1)} - p_{\sigma(1)}, \dots, v_{n\sigma(n)} - p_{\sigma(n)})$$

over all allocations in $F(e)$ that minimize the budget violation penalty.

We present the proof of this theorem in Appendix A. Specifically, we construct a polynomial-time algorithm (Algorithm 3) that produces an allocation that minimizes the budget violation penalty among all envy-free allocations. We find that if this allocation violates the budget constraints, it necessarily maximizes *any* additional criteria of justice among all envy-free allocations that minimize the maximum budget violation penalty (Lemma 12). If the allocation produced by Algorithm 3 satisfies the budget constraints, Algorithm 2 produces the desired outcome.

Going a step further, using the ideas of Section 3 and some additional machinery, we design, in Appendix B, a *purely combinatorial* algorithm that transforms the allocation produced by Algorithm 3 into a budget-constrained maximin solution in $O(n^3)$ time.

6 Discussion

The need for computational efficiency. From a computational complexity viewpoint, all of our algorithms are prac-

ticable. One may wonder, though, why computational complexity is even an issue, given that typical rent division instances are small (the vast majority of real-world instances have up to 4 agents). The reason is that the number of instances is relatively large. Spliddit, for example, employs Amazon Web Services, so running its algorithms comes at a cost. Informally speaking, if algorithm A is, say, twice as computationally efficient as Algorithm B , then the cost of solving tens of thousands of instances using algorithm A would (ideally) be half that of Algorithm B .

An alternative approach (which does not work). In a masterpiece of mathematical exposition, Su (1999) presents an algorithm for envy-free rent division that relies on Sperner’s Lemma; this Algorithm was actually implemented in 2014 by the New York Times.⁴ The algorithm asks participants a sequence of queries (“which room would you prefer at these prices?”), until it converges to an (approximately) envy-free allocation. Even though Su’s assumptions on utility are relatively weak, his approach has several shortcomings. First, preference elicitation is cumbersome, as it requires repeated interaction with users, and, moreover, that interaction may be lengthy as the rate of convergence can be exponential in the tolerance for envy. Second, the approach makes it impossible to select among envy-free allocations (or, at least, it is unknown how to do so), and, therefore, it may lead to envy-free allocations that are intuitively unfair. Even more importantly for the present paper, it seems inherently impossible to adapt Su’s algorithm to the budget-constrained case, as a participant must be able to choose a room under every rent division, including those where the prices of all rooms violate his budget constraints.

Incentives. A natural question is whether our algorithms are manipulable, that is, whether an agent can benefit by misrepresenting his preferences; the answer is “yes”, not only for our algorithms, but for any algorithm that produces an envy-free allocation for the utility reports (Alkan, Demange, and Gale 1991). However, as Gal et al. (2017) argue, the incentive problem may be mitigated in practice by the fact that participants are typically unaware of how the algorithm works (apart from its fairness guarantees), which makes strategic manipulation difficult.

But we can say even more. Velez (2015) shows that, in the non-budget-constrained setting, each envy-free algorithm has the remarkable property that the set of envy-free allocations, for true utilities, coincides with the set of limit equilibrium outcomes (Radner 1980) of their induced complete information games. The budget-constrained setting creates more opportunities for manipulation, as now agents report their values for the rooms (as before) and, additionally, their budget constraints. However, it is possible to show that the result of Velez (2015) goes through, by extending his analysis. Therefore, when agents know each other well and strategically engage in manipulating the algorithm, we can still expect to obtain an envy-free allocation.

⁴<https://www.nytimes.com/2014/04/29/science/to-divide-the-rent-start-with-a-triangle.html>

References

- Aleksandrov, M., and Walsh, T. 2017. Pure Nash equilibria in online fair division. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 42–48.
- Alkan, A.; Demange, G.; and Gale, D. 1991. Fair allocation of indivisible goods and criteria of justice. *Econometrica* 59(4):1023–1039.
- Andersson, T., and Ehlers, L. 2013. An algorithm for identifying agent- k -linked allocations in economies with indivisibilities. Cahier 17-2013, CIREQ.
- Andersson, T.; Ehlers, L.; and Svensson, L.-G. 2014. Budget-balance, fairness and minimal manipulability. *Theoretical Economics* 9(3):753–777.
- Aragones, E. 1995. A derivation of the money Rawlsian solution. *Social Choice and Welfare* 12:267–276.
- Aziz, H.; Rauchecker, G.; Schryen, G.; and Walsh, T. 2017. Algorithms for max-min share fair allocation of indivisible chores. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, 335–341.
- Bei, X.; Qiao, Y.; and Zhang, S. 2017. Networked fairness in cake cutting. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 3632–3638.
- Bliem, B.; Bredereck, R.; and Niedermeier, R. 2016. Complexity of efficient and envy-free resource allocation: Few agents, resources, or utility levels. In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, 102–108.
- Bouveret, S.; Cechlárová, K.; Elkind, E.; Igarashi, A.; and Peters, D. 2017. Fair division of a graph. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 135–141.
- Brânzei, S.; Caragiannis, I.; Kurokawa, D.; and Procaccia, A. D. 2016. An algorithmic framework for strategic fair division. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, 418–424.
- Gal, Y.; Mash, M.; Procaccia, A. D.; and Zick, Y. 2017. Which is the fairest (rent division) of them all? *Journal of the ACM*. Forthcoming.
- Goldman, J., and Procaccia, A. D. 2014. Spliddit: Unleashing fair division algorithms. *SIGecom Exchanges* 13(2):41–46.
- Herreiner, D. K., and Puppe, C. D. 2009. Envy freeness in experimental fair division problems. *Theory and decision* 67(1):65–100.
- Kurokawa, D.; Procaccia, A. D.; and Wang, J. 2016. When can the maximin share guarantee be guaranteed? In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, 523–529.
- Menon, V., and Larson, K. 2017. Deterministic, strategyproof, and fair cake cutting. In *Proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI)*, 352–358.
- Radner, R. 1980. Collusive behavior in noncooperative epsilon-equilibria of oligopolies with long but finite lives. *Journal Economic Theory* 22(2):136–154.

- 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.
- Svensson, L.-G. 2009. Coalitional strategy-proofness and fairness. *Economic Theory* 40(2):227–245.
- Velez, R. A. 2015. Sincere and sophisticated players in an equal-income market. *Journal Economic Theory* 157(0):1114–1129.
- Velez, R. A. 2017. Equitable rent division. Mimeo, Department of Economics, Texas A&M University.

Appendix

A Minimizing the Max Budget Violation

In order to prove Theorem 4 we first modify Algorithm 2 to calculate an allocation at which the agents, in aggregate, pay the maximal rent they could pay in an envy-free allocation while satisfying budget constraints. Then we prove two key observations that essentially complete the proof. First, a “horizontal” translation of this allocation minimizes the maximal budget violation. Second, whenever no envy-free allocation satisfies budget constraints, minimizing the maximal budget violation penalty selects a unique level of welfare.

For $e \equiv (N, A, v, r)$ let μ^v be the allocation computed in Step 6 of Algorithm 2 (note that this allocation is invariant to the value of r); for each $p \in \mathbb{R}^A$ for which there is σ such that $(\sigma, p) \in F(e)$, let $\Gamma_C(p)$ be the graph $(\mathcal{C}(N, A, v), E)$ where $(C, D) \in E$ if and only if there is $i \in C$ and $j \in D$ such that $(i, j) \in \Gamma(\mu^v, p)$. One can easily check that $(C, D) \in \Gamma_C(p)$ if and only if for each $(\sigma, p) \in F(u)$, there are $i \in C$ and $j \in D$ such that $(i, j) \in \Gamma(\sigma, p)$.⁵

Lemma 8. *Let (μ, p) and (σ, p) be two allocations in $F(u)$, and $(C, D) \in \Gamma_C(p)$ such that $\sum_{a \in \mu(C)} q_a > \sum_{a \in \mu(C)} p_a$. Thus, $\sum_{a \in \mu(D)} q_a > \sum_{a \in \mu(D)} p_a$.*

Proof. Since $(C, D) \in \Gamma_C(p)$ there are $i \in C$ and $j \in D$ such that $(i, j) \in \Gamma(\mu, p)$. By Lemma 4, $\sum_{a \in \mu(C)} q_a > \sum_{a \in \mu(C)} p_a$, implies $q_{\mu(i)} > p_{\mu(i)}$. Thus,

$$\begin{aligned} u_i(p_{\mu(j)}, \mu(j)) &= u_i(p_{\mu(i)}, \mu(i)) \\ &> u_i(q_{\sigma(i)}, \sigma(i)) \\ &\geq u_i(q_{\mu(j)}, \mu(j)), \end{aligned}$$

where the strict inequality holds by Lemma 1, and the weak inequality holds because $(\sigma, p) \in F(u)$. Thus, $q_{\mu(j)} > p_{\mu(j)}$. By Lemma 4, $\sum_{a \in \mu(D)} q_a > \sum_{a \in \mu(D)} p_a$. \square

⁵Let (σ, p) and (μ, p) be two allocations in $F(u)$, $\{C, D\} \subseteq \mathcal{C}$, $i \in C$, and $j \in D$ such that $u_i(p_{\sigma(i)}, \sigma(i)) = u_i(p_{\sigma(j)}, \sigma(j))$. By Lemma 5, $k \equiv \mu^{-1}(\sigma(j)) \in D$. By Lemma 1, $u_i(p_{\sigma(i)}, \sigma(i)) = u_i(p_{\mu(i)}, \mu(i))$. Thus, $u_i(p_{\mu(i)}, \mu(i)) = u_i(p_{\mu(k)}, \mu(k))$.

Lemma 9. Let $e \equiv (N, A, v, r)$, $b \equiv (b_i)_{i \in N}$, and $\mathcal{C} \equiv \mathcal{C}(N, A, v)$. Suppose that $(\mu^v, p) \in F_b(e)$ is such that for each $C \in \mathcal{C}$ there is $D \in \mathcal{C}$ such that there is a path from C to D in $\Gamma_{\mathcal{C}}(p)$ and $\sum_{a \in D} p_a = \delta_D$. Then, for each $s > r$, $F_b(N, A, v, s) = \emptyset$.

Proof. Let $s > r$ and $(\sigma, q) \in F(N, A, v, s)$. Thus, there is $C \in \mathcal{C}$ such that $\sum_{a \in C} q_a > \sum_{a \in C} p_a$. Let $D \in \mathcal{C}$ be such that there is a path from C to D in $\Gamma_{\mathcal{C}}(p)$ and $\sum_{a \in D} p_a = \delta_D$. By Lemma 8, $\sum_{a \in D} q_a > \sum_{a \in D} p_a$. Thus, $\sum_{a \in D} q_a > \delta_D$ and $(\sigma, q) \notin F_b(N, A, v, s)$. \square

The following is the basic step in our next algorithm. Starting from an allocation $(\mu^v, p) \in F_b(N, A, v, r)$, we extract the maximal amount of money possible from the set of agents who are not directly or indirectly budget constrained, while preserving no-envy and budget constraints.⁶ More precisely, let the sets of agents $L \neq \emptyset$ and $M \neq \emptyset$ be a partition of N . Let $q(p, L, M)$ be the vector $q \in \mathbb{R}^A$ defined in 1 below for the maximal Δ for which 2-3 are satisfied.

1. For each $i \in L$, $q_{\mu^v(i)} = p_{\mu^v(i)} + \Delta$; and for each $i \in M$, $q_{\mu^v(i)} = p_{\mu^v(i)}$.
2. For each $i \in L$ and each $j \in M$, $v_{i\mu(i)} - q_{\mu^v(i)} \geq v_{i\mu(j)} - q_{\mu^v(j)}$.
3. For each $i \in L$, $q_{\mu^v(i)} \leq b_{\mu^v(i)}$.

The Δ associated with $q(p, L, M)$ can be directly computed in $O(n^2)$ as the maximum satisfying conditions 2 and 3 above. By 2, $(\mu^v, q(p, L, M)) \in F(N, A, v, r - |L|\Delta)$. By 3, $(\mu^v, q(p, L, M)) \in F_b(N, A, v, r - |L|\Delta)$.

Lemma 10. Let N, A, v , and $b \equiv (b_i)_{i \in N}$. Then, Algorithm 3 stops in polynomial time. Let r and (μ^v, p) be the outcome of the algorithm. Then, $(\mu^v, p) \in F_b(N, A, v, r)$ and for each $s > r$, $F_b(N, A, v, s) = \emptyset$.

Proof. The algorithm coincides with Algorithm 2 up to Line 6. The set L decreases each time Line 13 holds. Thus, Line 16 holds at most n times. Lines 12 and 17 can be completed each in $O(n^2)$. Since (σ, p) in Line 1 is envy-free, (μ^v, p) defined in Line 8 belongs to $F(N, A, v, r)$ where r is defined in Line 9. By definition of Δ in Line 7, for each $C \in \mathcal{C}$, $\sum_{a \in A(C)} p_a \geq \delta_C$. Thus, (μ^v, p) defined in Line 8 belongs to $F_b(N, A, v, r)$ where r is defined in Line 9. Since Line 17 preserves no-envy and budget constraints, r and (μ^v, p) , the outcome of the algorithm, are such that $(\mu^v, p) \in F_b(N, A, v, r)$. Since $L = \emptyset$ just before the algorithm stops, the agents who belong to some $C \in \mathcal{C}$ such that there is a path in $\Gamma_{\mathcal{C}}(p)$ from C to some $D \in \mathcal{C}$ such that $\sum_{a \in D} p_a = \delta_D$ is N . By Lemma 9, for each $s > r$, $F_b(N, A, v, s) = \emptyset$. \square

Lemma 11. Let N, A, v , and $b \equiv (b_i)_{i \in N}$. Let r and (μ^v, p) be the outcome of Algorithm 3. For each $s \in \mathbb{R}$, let $p^s \equiv (p_a + (s - r)/n)_{a \in A}$. Then,

$$(\mu^v, p^s) \in \arg \min_{(\sigma, q) \in F(N, A, v, s)} \max_{i \in N} \{0, q_{\sigma(i)} - b_i\}.$$

⁶A similar balancing redistribution was proposed by Andersson and Ehlers (2013) to calculate, in polynomial time, an agent optimal envy-free allocation.

```

1: compute  $(\sigma, p) \in F(N, A, v, r)$  for some  $r \in \mathbb{R}$ 
2:  $\mathcal{C} \leftarrow$  strongly connected components of  $\Gamma(\sigma, p)$ 
3: for each  $C \in \mathcal{C}$  do
4:    $\delta_C$  and  $(\mu_C, p_C) \leftarrow$  output of Algorithm 1 on  $C$ ,
      $A(C)$ ,  $v_C$ , and  $b_C$ 
5: end for
6: let  $\mu^v : N \rightarrow A$  s.t. for all  $i \in N$ ,  $\mu^v(i) = \mu_C(i)$  for
    $C \in \mathcal{C}$  s.t.  $i \in C$ 
7:  $\Delta \leftarrow \min_{C \in \mathcal{C}} (\delta_C - p_C) / |C|$ 
8:  $p \leftarrow (p_a + \Delta)_{a \in N}$ 
9:  $r \leftarrow r + n\Delta$ 
10:  $L = N$ 
11: while  $L \neq \emptyset$  do
12:    $M \leftarrow$  set of agents who belong to some  $C \in \mathcal{C}$  such
     that there is a path in  $\Gamma_{\mathcal{C}}(p)$  from  $C$  to some  $D \in \mathcal{C}$ 
     such that  $\sum_{a \in D} p_a = \delta_D$ 
13:    $L \leftarrow N \setminus M$ 
14:   if  $L = \emptyset$  then
15:     return  $r$  and  $(\mu^v, p)$ 
16:   else
17:      $p \leftarrow q(p, L, M)$ 
18:      $r \leftarrow \sum_{a \in A} p_a$ 
19:   end if
20: end while

```

Algorithm 3: Maximum-rent envy-free allocation.

Proof. Since $(\mu^v, p) \in F_b(N, A, v, r)$, the result is obvious for $s < r$. Since $M = N$ before the algorithm stops, the set of $C \in \mathcal{C}(N, A, v)$ such that $\sum_{i \in C} p_{\sigma(i)} = \delta_C$ is non-empty. Thus, there is $i \in N$ such that $p_{\sigma(i)} = b_i$. Let $s \geq r$. Since for each $i \in N$, $p_{\sigma(i)} \leq b_i$,

$$\max_{i \in N} \max\{0, p_i^s - b_i\} = \frac{s - r}{n}.$$

Since $(\mu^v, p) \in F(N, A, v, r)$, we have that $(\mu^v, p^s) \in F(N, A, v, s)$. Let $(\sigma, q) \in F(N, A, v, s)$ and $\Delta \equiv \max_{i \in N} \max\{0, q_{\sigma(i)} - b_i\}$. Since $(\sigma, (q_a - \Delta)_{a \in A}) \in F_b(N, A, v, s)$, by Lemma 10, $s - n\Delta \leq r$. Thus, $\Delta \geq (s - r)/n$. \square

Lemma 12. Let N, A, v , and $b \equiv (b_i)_{i \in N}$. Let r and (μ^v, p) be the outcome of Algorithm 3. For each $s \in \mathbb{R}$, let $p^s \equiv (p_a + (s - r)/n)_{a \in A}$. Then, for each $s \geq r$, each

$$(\sigma, q) \in \arg \min_{(\sigma, q) \in F(N, A, v, s)} \max_{i \in N} \{0, q_{\sigma(i)} - b_i\},$$

and each $i \in N$, $v_{i\sigma(i)} - q_{\sigma(i)} = v_{i\mu^v(i)} - p_{\mu^v(i)}^s$.

Proof. Let $(\sigma, q) \in F(N, A, v, s)$ in the statement of the lemma. By Lemma 1 it is enough to prove that $q = p^s$. Suppose by contradiction that $q \neq p^s$. Since both q and p^s add up to s , there is $a \in A$ such that $q_a > p_a^s$. Let $C \in \mathcal{C}(N, A, v)$ be such that $a \in A(C)$. By Lemma 4, for each $a \in A(C)$, $q_a > p_a^s$. Since $M = N$ when Algorithm 3 terminates, there is a path in $\Gamma_{\mathcal{C}}(p)$ from C to some $D \in \mathcal{C}$ such that $\sum_{a \in A(D)} p_a = \delta_D$. Since for each $a \in A(C)$, $q_a > p_a^s$, by Lemma 8, $\sum_{a \in A(D)} q_a > \sum_{a \in A(D)} p_a^s$. By Lemma 4, for each $a \in A(D)$, $q_a > p_a^s$. By definition of p^s ,

$\sum_{a \in A(D)} p_a^s = \delta_D + |D|(s-r)/n$. Thus, $\sum_{a \in A(D)} q_a - |D|(s-r)/n > \delta_D$. From our proof of Lemma 11 we know that,

$$\min_{(\sigma, q) \in F(N, A, v, s)} \max_{i \in N} \{0, q_{\sigma(i)} - b_i\} = (s-r)/n.$$

Thus, $(\sigma, (q_a - (s-r)/n)_{a \in A}) \in F_b(N, A, v, r)$. Thus, $(\sigma_D, q_{A(D)}) \in F_{b_D}(D, A(D), v_D, t)$ for $t > \delta_D$. This is a contradiction. \square

1: \hat{r} and $(\mu^v, p) \leftarrow$ outcome of Algorithm 3 on N, A, v
2: $p \leftarrow (p_a + (r - \hat{r})/n)_{a \in A}$
3: **if** $r < \hat{r}$ **then**
4: $(\mu^v, p) \leftarrow$ outcome of Algorithm 2 on (N, A, v, r)
5: **end if**
6: **return** (μ^v, p)

Algorithm 4: Optimal envy-free allocation among minimizers of maximal budget violation.

Proof of Theorem 4. Each step in Algorithm 4 runs in polynomial time. Let (μ^v, p) be its outcome on $e \equiv (N, A, v, r)$. If Line 4 does not hold, by Lemma 12 the subset of allocations in $F(e)$ that minimize the maximal budget violation penalty are all welfare equivalent. Thus, (μ^v, p) trivially maximizes

$$\min_{s \in [t]} f_s(v_{1\sigma(1)} - p_{\sigma(1)}, \dots, v_{n\sigma(n)} - p_{\sigma(n)})$$

on this set. If Line 4 holds, the result follows from Theorem 1. \square

B Maxmin envy-free allocations

When the additional criteria of justice is to maximize the minimum utility among agents, Line 4 in Algorithm 4 can be substituted with a combinatorial subroutine that dispenses altogether of linear programming and runs in $O(n^3)$.

Lemma 13. *Let $e \equiv (N, A, v, r)$, $b \equiv (b_i)_{i \in N}$ such that $F_b(e) \neq \emptyset$, and $\mathcal{C} \equiv \mathcal{C}(N, A, v)$. Suppose that $(\mu^v, p) \in F_b(e)$ is such that for each $C \in \mathcal{C}$ there is a path in $\Gamma_{\mathcal{C}}(p)$ from C to some $D \in \mathcal{C}$ such that either $\sum_{a \in D} p_a = \delta_D$ or there is $i \in D$ with $i \in \arg \min_{i \in N} v_{i\mu^v(i)} - p_{\mu^v(i)}$. Then,*

$$(\mu^v, p) \in \arg \max_{(\sigma, q) \in F_b(e)} \min_{i \in N} v_{i\sigma(i)} - q_{\sigma(i)}.$$

Proof. Let $(\mu^v, p) \in F_b(e)$ satisfy the conditions in the lemma. Suppose by contradiction that there is $(\sigma, q) \in F_b(e)$ such that $\min_{i \in N} v_{i\sigma(i)} - q_{\sigma(i)} > \min_{i \in N} v_{i\mu^v(i)} - p_{\mu^v(i)}$. Let D be such that there is $i \in D$ such that $i \in \arg \min_{i \in N} v_{i\mu^v(i)} - p_{\mu^v(i)}$. Then,

$$\min_{i \in D} v_{i\mu^v(i)} - p_{\mu^v(i)} < \min_{i \in D} v_{i\mu^v(i)} - q_{\mu^v(i)}.$$

By Lemma 4, $\sum_{a \in D} p_a > \sum_{a \in D} q_a$. Thus, there is $H \in \mathcal{C}$ such that $\sum_{a \in H} p_a < \sum_{a \in H} q_a$. Let $F \in \mathcal{C}$ be such that there is a path from H to F in $\Gamma_{\mathcal{C}}(p)$ and either $\sum_{a \in F} p_a = \delta_F$ or there is $i \in F$ with $i \in \arg \min_{i \in N} v_{i\mu^v_D(i)} - p_{\mu^v(i)}$.

1: $L = N$
2: **while** $L \neq \emptyset$ **do**
3: $M \leftarrow$ the set of agents in a $C \in \mathcal{C}$ s.t. there is a path in $\Gamma_{\mathcal{C}}(p)$ from C to some $D \in \mathcal{C}$
s.t. $D \cap \arg \min_{i \in N} v_{i\mu^v(i)} - p_{\mu^v(i)} \neq \emptyset$
4: $K \leftarrow$ the set of agents, not in M , and in a $C \in \mathcal{C}$
s.t. there is a path in $\Gamma_{\mathcal{C}}(p)$ from C to some $D \in \mathcal{C}$ s.t. $\sum_{a \in D} p_a = \delta_D$
5: $L \leftarrow N \setminus (M \cup K)$
6: **if** $L = \emptyset$ **then**
7: **return** (μ^v, p)
8: **else**
9: $p \leftarrow q_m(p, L, K, M)$
10: **end if**
11: **end while**

Algorithm 5: Optimal envy-free allocation subject to budget constraints.

A standard argument (Velez 2017, Lemma 3) shows that $\sum_{a \in F} q_a > \sum_{a \in F} p_a = \delta_F$. Thus, either

$$\min_{i \in N} v_{i\sigma(i)} - q_{\sigma(i)} \leq \min_{i \in F} v_{i\sigma(i)} - q_{\sigma(i)} \leq \min_{i \in N} v_{i\mu^v(i)} - p_{\mu^v(i)}$$

or $(\sigma, q) \notin F_b(e)$. This is a contradiction. \square

Our basic step is, starting from an allocation $(\mu^v, p) \in F_b(N, A, v, r)$, to extract the maximal possible amount of money from a set of agents who are not directly or indirectly budget constrained, and give it to another set of agents preserving no-envy and increasing the minimum utility among all agents. More precisely, let the sets of agents $L \neq \emptyset$, K , and $M \neq \emptyset$ be a partition of N such that $\arg \min_{j \in N} v_{j\mu^v(j)} - p_j \subseteq M$. Let $q_m(p, L, K, M)$ be the allocation q defined in 1 below for the maximal Δ for which 2-4 are satisfied.

1. For each $i \in L$, $q_{\mu^v(i)} = p_{\mu^v(i)} + \Delta$; for each $i \in K$, $q_{\mu^v(i)} = p_{\mu^v(i)}$; and for each $i \in M$, $q_{\mu^v(i)} = p_{\mu^v(i)} - \Delta|L|/|M|$.
2. For each $i \in L$,
 - (a) for each $j \in K \cup M$, $v_{i\mu^v(i)} - q_{\mu^v(i)} \geq v_{i\mu^v(j)} - q_{\mu^v(j)}$;
 - (b) $v_{i\mu^v(i)} - q_{\mu^v(i)} \geq \min_{j \in M} v_{j\mu^v(j)} - q_{\mu^v(j)}$.
3. For each $i \in K$,
 - (a) for each $j \in M$, $v_{i\mu^v(i)} - q_{\mu^v(i)} \geq v_{i\mu^v(j)} - q_{\mu^v(j)}$.
 - (b) $v_{i\mu^v(i)} - q_{\mu^v(i)} \geq \min_{j \in M} v_{j\mu^v(j)} - q_{\mu^v(j)}$.
4. For each $i \in L$, $q_{\mu^v(i)} \leq b_i$.

The Δ associated with $q_m(p, L, K, M)$ can be directly computed in $O(n^2)$ as the maximum satisfying conditions 2, 3, and 4 above.⁷ By 2(a) and 3(a), $(\mu^v, q_m(p, L, K, M)) \in F(N, A, v, r)$. By 4,

$$(\mu^v, q_m(p, L, K, M)) \in F_b(N, A, v, r). \quad (2)$$

By 2(b) and 3(b),

$$\arg \min_{j \in N} v_{j\mu^v(j)} - p_{\mu^v(j)} \subseteq \arg \min_{j \in N} v_{j\mu^v(j)} - q_m(p, L, K, M)_{\mu^v(j)}. \quad (3)$$

⁷Note that conditions 2(b) and 3(b) can be written as $v_{i\mu^v(i)} - p_{\mu^v(i)} \geq \min_{j \in M} v_{j\mu^v(j)} - p_{\mu^v(j)} + \Delta(1 + |L|/|M|)$.

If Line 4 in Algorithm 4 holds, the allocation (μ^v, p) calculated in Line 2 is envy-free and satisfies budget constraints. Algorithm 5 transforms (μ^v, p) into a maximum budget-constrained envy-free allocation.

Lemma 14. *Let $e \equiv (N, A, v, r)$, $b \equiv (b_i)_{i \in N}$, $\mathcal{C} \equiv \mathcal{C}(N, A, v)$, $(\delta_C)_{C \in \mathcal{C}}$, and $(\mu^v, p) \in F_b(e)$. Then, Algorithm 5 runs in $O(n^3)$. Its outcome maximizes the minimum utility among agents among all allocations in $F_b(e)$.*

Proof. By (3) the set M weakly increases every time Line 3 holds. Thus, the set L weakly decreases every time Line 5 holds. If the set L does not decrease when Line 5 holds, then the set M increases. Thus, Line 2 holds $O(n)$ times. Since each step runs in $O(n^2)$, the algorithm runs in $O(n^3)$. By (2), (μ^v, p) calculated in Line 9 belongs to $F_b(e)$. Since $L = \emptyset$, i.e., $M \cup K = N$, when the algorithm terminates, Lemma 13 guarantees its outcome maximizes the minimum utility among agents among all allocations in $F_b(e)$. \square