Better Redistribution with Inefficient Allocation in Multi-Unit Auctions*

Mingyu Guo
University of Adelaide
School of Computer Science
Adelaide, Australia
mingyu.guo@adelaide.edu.au

Vincent Conitzer[†]
Duke University
Department of Computer Science
Durham, NC, USA
conitzer@cs.duke.edu

Abstract

For the problem of allocating one or more items among a group of competing agents, the Vickrey-Clarke-Groves (VCG) mechanism is strategy-proof and efficient. However, the VCG mechanism is not strongly budget balanced: in general, value flows out of the system of agents in the form of VCG payments, which reduces the agents' utilities. In many settings, the objective is to maximize the sum of the agents' utilities (taking payments into account). For this purpose, several VCG redistribution mechanisms have been proposed that redistribute a large fraction of the VCG payments back to the agents, in a way that maintains strategy-proofness and the non-deficit property. Unfortunately, sometimes even the best VCG redistribution mechanism fails to redistribute a substantial fraction of the VCG payments. This results in a low total utility for the agents, even though the items are allocated efficiently. In this paper, we study strategy-proof allocation mechanisms that do not always allocate the items efficiently. It turns out that by allocating inefficiently, more payment can sometimes be redistributed, so that the net effect is an increase in the sum of the agents' utilities.

Our objective is to design mechanisms that are *competitive* against the first-best mechanism in terms of the agents' total utility. We first

study multi-unit auctions with unit demand. We characterize the family of *linear* allocation mechanisms. We propose an optimization technique for simultaneously finding a linear allocation mechanism and a payment redistribution rule, which together are optimal. With the help of this technique, we also analytically characterize several competitive mechanisms, which are based on burning units and partitioning the agents into groups. Finally, we extend the definition of linear allocation mechanisms and the optimization technique to general multi-unit auctions.

1   Introduction

Many problems in multiagent systems involve resource allocation. This paper studies the fundamental problem of using economic mechanisms to allocate resources among autonomous agents, with the aim of maximizing the agents' total utility. We focus on settings where the resource consists of multiple indistinguishable units. Example real-life scenarios include the allocation of cluster nodes to supercomputing clients, where the cluster nodes are identical. Mechanisms proposed in this paper can be applied to scenarios where each agent requires only one unit of the resource (*e.g.*, the client application contains a single thread, thus requiring only one node) or each agent's marginal valuations for the units are nonincreasing (*e.g.*, the client application runs faster with more nodes, but the marginal benefit is nonincreasing). Besides proposing mechanisms that address a fundamental problem in multiagent systems, our results also demonstrate the effectiveness of using computing techniques to deliver new results in economic theory.

For the problem of allocating one or more items among a group of competing agents, the well-known Vickrey-Clarke-Groves (VCG) mechanism [30, 5, 12] is strategy-proof and efficient.[1] That is, under the VCG mechanism, it is a dominant strategy for the agents to report their true valuations for the items, and the mechanism allocates the items in a way that maximizes the sum of the agents' valuations. In the context of resource allocation, the VCG mechanism also satisfies the non-deficit property: the sum of the agents' payments is nonnegative. However, the VCG mechanism is not strongly budget balanced, that is, the agents' payments do not sum to 0. Hence, in general, value flows out of the system of agents in the form of VCG payments, which reduces the agents' utilities. Actually, for sufficiently general settings, no mechanism satisfies strategy-proofness, efficiency, and strong budget balance at the same time [22, 11, 10, 24].[2]

In many cases, the objective of the allocation process is to maximize the welfare of the agents, that is, the sum of the agents' utilities, *taking payments into account*. For this purpose, several *VCG redistribution mechanisms* have been proposed. These mechanisms first allocate the items efficiently and charge the VCG payments. Then, a large fraction of the VCG revenue is redistributed

---

[1]As is often done, we use "the VCG mechanism" to refer to the Clarke mechanism, and not to any other Groves mechanism.

[2]By sacrificing one of these properties, the others can be achieved [2, 8, 28, 7].

back to the agents, in a way that maintains the desirable properties of the original VCG mechanism, including strategy-proofness, the non-deficit property, and (sometimes) individual rationality [4, 16, 23, 18, 17, 14, 13, 15, 19, 25].[3]

However, in some cases, even the best redistribution mechanism fails to redistribute a substantial amount of the VCG revenue. That is, even though the VCG redistribution mechanisms maximize *efficiency* (the sum of the agents' valuations), the total *welfare* (the sum of the agents' utilities, taking payments into account) can be very low (in fact, zero), as a result of poor redistribution. Still, the previously proposed VCG redistribution mechanisms are optimal in various senses—but only under the constraint that allocation is efficient.

In this paper, we consider the natural next step of allowing for inefficient allocation. It turns out that even though inefficient allocation reduces efficiency, it sometimes allows for greater redistributions, so that the net effect is an increase in the sum of the agents' utilities. Moulin [23] already provided an example where inefficient allocation can lead to better results, but left a more thorough investigation for future research. As we will see, the example mechanism that he proposed will turn out to be useful for us. In work subsequent to the conference version of our paper, de Clippel *et al.* [6] also study the problem of allocating inefficiently in order to increase the amount that is redistributed. We will discuss their work in more detail below.

In Section 2, we cover the preliminaries. In Section 3, we briefly review the worst-case optimal (WCO) VCG redistribution mechanism [16, 23]. For our setting and objective, the WCO mechanism is optimal among all strategy-proof mechanisms that are based on efficient allocation. We give an example where the WCO mechanism performs poorly, and its performance actually improves when we burn (throw away) one item.[4] This shows that inefficient allocation can indeed (sometimes) lead to better results. In Section 4, we characterize the *linear allocation mechanisms* for multi-unit auctions with unit demand, and propose an optimization technique for simultaneously finding a linear allocation mechanism and a payment redistribution rule, which together are optimal. With the help of the technique developed in Section 4, in Section 5, we analytically characterize several mechanisms that are guaranteed to perform well for all type profiles, which are based on burning units and partitioning the agents into groups. In Section 6, we compare the performances of the obtained mechanisms. Finally, in Section 7, we extend the definition of linear allocation mechanisms and the optimization technique to general multi-unit auctions.

---

[3]Recently, Tsuruta *et al.* [29] also studied how to redistribute while maintaining the *false-name-proofness* of the original VCG mechanism in single-item auctions. A mechanism is said to be false-name-proof if no agent can benefit by using multiple fake identities to participate in the mechanism. This paper does not consider false-name manipulation. This requirement restricts the space of redistribution mechanisms significantly, ruling out the standard strategy-proof redistribution mechanisms.

[4]We assume *free disposal* in this paper. That is, it is free to burn (throw away, or allocate to an external agent – but only to one about whose welfare the participating agents do not care one way or another) items. In some markets, payments are needed in order to get rid of extra resources. For such markets, mechanisms involving burning items perform worse or become infeasible.

2 Preliminaries

From this section to Section 6, we restrict our attention to multi-unit auctions with unit demand. In such auctions, multiple indistinguishable units are for sale at the same time, and each agent is interested in obtaining one and only one unit. In Section 7, we drop the unit demand restriction. To simplify the presentation, we defer everything beyond the unit demand setting to Section 7.

Let the number of units be m. Let the number of agents be n. We consider only cases where $n > m$; otherwise, it is clearly optimal to give every agent a unit and charge nothing. (This reasoning is not valid without the unit demand assumption, so in Section 7 we also consider cases where $m \geq n$.)

In the unit demand setting, an agent's type can be denoted by a single nonnegative value, which represents her valuation for winning (one unit). For the i-th agent, we denote her type by v_i. We focus on strategy-proof mechanisms in this paper. Therefore, we do not differentiate between "true types" and "reported types." Without losing generality, we assume that $v_1 \geq v_2 \geq \ldots \geq v_n \geq 0$. Vector $(v_1, v_2, \ldots, v_n)$ is called the *type profile*.

An allocation mechanism determines how we allocate the units, and how much each agent pays (or receives). We allow for randomized allocation mechanisms. In this paper, we assume that the agents are maximizing their expected (quasilinear) utilities, and random bits are not chosen adversarially. Please refer to Section 8 for more discussion on random bits. We use $a_i^{M,V}$ to denote agent i's probability of winning (one unit) under allocation mechanism M and for type profile V. When M and V are clear from the context, we simply use a_i to denote agent i's probability of winning. We denote the overall allocation by $(a_1, a_2, \ldots, a_n)$. Since there are only m units, we must have $\sum_{i=1}^{n} a_i \leq m$. The set of all feasible allocations is $\{(a_1, a_2, \ldots, a_n) | 0 \leq a_1, a_2, \ldots, a_n \leq 1, \sum_{i=1}^{n} a_i \leq m\}$.[5] Similarly, we use $t_i^{M,V}$ to denote agent i's payment (expected payment if the allocation mechanism is randomized) under allocation mechanism M and for type profile V. When M and V are clear from the context, we simply use t_i to denote agent i's payment. For example, under the VCG mechanism, for a given type profile $(v_1, v_2, \ldots, v_n)$, if $i \leq m$, then $a_i = 1$ and $t_i = v_{m+1}$, and if $i > m$, then $a_i = 0$ and $t_i = 0$. (Under the VCG mechanism, the agents with the m highest types each win one unit and each pay the value of the $(m+1)$-th highest type.) As a second example, under the mechanism in which the win-

---

[5]The reader may be concerned that the resulting mechanism may nevertheless not be feasible, because all the constraint implies is that the *expected* number of allocated items does not exceed m. In this footnote, we show that if $\sum_{i=1}^{n} a_i \leq m$, then $(a_1, a_2, \ldots, a_n)$ is feasible (able to be achieved using a randomized mechanism that allocates at most m units under any random realization). We first consider the case where $\sum_{i=1}^{n} a_i = m$. In this case, $(a_1, a_2, \ldots, a_n)$ is a convex combination of binary allocations, where a binary allocation is an allocation with m coordinates equal to 1 and the rest equal to 0. Every binary allocation is feasible. Therefore, any convex combination $(a_1, a_2, \ldots, a_n)$ is also feasible because it is a distribution of feasible allocations. If $\sum_{i=1}^{n} a_i < m$, then there exist $(a_1', a_2', \ldots, a_n')$ where $a_i \leq a_i' \leq 1$ and $\sum_{i=1}^{n} a_i' = m$. By the preceding, we know that $(a_1', a_2', \ldots, a_n')$ is feasible. We use $(a_1', a_2', \ldots, a_n')$, and we remove an agent from the allocation (not allowing her to win) with probability $1 - \frac{a_i}{a_i'}$. As a result, $(a_1, a_2, \ldots, a_n)$ is achieved, and therefore it is feasible.

ners are picked uniformly at random (without payments), for all i, $a_i = \frac{m}{n}$ and $t_i = 0$. We assume that an agent wishes to maximize her *expected* utility. That is, agent i wishes to maximize $a_i v_i - t_i$.

Given a strategy-proof allocation mechanism, we can use the following approach to attempt to increase the agents' utilities while maintaining strategy-proofness: redistribute (pay) to each agent an amount that is independent of her own type. When an agent's redistribution is independent of her own type, no matter what she reports, she receives the same amount of redistribution. This implies that the redistribution is irrelevant to the agent from the perspective of strategic reporting; and hence, if the allocation mechanism without redistribution is strategy-proof, then so is the mechanism with redistribution. In this paper, we use "allocation mechanism" to refer to the mechanism before redistribution – for example, in a VCG redistribution mechanism, the VCG mechanism is the allocation mechanism, whereas the complete mechanism also includes the redistributions. We use the function R to denote the redistribution rule. $R(v_{-i})$ is the amount of redistribution received by agent i, where v_{-i} are the types of the agents other than i.[6] Not all redistribution rules are feasible. We require that the complete mechanism (after redistribution) satisfies two properties. First, we require that the complete mechanism satisfies the *non-deficit* property: the total redistribution should never exceed the total payment collected by the allocation mechanism ($\sum_{i=1}^{n} R(v_{-i}) \leq \sum_{i=1}^{n} t_i$). Otherwise, we need an external subsidy to run the mechanism. We also require that the complete mechanism is *individually rational*: every agent's expected utility after redistribution must be nonnegative ($a_i v_i - t_i + R(v_{-i}) \geq 0$).

With redistribution, the agents' welfare is $\sum_{i=1}^{n} a_i v_i - \sum_{i=1}^{n} t_i + \sum_{i=1}^{n} R(v_{-i})$. Our goal is to find a complete mechanism (a strategy-proof allocation mechanism M and a feasible redistribution function R) that maximizes this expression. However, because the value of the above expression depends on the type profile, it could be that one complete mechanism (M, R) maximizes the expression for some type profiles, while another complete mechanism (M, R) maximizes the expression for some other type profiles. In this paper, we pursue a worst-case analysis approach. Such approaches have been commonly used in the literature on redistribution mechanisms [16, 23, 13, 14, 15, 19, 25][7] as well as the literature on digital goods auctions [1, 9, 20]. Specifically, consider the *first-best mechanism* that always *correctly* identifies the agents with the m highest true valuations, without asking for their types, and allocates the units to these agents at no charge. Clearly this mechanism obtains the largest welfare that we could hope for without external subsidy. Of course, the first-best mecha-

---

[6]Since the agents' type spaces are identical, and we allow for randomized mechanisms, for our worst-case objective, it is without loss of generality to consider only anonymous mechanisms: given a nonanonymous mechanism, we can always modify it by permuting the roles of the agents, and then run each permuted mechanism with equal probability. That is, it is without loss of generality to only consider redistribution rules under which all agents share the same redistribution function.

[7]However, other objectives have been pursued as well, for example, maximizing the expected redistribution with respect to a prior [17], or making sure that there is no feasible mechanism that always redistributes at least as much and sometimes more [18].

nism is not strategy-proof, so we only use it as a benchmark. Our objective is to design strategy-proof mechanisms that are *competitive* against the first-best mechanism. We say a complete mechanism (M, R) is α-*competitive* against the first-best mechanism if the agents' welfare under (M, R) is at least $\alpha \sum_{i=1}^{m} v_i$, for all type profiles. ($\sum_{i=1}^{m} v_i$ is the agents' welfare under the first-best mechanism.) Our objective is to find the complete mechanism (M, R) that is the most competitive, that is, that maximizes α, under the constraints that (M, R) must satisfy strategy-proofness, individual rationality, and the non-deficit property.[8]

3 VCG redistribution mechanisms

In this section, we first review the worst-case optimal (WCO) VCG redistribution mechanism. For our setting and objective, the WCO mechanism is optimal among all strategy-proof mechanisms that are based on efficient allocation. Then, to motivate the rest of the paper, we construct an example mechanism that allocates inefficiently and is more competitive than the WCO mechanism. This shows that inefficient allocation can indeed (sometimes) lead to higher competitive ratios.

In previous work [16], we characterized the worst-case optimal VCG redistribution mechanism, which maximizes the fraction of VCG payment that is redistributed in the worst case, among all VCG redistribution mechanisms that are individually rational and satisfy the non-deficit property. The WCO mechanism is defined for all multi-unit auctions where each agent has nonincreasing marginal values for units.[9] In the more restricted unit demand setting, the WCO mechanism uses the following redistribution function:

$$R(v_{-j}) = \sum_{i=m+1}^{n-1} c_i v_{-j,i}$$

$$c_i = \frac{(-1)^{i+m-1}(n-m)\binom{n-1}{m-1}}{i\sum_{j=m}^{n-1}\binom{n-1}{j}} \frac{1}{\binom{n-1}{i}} \sum_{j=i}^{n-1} \binom{n-1}{j} \quad \text{for } i = m+1, \ldots, n-1$$

Here, $v_{-j,i}$ is the i-th highest type among v_{-j}.

Around the same time, Moulin [23] independently characterized the same mechanism based on a different objective.[10] In fact, Moulin's characterization

---

[8]It should be noted that this paper allows randomized mechanisms. A randomized mechanism is α-competitive if the agents' welfare under the mechanism is at least α times the agents' welfare under the first-best mechanism *in expectation over the random choices made by the mechanism* (but for any valuation profile). Similarly, for randomized mechanisms, strategy-proofness and individual rationality also only need to be satisfied in expectation over the random choices made by the mechanism (but for any valuation profile).

[9]WCO was later generalized to heterogeneous-item auctions with unit demand [13, 15]. The generalized version was proven to be still optimal.

[10]Moulin's mechanism is defined only for the unit demand setting. Moulin tries to minimize the worst-case ratio between the VCG payment that fails to be redistributed and the total efficiency, also among all VCG redistribution mechanisms that are individually rational and satisfy the non-deficit property. Moulin's objective and WCO's objective result in the same mechanism only if individual rationality is required. If individual rationality is not required, then WCO's objective still results in the WCO mechanism, but Moulin's does not.

of WCO implies that the WCO mechanism has the highest competitive ratio (in the sense of this paper) among all VCG redistribution mechanisms with the required properties.

Claim 1. *[23]: The WCO mechanism has the highest competitive ratio against the first-best mechanism, among all VCG redistribution mechanisms that are individually rational and satisfy the non-deficit property.*

It should be noted that, by the result of Holmstrom [21], in multi-unit auctions with unit demand, all efficient and strategy-proof mechanisms are VCG redistribution mechanisms. That is, the above claim essentially says that WCO has the highest competitive ratio among all efficient mechanisms that are strategy-proof, individually rational, and satisfy the non-deficit property.

Given the number of agents n and the number of units m ($n > m$), the WCO mechanism's competitive ratio $\alpha_{WCO}(n, m)$ is characterized by the following equation:

$$\alpha_{WCO}(n, m) = 1 - \frac{\binom{n-1}{m}}{\sum_{j=m}^{n-1} \binom{n-1}{j}}$$

When $n = 3$ and $m = 2$, the WCO mechanism (the optimal efficient mechanism) is not competitive at all: $\alpha_{WCO}(3, 2) = 0$. In contrast, the following simple mechanism that allocates inefficiently is somewhat competitive:

- Burn (throw away) one unit.

- Allocate the remaining unit according to the WCO mechanism for $n = 3$ and $m = 1$.

The new mechanism satisfies all the required properties, because it is equivalent to the WCO mechanism for $n = 3$ and $m = 1$. It is not efficient because one unit is burned. Since $\alpha_{WCO}(3, 1) = \frac{1}{3}$, the new mechanism is $\frac{1}{3}$-competitive against the first-best mechanism for one unit ($m = 1$). That is, the new mechanism guarantees a welfare of $\frac{1}{3}v_1$ for any type profile. Since $v_1 \geq v_2$, it also guarantees a total utility of $\frac{1}{6}(v_1 + v_2)$ for all type profiles. Hence, the competitive ratio of the new mechanism against the first-best mechanism for two units ($m = 2$) is at least $\frac{1}{6}$. That is, this new mechanism has a higher competitive ratio. Essentially, if we keep both items, then no redistribution is possible. On the other hand, if we throw away one item, then we are able to redistribute some payments, which results in a higher competitive ratio. So, ironically, in some cases, the agents are happier if one unit is burned. Motivated by this example, in the rest of the paper, we study mechanisms that allocate inefficiently (and in Section 5.1, we specifically study mechanisms that are based on burning units).

4   Linear allocation mechanisms

In this section, we define and characterize a family of mechanisms that we call *linear allocation mechanisms*.[11] We then propose an optimization technique for simultaneously finding a linear allocation mechanism and a redistribution rule, which together have the highest competitive ratio among all mechanisms whose allocation part is linear.[12]

Definition 1. An allocation mechanism is *linear* if the following two conditions are satisfied:

- An agent's probability of winning depends only on her rank. That is, under a linear allocation mechanism, the a_i are constants (a_i only depends on the rank i, but not the value of v_i). (We recall that agent i is the agent with the i-th highest type, and a_i is agent i's probability of winning.) That is, a linear allocation mechanism corresponds to a constant allocation vector $(a_1, a_2, \ldots, a_n)$.

- An agent's payment is always 0 if her type is 0.

Example 1. The VCG mechanism is linear, for the following reasons. In the VCG mechanism, the agents with the highest m types each win one unit. That is, for all type profiles, we have $a_i = 1$ for $i \leq m$ and $a_i = 0$ for $i > m$. Under the VCG mechanism, when an agent's type is 0, her payment is 0.

Example 2. The *random allocation mechanism* in which the winners are picked uniformly at random (without replacement), and there are no payments, is linear, for the following reasons. Under this mechanism, for all type profiles, we have $a_i = m/n$ for all i, and every agent pays 0.

We are only interested in linear allocation mechanisms that are strategy-proof. The next claim gives a standard type of monotonicity characterization of all strategy-proof linear allocation mechanisms.

Claim 2. *There is a one-to-one correspondence between strategy-proof linear allocation mechanisms and the following set of allocation vectors:*

$$\{(a_1, a_2, \ldots, a_n) | 1 \geq a_1 \geq a_2 \geq \ldots \geq a_n \geq 0, \sum_{i=1}^{n} a_i \leq m\}$$

That is, a strategy-proof linear allocation mechanism is characterized by its corresponding allocation vector $(a_1, a_2, \ldots, a_n)$ where $1 \geq a_1 \geq a_2 \geq \ldots \geq a_n \geq 0$.

---

[11]The linear allocation mechanisms are randomized mechanisms characterized by n parameters. In de Clippel *et al.* [6], the authors also studied multi-unit auctions with unit demand, and proposed a family of *nonlinear* and *deterministic* mechanisms that are characterized by exactly two parameters. Optimization within their mechanism family is easy because there are only two parameters. The optimal linear allocation mechanism performs better, but it is not deterministic.

[12]It is not clear whether there exist nonlinear allocation mechanisms with even higher competitive ratios. Nevertheless, as we will show, the optimal linear allocation mechanisms generally achieve competitive ratios that are close to 1 (we recall that 1 is the trivial upper bound on competitive ratio of any mechanism, linear or nonlinear).

From now on, we will simply use allocation vector $(a_1, a_2, \ldots, a_n)$ to refer to the (only) corresponding strategy-proof linear allocation mechanism. We will only consider strategy-proof linear allocation mechanisms.

Under linear allocation mechanism $(a_1, a_2, \ldots, a_n)$, the mechanism's efficiency equals $\sum_{i=1}^{n} a_i v_i$. Agent i's VCG-based payment equals the following:

$$\sum_{j=1}^{i-1} a_j v_j + \sum_{j=i+1}^{n} a_{j-1} v_j - \sum_{j=1}^{i-1} a_j v_j - \sum_{j=i+1}^{n} a_j v_j = \sum_{j=i+1}^{n} (a_{j-1} - a_j) v_j$$

The agents' total payment equals

$$\sum_{i=1}^{n} \sum_{j=i+1}^{n} (a_{j-1} - a_j) v_j$$

Under a given linear allocation mechanism, the mechanism's efficiency and the agents' total payment are both linear functions of the types, and the coefficients are constants. This is why we call these *linear* allocation mechanisms.

Next, we present more example linear allocation mechanisms.

Example 3. Let us consider the case where $m > 1$. The mechanism that burns one unit and allocates the remaining units efficiently is linear. It corresponds to the allocation vector $(\underbrace{1, 1, \ldots, 1}_{m-1}, 0, \ldots, 0)$. (The agents' payments are VCG-based.)

Example 4. The mechanism that excludes one agent uniformly at random and allocates the units among the remaining agents efficiently is linear. Agents 1 to m will win unless excluded, so $a_i = \frac{n-1}{n}$ for $i = 1, 2, \ldots, m$. Agent $m + 1$ will win only when one of the first m agents gets excluded, so $a_{m+1} = \frac{m}{n}$. The other agents never win. The mechanism corresponds to the allocation vector $(\underbrace{\frac{n-1}{n}, \frac{n-1}{n}, \ldots, \frac{n-1}{n}}_{m}, \frac{m}{n}, 0, \ldots, 0)$. (The agents' payments are VCG-based.)

Example 5. It is easy to see that the set $\{(a_1, a_2, \ldots, a_n) | 1 \geq a_1 \geq a_2 \geq \ldots \geq a_n \geq 0, \sum_{i=1}^{n} a_i \leq m\}$ is convex. That is, any probability mixture over linear allocation mechanisms is also linear.

In what follows, we introduce an optimization technique for simultaneously finding a linear allocation mechanism and a payment redistribution rule, which together have the highest competitive ratio among all mechanisms whose allocation part is linear.

Since there is a one-to-one correspondence between linear allocation mechanisms and allocation vectors in $\{(a_1, a_2, \ldots, a_n) | 1 \geq a_1 \geq a_2 \geq \ldots \geq a_n \geq 0, \sum_{i=1}^{n} a_i \leq m\}$, essentially, we are solving for the optimal values of the a_i, as well as an optimal redistribution function R, so that the linear allocation mechanism $(a_1, a_2, \ldots, a_n)$ and R together have the highest competitive ratio. This comes down to the following optimization model:

Variable function: $R : [0, \infty)^{n-1} \to \mathbb{R}$
Variables: $\alpha, a_1, a_2, \ldots, a_n$
Maximize α
Subject to:
$1 \geq a_1 \geq a_2 \geq \ldots \geq a_n \geq 0$
$\sum_{i=1}^{n} a_i \leq m$
For every type profile $v_1 \geq v_2 \geq \ldots \geq v_n \geq 0$
$\quad R(v_{-n}) \geq 0$ (individual rationality)
$\quad \sum_{i=1}^{n} \sum_{j=i+1}^{n} (a_{j-1} - a_j) v_j \geq \sum_{i=1}^{n} R(v_{-i})$ (non-deficit)
$\quad \sum_{i=1}^{n} a_i v_i - \sum_{i=1}^{n} \sum_{j=i+1}^{n} (a_{j-1} - a_j) v_j + \sum_{i=1}^{n} R(v_{-i}) \geq \alpha \sum_{i=1}^{m} v_i$
$\quad$ (competitive ratio constraint)

The non-deficit constraint and the competitive ratio constraint are both self-explanatory. (We recall that $\sum_{i=1}^{n} \sum_{j=i+1}^{n} (a_{j-1} - a_j) v_j$ is the agents' total payment, $\sum_{i=1}^{n} a_i v_i$ is the mechanism's efficiency, and $\sum_{i=1}^{m} v_i$ is the welfare under the first-best mechanism.) What remains to be explained is why the individual rationality constraint is equivalent to $R(v_{-n}) \geq 0$. First, $R(v_{-n}) \geq 0$ is necessary for individual rationality, because we need to make sure that every agent's final utility is nonnegative for every type profile. When agent n's type is 0, she pays 0 under any linear allocation mechanism. Because she does not care for receiving any unit, her final utility is just her redistribution, which is $R(v_{-n})$. Hence, $R(v_{-n}) \geq 0$ is necessary for individual rationality. On the other hand, when $R(v_{-n})$ is nonnegative for every v_{-n}, we have that for every collection of $n-1$ types, denoted by S, $R(S) \geq 0$. That is, $R(v_{-i}) \geq 0$ for all i. That is, redistribution makes every agent (weakly) better off. Even without redistribution, every agent's utility is at least 0. (An agent can always report 0 and pay 0.) Therefore, $R(v_{-n}) \geq 0$ is also sufficient for individual rationality.

The above formulation is a functional optimization problem, which may appear difficult to solve. In what follows, we present a linear program formulation. Based on the solution to this linear program, we can construct an optimal solution to the above functional optimization problem.

If we only consider the type profiles in which the v_i are binary, then the above functional optimization problem turns into a value optimization problem, as follows. When we only consider binary type profiles, we do not have to consider the values of R for other profiles. We use r_i ($i = 0, 1, \ldots, n-1$) to denote $R(\underbrace{1, \ldots, 1}_{i}, \underbrace{0, \ldots, 0}_{n-1-i})$: the redistribution received by an agent when i other agents report 1, and the remaining $n - 1 - i$ agents report 0. In the original functional optimization problem, we are optimizing over the a_i and the function R. When we only consider binary type profiles, we only need to optimize over the a_i and the r_i, which corresponds to the following value optimization problem:

> **Variables:** $\alpha, a_1, a_2, \ldots, a_n, r_0, r_1, \ldots, r_{n-1}$
> **Maximize** α
> **Subject to:**
> $1 \geq a_1 \geq a_2 \geq \ldots \geq a_n \geq 0$
> $\sum_{i=1}^{n} a_i \leq m$
> For every type profile $v_1 \geq v_2 \geq \ldots \geq v_n \geq 0$, where the v_i are binary
> $\quad R(v_{-n}) \geq 0$ (individual rationality)
> $\quad \sum_{i=1}^{n} \sum_{j=i+1}^{n} (a_{j-1} - a_j) v_j \geq \sum_{i=1}^{n} R(v_{-i})$ (non-deficit)
> $\quad \sum_{i=1}^{n} a_i v_i - \sum_{i=1}^{n} \sum_{j=i+1}^{n} (a_{j-1} - a_j) v_j + \sum_{i=1}^{n} R(v_{-i}) \geq \alpha \sum_{i=1}^{m} v_i$
> $\quad$ (competitive ratio constraint)

The optimal objective value of the above value optimization problem is greater than or equal to the optimal objective value of the original functional optimization problem, because there are fewer constraints (here, we are only enforcing the constraints on binary type profiles).

We now show that this is indeed a linear program. The objective is clearly linear in the variables. The two constraints involving only the a_i are clearly linear inequalities. The individual rationality constraint is just "for all i from 0 to $n-1$, $r_i \geq 0$." The non-deficit constraint is also equivalent to a set of linear inequalities, as follows. When $v_i = 0$ for all i, the non-deficit constraint is $0 \geq n r_0$.[13] When $v_1 = 1$ and $v_i = 0$ for all $i > 1$, the non-deficit constraint is $0 \geq (n-1)r_1 + r_0$. When $v_1 = v_2 = 1$ and $v_i = 0$ for all $i > 2$, the non-deficit constraint is $(a_1 - a_2)v_2 = a_1 - a_2 \geq (n-2)r_2 + 2r_1$. Similarly, for any other binary type profile, the non-deficit constraint also corresponds to a linear inequality. In total, the v_i can take $n+1$ sets of values (from all 0s to all 1s). Therefore, the non-deficit constraint is equivalent to $n+1$ linear inequalities. By similar reasoning, the competitive ratio constraint is also equivalent to $n+1$ linear inequalities. We conclude that the above value optimization problem is indeed a linear program.

We use the a_i^*, the r_i^*, and α^* to denote an optimal solution to the linear program. As we noted, α^*, the optimal objective value of the linear program, is greater than or equal to the optimal objective value of the original functional optimization problem. That is, α^* is an upper bound on the competitive ratio of any complete mechanism whose allocation part is linear. In the following theorem, we show that this upper bound can be achieved: we construct a complete mechanism whose allocation part is linear, and prove that this mechanism has a competitive ratio of α^*. It follows immediately that the mechanism we construct is optimal among all complete mechanisms whose allocation part is linear.

Theorem 1. *For all i and all v_{-i}, let $R^*(v_{-i}) = \sum_{j=1}^{n-1} (r_j^* - r_{j-1}^*) v_{-i,j}$. (We recall that $v_{-i,j}$ is the j-th highest type among the agents other than i.) Complete mechanism $((a_1^*, a_2^*, \ldots, a_n^*), R^*)$ is α^*-competitive. That is, this mechanism has the highest competitive ratio among all complete mechanisms whose allocation*

---

[13]Combined with the individual rationality constraint, we have that r_0 should always be 0.

part is linear.[14]

According to Footnote 13, we have $r_0^* = 0$. We note that when the v_i are binary, from the perspective of an agent, if all the other agents report 0, then her redistribution equals $\sum_{j=1}^{n-1}(r_j^* - r_{j-1}^*)0 = 0 = r_0^*$. If among the other agents, k $(k > 0)$ agents report 1, and the remaining agents report 0, then her redistribution equals $\sum_{j=1}^{k}(r_j^* - r_{j-1}^*) = r_k^*$. That is, in a sense, R^* is a linear interpolation of the r_i^*. Before proving the theorem, we give the following lemma. A similar lemma appeared in Guo and Conitzer [16].

Lemma 1. *When the c_i do not depend on the x_i, the following two systems of inequalities are equivalent:*

(a) $c_1 x_1 + c_2 x_2 + \ldots + c_s x_s \geq 0$ *for all* $x_1 \geq x_2 \geq \ldots \geq x_s \geq 0$.

(b) $c_1 x_1 + c_2 x_2 + \ldots + c_s x_s \geq 0$ *for all* $x_1 \geq x_2 \geq \ldots \geq x_s$, *where the x_i are binary.*

Proof. (a) $\Rightarrow$ (b) is trivial. We now prove (b) $\Rightarrow$ (a). (b) implies that $\sum_{i=1}^{j} c_i \geq 0$ for all $1 \leq j \leq s$. We have $c_1 x_1 + c_2 x_2 + \ldots + c_s x_s = \sum_{j=1}^{s-1}(\sum_{i=1}^{j} c_i)(x_j - x_{j+1}) + (\sum_{i=1}^{s} c_i)x_s$. For all $x_1 \geq x_2 \geq \ldots \geq x_s \geq 0$, each term of the above expression is nonnegative, hence the whole expression is nonnegative. So (b) $\Rightarrow$ (a). $\square$

We are now ready to prove the theorem.

Proof. We only need to prove that α^*, the a_i^*, and the function R^* form a feasible solution to the original functional optimization problem. If they form a feasible solution, then $((a_1^*, a_2^*, \ldots, a_n^*), R^*)$ is α^*-competitive, implying that it is optimal. That is, we need to prove that the following constraints hold (it is obvious that the constraints involving only the a_i hold):

For every type profile $v_1 \geq v_2 \geq \ldots \geq v_n \geq 0$,
$R^*(v_{-n}) \geq 0$ (individual rationality)
$\sum_{i=1}^{n} \sum_{j=i+1}^{n}(a_{j-1}^* - a_j^*)v_j \geq \sum_{i=1}^{n} R^*(v_{-i})$ (non-deficit)
$\sum_{i=1}^{n} a_i^* v_i - \sum_{i=1}^{n} \sum_{j=i+1}^{n}(a_{j-1}^* - a_j^*)v_j + \sum_{i=1}^{n} R^*(v_{-i}) \geq \alpha^* \sum_{i=1}^{m} v_i$
(competitive ratio constraint)

We notice that all the above inequalities are linear inequalities of the v_i, which must hold for all $v_1 \geq v_2 \geq \ldots \geq v_n \geq 0$, where the coefficients only depend on the r_i^*, the a_i^*, and α^*. For example, the individual rationality constraint can be rewritten as "for every type profile $v_1 \geq v_2 \geq \ldots \geq v_n \geq 0$, $\sum_{j=1}^{n-1}(r_j^* - r_{j-1}^*)v_j \geq 0$." Similarly, the non-deficit constraint is "for every type profile $v_1 \geq v_2 \geq \ldots \geq v_n \geq 0$, $\sum_{i=1}^{n} \sum_{j=i+1}^{n}(a_{j-1}^* - a_j^*)v_j \geq \sum_{i=1}^{n} \sum_{j=1}^{n-1}(r_j^* - r_{j-1}^*)v_{-i,j}$. Here, the left-hand side of $\geq$ is a linear function of the v_i, where the

---

[14]By generalizing our theorem, Naroditskiy *et al.* [26] proposed a technique for finding optimal payments for given non-linear allocation mechanisms in single-parameter domains.

linear coefficients only depend on the a_i^*. We have that if $j < i$, then $v_{-i,j} = v_j$, otherwise $v_{-i,j} = v_{j+1}$. The right-hand side of $\geq$ is also a linear function of the v_i, where the linear coefficients only depend on the r_i^*. Similarly, the competitive ratio constraint is "for every type profile $v_1 \geq v_2 \geq \ldots \geq v_n \geq 0$, $\sum_{i=1}^n a_i^* v_i - \sum_{i=1}^n \sum_{j=i+1}^n (a_{j-1}^* - a_j^*) v_j + \sum_{i=1}^n \sum_{j=1}^{n-1} (r_j^* - r_{j-1}^*) v_{-i,j} \geq \alpha^* \sum_{i=1}^m v_i$. That is, the competitive ratio constraint is also a linear inequality of the v_i, which must hold for all $v_1 \geq v_2 \geq \ldots \geq v_n \geq 0$, and the linear coefficients only depend on the r_i^*, the a_i^*, and α^*. If we apply Lemma 1, then we only need to prove that these constraints hold when the v_i are binary. (When the v_i are binary, R^* reduces to the r_i^*.) But, we already know that linear allocation mechanism $(a_1^*, a_2^*, \ldots, a_n^*)$, combined with the redistribution rule characterized by the r_i^*, satisfies all the constraints when the v_i are binary. □

Using Theorem 1, for different numbers of agents n and different numbers of units m, we solve for the most competitive complete mechanisms whose allocation part is linear. The obtained mechanisms' competitive ratios will be presented in Section 6.

5 Analytical competitive mechanisms

One drawback of the technique proposed in the previous section is that we need to solve a linear program for every n and m, and we do not have an analytical characterization of the resulting mechanisms. In this section, we study two restricted families of linear allocation mechanisms. One family is based on burning units. The other is based on partitioning the agents into groups. We will use the technique introduced in the previous section to optimize within both families (adding more constraints on the a_i, as we are looking at more restricted families of linear allocation mechanisms). By observing the patterns of the obtained mechanisms, we have analytically characterized several competitive mechanisms. In Section 6, we compare the competitive ratio of the optimal mechanism that is based on linear allocation, and the competitive ratios of the analytical mechanisms.

5.1 Burning units

In this subsection, we study linear allocation mechanisms that are based on (sometimes) burning units. We start with the characterization. First, we can construct m linear allocation mechanisms that are based on burning a deterministic number of units. Let M_i $(i = 1, \ldots, m)$ be the allocation mechanism in which $m - i$ units are burned, and the remaining i units are allocated efficiently according to the VCG mechanism. M_m is just the original VCG mechanism. We note that it makes no sense to burn all units, hence $i > 0$. We call the M_i *deterministic burning allocation mechanisms*. M_i is the linear allocation mechanism $(\underbrace{1, 1, \ldots, 1}_{i}, 0, 0, \ldots, 0)$. We can also construct mechanisms in which
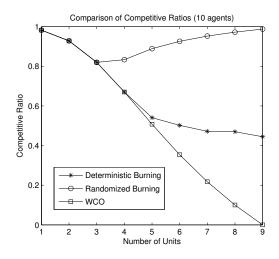
Figure 1: A comparison of α_D^*, α_R^* and α_{WCO}^*.

a random number of units is burned, by randomizing over the M_i. Let M be a mixture of the M_i, where mechanism M_i is chosen with probability p_i. That is, M is the mechanism in which with probability p_i, exactly $m - i$ units are burned. We call such mixtures over the M_i *randomized burning alloca-tion mechanisms.* Under M, for $i = 1, 2, \ldots, m$, agent i wins with probability $\sum_{j=i}^m p_j$, and the other agents never win. M is the linear allocation mechanism $(\sum_{i=1}^m p_i, \sum_{i=2}^m p_i, \sum_{i=3}^m p_i, \ldots, \sum_{i=m}^m p_i, 0, 0, \ldots, 0)$. The family of linear allo-cation mechanisms that are based on the idea of burning a random number of units is then the linear allocation mechanisms whose allocation vectors are in $\{(a_1, a_2, \ldots, a_n) | a_{m+1} = a_{m+2} = \ldots = a_n = 0, 1 = a_1 \geq a_2 \geq \ldots \geq a_m \geq 0\}$.

Using Theorem 1, we can find the most competitive burning-based mecha-nism. To solve for the most competitive randomized burning-based mechanism, we need to add two more constraints to the linear program: $a_{m+1} = 0$ and $a_1 = 1$. To solve for the most competitive deterministic burning-based mecha-nism, we need to add n more constraints to the linear program and turn it into a mixed integer program: for all i, a_i is binary. We will present more results in Section 6. Here, for the case of $n = 10$, $m = 1, \ldots, 9$, we compare the values of α_{WCO}^* (competitive ratio of the WCO mechanism, which burns nothing), α_D^* (competitive ratio of the optimal deterministic burning-based mechanism) and α_R^* (competitive ratio of the optimal randomized burning-based mechanism) in Figure 1. When m is small, the three values are the same. As m gets large, the value of α_{WCO}^* decreases all the way to 0; the value of α_D^* also decreases but it gets stable when its value goes down to around 0.5; the value of α_R^* first decreases, but then increases, and at the end almost reaches 1.

Of course, $\alpha_{WCO}^* \leq \alpha_D^* \leq \alpha_R^*$; it turns out that all of these inequalities are sometimes strict. Therefore, in general, we need to burn a random number of units to get the most competitive burning-based mechanism.

14

While we can use Theorem 1 to numerically solve for the most competitive burning-based mechanism for given n and m, it is still desirable to have a general analytical characterization for all n and m. The following claim analytically characterizes a burning-based mechanism for all n and m, and gives its competitive ratio. We obtained this analytical competitive mechanism by observing the patterns of the (numerically obtained) optimal burning-based mechanisms for all $n < 10$. We do not know whether this mechanism is still optimal for larger n. However, we do know that this mechanism is feasible and has the specified competitive ratio for all n and m.

The intuition behind the proposed mechanism is as follows. We aim to redistribute more by allocating inefficiently, in order to achieve a net benefit. It turns out that we can guarantee 100% redistribution by burning at most one unit. For any type profile, the total VCG payment is mv_{m+1} when nothing is burned, and it is $(m-1)v_m$ when one unit is burned. If we burn one unit with probability $\frac{n-m}{n-1}$, then in expectation, the total payment collected by the allocation mechanism is $\frac{n-m}{n-1}(m-1)v_m + (1 - \frac{n-m}{n-1})mv_{m+1} = \frac{m-1}{n-1}(mv_{m+1} + (n-m)v_m)$. Throughout this paper, for randomized mechanisms, we simply charge agents the expected payments, instead of charging a different payment depending on the realization of the mechanism (*e.g.*, whether a unit was burned). The agents' incentives will not be affected, because agents are assumed to be maximizing their expected (quasilinear) utility. By redistributing to every agent $\frac{m-1}{n-1}$ times the m-th highest type among the others, the total redistribution equals exactly the total payment. When m is not too small, it is an acceptable sacrifice of allocative efficiency to burn at most one unit to achieve strong budget balance. When m is very small, we cannot afford to burn any units, so it is better to use the original worst-case optimal VCG redistribution mechanism (nothing is burned).

Claim 3. *For any n and m, we can achieve the following competitive ratio with a burning-based mechanism:*

$$\max\{1 - \frac{\binom{n-1}{m}}{\sum_{j=m}^{n-1}\binom{n-1}{j}}, \frac{mn-n}{mn-m}\}$$

If the first expression is greater (or equal), then the mechanism achieving the above ratio is the worst-case optimal VCG redistribution mechanism (nothing is burned).

If the second expression is greater, then the mechanism achieving the above ratio is the following:

- *Burn (throw away) one unit with probability $\frac{n-m}{n-1}$. The remaining units are allocated efficiently. Every agent pays the VCG-based payment. (This is linear allocation mechanism $(1, 1, \ldots, 1, \frac{m-1}{n-1}, 0, 0, \ldots, 0)$.)*

- *Every agent receives a redistribution of $\frac{m-1}{n-1}$ times the m-th highest type among types other than this agent's own type.*

15

5.2 Partitioning units and agents

In this subsection, we study linear allocation mechanisms that are based on partitioning the agents into groups. This is an idea that has previously been proven effective in mechanism design [3, 9]. Based on this idea, we characterize a family of strongly budget balanced mechanisms. Some of the mechanisms in this class have been proposed previously [7, 23]. We focus on finding the most competitive mechanism in this family. Because all the mechanisms in this family are strongly budget balanced, no redistributions are ever needed.

We start with two example mechanisms. They are both based on excluding one agent from the set of all agents. The first one is due to Moulin [23], and the second one is due to Faltings [7].

Example Mechanism 1

- Exclude one agent uniformly at random.

- Assign one unit to the excluded agent at no charge.

- The remaining units are allocated to the remaining agents according to the VCG mechanism.

- Transfer all the VCG revenue to the excluded agent.

Example Mechanism 2

- Exclude one agent uniformly at random.

- Units are allocated to the remaining agents according to the VCG mechanism.

- Transfer all the VCG revenue to the excluded agent.

Both example mechanisms are strategy-proof, individually rational, and strongly budget balanced. In the first mechanism, one agent is excluded and assigned one unit. In the second mechanism, one agent is excluded and assigned zero units. It should be noted that if we omit the last step in Example Mechanism 2, that is, if we do not transfer all the VCG revenue to the excluded agent, then the mechanism becomes exactly the linear allocation mechanism in Example 4. That is, it becomes linear allocation mechanism $(\underbrace{\frac{n-1}{n}, \frac{n-1}{n}, \ldots, \frac{n-1}{n}}_{m}, \frac{m}{n}, 0, \ldots, 0)$. Similarly, if we omit the last step in Example Mechanism 1, then the mechanism becomes linear allocation mechanism $(\underbrace{1, 1, \ldots, 1}_{m-1}, \frac{m}{n}, \frac{1}{n}, \frac{1}{n}, \ldots, \frac{1}{n})$. (The first $m-1$ agents always win, excluded or not. The m-th agent wins only if she is excluded or one agent from 1 to $m-1$ is excluded. The other agents win only when excluded.)

We now introduce the family of partitioning-based mechanisms; this family generalizes both of the previous two mechanisms.

Definition 2. Given n and m, for $n_1 \in \{1, \ldots, \lfloor \frac{n}{2} \rfloor\}$, $m_1 \in \{0, \ldots, \min\{n_1, m\}\}$, we define the following mechanism:

- Pick n_1 agents to form one group, uniformly at random. The other $n - n_1$ agents form the second group.

- Allocate m_1 units among the first group, according to the VCG mechanism.

- Allocate the remaining $m - m_1$ units among the second group, according to the VCG mechanism.

- Transfer the VCG revenue from the first group to the second group, in any predetermined way.

- Transfer the VCG revenue from the second group to the first group, in any predetermined way.

We call this mechanism the (n_1, m_1)-*partition mechanism*.

We note that Example Mechanisms 1 and 2 are the $(1, 1)$-partition mechanism and the $(1, 0)$-partition mechanism, respectively. Under any partitioning-based mechanism, whether an agent wins depends on which group she is assigned to, and her rank within her group. Since groups are randomly formed, independent of the type profile, an agent's probability of winning depends only on her rank. If we do not transfer the VCG revenue between the groups, then an agent pays 0 when her type is 0. That is, without transferring the VCG revenue, every partitioning-based mechanism is linear.

Claim 4. *The partitioning-based mechanisms are strategy-proof, individually rational, and strongly budget balanced.*

Proof. Without transferring the VCG revenue, every agent is participating in a VCG mechanism, which is strategy-proof. For each agent, the transfer payment she receives depends only on the types in the other group of agents, hence it does not affect her incentives. Therefore, the mechanisms are strategy-proof. Similarly, without transferring the VCG revenue, every agent is participating in a VCG mechanism, which is individually rational. With transferring, the agents' utilities become higher or stay the same. Therefore, the mechanisms are individually rational. Finally, the strong budget balance property follows from the fact that the entire VCG revenue is transferred. $\square$

Since the partitioning-based mechanisms are strongly budget balanced, welfare equals efficiency for them. Hence, for our objective of finding the most competitive partitioning-based mechanism, we can completely ignore the VCG payments and the revenue transferring process. That is, for the analysis that follows, we pretend that there are no payments of any kind; when we use the mechanism, we add the VCG payments and transfers back to achieve strategy-proofness.

We can use the technique in Theorem 1 to solve for the most competitive partitioning-based mechanism (ignoring all payments and redistributions). However, we now present a simpler solution technique based on the special structure of the family of partitioning-based mechanisms. The following claim characterizes the competitive ratio of a given partitioning-based mechanism.

Claim 5. *Given n and m, the competitive ratio of the (n_1, m_1)-partition mechanism equals*

$$\frac{U_{M_{n_1, m_1}}(m)}{m}$$

Here, $U_{M_{n_1, m_1}}(m)$ is the expected efficiency (welfare) under the (n_1, m_1)-partition mechanism when m agents report 1 and the remaining agents report 0. This competitive ratio is equal to

$$\frac{\sum_{x \in X} \binom{n_1}{x}\binom{n - n_1}{m - x}(\min\{x, m_1\} + \min\{m - x, m - m_1\})}{m\binom{n}{m}}$$

Here, $X = \{x | x \in \{0, 1, \ldots, n_1\}, 0 \le m - x \le n - n_1\}$. We will call this competitive ratio $\alpha^*_{n_1, m_1}$.

So far, we have not considered mixtures over partitioning-based mechanisms. One might think that, by taking such mixtures, we can obtain more competitive mechanisms. However, the following claim rules out this idea.

Claim 6. *If M is a mixture over $M_1, M_2, \ldots, M_t$, where the M_i are partitioning-based mechanisms for different values of n_1 and m_1, and M_i is chosen with probability p_i, then there exists $1 \le j \le t$ so that M_j attains at least the competitive ratio of M.*

Proof. By the same argument as in Claim 5, the competitive ratio of M is at most $\frac{U_M(m)}{m}$. We have that $\frac{U_M(m)}{m} = \frac{\sum_{i=1}^{t} p_i U_{M_i}(m)}{m} \le \max_j \frac{U_{M_j}(m)}{m}$. But $\frac{U_{M_j}(m)}{m}$ is the competitive ratio for M_j by Claim 5. Hence, taking $j \in \arg\max_j \frac{U_{M_j}(m)}{m}$, M_j is at least as competitive as M. $\qquad\square$

By Claim 5, for given n and m, by maximizing

$$\frac{\sum_{x \in X} \binom{n_1}{x}\binom{n - n_1}{m - x}(\min\{x, m_1\} + \min\{m - x, m - m_1\})}{m\binom{n}{m}}$$

over n_1 and m_1, we obtain the optimal (n_1, m_1)-partition mechanism. This mechanism is also optimal among all mixtures of partitioning-based mechanisms by Claim 6. It is still desirable to have a general analytical characterization of the optimal n_1 and m_1 for all n and m. The following claim analytically characterizes three partitioning-based mechanisms for all n and m, and gives their corresponding competitive ratios. We identified these three mechanisms by observing the patterns of the (numerically obtained) optimal partitioning-based mechanisms for all $n < 10$. We do not know whether there are other mechanisms

18

that obtain higher competitive ratios for larger n. However, we do know that all three mechanisms are feasible and achieve the specified competitive ratios for all n and m.

The intuition behind the following mechanism is as follows. When m is small, given a random agent, most likely this agent does not belong to the m highest agents. Therefore, when m is small, it is an acceptable sacrifice of allocative efficiency to randomly remove one agent and deny her any units (run the $(1,0)$-partition mechanism). By doing so, we achieve strong budget balance, as well as efficiency for the rest of the agents. Similarly, when m is large relative to n, given a random agent, most likely this agent does belong to the m highest agents. Therefore, when m is large relative to n, it is an acceptable sacrifice of allocative efficiency to randomly remove one agent and give her one unit (run the $(1,1)$-partition mechanism). Finally, when m is neither small nor large (*e.g.*, m is close to $n/2$), neither $(1,0)$ nor $(1,1)$ works the best. When m is close to $n/2$, given two random agents, most likely exactly one of them belongs to the m highest agents. Therefore, in this scenario, it is more reasonable to exclude two agents at random, and give them only one unit (run the $(2,1)$-partition mechanism).

Claim 7. *For any n and m, we can achieve the following competitive ratio with a partitioning-based mechanism:*

$$\max\{\frac{n-1}{n}, \frac{nm+m-n}{nm},$$

$$\frac{\sum_{x\in\{x|0\leq x\leq 2, 0\leq m-x\leq n-2\}}\binom{2}{x}\binom{n-2}{m-x}(\min\{x,1\}+\min\{m-x,m-1\})}{m\binom{n}{m}}\}$$

If the first expression is the largest, then the mechanism achieving the above ratio is the $(1,0)$-partition mechanism.

If the second expression is the largest, then the mechanism achieving the above ratio is the $(1,1)$-partition mechanism.

If the third expression is the largest, then the mechanism achieving the above ratio is the $(2,1)$-partition mechanism.

6   Numerical results

In the following table, for different numbers of agents and different numbers of units, we present the competitive ratios of the WCO mechanism (α^*_{WCO}), the competitive ratios of the optimal deterministic burning-based mechanism (α^*_D), the competitive ratios of the optimal randomized burning-based mechanism (α^*_R), the competitive ratios of the optimal partitioning-based mechanism ($\alpha^*_{n_1,m_1}$), and the competitive ratios of the optimal mechanism whose allocation part is linear (α^*). It should be noted that in all the experiments, the optimal randomized burning-based mechanism matches the analytical characterization

in Claim 3, and the optimal partitioning-based mechanism matches the analytical characterization in Claim 7. The numbers below α_D^* are the numbers of units burned to achieve α_D^*. The numbers below α_R^* are the probabilities of having one unit burned to achieve α_R^*. The numbers below α_{n_1,m_1}^* are the values of n_1 and m_1 to achieve α_{n_1,m_1}^*.

	α_{WCO}^*	α_D^*	α_R^*	α_{n_1,m_1}^*	α^*
n=4,m=1	0.571	0.571	0.571	0.750	0.842
		0	0	(1,0)	
n=4,m=2	0.250	0.286	0.667	0.833	0.864
		1	0.67	(2,1)	
n=4,m=3	0	0.267	0.889	0.917	0.923
		2	0.33	(1,1)	
n=6,m=1	0.839	0.839	0.839	0.833	0.918
		0	0	(1,0)	
n=6,m=3	0.375	0.410	0.800	0.867	0.914
		1	0.6	(2,1)	
n=6,m=5	0	0.356	0.960	0.967	0.968
		3	0.2	(1,1)	
n=8,m=1	0.945	0.945	0.945	0.875	0.962
		0	0	(1,0)	
n=8,m=3	0.646	0.646	0.762	0.875	0.929
		0	0.71	(1,0)	
n=8,m=5	0.276	0.452	0.914	0.925	0.950
		2	0.43	(1,1)	
n=8,m=7	0	0.422	0.980	0.982	0.982
		4	0.14	(1,1)	

Example 6. For illustration, for $n = 8$ and $m = 5$, for type profile $(8, 7, 6, 5, 4, 3, 2, 1)$, we present the specific optimal mechanisms:

- Optimal deterministic burning-based mechanism: Two units are burned. Agents compete for the remaining 3 units under the original VCG mechanism. Agents pay the VCG payments. On top of this, for all i, agent i receives back a redistribution amount that equals $0.4518v_{-i,4} - 0.2047v_{-i,5} + 0.1412v_{-i,6} - 0.1059v_{-i,7}$, where $v_{-i,j}$ is the j-th highest type among the agents other than i. For the example type profile, the total valuation is 21. The total VCG-based payment equals 15 while the total redistribution equals 12.3247. The efficiency ratio equals 0.6108 (slightly higher than the worst-case ratio $\alpha_D^* = 0.422$).

- Optimal randomized burning-based mechanism: One unit is burned with probability 0.4286. That is, the four highest agents win for sure. The fifth highest agent wins with probability 0.5714. Agents pay the VCG-based payments. On top of this, for all i, agent i receives back a redistribution amount that equals $0.5714v_{-i,5}$, where $v_{-i,5}$ is the fifth highest type among

the agents other than i. For the example type profile, the total valuation is 28.2857. The total VCG-based payment equals 15.4285 while the total redistribution is the same. The efficiency ratio equals 0.9429 (slightly higher than the worst-case ratio $\alpha_R^* = 0.914$).

- Optimal partitioning-based mechanism: One agent is chosen uniformly at random. This agent receives one unit for free and is excluded from the auction. The remaining seven agents compete for the remaining four units under the original VCG mechanism. The VCG payments are redistributed to the excluded agent. Redistribution is always 100% under this mechanism. For the example type profile, the total valuation is 28.1250. The efficiency ratio equals 0.9375 (slightly higher than the worst-case ratio $\alpha_{1,1}^* = 0.925$).

- Optimal linear mechanism: The four highest agents win for sure. The fifth highest agent wins with probability 0.7506 and the sixth highest agent wins with probability 0.2494. Agents pay the VCG-based payments. On top of this, for all i, agent i receives back a redistribution amount that equals $0.3325v_{-i,5} + 0.2969v_{-i,6} - 0.0356v_{-i,7}$, where $v_{-i,j}$ is the j-th highest type among the agents other than i. For the example type profile, the total valuation is 29.7596. The total VCG-based payment equals 14.5012 while the total redistribution equals 14.0024. The efficiency ratio equals 0.9751 (slightly higher than the worst-case ratio $\alpha^* = 0.950$).

7   General multi-unit auctions

In this section, we drop the unit demand restriction. That is, we study multi-unit auctions in which an agent may demand more than one unit.

We no longer require that $m < n$. However, in this section, we only consider cases where $m > 1$ (cases where $m = 1$ reduce to the unit demand setting). An agent's type is no longer a single value, but a vector of length m. We use $v_{i,j}$ to denote agent i's marginal value for winning her j-th unit. That is, if agent i wins k units, then her valuation equals $\sum_{j=1}^{k} v_{i,j}$. A type profile is then an $m \times n$ matrix:

$$
\begin{bmatrix}
v_{1,m} & v_{2,m} & \cdots & v_{n,m} \\
\vdots & \vdots & \vdots & \vdots \\
v_{1,2} & v_{2,2} & \cdots & v_{n,2} \\
v_{1,1} & v_{2,1} & \cdots & v_{n,1}
\end{bmatrix}
$$

We use e_i $(i = 1, 2, \ldots, mn)$ to denote the i-th highest element in the type profile. We still use v_{-i} to denote the types of the agents other than i. That is, v_{-i} is an $m \times (n-1)$ matrix:

$$
\begin{bmatrix}
v_{1,m} & \cdots & v_{i-1,m} & v_{i+1,m} & \cdots & v_{n,m} \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
v_{1,2} & \cdots & v_{i-1,2} & v_{i+1,2} & \cdots & v_{n,2} \\
v_{1,1} & \cdots & v_{i-1,1} & v_{i+1,1} & \cdots & v_{n,1}
\end{bmatrix}
$$

We still use R to denote the redistribution function, where $R : [0,\infty)^{m(n-1)} \to \mathbb{R}$. Agent i's redistribution equals $R(v_{-i})$. If for all possible type profiles, $v_{i,j} \geq v_{i,j+1}$ for all i and all j ($j < m$), then we say we are in the setting of multi-unit auctions with *nonincreasing* marginal values.

In general multi-unit auctions, the first-best mechanism is defined as the mechanism that always *correctly* identifies the allocation that maximizes the agents' total efficiency, without asking for their types, and allocates the units among the agents accordingly at no charge. Again, clearly this mechanism obtains the largest welfare that we could hope for without an external subsidy. We will only use the first-best mechanism as a benchmark, because it is not strategy-proof. Again, our objective is to design strategy-proof complete mechanisms (consisting of both an allocation mechanism and a redistribution rule) that are competitive against the first-best mechanism, in terms of the agents' welfare.

The welfare under the first-best mechanism is

$$
\max_{\substack{x_1, x_2, \ldots, x_n \in \{0, 1, \ldots, m\} \\ \sum_{i=1}^{n} x_i = m}} \sum_{i=1}^{n} \sum_{1 \leq j \leq x_i} v_{i,j}
$$

If the marginal values are nonincreasing, then the welfare under the first-best mechanism is simply $\sum_{i=1}^{m} e_i$.

Our objective is to simultaneously find an allocation mechanism and a redistribution rule, which together form a mechanism with a high competitive ratio. In the unit demand setting, we had a useful set of candidate allocation mechanisms to choose from – the family of linear allocation mechanisms. The following properties made this set useful. First, the linear allocation mechanisms are characterized by n parameters (the a_i). Therefore, it is a value optimization problem to find an optimal linear allocation mechanism (under the premise that we know how to design a corresponding redistribution rule). Second, under a linear allocation mechanism, both the mechanism's efficiency and the agents' total payment are linear functions of the agents' types (ordered types, to be more precise). This fact allows us to consider only binary type profiles when designing a redistribution rule for a given linear allocation mechanism, as shown in Theorem 1.

As we will show, we can generalize the definition of linear allocation mechanisms to general multi-unit auctions. However, we no longer have a parameterized characterization of the whole linear family, like that in Claim 2. We can also generalize the optimization technique behind Theorem 1, *but only when the marginal values are nonincreasing*. Therefore, when the marginal values are nonincreasing, *for a given parameterized family of linear allocation mechanisms,*

22

we are able to simultaneously find both a linear allocation mechanism *within the given family*, and a redistribution rule, which together maximize the competitive ratio. As we will see, both the burning allocation mechanisms and the partitioning allocation mechanisms, when generalized to multi-unit auctions with nonincreasing marginal values, can be parameterized and are linear. This allows us to use our technique to solve for the most competitive burning/partitioning-based mechanism, when the marginal values are nonincreasing. Actually, it is easy to prove that the $(1,0)$-partition mechanism (Faltings' mechanism [7]) has the highest competitive ratio among all partitioning-based mechanisms, so we do not actually need our general technique for this. When the marginal values are not necessarily nonincreasing, our technique no longer works, but still, both the burning allocation mechanisms and the partitioning allocation mechanisms are well defined. It turns out that the $(1,0)$-partition mechanism is still the most competitive partitioning-based mechanism. We conclude with a negative result saying that no burning-based mechanism can be as competitive as the $(1,0)$-partition mechanism when the marginal values are not necessarily nonincreasing.

The key idea behind Theorem 1 is that if the allocation is linear, then when designing the redistribution rule, we can focus on binary type profiles (type profiles in which the v_i are binary). The generalized optimization technique in this section builds on the same idea. That is, in multi-unit auctions with nonincreasing marginal values, if the allocation is linear (according to the generalized definition of linearity), then when we design the redistribution rule, we can focus on binary marginal values (type profiles in which the e_i are binary). It should be noted that in previous work [17], when designing VCG redistribution mechanisms, also in the setting of multi-unit auctions with nonincreasing marginal values, we have implicitly used an almost identical idea. In that paper, the objective is to design a redistribution rule that maximizes the expected total redistribution, and the non-deficit constraint is the only constraint. In contrast, in this paper, we adopt a prior-free worst-case approach, and besides the non-deficit property, we also require individual rationality. Despite all these differences, the technique in this section can be interpreted as a slight extension of the technique in our earlier work [17]. There, the basic idea is that, if the allocation mechanism is the VCG mechanism, then for the objective and constraint in that paper, when designing the redistribution rule, it is enough to focus on binary marginal values. Here, we show that as long as the allocation mechanism is linear (VCG is a special case of the linear allocation mechanisms), then when designing the redistribution rule for our objective and constraints, we can focus on binary marginal values.

We now generalize the definition of linearity to multi-unit auctions with nonincreasing marginal values.

Definition 3. A type profile is *tie-free* if no two marginal values in it are equal. That is, for any i, i', j, j', $v_{i,j} = v_{i',j'}$ implies $i = i'$ and $j = j'$.

Definition 4. Let V and V' be two tie-free type profiles. We use the $v_{i,j}$ to denote the elements in V, and we use the $v'_{i,j}$ to denote the elements in V'. We

say V and V' are *order-consistent* [15] if for any i_1, i_2, j_1, j_2, $v_{i_1,j_1} > v_{i_2,j_2}$ implies $v'_{i_1,j_1} > v'_{i_2,j_2}$.

The set of all tie-free type profiles can be divided into a finite number of *maximal order consistent classes*, where tie-free type profiles that are order consistent with each other all belong to the same class, and tie-free type profiles that are not order consistent belong to different classes.

For example, let us consider the case with three agents and two units ($n = 3$ and $m = 2$). In this case, a type profile is a 2×3 matrix:

$$\left[\begin{array}{ccc} v_{1,2} & v_{2,2} & v_{3,2} \\ v_{1,1} & v_{2,1} & v_{3,1} \end{array} \right]$$

An example maximal order consistent class is the set of all type profiles that satisfy $v_{1,1} > v_{2,1} > v_{3,1} > v_{1,2} > v_{2,2} > v_{3,2}$. Actually, each ordering of the mn marginal values (satisfying the constraint that marginal values of each agent are non-increasing) corresponds to a maximal order consistent class.

We use $E^M(V)$ to denote the mechanism's efficiency under allocation mechanism M for type profile V. We use $P^M(V)$ to denote the agents' total payment under allocation mechanism M for type profile V. (For redistribution functions, we will omit the M superscript for simplicity.)

Now we are ready to define the linear allocation mechanisms in multi-unit auctions with nonincreasing marginal values.

Definition 5. In multi-unit auctions with nonincreasing marginal values, an allocation mechanism M is *linear* if the following conditions are satisfied:

- When we restrict our attention to type profiles in a particular maximal order consistent class S, both $E^M(V)$ and $P^M(V)$ can be written as linear functions of the elements in V, and the coefficients are constants.

- $E^M(V)$ and $P^M(V)$ are continuous in the elements in V.

- An agent's payment is always 0 if all her marginal values are 0.

As in the unit demand setting, the family of linear allocation mechanisms is convex: any probability mixture over linear allocation mechanisms is also linear.

When the marginal values are nonincreasing, the VCG mechanism is linear. For the earlier example with three agents and two items, we revisit the example maximal order consistent class containing all type profiles that satisfy $v_{1,1} > v_{2,1} > v_{3,1} > v_{1,2} > v_{2,2} > v_{3,2}$. When we only consider type profiles in this class, the mechanism's efficiency under the VCG mechanism is always $v_{1,1} + v_{2,1}$ (agents 1 and 2 each win one unit). That is, the mechanism's efficiency is a linear function of the $v_{i,j}$, with constant coefficients. Similarly, when we only consider type profiles in this class, the agents' total payment under the VCG mechanism is always $2v_{3,1}$ (agents 1 and 2 each pay $v_{3,1}$). That is, the agents' total payment

---

[15]The definition of order-consistency in our earlier work [17] is slightly different.

is also a linear function of the $v_{i,j}$, with constant coefficients. The same holds for other maximal consistent classes.

Next, we show that, under the VCG mechanism, the mechanism's efficiency and total payment are both continuous in the $v_{i,j}$. We use $\max^i(S)$ to denote the sum of the i highest values in set S. For any type profile, the mechanism's efficiency equals $\max^2(\{v_{1,1}, v_{1,2}, v_{2,1}, v_{2,2}, v_{3,1}, v_{3,2}\})$, which is continuous in the $v_{i,j}$. For any type profile, the agents' total payment equals $\max^2(\{v_{1,1}, v_{1,2}, v_{2,1}, v_{2,2}\}) + \max^2(\{v_{1,1}, v_{1,2}, v_{3,1}, v_{3,2}\}) + \max^2(\{v_{2,1}, v_{2,2}, v_{3,1}, v_{3,2}\}) - 2\max^2(\{v_{1,1}, v_{1,2}, v_{2,1}, v_{2,2}, v_{3,1}, v_{3,2}\})$, which is also continuous in the $v_{i,j}$.

Finally, it is clear that when an agent's marginal values are all 0, she pays 0 under the VCG mechanism. Therefore, for $n = 3$ and $m = 2$, the VCG mechanism is linear. By similar reasoning, it can be seen that the VCG mechanism is linear for all n and m, when the marginal values are nonincreasing.

We can straightforwardly generalize the burning allocation mechanisms and the partitioning allocation mechanisms to general multi-unit auctions. Just as in the unit demand setting, in a general multi-unit auction, under a burning allocation mechanism, we burn a deterministic/random number of units, and then allocate the remaining units according to the VCG mechanism. Again, a burning allocation mechanism is characterized by the p_i: with probability p_i, we burn $m - i$ units. (If one of the p_i is 1, then we are burning a deterministic number of units.) In general multi-unit auctions, under a partitioning allocation mechanism, we partition both the agents and the units into two groups, and run the VCG mechanism on each group. Again, a partitioning allocation mechanism is characterized by the number of agents in the smaller group (n_1) and the number of units assigned to the smaller group (m_1). Both the burning allocation mechanisms and the partitioning allocation mechanisms are based on VCG mechanisms. The burning allocation mechanisms are mixtures over VCG mechanisms for different numbers of units, and under a partitioning allocation mechanism, for each subgroup of agents that results from partitioning, we use the VCG mechanism to allocate the units assigned to this subgroup. We can use the same technique that we used to prove the linearity of the VCG mechanism to prove that both the burning allocation mechanisms and the partitioning allocation mechanisms (after ignoring payments) are linear in multi-unit auctions with nonincreasing marginal values.

Let us consider a parameterized family of linear allocation mechanisms (*e.g.*, the family of burning allocation mechanisms). We denote the parameters by the p_i ($i = 1, 2, \ldots, k$). The p_i may need to satisfy certain constraints for the corresponding mechanism to be feasible. For example, for the family of burning allocation mechanisms, the p_i must sum to 1, because they represent probabilities of burning different numbers of units ($k = m$ in this case). The following optimization model solves for the optimal values of the p_i, as well as a redistribution function R, so that together, the linear allocation mechanism characterized by the p_i and the redistribution function R have the highest competitive ratio. (We use F to denote the first-best mechanism.)

> **Variable function:** $R : [0, \infty)^{m(n-1)} \to \mathbb{R}$
> **Variables:** $\alpha, p_1, p_2, \ldots, p_k$
> **Maximize** α
> **Subject to:**
> The p_i characterize a feasible linear allocation mechanism M.
> For every type profile V in which the marginal values, denoted by the $v_{i,j}$,
> are nonincreasing,
>   For all i, $R(v_{-i}) \geq 0$ (individual rationality)
> $P^M(V) \geq \sum_{i=1}^n R(v_{-i})$ (non-deficit)
> $E^M(V) - P^M(V) + \sum_{i=1}^n R(v_{-i}) \geq \alpha E^F(V)$ (competitive ratio)

Once again, this is a functional optimization problem, which appears difficult to solve. We use the vector $(x_1, x_2, \ldots, x_{n-1})$ to denote the matrix v_{-i} in which the j-th column contains x_j ones (the x_i are integers). For example, $\begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \end{bmatrix}$ is denoted by the vector $(2, 1, 0)$. We use $r(x_1, x_2, \ldots, x_{n-1})$ to denote the redistribution received by an agent when the other agents' types are $(x_1, x_2, \ldots, x_{n-1})$. Just as in the unit demand setting, if we only consider binary marginal values, then the redistribution function R can be represented by the $r(\cdot)$, where $0 \leq x_i \leq m$. In the original functional optimization problem, we are optimizing over the p_i and the redistribution function R. When we only consider binary marginal values, we are optimizing over the p_i and the $r(\cdot)$. Hence, we have a value optimization problem. Under some conditions, this optimization problem corresponds to a linear program (*e.g.*, when we focus on the family of burning allocation mechanisms).

> **Variables:** $\alpha, p_1, p_2, \ldots, p_k, r(x_1, x_2, \ldots, x_{n-1})$ $(0 \leq x_1, x_2, \ldots, x_{n-1} \leq m)$
> **Maximize** α
> **Subject to:**
> The p_i characterize a feasible linear allocation mechanism M.
> For every type profile V in which the marginal values, denoted by the $v_{i,j}$,
> are nonincreasing and binary,
>   For all i, $R(v_{-i}) \geq 0$ (individual rationality)
> $P^M(V) \geq \sum_{i=1}^n R(v_{-i})$ (non-deficit)
> $E^M(V) - P^M(V) + \sum_{i=1}^n R(v_{-i}) \geq \alpha E^F(V)$ (competitive ratio)

We use the p_i^*, the $r^*(x_1, x_2, \ldots, x_n)$, and α^* to denote an optimal solution of the above value optimization program. We always have that $r^*(0, 0, \ldots, 0) = 0$. (For the type profile in which all the marginal values are 0, individual rationality implies $r^*(0, 0, \ldots, 0) \geq 0$, and the non-deficit constraint implies $r^*(0, 0, \ldots, 0) \leq 0$.) Again, α^*, the optimal objective value of the above linear program, must be greater than or equal to the optimal objective value of the original functional optimization problem. That is, α^* is an upper bound on the competitive ratio of any complete mechanism whose allocation part is within the family under consideration. It turns out that we can achieve this upper bound. In what follows, we construct a complete mechanism whose allocation

part is within the family under consideration, and prove that this mechanism has a competitive ratio of α^*. Hence, the constructed mechanism is an optimal solution to the original functional optimization problem.

We define function R^* as follows:

$$R^*(v_{-i}) = \sum_{j=1}^{m(n-1)} (r^*(T_{-i,j}) - r^*(T_{-i,j-1}))e_{-i,j} \tag{1}$$

$e_{-i,j}$ is the j-th highest element in v_{-i}. $T_{-i,j}$ is obtained from v_{-i} by replacing the j highest marginal values in v_{-i} by 1, and by replacing the remaining marginal values by 0.[16]

R^* is defined in such a way that when v_{-i} is binary, $R^*(v_{-i})$ matches $r^*(v_{-i})$. Therefore, R^* can be interpreted as a linear interpolation of the $r^*(\cdot)$.

R^*, as defined above, is continuous in the elements in v_{-i}, for the following reasons. If we perturb an element x in v_{-i} by δ, then some of the $e_{-i,j}$ are perturbed. Let k_1 be the smallest number for which e_{-i,k_1} is perturbed. Let k_2 be the largest number for which e_{-i,k_2} is perturbed. We must have that before the perturbation, for all $k_1 \leq j \leq k_2$, $e_{-i,j}$ differs from x by at most δ, and after the perturbation, for all $k_1 \leq j \leq k_2$, $e_{-i,j}$ differs from x by at most δ. We also have that $T_{-i,j}$ is unchanged for $j < k_1$ and $j \geq k_2$. Therefore, before the perturbation, $R^*(v_{-i})$ is close to $\sum_{j=1}^{k_1-1}(r^*(T_{-i,j}) - r^*(T_{-i,j-1}))e_{-i,j} + \sum_{j=k_1}^{k_2}(r^*(T_{-i,j}) - r^*(T_{-i,j-1}))x + \sum_{j=k_2+1}^{n}(r^*(T_{-i,j}) - r^*(T_{-i,j-1}))e_{-i,j} = \sum_{j=1}^{k_1-1}(r^*(T_{-i,j}) - r^*(T_{-i,j-1}))e_{-i,j} + r^*(T_{-i,k_2})x - r^*(T_{-i,k_1-1})x + \sum_{j=k_2+1}^{n}(r^*(T_{-i,j}) - r^*(T_{-i,j-1}))e_{-i,j}$. (The difference between $R^*(v_{-i})$ and this expression is bounded above by a multiple of δ.) It should be noted that the value of this expression is unchanged after the perturbation, because $T_{-i,j}$ is unchanged for $j < k_1$ and $j \geq k_2$, and $e_{-i,j}$ is unchanged for $j < k_1$ and $j > k_2$. After the perturbation, with the new values of the $T_{-i,j}$ and the $e_{-i,j}$, we still have that $R^*(v_{-i})$ is close to $\sum_{j=1}^{k_1-1}(r^*(T_{-i,j}) - r^*(T_{-i,j-1}))e_{-i,j} + r^*(T_{-i,k_2})x - r^*(T_{-i,k_1-1})x + \sum_{j=k_2+1}^{n}(r^*(T_{-i,j}) - r^*(T_{-i,j-1}))e_{-i,j}$. (The difference between $R^*(v_{-i})$ after the perturbation and this expression is also bounded above by a multiple of δ.) That is, with or without the perturbation, $R^*(v_{-i})$ is close to the same expression (which has the same value with or without the perturbation). Therefore, R^*, as defined above, is continuous in the elements in v_{-i}.

Theorem 2. *Within the family of allocation mechanisms that we consider, let $M(p_1^*, p_2^*, \ldots, p_k^*)$ denote the allocation mechanism characterized by the p_i^*. $(M(p_1^*, p_2^*, \ldots, p_k^*), R^*)$ is α^*-competitive. That is, it is optimal among all complete mechanisms whose allocation part is within the family under consideration.*

Before proving the theorem, we give the following lemma, which is similar to Lemma 1.

---

[16]Here, ties can be broken arbitrarily (we will show later that R^*, as defined above, is continuous in the elements in v_{-i}).

Lemma 2. *Let the c_i be constants. The following two systems of inequalities are equivalent:*

(a) $c_1x_1 + c_2x_2 + \ldots + c_sx_s \geq 0$ *for all* $x_1 > x_2 > \ldots > x_s > 0$.

(b) *For any $\delta > 0$, there exists $\epsilon > 0$, so that $c_1x_1 + c_2x_2 + \ldots + c_sx_s > -\delta$ for all $(x_1, x_2, \ldots, x_s) \in \bigcup_{k=1,2,\ldots,s}\{(x_1, x_2, \ldots, x_s)|1 > x_1 > x_2 > \ldots > x_s > 0, x_k > 1 - \epsilon, x_{k+1} < \epsilon\}$.*

Proof. (a) $\Rightarrow$ (b) is trivial. We now prove (b) $\Rightarrow$ (a). By Lemma 1, we only need to prove (b) implies the following: $c_1x_1 + c_2x_2 + \ldots + c_sx_s \geq 0$ for all $x_1 \geq x_2 \geq \ldots \geq x_s$, where the x_i are binary. For any δ and ϵ, if we have that $c_1x_1 + c_2x_2 + \ldots + c_sx_s > -\delta$ for all $(x_1, x_2, \ldots, x_s) \in \bigcup_{k=1,2,\ldots,s}\{(x_1, x_2, \ldots, x_s)|1 > x_1 > x_2 > \ldots > x_s > 0, x_k > 1 - \epsilon, x_{k+1} < \epsilon\}$, then we have $c_1x_1 + c_2x_2 + \ldots + c_sx_s \geq -\delta$ for all $(x_1, x_2, \ldots, x_s) \in \bigcup_{k=1,2,\ldots,s}\{(x_1, x_2, \ldots, x_s)|1 \geq x_1 \geq x_2 \geq \ldots \geq x_s \geq 0, x_k \geq 1 - \epsilon, x_{k+1} \leq \epsilon\}$. This implies that $c_1x_1 + c_2x_2 + \ldots + c_sx_s \geq 0$ for all $x_1 \geq x_2 \geq \ldots \geq x_s$, where the x_i are binary. $\square$

The above lemma is basically a variant of Lemma 1, where part (b) has been changed slightly (from binary to almost binary and from nonnegative to almost nonnegative). As we mentioned earlier, the technique proposed in this section is a slight generalization of the result in [17]. Therefore, we only present the following proof sketch.

Proof. We only need to prove that α^*, the p_i^*, and the function R^* form a feasible solution to the original functional optimization problem. If they form a feasible solution, then $(M(p_1^*, p_2^*, \ldots, p_k^*), R^*)$ is α^*-competitive, implying that it is optimal. That is, we need to prove that the constructed mechanism satisfies individual rationality, the non-deficit constraint, and the competitive ratio constraint (it is obvious that the constraints involving only the p_i hold).

According to the definition of the linear allocation mechanisms, both $E^M(V)$ and $P^M(V)$ are continuous in the elements in V. We have also proved that the constructed R^* is continuous in the elements in V. We now consider the constraints in the optimization model that are related to the redistribution rule. The individual rationality constraint is that for all type profiles V in which the elements are the $v_{i,j}$, we must have $R^*(v_{-i}) \geq 0$ for all i and all v_{-i}. Because the set of tie-free type profiles is dense in the set of all type profiles, it suffices to prove that for all tie-free type profiles V in which the elements are the $v_{i,j}$, we must have $R^*(v_{-i}) \geq 0$ for all i and all v_{-i}. Let V be an arbitrary tie-free type profile. Let S be the maximal order consistent class containing V. When we focus only on type profiles in S, for any type profile/matrix V, the i-th highest element is always in a fixed position (j, k). To prove that all type profiles in S satisfy $R^*(v_{-i}) \geq 0$, we need to prove a linear inequality in the form of part (a) in Lemma 2, because R^* is linear in the e_i, and the e_i can take any values as long as we have $e_1 > e_2 > \ldots > e_{mn} > 0$. When the type profiles are binary, the individual rationality constraint is satisfied, by the definition of R^*. Because R^* is continuous, when we perturb a binary type profile, the

individual rationality constraint can at most be violated by a small amount. It follows that the aforementioned linear inequality in the e_i can at most be violated by a small amount if the e_i are close to binary. This gives us part (b) in Lemma 2. Therefore, we have that individual rationality is satisfied for all tie-free type profiles (and hence, as we argued before, for all profiles) in an arbitrary maximal order consistent class. Similarly, we can prove that the non-deficit constraint and the competitive ratio constraint are satisfied for all type profiles. □

Based on the technique proposed above, we can solve for an optimal burning-based mechanism. For $n = 5$ and different m, the following table presents the optimal competitive ratios of burning-based mechanisms in multi-unit auctions with nonincreasing marginal values.

$m = 1$	$m = 3$	$m = 5$	$m = 7$	$m = 9$
0.733	0.643	0.610	0.595	0.587

As a comparison, for the unit demand setting, when $n = 5$ and $m = 1$, the competitive ratio of the optimal randomized burning-based mechanism is also 0.733. This is not surprising, because when there is only one item, the two settings are the same. When $n = 5$ and $m = 3$, for the unit demand setting, the competitive ratio of optimal randomized burning-based mechanism is 0.833 (higher than the value of 0.643 for the multi-demand setting). For the cases where $n = 5$ and $m \geq 5$, in the unit demand setting we naturally obtain a competitive ratio of 1 because we can satisfy all agents trivially.

Example 7. For illustration, for $n = 3$ and $m = 2$, we present the specific optimal burning-based mechanism.

With probability 0.8, nothing is burned. With probability 0.2, one unit is burned. Optimal redistributions are characterized by the following r^* function:

$r^*(0,0) = r^*(1,0) = r^*(2,0) = 0$
$r^*(1,1) = r^*(2,1) = 0.2$
$r^*(2,2) = 0.6$

Let us consider a specific type profile:

$$\begin{bmatrix} 5 & 3 & 1 \\ 6 & 4 & 2 \end{bmatrix}$$

Agent 1's marginal value for winning the first unit equals 6, and her marginal value for winning the second unit equals 5. Agent 2's marginal value for winning the first unit equals 4, and her marginal value for winning the second unit equals 3. Finally, agent 3's marginal value for winning the first unit equals 2, and her marginal value for winning the second unit equals 1.

For this specific type profile, agent 1 wins one unit for sure, and she wins a second unit with probability 0.8. The total valuation equals $6 + 5 \cdot 0.8 = 10$. Agent 1 is the only agent who pays positive VCG payment. The total VCG payment equals $4 + 3 \cdot 0.8 = 6.4$.

According to Equation 1, agent 1's redistribution equals

$$R(\begin{bmatrix} 3 & 1 \\ 4 & 2 \end{bmatrix}) = 4(r^*(1,0) - r^*(0,0)) + 3(r^*(2,0) - r^*(1,0))$$

$$+2(r^*(2,1) - r^*(2,0)) + 1(r^*(2,2) - r^*(2,1))$$

$$= 2 \cdot 0.2 + 1 \cdot 0.4 = 0.8$$

Similarly, agent 2's redistribution and agent 3's redistribution equal 0.8 and 2, respectively.

The total redistribution is then 3.6. The efficiency ratio equals $\frac{10-6.4+3.6}{11} = 0.6545$. The worst-case ratio is only 0.3. As a comparison, for $n = 3$ and $m = 2$, in the unit demand setting, the optimal randomized burning-based mechanism has a worst-case ratio of 0.75.

We can also use the proposed technique to solve for the optimal partitioning-based mechanism. Actually, it turns out that this is not necessary, as it is easy to prove that the $(1,0)$-partition mechanism is the most competitive partitioning-based mechanism, whether the marginal values are assumed to be nonincreasing or not.

Theorem 3. *In general multi-unit auctions, the $(1,0)$-partition mechanism is the most competitive partitioning-based mechanism, whether the marginal values are assumed to be nonincreasing or not.*

Example 8. For the type profile in Example 7, if we run the $(1,0)$-partition mechanism, then the valuation equals $\frac{7+11+11}{3} = 9.6667$. All payments are redistributed. The efficiency ratio equals $\frac{9.6667}{11} = 0.8788$ (higher than the worst-case ratio $\frac{n-1}{n} = 0.6667$).

Finally, we conclude with a negative result saying that no burning-based mechanism can be more competitive than the $(1,0)$-partition mechanism when the marginal values are not necessarily nonincreasing.

Claim 8. *In multi-unit auctions, if there are at least two agents,[17] and the marginal values are not necessarily nonincreasing, then no burning-based mechanism's competitive ratio can be strictly more than $\frac{n-1}{n}$.*

8   Future Research

Future research on maximizing the agents' welfare under the constraint of strategy-proofness can take a number of directions. Are there other families of inefficient mechanisms that result in a high welfare for the agents? Among what more general classes of mechanisms are the mechanisms proposed in this paper optimal? Can we generalize some of these results to wider settings, for

---

[17]If there is only one agent, we can simply allocate all the units to the only agent at no charge.

example, settings with heterogeneous objects? What happens if we change the objective, for example, if we have a prior distribution over type profiles and we wish to maximize expected welfare?

Finally, we have assumed in this paper that random bits are not chosen adversarially, that is, the worst-case adversary controls the agents' valuations but not our random choices. If the adversary also controls our random choices (which would correspond to a different notion of worst-case optimality), then it never hurts to use a deterministic mechanism. To illustrate how this affects our results, let us consider Example Mechanism 2 of Subsection 5.2 (the Faltings mechanism), in which we (uniformly) randomly choose one agent to be excluded from the auction and assign this agent zero units. This mechanism is competitive in the sense of this paper. (An agent is excluded with probability $\frac{1}{n}$, hence the competitive ratio is $\frac{n-1}{n}$.) However, if we use the alternative notion of worst-case optimality where the adversary controls the random choices, then the agent with the highest valuation will always be excluded, so that the mechanism is not competitive. It remains to be seen what we can do for this alternative notion of worst-case optimality. There has already been work in this direction. de Clippel *et al.* [6] study deterministic inefficient mechanisms that are competitive in multi-unit auctions with unit demand. Specifically, their mechanisms are based on burning deterministic numbers of units. The difference between their mechanisms and the deterministic burning-based mechanisms studied in this paper is that in their mechanisms, different numbers of units may be burned for different type profiles, while in our mechanisms, we only consider burning a fixed number of units for all type profiles.

9   Conclusion

The VCG mechanism is not strongly budget balanced: in general, value flows out of the system of agents in the form of VCG payments, which reduces the agents' utilities. In many settings, the objective is to maximize the sum of the agents' utilities (taking payments into account). For this purpose, several VCG redistribution mechanisms have been proposed that redistribute a large fraction of the VCG payments back to the agents, in a way that maintains strategy-proofness and the non-deficit property. Unfortunately, sometimes even the best VCG redistribution mechanism fails to redistribute a substantial fraction of the VCG payments. This results in a low welfare for the agents, even though the items are allocated efficiently. In this paper, we studied strategy-proof allocation mechanisms that do not always allocate the items efficiently. It turns out that by allocating inefficiently, more payment can sometimes be redistributed, so that the net effect is an increase in the sum of the agents' utilities. The objective that we pursued is to design mechanisms that are competitive in terms of welfare against the first-best mechanism. The results obtained in our paper are summarized as follows.

- Multi-Unit Auctions with Unit Demand:

- We characterized the family of linear allocation mechanisms (Claim 2). We proposed an optimization technique for simultaneously finding a linear allocation mechanism and a payment redistribution rule, which together are optimal (Theorem 1).

- Given n and m, we analytically characterized a competitive mechanism based on burning units (Claim 3). The mechanism burns at most one unit. For some values of n and m, the mechanism does not burn any units, in which case it is identical to the previously known WCO mechanism [16]. For other n and m values, the mechanism burns one unit with a certain probability, in which case all payments can be redistributed, achieving strong budget balance.

- Given n and m, we analytically characterized a competitive mechanism based on partitioning the agents into groups (Claim 7). Depending on the values of n and m, the mechanism is one of the following mechanisms: the $(1,0)$-partition mechanism, the $(1,1)$-partition mechanism, and the $(2,1)$-partition mechanism. Here, both $(1,0)$ and $(1,1)$ are previously known mechanisms [7, 23].

- Multi-Unit Auctions with Nonincreasing Marginal Values:

  - We generalized the definition of linear allocation mechanisms to this setting and also generalized the corresponding optimization technique (Theorem 2).

- General Multi-Unit Auctions:

  - We showed that the previously known Faltings mechanism (the $(1,0)$-partition mechanism) is optimal among all partitioning-based mechanisms, whether the marginal values are nonincreasing or not (Theorem 3).

  - We showed that for general multi-unit auctions, no burning-based mechanism can be better than Faltings' mechanism (Claim 8).

References

[1] G. Aggarwal, A. Fiat, A. Goldberg, J. Hartline, N. Immorlica, and M. Sudan. Derandomization of auctions. In *Proceedings of the Annual Symposium on Theory of Computing (STOC)*, pages 619–625, 2005.

[2] M. J. Bailey. The demand revealing process: to distribute the surplus. *Public Choice*, 91:107–126, 1997.

[3] M.-F. Balcan, A. Blum, J. D. Hartline, and Y. Mansour. Mechanism design via machine learning. In *Proceedings of the Annual Symposium on Foundations of Computer Science (FOCS)*, pages 605–614, 2005.

[4] R. Cavallo. Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In *Proceedings of the International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pages 882–889, Hakodate, Japan, 2006.

[5] E. H. Clarke. Multipart pricing of public goods. *Public Choice*, 11:17–33, 1971.

[6] G. de Clippel, V. Naroditskiy, and A. Greenwald. Destroy to save. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 207–214, Stanford, CA, USA, 2009.

[7] B. Faltings. A budget-balanced, incentive-compatible scheme for social choice. In *Agent-Mediated Electronic Commerce (AMEC), LNAI, 3435*, pages 30–43, 2005.

[8] J. Feigenbaum, C. H. Papadimitriou, and S. Shenker. Sharing the cost of muliticast transmissions. *Journal of Computer and System Sciences*, 63:21–41, 2001.

[9] A. Goldberg, J. Hartline, A. Karlin, M. Saks, and A. Wright. Competitive auctions. *Games and Economic Behavior*, 2006.

[10] J. Green and J.-J. Laffont. Characterization of satisfactory mechanisms for the revelation of preferences for public goods. *Econometrica*, 45:427–438, 1977.

[11] J. Green and J.-J. Laffont. *Incentives in Public Decision Making*. Amsterdam: North-Holland, 1979.

[12] T. Groves. Incentives in teams. *Econometrica*, 41:617–631, 1973.

[13] S. Gujar and Y. Narahari. Redistribution mechanisms for assignment of heterogeneous objects. *J. Artif. Intell. Res. (JAIR)*, 41:131–154, 2011.

[14] M. Guo. VCG redistribution with gross substitutes. In *Proceedings of the National Conference on Artificial Intelligence (AAAI)*, San Francisco, CA, USA, 2011.

[15] M. Guo. Worst-case optimal redistribution of VCG payments in heterogeneous-item auctions with unit demand. In *Proceedings of the Eleventh International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, Valencia, Spain, 2012.

[16] M. Guo and V. Conitzer. Worst-case optimal redistribution of VCG payments in multi-unit auctions. *Games and Economic Behavior*, 67(1):69–98, 2009.

[17] M. Guo and V. Conitzer. Optimal-in-expectation redistribution mechanisms. *Artificial Intelligence*, 174(5-6):363–381, 2010.

[18] M. Guo, E. Markakis, K. R. Apt, and V. Conitzer. Undominated groves mechanisms. *Journal of Artificial Intelligence Research*, 46:129–163, 2013.

[19] M. Guo, V. Naroditskiy, V. Conitzer, A. Greenwald, and N. R. Jennings. Budget-balanced and nearly efficient randomized mechanisms: Public goods and beyond. In *Proceedings of the Seventh Workshop on Internet and Network Economics (WINE)*, Singapore, 2011.

[20] J. Hartline and R. McGrew. From optimal limited to unlimited supply auctions. In *Proceedings of the ACM Conference on Electronic Commerce (EC)*, pages 175–182, Vancouver, BC, Canada, 2005.

[21] B. Holmström. Groves' scheme on restricted domains. *Econometrica*, 47(5):1137–1144, 1979.

[22] L. Hurwicz. On the existence of allocation systems whose manipulative Nash equilibria are Pareto optimal, 1975. Presented at the 3rd World Congress of the Econometric Society.

[23] H. Moulin. Almost budget-balanced VCG mechanisms to assign multiple objects. *Journal of Economic Theory*, 144(1):96–119, 2009.

[24] R. Myerson and M. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 28:265–281, 1983.

[25] V. Naroditskiy, M. Guo, L. Dufton, M. Polukarov, and N. R. Jennings. Redistribution of VCG payments in public project problems. In *Proceedings of the Eighth Workshop on Internet and Network Economics (WINE)*, Liverpool, 2012.

[26] V. Naroditskiy, M. Polukarov, and N. R. Jennings. Optimization in payments in dominant strategy mechanisms for single-parameter domains. In *Proceedings of the Third International Workshop on Computational Social Choice (COMSOC)*, Dusseldorf, Germany, 2010.

[27] N. Nisan and A. Ronen. Computationally feasible VCG mechanisms. *Journal of Artificial Intelligence Research*, 29(1):19–47, 2007.

[28] D. Parkes, J. Kalagnanam, and M. Eso. Achieving budget-balance with Vickrey-based payment schemes in exchanges. In *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1161–1168, Seattle, WA, USA, 2001.

[29] S. Tsuruta, M. Oka, T. Todo, Y. Kawasaki, M. Guo, Y. Sakurai, and M. Yokoo. Optimal false-name-proof single-item redistribution mecanisms. In *Proceedings of the Thirteenth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, Paris, France, 2014.

[30] W. Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16:8–37, 1961.

Appendix

Proof of Claim 2[18]:

Proof. Let us consider a specific strategy-proof linear allocation mechanism, whose allocation vector is $(a_1, a_2, \ldots, a_n)$. Let us reason from the perspective of a particular agent Alice. (We use "agent Alice" instead of "agent i" because agent i specifically refers to the agent whose type is the i-th highest.) We use x_i ($i = 1, 2, \ldots, n-1$) to denote the i-th highest type among the agents other than Alice. We consider a specific type profile in which the x_i are all distinct, and $x_{n-1} > 0$.

First of all, we observe that (holding the types of the other agents fixed), Alice's payment depends only on her rank. This is because if there were two types that resulted in the same rank but a different payment, then if she had the type that resulted in the higher payment, she would have an incentive to report the one that results in the lower payment instead, because, by assumption, the allocation will be the same. In particular, if she is ranked n-th, she will pay 0, because an agent bidding 0 always pays 0.

We use z_i to denote Alice's payment when she is the i-th highest agent (so $z_n = 0$). Since we are considering a strategy-proof mechanism, for any $1 \leq i \leq n-1$, we have the following:

$$a_{i+1}(x_i - \epsilon) - z_{i+1} \geq a_i(x_i - \epsilon) - z_i$$
$$a_{i+1}(x_i + \epsilon) - z_{i+1} \leq a_i(x_i + \epsilon) - z_i$$

Here, ϵ is a sufficiently small positive value. The first inequality says that when Alice's type is just below x_i, she prefers winning with probability a_{i+1} and paying z_{i+1} (reporting truthfully) to winning with probability a_i and paying z_i (misreporting to be ranked i-th). The second inequality says that this is reversed when when Alice's type is just above x_i. It is straightforward to verify that if $a_{i+1} > a_i$, then the above two inequalities cannot be simultaneously satisfied. That is, for a linear allocation mechanism to be strategy-proof, we must have that $a_i \geq a_{i+1}$ for all $1 \leq i \leq n-1$. That is, every strategy-proof linear allocation mechanism corresponds to an allocation vector in $\{(a_1, a_2, \ldots, a_n) | 1 \geq a_1 \geq a_2 \geq \ldots \geq a_n \geq 0, \sum_{i=1}^{n} a_i \leq m\}$.

We now consider Alice's payment for an arbitrary type profile. We still use x_i ($i = 1, 2, \ldots, n-1$) to denote the i-th highest type from the agents other than Alice. We no longer require the x_i to be distinct or $x_{n-1} > 0$. We still use z_i to denote Alice's payment when she is the i-th highest agent. We know that $z_n = 0$. When Alice's type is exactly x_{n-1}, she should be indifferent between being the $(n-1)$-th highest agent and the n-th highest agent. That is, we have $a_{n-1}x_{n-1} - z_{n-1} = a_n x_{n-1}$, or equivalently, $z_{n-1} = (a_{n-1} - a_n)x_{n-1}$. When Alice's type is exactly x_{n-2}, she should be indifferent between being

---

[18]It should be noted that in the proof, to address the issue of tie-breaking, we assume that a strategic agent can choose the way that the tie is broken for her—that is, she can choose any rank consistent with her bid. To avoid conflict during tie-breaking, we also require that if multiple agents have the same bid, then they are all indifferent about how the tie is broken.

the $(n-2)$-th highest agent and the $(n-1)$-th highest agent. That is, we have $a_{n-2}x_{n-2}-z_{n-2} = a_{n-1}x_{n-2}-z_{n-1}$. Using the previous expression for z_{n-1}, we have $z_{n-2} = (a_{n-2}-a_{n-1})x_{n-2}+(a_{n-1}-a_n)x_{n-1}$. Continuing this reasoning, we obtain $z_i = \sum_{j=i}^{n-1}(a_j-a_{j+1})x_j$. Since we are considering an arbitrary agent and an arbitrary type profile, we have that the a_i characterize the agents' payments, and thus characterize the entire mechanism (because the a_i also characterize the allocation). When the a_i are nonincreasing and sum to at most m, the linear allocation mechanism characterized by the a_i is exactly the maximal-in-range [27] allocation mechanism under which the allowable allocations are those in which the agents win with probabilities $a_1, a_2, \ldots, a_n$. Moreover, the payments described above are exactly the corresponding VCG-based payments. We know that maximal-in-range allocations combined with VCG-based payments result in strategy-proof mechanisms. Therefore, we have that each allocation vector in $\{(a_1, a_2, \ldots, a_n)|1 \geq a_1 \geq a_2 \geq \ldots \geq a_n \geq 0, \sum_{i=1}^{n} a_i \leq m\}$ corresponds to a strategy-proof linear allocation mechanism. $\square$

Proof of Claim 3:

Proof. We already know that WCO has competitive ratio $1-\dfrac{\binom{n-1}{m}}{\sum_{j=m}^{n-1}\binom{n-1}{j}}$. Hence, we only need to show that the other mechanism proposed in the claim is feasible and has competitive ratio $\frac{mn-n}{mn-m}$, for the values of n and m for which this mechanism outperforms WCO.

If $m = 1$, then $\frac{mn-n}{mn-m} = 0$, which cannot be greater than the competitive ratio of WCO. So we only need to consider $m > 1$. We first prove that the mechanism is feasible. That is, it is individually rational and satisfies the non-deficit property. Individual rationality is satisfied because the redistribution is always nonnegative. For any type profile, the total VCG payment is mv_{m+1} when nothing is burned, and it is $(m-1)v_m$ when one unit is burned. In expectation, the total payment collected by the allocation mechanism is $\frac{n-m}{n-1}(m-1)v_m + (1 - \frac{n-m}{n-1})mv_{m+1} = \frac{m-1}{n-1}(mv_{m+1} + (n-m)v_m)$. For the agents reporting $v_1, \ldots, v_m$, the redistribution received is $\frac{m-1}{n-1}v_{m+1}$. For the other agents, the redistribution received is $\frac{m-1}{n-1}v_m$. Therefore, the total redistribution equals the total payment, so the non-deficit constraint is satisfied. We conclude that the mechanism is feasible.

Now we show that the mechanism has competitive ratio $\frac{mn-n}{mn-m}$. With probability $\frac{n-m}{n-1}$, the total efficiency is $\sum_{i=1}^{m-1} v_i$ (one unit is burned). When nothing is burned, the total efficiency is $\sum_{i=1}^{m} v_i$. In expectation, the total efficiency is $\frac{n-m}{n-1}\sum_{i=1}^{m-1} v_i + (1 - \frac{n-m}{n-1})\sum_{i=1}^{m} v_i$. This is greater than or equal to $\frac{n-m}{n-1}\frac{m-1}{m}\sum_{i=1}^{m} v_i + (1-\frac{n-m}{n-1})\sum_{i=1}^{m} v_i = \frac{mn-n}{mn-m}\sum_{i=1}^{m} v_i$ (we have equality when all the v_i are equal). Since the total payment equals the total redistribution, efficiency is equal to welfare, so we conclude that we obtain the competitive ratio $\frac{mn-n}{mn-m}$. $\square$

Proof of Claim 5:

Proof. For the type profile in which m agents report 1 and the remaining agents report 0, the first-best mechanism achieves an efficiency of m. Hence, $\frac{U_{M_{n_1,m_1}}(m)}{m}$ is an upper bound on $\alpha^*_{n_1,m_1}$.

We now show that $\frac{U_{M_{n_1,m_1}}(m)}{m}$ is equal to the second expression in the claim; then, we will show that M_{n_1,m_1} does in fact attain this competitive ratio. In the (n_1, m_1)-partition mechanism, n_1 agents are randomly picked to form one group, and the remaining $n - n_1$ agents form a second group. If m agents report 1 and the remaining agents report 0, then the probability of having x agents that report 1 in the group of size n_1 is $\frac{\binom{n_1}{x}\binom{n-n_1}{m-x}}{\binom{n}{m}}$. The corresponding total welfare is $(\min\{x, m_1\} + \min\{m - x, m - m_1\})$. The set of possible values of x is X. It follows that $U_{M_{n_1,m_1}}(m)$ is equal to $\frac{\sum_{x \in X} \binom{n_1}{x}\binom{n-n_1}{m-x}(\min\{x,m_1\}+\min\{m-x,m-m_1\})}{\binom{n}{m}}$.

All that is left to show is that M_{n_1,m_1} does in fact attain this competitive ratio. Let us consider the following allocation mechanism, which is never better than M_{n_1,m_1}:

- Pick n_1 agents to form one group, uniformly at random. The other $n - n_1$ agents form the second group.

- Remove the agents with the lowest $n - m$ types.

- For the first group, if there are more than m_1 agents left, allocate m_1 units uniformly at random among the remaining agents in group one. Otherwise, allocate one unit to every remaining agent in group one.

- For the second group, if there are more than $m - m_1$ agents left, allocate $m - m_1$ units uniformly at random among the remaining agents in group two. Otherwise, allocate one unit to every remaining agent in group two.

For any type profile, the above mechanism results in (weakly) lower efficiency than the (n_1, m_1)-partition mechanism, because in the partitioning-based mechanism, the units are assigned efficiently within each group, and in the modified mechanism they are not because of agent removal and random assignment.

Under the modified mechanism, only the agents reporting $v_1, \ldots, v_m$ possibly win any units, and the probability of winning is the same for each of them. For the type profile in which m agents report 1 and the remaining agents report 0, the modified mechanism results in the same efficiency as the partitioning-based mechanism. Therefore, because in this case, a winning agent's utility is 1, the expected number of winners under the modified mechanism is $U_{M_{n_1,m_1}}(m)$. But this number must be the same for all type profiles. So, using the fact that each of the top m agents is equally likely to win, for a general type profile, the expected efficiency under the modified mechanism is $\frac{U_{M_{n_1,m_1}}(m)}{m} \sum_{i=1}^{m} v_i$; and we know that this is (weakly) lower than the expected efficiency under the (n_1, m_1)-partition mechanism. Hence, the (n_1, m_1)-partition mechanism has a competitive ratio of at least $\frac{U_{M_{n_1,m_1}}(m)}{m}$. We have already proved that

37

$\frac{U_{M_{n_1,m_1}}(m)}{m}$ is an upper bound on $\alpha^*_{n_1,m_1}$, so this expression must be exactly equal to the competitive ratio. □

Proof of Theorem 3:

Proof. Let us consider the type profile in which agent 1's marginal value is 1 for every additional unit, and all the other agents have only marginal values of 0. For this type profile, under the (n_1, m_1)-partition mechanism, the agents' total welfare equals $\frac{n_1}{n}m_1 + \frac{n-n_1}{n}(m - m_1)$. (With probability $\frac{n_1}{n}$, agent 1 is assigned to the group of size n_1, in which case the welfare equals m_1. With probability $\frac{n-n_1}{n}$, agent 1 is assigned to the group of size $n - n_1$, in which case the welfare equals $m - m_1$.) This expression is at most $\frac{n-1}{n}m$ for any n_1 and m_1. The agents' total welfare under the first-best mechanism is simply m. Therefore, whether the marginal values are assumed to be nonincreasing or not, any partitioning-based mechanism is at most $\frac{n-1}{n}$-competitive.

Given any type profile, let S denote the set of agents that win something under the first-best mechanism. If we apply the $(1, 0)$-partition mechanism, then every agent is excluded with probability $\frac{1}{n}$. Suppose that, when a winner is excluded, we do not change the allocation, but simply throw away the units that the excluded agent would have won; the resulting welfare is exactly $\frac{n-1}{n}$ times the welfare under the first-best mechanism; moreover, it is always smaller than or equal to the agents' welfare under the $(1, 0)$-partition mechanism (which allocates efficiently among the non-excluded bidders). Therefore, the $(1, 0)$-partition mechanism is at least $\frac{n-1}{n}$-competitive, which means that it is optimal among all partitioning-based mechanisms, whether the marginal values are assumed to be nonincreasing or not. □

Proof of Claim 8:

Proof. To ensure individual rationality, any agent's redistribution should be nonnegative. To ensure the non-deficit property, an agent's redistribution should never exceed the minimal total payment that can result when we hold the other agents' types fixed and change her own type. Otherwise, when her type is one that minimizes the total payment, the total redistribution, which is at least the redistribution received by this agent, is more than the total payment.[19]

We first prove the result for the case of two agents. Let M be a burning-based mechanism under which $m - i$ units are burned with probability p_i. Let us consider the type profile in which agent 1's marginal values are $(0, 0, \ldots, 0, 1)$. That is, agent 1's valuation equals 1 if she wins all the units, and 0 otherwise. Let agent 2's type be the same. For this type profile, with probability p_m, the efficiency equals 1, and the payment equals 1. With probability $1 - p_m$, the efficiency equals 0, and the payment equals 0. That is, the expected efficiency is p_m, and we try to redistribute as much as possible from the expected payment p_m. The welfare is maximized when $p_m = 1$ (nothing is burned). That is, among all burning-based mechanisms, nothing can outperform the VCG mechanism

---

[19]This observation is due to [4].

with redistribution on this profile. However, when there are only two agents, neither agent can collect any redistribution: for either agent, the minimal total payment is always 0 (when we fix the other agent's type and change her own type), because if her type is $(0, 0, \ldots, 0)$, then the total payment is 0. Therefore, no burning-based mechanism can result in greater welfare on this profile than the VCG mechanism itself. However, under the VCG mechanism, the agents' total welfare is 0 for this type profile. Thus, when there are two agents, no burning-based mechanism can be competitive at all.

Next, we prove the result for the case of at least three agents. Let M be a burning-based mechanism under which $m - i$ units are burned with probability p_i, and suppose that its competitive ratio is strictly more than $\frac{n-1}{n}$. Let us consider the type profile in which agent 1's marginal values are $(0, 0, \ldots, 0, 1)$. That is, agent 1's valuation equals 1 if she wins all the units, and 0 otherwise. Let all the other agents have only marginal values of 0. That is, the other agents are not interested in any units. For this type profile, it is clear that if $p_m \leq \frac{n-1}{n}$, then the mechanism is at most $\frac{n-1}{n}$-competitive. That is, we must have $p_m > \frac{n-1}{n}$. Now, let us consider another type profile in which agent 1's type is the same, but we let agent 2's marginal values be $(0, 0, \ldots, 0, 1 - \epsilon, 0)$, where ϵ is a very small positive value. That is, agent 2's valuation equals $1 - \epsilon$ if she wins at least $m - 1$ units. Let all the other agents' marginal values all be 0. For this type profile, burning more than one unit results in a welfare of 0, and as we have proved, the probability of burning one unit is less than $\frac{1}{n}$. If agent 1's type were $(0, 0, \ldots, 0)$ instead, then the total payment would be 0 under M. (The total payment would be 0 both under the VCG mechanism, and under the VCG mechanism with one unit burned.) It follows that agent 1's redistribution is 0 for the type profile considered. Similarly, agent 2's redistribution is also 0. For any other agent, if we were to replace her type by $(1 + \epsilon, 0, 0, \ldots, 0)$, then the total payment under the VCG mechanism would be ϵ, and the total payment under the VCG mechanism with one unit burned would be $1 - \epsilon$. Since $p_m > \frac{n-1}{n}$, the expected total payment would be at most $\frac{n-1}{n}\epsilon + \frac{1}{n}(1 - \epsilon) = \frac{n-2}{n}\epsilon + \frac{1}{n}$. Therefore, each agent other than agent 1 and 2 at most receives a redistribution of $\frac{n-2}{n}\epsilon + \frac{1}{n}$. Thus, the total redistribution is at most $\frac{(n-2)^2}{n}\epsilon + \frac{n-2}{n}$. The total efficiency for this type profile is at most 1. The VCG payment for this type profile is $1 - \epsilon$. The VCG payment with one unit burned is 0. Therefore, the total payment is at least $\frac{n-1}{n} - \frac{n-1}{n}\epsilon$. So, the welfare is at most $1 - \frac{n-1}{n} + \frac{n-1}{n}\epsilon + \frac{(n-2)^2}{n}\epsilon + \frac{n-2}{n} = \frac{n-1}{n} + \frac{(n-2)^2 + n - 1}{n}\epsilon$. The welfare under the first-best mechanism is 1. Since ϵ can be arbitrarily small, the burning-based mechanism's competitive ratio cannot be strictly more than $\frac{n-1}{n}$, and we have the desired contradiction. $\square$