Efficient Money Burning in General Domains*

Dimitris Fotakis[1], Dimitris Tsipras[2], Christos Tzamos[2], and Emmanouil Zampetakis[2]

[1] School of Electrical and Computer Engineering,
National Technical University of Athens, 157 80 Athens, Greece
[2] Computer Science and Artificial Intelligence Laboratory,
Massachusetts Institute of Technology, Cambridge, MA 02139, U.S.A.
Emails: fotakis@cs.ntua.gr, tsipras@mit.edu, tzamos@mit.edu, mzampet@mit.edu

Abstract. We study mechanism design where the objective is to maximize the *residual surplus*, i.e., the total value of the outcome minus the payments charged to the agents, by truthful mechanisms. The motivation comes from applications where the payments charged are not in the form of actual monetary transfers, but take the form of wasted resources. We consider a general mechanism design setting with m discrete outcomes and n multidimensional agents. We present two randomized truthful mechanisms that extract an $O(\log m)$ fraction of the maximum social surplus as residual surplus. The first mechanism achieves an $O(\log m)$-approximation to the social surplus, which is improved to an $O(1)$-approximation by the second mechanism. An interesting feature of the second mechanism is that it optimizes over an appropriately restricted space of probability distributions, thus achieving an efficient tradeoff between social surplus and the total amount of payments charged to the agents.

1 Introduction

The extensive use of monetary transfers in Mechanism Design is due to the fact that so little can be implemented truthfully in their absence (see e.g., [20]). On the other hand, if monetary transfers are available (and their use is acceptable and feasible in the particular application), the famous Vickrey-Clarke-Groves (VCG) mechanism succeeds in truthfully maximizing the *social surplus* (a.k.a. the social welfare, that is the total value generated for the agents), albeit with possibly very large monetary transfers from the agents to the mechanism. In many typical mechanism design settings (e.g., mechanisms for public good allocation or auctions for allocation of private goods), large monetary transfers are acceptable just because the revenue of the mechanism is not lost and may be either redistributed among the agents (see e.g., [11, 12]) or invested in favor of the society.

However, there are mechanism design settings where monetary transfers are not acceptable (due to practical or ethical reasons) and the payments required for truthful implementation take the form of wasted resources, a.k.a. *money burning*, instead of actual monetary transfers. One could think of the time wasted in "computational" challenges (e.g., captcha) or in waiting

* This research was partially supported by the project AlgoNow, co-financed by the European Union (European Social Fund - ESF) and Greek national funds, through the Operational Program "Education and Lifelong Learning"of the National Strategic Reference Framework (NSRF) – Research Funding Program: THALES, investing in knowledge society through the European Social Fund, by NSF Award CCF-1101491, and by an award for Graduate Students in Theoretical Computer Science by the Simons Foundation. Part of this work was done while D. Tsipras was with the School of Electrical and Computer Engineering, National Technical University of Athens, Greece.

queues or lists (e.g., in hospitals [2] or in popular events or places), where each agent's waiting time serves as an implicit proof of how much the agent values the service (see also [13, 4] for more examples in the same direction). Assuming that the value of the wasted resources is measured in the same unit as the agent valuations, the natural objective in such settings is to maximize the net gain of the agents. This is quantified by the social surplus minus the payments charged to the agents and is usually referred to as the *residual surplus* (a.k.a. the social utility).

Within the Algorithmic Game Theory community, the general idea of money burning and residual surplus maximization by truthful mechanisms was first considered by Hartline and Roughgarden [13]. They considered single-unit and k-unit (unit demand) auctions and presented a family of truthful prior-free mechanisms that guarantee at least a constant fraction of the optimal (w.r.t. the residual surplus) Bayesian mechanism. Their mechanisms randomize among a VCG auction and a randomized posted price mechanism. To show that these mechanisms achieve an $O(1)$-approximation to the residual surplus extracted by an optimal Bayesian mechanism with complete knowledge of the agents' distribution (under the i.i.d. assumption), Hartline and Roughgarden used Myerson's theorem and characterized the optimal Bayesian mechanism for single-parameter agents. They also proved that if we compare the residual surplus of a truthful mechanism to the maximum social surplus, the best possible approximation guarantee for k-unit unit-demand auctions is $\Theta(1+\log(n/k))$, where n denotes the number of agents participating in the auction.

1.1 Contribution and Techniques

In this work, we consider residual surplus maximization by truthful mechanisms in a general mechanism design setting with m discrete outcomes and n multidimensional agents with general nonnegative valuations over the outcomes. Due to the fact that residual surplus maximization is closely related to revenue maximization, establishing a characterization of the optimal (w.r.t. the residual surplus) truthful Bayesian mechanism, as in [13], for multidimensional agents is a daunting task and far beyond the scope of this work. Actually, such a characterization not only would allow for a strong approximation guarantee (e.g., a constant approximation ratio) for the residual surplus, but would also provide a much better understanding of revenue-optimal mechanisms for multidimensional agents. Characterizing revenue-optimal mechanisms, even in relatively simple domains with multiple goods, is an important and extensively studied problem in mechanism design (see e.g., the surveys [22, 16] for some of the previous work on the problem). Therefore, instead of characterizing the optimal Bayesian mechanism and comparing the residual surplus of our mechanisms against the optimal residual surplus, we evaluate the performance of our mechanisms by comparing their residual surplus against the maximum social surplus. In fact, we present mechanisms that achieve nontrivial approximation guarantees w.r.t. both the residual surplus and the social surplus. Our main contribution is two randomized truthful mechanisms that approximate residual surplus within a best possible factor of $O(\log m)$, thus extending [13, Theorem 5.2] to general multidimensional domains.

Probably the simplest candidate mechanisms for residual surplus maximization are the random allocation, where each outcome is selected with probability $1/m$, and the VCG mechanism. It is not hard to see that the approximation ratio of the random allocation for both the residual surplus and the social surplus is m. Moreover, VCG cannot approximate the residual surplus within a factor better than m even for the simple case of m uniform i.i.d. single-minded agents (see Proposition 1). A natural way to approximate residual surplus is through a careful tradeoff between VCG, which optimizes the social surplus, but may result in a poor residual surplus due to high payments, and the random allocation on appropriately selected subsets of outcomes, which is truthful without payments and thus, translates all the social surplus into residual surplus.

Exploiting this intuition and building on the mechanism of [13, Theorem 5.2], we present a randomized truthful mechanism that approximates both the residual surplus and the social surplus within a factor of $O(\log m)$ (Theorem 2). The idea of the mechanism is to draw a random integer j from 0 to $\log m$, select a random outcome i among the best (in total value) 2^j outcomes and apply VCG payments. Hartline and Roughgarden [13, Theorem 5.2] proved that in k-unit unit-demand auctions with n agents, this mechanism is truthful and that its social and residual surplus approximate the maximum social surplus within a factor of $\Theta(1 + \log(n/k))$. The key step in our analysis is to show that in terms of residual surplus maximization, the worst-case instances correspond to single item auctions (Lemma 1). Then, the upper bound of [13, Theorem 5.2] carries over to our multidimensional setting. Moreover, since the single item auction is a special case of our setting, the lower bound of [13, Proposition 5.1] implies that our approximation ratio is asymptotically tight.

For randomized mechanisms with m discrete possible outcomes, the space of all possible random allocations coincides with the space of all possible probability distributions with support size m. If a mechanism results in a particular random allocation over the outcomes, then each outcome is chosen and implemented with the corresponding probability. Hence, in general, one may regard the feasible region of a randomized mechanism as the m-dimensional unit simplex, which contains all probability distributions with support size m.

Building on this understanding, our second mechanism optimizes the social surplus (using VCG) over a carefully defined subspace of the m-dimensional unit simplex. Intuitively, if we optimized over the unit simplex, we could achieve an optimal social surplus, but with a poor residual surplus, due to the high payments when the two best outcomes are close in total value. So, we define a subspace that is slightly curved close to the vertices of the m-dimensional unit simplex (see also Figure 1), thus achieving a significant decrease in the payments if the best outcomes are close in total value. Due to this fact, our second mechanism is *partial*, in the sense that with probability $1 - \varepsilon$, it may not implement any outcome. For any $\varepsilon > 0$, the approximation ratio is $1 + \varepsilon$ for the social surplus and $O(\frac{(1+\varepsilon)^2}{\varepsilon} \log m)$ for the residual surplus (Theorem 3). Hence, this mechanism achieves an essentially best possible approximation ratio for the residual surplus and a constant approximation to the social surplus, significantly improving on our first mechanism. The main idea behind this mechanism is that by restricting the solution space appropriately, we can achieve a tradeoff between the social surplus and the total amount of payments charged to the agents. Moreover, for mechanisms that achieve an almost optimal social surplus, the payments required for truthfulness decrease significantly faster than the resulting social surplus. We believe that such mechanisms, which are based on carefully chosen restricted subspaces and provide smooth tradeoffs between approximation ratio and payments, are of independent interest and may find other applications in mechanism design settings with restricted payments.

Our mechanisms run in time polynomial in the total number of outcomes m and in the number of agents n. In domains that allow for succinct input representation (e.g., Combinatorial Auctions, Combinatorial Public Projects), m is usually exponential in the size of the input. This is not surprising, since our approximation guarantees are significantly better than known lower bounds on the polynomial-time approximability of several **NP**-hard optimization problems. In certain domains, we can combine our mechanisms with existing Maximal-in-Range mechanisms so that the combined mechanism runs in time polynomial in the number of agents n and in the cardinality of the ground set on which the set of outcomes is defined (e.g., this coincides with the number of items in Combinatorial Auctions and in Combinatorial Public Projects), even though the number of outcomes m may be exponential in these parameters. For subadditive Combinatorial Public Projects, for example, we can use the Maximal-in-Range mechanism of [23, Section 3.2] and obtain a randomized polynomial-time truthful mechanism with $O(\min\{k, \sqrt{u}\})$-approximation to the social surplus and $O(\min\{k, \sqrt{u}\} \log u)$-approximation to the residual surplus, where u is the number of items and k is the size of the project.

1.2 Related Work

There is much work on (mostly polynomial-time) truthful mechanisms with monetary transfers that seek to maximize (exactly or approximately) the social surplus. In this general agenda, our work is closest in spirit to mechanisms with frugal payments (see e.g., [1, 7]). Moreover, our partial allocation mechanism was inspired by the work of Cole, Gkatzelis and Goel [5], where they present a truthful partial allocation mechanism that does not resort to monetary transfers and achieves an $1/e$-approximate proportionally fair division of divisible items by wasting roughly an $1/e$ fraction of each agent's value in the optimal allocation. Interestingly, the wasted item value for each agent in [5] is equal to the corresponding VCG payment. Thus, a partial allocation of items is used to simulate VCG payments and the mechanism becomes truthful without monetary transfers. This, in turn, implies that the entire social surplus of the mechanism is translated into residual surplus. In our second mechanism, we use the idea of partial allocations to decrease the total amount of payments required for truthfulness, so that the residual surplus is within a logarithmic factor of the mechanism's (and the optimal) social surplus.

Prior to [13], Chakravarty and Kaplan [4] characterized the Bayesian mechanism of maximum residual surplus in multi-unit (unit demand) auctions. More recently, Braverman et al. [2] considered residual surplus optimization in health care service allocation, but they focused on the complexity of computing efficient equilibrium allocations, instead of approximately truthful mechanisms.

An orthogonal direction is that of revenue redistribution (see e.g, [3, 11, 12] and the references therein). Although most of the literature focuses on maximizing the amount of redistributed VCG payments, some positive results in this direction concern residual surplus optimization relaxing the requirement for social surplus maximization (see e.g., [12]). Our viewpoint and results are incomparable, both technically and conceptually, to those in the area of redistribution mechanisms. A crucial difference is that in any efficient redistribution mechanism, certain agents should receive payments (this is unavoidable if one insists on efficiency and individual rationality, see e.g., [15]). In our setting, where payments are in the form of wasted social surplus, such redistribution is infeasible.

1.3 Organization

We start, in Section 2, with formally introducing the notation, the model, and some well known technical facts. In Section 3, we show that VCG cannot achieve any nontrivial approximation ratio to the residual surplus, even in a very simple setting, and explain why the best possible approximation ratio to the residual surplus is at least logarithmic in the number of outcomes. In Section 4, we present a randomized truthful mechanism that approximates both the residual surplus and the social surplus within a logarithmic factor. In Section 5, we present a truthful randomized mechanism, based on the idea of smooth partial allocations, which achieves a logarithmic approximation ratio to the residual surplus and a constant approximation ratio to the social surplus. We conclude with some discussion and some directions for further research in Section 6. An extended abstract of this work appears in [9].

2 Notation and Preliminaries

For any integer m, we let $[m] := \{1, \ldots, m\}$. We denote the j-th coordinate of a vector $\boldsymbol{v}$ by v_j. For a vector $\boldsymbol{v} = (v_1, \ldots, v_m)$ and an index $i \in [m]$, $\boldsymbol{v}_{-i}$ denotes $\boldsymbol{v}$ without coordinate i. For a vector $\boldsymbol{v} \in \mathbb{R}^m$ and some $\ell \geq 1$, $\boldsymbol{v}^\ell := (v_1^\ell, \ldots, v_m^\ell)$ denotes the coordinate-wise ℓ-th power of $\boldsymbol{v}$ and $\|\boldsymbol{v}\|_\ell := (\sum_{j=1}^m v_j^\ell)^{1/\ell}$ is the ℓ-norm of $\boldsymbol{v}$. For convenience, we let $|\boldsymbol{v}| := \|\boldsymbol{v}\|_1$ denote the 1-norm and $\|\boldsymbol{v}\|_\infty := \max_{j\in[m]}\{v_j\}$ denote the infinity norm of $\boldsymbol{v}$.

2.1 The Setting

We consider a general mechanism design setting with a finite set of possible outcomes O. We denote the number of outcomes $|O|$ by m. There is a set of n strategic agents, each with a private non-negative value for every outcome. The *valuation* of each agent i is given by a vector $\boldsymbol{v}_i \in \mathbb{R}_+^m$, where v_{ij} is the valuation of agent i for outcome j. We refer to the vector of all valuations $\boldsymbol{v} = (\boldsymbol{v}_1, \ldots, \boldsymbol{v}_n)$ as a *valuation profile*. For a valuation profile $\boldsymbol{v}$, $\boldsymbol{w}(\boldsymbol{v}) := \boldsymbol{v}_1 + \ldots + \boldsymbol{v}_n$ is the vector with the total value (or simply, the *weight*) of each outcome. We usually write $\boldsymbol{w}$, instead of $\boldsymbol{w}(\boldsymbol{v})$, and $\boldsymbol{w}_{-i}$, instead of $\boldsymbol{w}(\boldsymbol{v}_{-i})$, when $\boldsymbol{v}$ is clear from the context. For a valuation profile $\boldsymbol{v}$ and valuation $\boldsymbol{z}$, $(\boldsymbol{v}_{-i}, \boldsymbol{z})$ denotes the profile where vector $\boldsymbol{v}_i$ is replaced by vector $\boldsymbol{z}$.

2.2 Allocation Rules and Mechanisms

For a finite set S, $\Delta(S)$ denotes the unit simplex over S, that is the set of vectors in $\mathbb{R}_+^{|S|}$ with non-negative coordinates summing up to 1. Throughout this work, $\Delta(S)$ should be understood as the set of all probability distributions over S. A (randomized) allocation rule is a function $f : (\mathbb{R}_+^m)^n \to \Delta(O)$, mapping valuation profiles to m-dimensional probability vectors over the outcomes. Then $f_j(\boldsymbol{v})$ denotes the probability that outcome j is chosen and implemented by the randomized allocation rule f on valuation profile $\boldsymbol{v}$. Hence, outcome j is selected with probability equal to the j-th coordinate of the m-dimensional probability vector output by the randomized allocation rule f. This notation is very convenient since it allows us to think of randomized allocation rules as m-dimensional vectors in the unit simplex $\Delta(O)$ over the set of outcomes O. Then, several quantities of interest, such as the expected utility, the expected social surplus, the payments and the expected residual surplus can be expressed simply as the dot-product of two m-dimensional vectors.

Specifically, the expected value of an agent i is equal to $\boldsymbol{v}_i \cdot f(\boldsymbol{v})$. Throughout this work, we consider allocation rules that are *strongly anonymous*, in the sense that $f(\boldsymbol{v})$ depends only on the total value $\boldsymbol{w}(\boldsymbol{v})$ of the outcomes. Hence, we sometimes write the allocation rules as functions of the weight vector $\boldsymbol{w}(\boldsymbol{v})$ (or simply, of $\boldsymbol{w}$, when $\boldsymbol{v}$ is clear from the context).

A *payment rule* is a function $p : (\mathbb{R}_+^m)^n \to \mathbb{R}^n$ mapping valuation profiles to a payment for each agent. A *mechanism* is a pair $\mathcal{M} = (f, p)$ which on a valuation profile $\boldsymbol{v}$, outputs a probability vector $f(\boldsymbol{v})$ over the outcomes and charges each agent i an amount of $p_i(\boldsymbol{v})$.

The expected *utility* of agent i on a valuation profile $\boldsymbol{v}$ under mechanism $\mathcal{M} = (f, p)$ is

$$u_i(\boldsymbol{v}) := \boldsymbol{v_i} \cdot f(\boldsymbol{v}) - p_i(\boldsymbol{v})\,.$$

We assume that each agent i aims to maximize his expected utility u_i.

We require that our mechanisms are truthful and individually rational in expectation. A randomized mechanism $\mathcal{M} = (f, p)$ is *truthful* (in expectation) if for every agent i, any valuation profile $\boldsymbol{v}$ and any possible valuation $\boldsymbol{v}'_i$ of agent i,

$$\boldsymbol{v}_i \cdot f(\boldsymbol{v}) - p_i(\boldsymbol{v}) \geq \boldsymbol{v}_i \cdot f(\boldsymbol{v}_{-i}, \boldsymbol{v}'_i) - p_i(\boldsymbol{v}_{-i}, \boldsymbol{v}'_i)\,,$$

For brevity, we usually write that a randomized mechanism is truthful, with the understanding that this means that the mechanism is truthful in expectation (unless stated otherwise).

A randomized mechanism $\mathcal{M} = (f, p)$ is *individually rational* (in expectation) if for every agent i and any valuation profile $\boldsymbol{v}$,

$$\boldsymbol{v}_i \cdot f(\boldsymbol{v}) - p_i(\boldsymbol{v}) \geq 0\,.$$

2.3 Objectives and Approximation

Let $\mathcal{M} = (f, p)$ be some mechanism and let $\boldsymbol{v}$ be any valuation profile. We denote the total amount of payments of $\mathcal{M}$ on input $\boldsymbol{v}$ by

$$P(\boldsymbol{v}) := \sum_i p_i(\boldsymbol{v}).$$

In this work, we are interested in maximizing the social surplus and the residual surplus. The *social surplus* of $\mathcal{M}$ on $\boldsymbol{v}$ is

$$S(\boldsymbol{v}) := \sum_i \boldsymbol{v}_i \cdot f(\boldsymbol{v}) = \boldsymbol{w} \cdot f(\boldsymbol{v}) .$$

The *residual surplus* of $\mathcal{M}$ on $\boldsymbol{v}$ is

$$R(\boldsymbol{v}) := S(\boldsymbol{v}) - P(\boldsymbol{v}) = \sum_i u_i(\boldsymbol{v}) .$$

The social surplus and residual surplus of any mechanism on valuation profile $\boldsymbol{v}$ is at most $\|\boldsymbol{w}(\boldsymbol{v})\|_\infty$. We say that a mechanism $\mathcal{M}$ is ρ-approximate for the social surplus (resp. the residual surplus) if for any valuation profile $\boldsymbol{v}$, $S(\boldsymbol{v}) \geq \|\boldsymbol{w}(\boldsymbol{v})\|_\infty/\rho$ (resp. $R(\boldsymbol{v}) \geq \|\boldsymbol{w}(\boldsymbol{v})\|_\infty/\rho$). We say that a mechanism $\mathcal{M}$ is (ρ_1, ρ_2)-approximate for the social surplus and the residual surplus if $\mathcal{M}$ is ρ_1-approximate for the social surplus and ρ_2-approximate for the residual surplus. For brevity, we usually write simply that $\mathcal{M}$ is (ρ_1, ρ_2)-approximate, without explicitly referring to the social surplus and to the residual surplus.

2.4 Maximal-in-Distributional-Range Mechanisms

We mostly consider randomized allocation rules f such that there is a subset $S \subseteq \Delta(O)$ of probability distributions over O (a.k.a. the *range* of f) over which f optimizes on any input $\boldsymbol{v}$. Namely, for any valuation profile $\boldsymbol{v}$, $f(\boldsymbol{v}) = \arg\max_{\boldsymbol{s} \in S} \boldsymbol{s} \cdot \boldsymbol{w}$. Then, any mechanism $\mathcal{M} = (f, p)$, which is based on such a randomized allocation rule f and uses VCG payments $p_i(\boldsymbol{v}) = \boldsymbol{w}_{-i} \cdot f(\boldsymbol{v}_{-i}) - \boldsymbol{w}_{-i} \cdot f(\boldsymbol{v})$, is called a *Maximal-in-Distributional-Range* (MIDR) mechanism. From the analysis of the VCG mechanism (see e.g., [20]), we know that MIDR mechanisms are truthful in expectation and individually rational in expectation. Mechanism 1, presented in Section 4, is a random selection of the MIDR mechanisms introduced in Definition 1. The mechanism presented in Section 5 is MIDR. These mechanisms achieve truthfulness in expectation through the use of VCG payments.

We observe that using a simple payment transformation described in [8], an MIDR mechanism (with VCG payments) can become truthful in expectation and individually rational in the universal sense. Specifically, let us consider some agent i and let $p_i(\boldsymbol{v}) = \boldsymbol{w}_{-i} \cdot f(\boldsymbol{v}_{-i}) - \boldsymbol{w}_{-i} \cdot f(\boldsymbol{v})$ denote the VCG payment for agent i on some valuation profile $\boldsymbol{v}$. When outcome j is selected, we charge agent i with $p_{ij} = \frac{p_i(\boldsymbol{v})}{\boldsymbol{v}_i \cdot f(\boldsymbol{v})} v_{ij}$. It is not hard to verify that the expected payments do not change, and thus, truthfulness is not affected. Moreover, since $p_i(\boldsymbol{v}) \leq \boldsymbol{v}_i \cdot f(\boldsymbol{v})$, because $\boldsymbol{w}_{-i} \cdot f(\boldsymbol{v}_{-i}) \leq \boldsymbol{w}_{-i} \cdot f(\boldsymbol{v}) + \boldsymbol{v}_i \cdot f(\boldsymbol{v}) = \boldsymbol{w} \cdot f(\boldsymbol{v})$, the mechanism is now individually rational. Hence, we no longer distinguish between mechanisms that are individually rational in expectation and mechanisms that are individually rational in the universal sense.

3 A Lower Bound on the Approximability of Residual Surplus

In social surplus maximization, monetary transfers can be used freely to truthfully elicit the agents' preferences. In residual surplus maximization, on the other hand, the transfers needed for the truthful implementation of some mechanisms may comprise a significant portion of the social surplus, thus prohibiting any non-trivial approximation guarantees.

To obtain a lower bound on the approximability of the residual surplus by truthful mechanisms, we observe that the single-item auction can be easily cast as a special case of our general mechanism design setting. To show this, we restrict our attention to m outcomes and m agents, with each agent i having a value $v_i \geq 0$ for outcome i and a value of 0 for each of the remaining outcomes. For brevity, we call such agents *single minded*. Hence, it suffices to show a lower bound on the approximability of the residual surplus by truthful mechanisms for the special case of single-minded agents. However, even in this special case, such a lower bound cannot be simply based on a single valuation profile where every mechanism would perform poorly. The reason is that the trivial dictatorial mechanism can output the optimal allocation on any fixed valuation profile and charge zero payments. We therefore need to evaluate mechanisms over a large collection of valuation profiles. To this end, we use Myerson's characterization of truthful single-item auctions.

Theorem 1 (Myerson [19]). *Let $\mathcal{M} = (f, p)$ be any truthful mechanism and let $\boldsymbol{v}$ any valuation profile, where each agent i has some value $v_i > 0$ only for outcome i and v_i is drawn independently from a probability distribution $\mathcal{D}_i$ with cumulative distribution function $D_i(v)$ and probability density function $d_i(v)$. Then,*

$$\mathbb{E}[P(\boldsymbol{v})] = \mathbb{E}[\boldsymbol{\phi} \cdot f(\boldsymbol{v})]\,, \quad \textit{with} \ \ \phi_i = v_i - \frac{1 - D_i(v_i)}{d_i(v_i)}\,.$$

Theorem 1 determines the expected amount of payments charged by any truthful allocation rule. This, in turn, determines the expected residual surplus in terms of the allocation rule. By plugging in an appropriate distribution, we come up with lower bounds on the residual surplus of truthful mechanisms.

Proposition 1. *The Vickrey Auction cannot approximate residual surplus within a factor better than m.*

Proof. Assume that the agent values are drawn from the uniform distribution over $[0, 1]$. For the uniform distribution $\phi_i = 2v_i - 1$. By direct calculations, we obtain that $\mathbb{E}[R(\boldsymbol{v})] = 1 - \mathbb{E}[\|\boldsymbol{w}(\boldsymbol{v})\|_\infty]$. The expected maximum value of m independent and identically distributed (i.i.d.) uniform random variables in $[0, 1]$ is $\frac{m}{m+1}$. Therefore, $\mathbb{E}[R(\boldsymbol{v})] = \mathbb{E}[\|\boldsymbol{w}(\boldsymbol{v})\|_\infty]/m$. By the probabilistic method, there exists a profile $\boldsymbol{v}$ for which $R(\boldsymbol{v}) \leq \|\boldsymbol{w}(\boldsymbol{v})\|_\infty/m$. □

The proof of Proposition 1 shows that the VCG mechanism for the natural case of uniform i.i.d. single-minded agents approximates the residual surplus within a factor no better than m. Namely, VCG on this simple class of instances achieves an approximation ratio no better than the worst case approximation ratio of the simple random allocation. The reason for the poor performance of VCG on these instances is that the expectation of the second largest value of m i.i.d. uniform random variables in $[0, 1]$ is $\frac{m-1}{m+1}$, i.e., very close to their expected maximum value. So, the expected payment is very large and the expected residual surplus is only $1/(m + 1)$. On the other hand, the random allocation on this simple class of instances achieves an approximation ratio of roughly 2 for both the social surplus and the residual surplus, since the expected value of any agent is 1/2, while the maximum value is less than 1 and no payments are required. However, the standard trick of combining VCG and the random allocation fails to achieve a good approximation ratio to the residual surplus on simple instances with single-minded agents[3].

More generally, the reason that VCG fails to achieve a good approximation ratio is that aiming to maximize the social surplus, it has to charge every agent his critical value. This results in a large total amount of payments and does not allow the high social surplus of the VCG mechanism to be translated into an almost as high residual surplus. We therefore need to come up with mechanisms that instead of maximizing the social surplus, employ suboptimal allocations to reduce payments, while preserving a significant amount of social surplus. Our goal is to achieve the best possible worst-case guarantee for residual surplus maximization. A lower bound on the best approximation ratio in our setting can be obtained from [13, Proposition 5.1]. Since the proof is simple and informative, we include it here for completeness.

Proposition 2 (Hartline and Roughgarden [13]). *No truthful mechanism can approximate residual surplus within a factor of $o(\log m)$.*

Proof. We consider m agents drawn from the exponential distribution. Then, in Theorem 1, we have that $d_i(v) = e^{-v}$, $D_i(v) = 1 - e^{-v}$ and $\phi_i = v_i - 1$ for all agents i. Hence, we obtain that

$$\mathbb{E}[P(\boldsymbol{v})] = \mathbb{E}[\boldsymbol{w} \cdot f(\boldsymbol{v}) - |f(\boldsymbol{v})|] = \mathbb{E}[S(\boldsymbol{v}) - 1] .$$

By linearity of expectation, $\mathbb{E}[R(\boldsymbol{v})] = 1$. It is not hard to verify that the expected maximum value of m i.i.d. exponential random variables is equal to H_m, where H_m denotes the m-th harmonic number. Then

$$\mathbb{E}[R(\boldsymbol{v})] = \mathbb{E}\left[\frac{\|\boldsymbol{w}\|_\infty}{H_m}\right] .$$

Therefore, by the probabilistic method, there is a valuation profile $\boldsymbol{v}$ for which the approximation ratio for the residual surplus is at least logarithmic. □

4 Achieving Best-Possible Residual Surplus Guarantees

In this section, we present a randomized truthful mechanism that approximates both the social surplus and the residual surplus within a logarithmic factor for multidimensional agents. Our mechanism builds on the mechanism of [13, Theorem 5.2], which is truthful and approximates both the social surplus and the residual surplus of a k-unit unit-demand auction with m agents within a factor of $\Theta(1 + \log(m/k))$. For single-item auctions, the mechanism restricts

[3] For example, we consider an instance with m agents and m outcomes, where each agent i has some value $v_i > 0$ only for outcome i and $v_1 = 1$, $v_2 = 1 - \varepsilon$ and $v_3 = \cdots = v_m = \varepsilon$, for any small $\varepsilon > 0$. Then, the residual surplus of VCG is ε and the expected residual surplus of the random allocation is $\varepsilon + \frac{2-\varepsilon}{m}$. Since the optimal social surplus is 1, any randomization between VCG and the random allocation yields an approximation ratio of $\Omega(1/\varepsilon)$.

Mechanism 1 An $(O(\log m), O(\log m))$-approximate mechanism.

Choose j uniformly at random from $\{0, 1, 2, \ldots, \log m\}$ and let $k \leftarrow 2^j$

Output the probability vector $\mathrm{Top}_k(\boldsymbol{v})$ over outcomes

Charge each agent i with $p_i^k(\boldsymbol{v}) = \boldsymbol{w}_{-i} \cdot \mathrm{Top}_k(\boldsymbol{v}_{-i}) - \boldsymbol{w}_{-i} \cdot \mathrm{Top}_k(\boldsymbol{v})$

the allocation to the agents with the 2^j highest bids, for an integer j, $0 \leq j \leq \log m$, selected uniformly at random. Then, an agent selected randomly among them gets the item and is charged with the $(2^j + 1)$-th highest bid (this is the critical bid for this group of agents, the payment is 0 if $j = \log m$). Intuitively, the mechanism performs well because for at least one integer j, the difference between the 2^{j-1}-th highest bid and the $(2^j + 1)$-th highest bid is within a constant factor of the maximum bid, which is equal to the maximum social surplus (see also the proof of Theorem 2).

Next, we extend this mechanism and its analysis to the general setting of m discrete outcomes and n multidimensional agents. We start with defining the following class of allocation rules, which generalize the allocation rules used in the mechanism of [13, Theorem 5.2].

Definition 1. *For some integer $k \in [m]$, the* Top_k *allocation rule on input $\boldsymbol{v}$, orders outcomes in non-increasing order of their weight, $w_1 \geq \ldots \geq w_m$ (breaking ties lexicographically) and assigns probability $1/k$ to the first k outcomes. Formally, $\mathrm{Top}_k(\boldsymbol{v}) = \arg\max_{\boldsymbol{s} \in S_k} \boldsymbol{s} \cdot \boldsymbol{w}$, where the range S_k is the set of all vectors in $\Delta(O)$ with k coordinates equal to $1/k$ and $m - k$ coordinates equal to 0 (in case of ties, the lexicographically smaller vector $\boldsymbol{s}$ is returned).*

By definition, the class of Top_k allocation rules maximizes the social surplus over the range $S_k \subseteq \Delta(O)$ consisting of all probability vectors with exactly k coordinates equal to $1/k$. Therefore, they can be turned into truthful and individually rational mechanisms using the VCG payment scheme, as explained in Section 2.4. We denote mechanisms of this family by $\mathcal{M}_k = (\mathrm{Top}_k, p^k)$. Each such mechanism achieves a different approximation guarantee, as a function of k, with respect to the social surplus and to the residual surplus. By randomizing over k, we can combine these guarantees, while maintaining truthfulness and individual rationality. In Mechanism 1, we use this approach to achieve a best possible approximation ratio for the residual surplus by randomizing over exponentially increasing values of k (as in the mechanism of [13, Theorem 5.2]). For simplicity we assume that m is a power of 2 (as we can always pad valuation vectors $\boldsymbol{v}_i$ with zero value outcomes).

Mechanism 1 is a randomization over an appropriate selection of $\mathcal{M}_k$ mechanisms, where the indices k are independent of the input $\boldsymbol{v}$. As a result, Mechanism 1 is truthful and individually rational. In order to quantify its efficiency w.r.t. the residual surplus, we first show that the worst-case instances for Mechanism 1 are those corresponding to the simple single-item auction, where there is a single agent with positive valuation for each outcome.

Lemma 1. *For any valuation profile $\boldsymbol{v} = (\boldsymbol{v_1}, \ldots, \boldsymbol{v_n})$, the residual surplus of Mechanism 1 on $\boldsymbol{v}$ is no less than its residual surplus on the valuation profile $\boldsymbol{y} = (\boldsymbol{y}_1, \ldots, \boldsymbol{y_m})$, where each $\boldsymbol{y_i} = (0, \ldots, w_i(\boldsymbol{v}), \ldots, 0)$ is the valuation of a single-minded agent i with $y_{ii} = w_i(\boldsymbol{v})$ for outcome i and $y_{ij} = 0$ for any other outcome $i \neq j$.*

Before we proceed with the formal proof of Lemma 1, let us provide the intuition behind the proof and describe its main steps. Since Mechanism 1 is a randomization over mechanisms $\mathcal{M}_k$, it suffices to show the lemma for each mechanism $\mathcal{M}_k$ separately. The proof of the lemma consists of two key steps:

- First, we show that if an agent has positive value for multiple outcomes, splitting this agent into single-minded agents (one for each outcome) can only decrease the residual surplus of the mechanism. This holds since the "competition" between agents increases, and as a result, so do the payments, thus decreasing the residual surplus. We note here that the social surplus is not affected, since the mechanism $\mathcal{M}_k$ is strongly anonymous and depends only on the weight of each outcome, which is not affected by this transformation. By induction, we transform any valuation profile to one with single-minded agents without increasing the residual surplus.
- Next, we show that if there are multiple single-minded agents for the same outcome, joining their values into a single agent can only decrease the residual surplus. The reason is that the value that the agents should "prove" (in the form of payments) to the mechanism is initially split among them, and can only increase as they aggregate their values. A single agent with high value is more critical for the auction than many agents with small values. Again by induction we can transform any valuation profile with single-minded agents to a valuation profile with one single-minded agent per outcome, without increasing the residual surplus.

Proof. We proceed to formally state and prove the two claims above. We fix an integer $k \in [m]$ and consider the corresponding mechanism $\mathcal{M}_k = (\text{Top}_k, p^k)$. For simplicity of notation, we let f denote the allocation of Top_k and let p denote the corresponding VCG payments throughout the proof. Let $\boldsymbol{v}$ be an arbitrary valuation profile and let i be any agent. The utility of agent i under mechanism $\mathcal{M}_k$ is

$$\begin{aligned} u_i(\boldsymbol{v}) &= \boldsymbol{v}_i \cdot f(\boldsymbol{v}) - p_i(\boldsymbol{v}) \\ &= \boldsymbol{v}_i \cdot f(\boldsymbol{v}) - (\boldsymbol{w}_{-i} \cdot f(\boldsymbol{v}_{-i}) - \boldsymbol{w}_{-i} \cdot f(\boldsymbol{v})) \\ &= \boldsymbol{w} \cdot f(\boldsymbol{v}) - \boldsymbol{w}_{-i} \cdot f(\boldsymbol{v}_{-i}) \,. \end{aligned} \tag{1}$$

If agent i is not single-minded, there are multiple coordinates of $\boldsymbol{v}_i$ that are strictly positive, let j be one of them. We split $\boldsymbol{v}_i$ into $\boldsymbol{y}_1 = v_{ij}\boldsymbol{e}_j$ and $\boldsymbol{y}_2 = \boldsymbol{v}_i - \boldsymbol{y}_1$, where $\boldsymbol{e}_j$ the unit vector in the direction j. We denote these two agents with indices i_1 and i_2 respectively. Clearly, splitting agent i into agents i_1 and i_2 does not affect the weight vector $\boldsymbol{w}$ of the outcomes (and thus the allocation of the mechanism). Moreover, for any agent different from i, splitting $\boldsymbol{v}_i$ into $\boldsymbol{y}_1$ and $\boldsymbol{y}_2$ does not change the weight vector $\boldsymbol{w}_{-i}$ of the other agents. Therefore, the utility of any agent different from i does not change. Hence, for the first claim, it suffices to show that the combined utility of agents i_1 and i_2 is no more than the utility of agent i. Formally, it suffices to show that

$$u_i(\boldsymbol{v}) \geq u_{i_1}(\boldsymbol{v}_{-i}, \boldsymbol{y}_1, \boldsymbol{y}_2) + u_{i_2}(\boldsymbol{v}_{-i}, \boldsymbol{y}_1, \boldsymbol{y}_2) \,. \tag{2}$$

We first observe that

$$\begin{aligned}
u_{i_2}(\boldsymbol{v}_{-i}, \boldsymbol{y}_1, \boldsymbol{y}_2) &= \boldsymbol{w} \cdot f(\boldsymbol{v}) - \boldsymbol{w}(\boldsymbol{v}_{-i}, \boldsymbol{y}_1) \cdot f(\boldsymbol{v}_{-i}, \boldsymbol{y}_1) \\
&= \boldsymbol{w} \cdot f(\boldsymbol{v}) - \boldsymbol{w}_{-i} \cdot f(\boldsymbol{v}_{-i}) - (\boldsymbol{w}(\boldsymbol{v}_{-i}, \boldsymbol{y}_1) \cdot f(\boldsymbol{v}_{-i}, \boldsymbol{y}_1) - \boldsymbol{w}_{-i} \cdot f(\boldsymbol{v}_{-i})) \\
&= u_i(\boldsymbol{v}) - u_{i_1}(\boldsymbol{v}_{-i}, \boldsymbol{y}_1)\,.
\end{aligned}$$

Then, we can rewrite (2) as

$$u_{i_1}(\boldsymbol{v}_{-i}, \boldsymbol{y}_1, \boldsymbol{y}_2) \leq u_{i_1}(\boldsymbol{v}_{-i}, \boldsymbol{y}_1)\,.$$

So, it suffices to show that the utility of a single-minded agent does not decrease when the total value of the competing outcomes decreases.

Since we break ties between outcomes in a consistent (lexicographic) way, tie-breaking is irrelevant to the analysis of the mechanism. Hence, we assume, for simplicity, that no ties occur. The mechanism allocates probability $1/k$ to the top k outcomes and 0 to the rest of them. Let w_{crit} denote the minimum weight amongst the outcomes with positive probability on valuation profile $\boldsymbol{v}_{-i}$. Then, applying (1) for the utility of the single-minded agent i_1, we obtain that

$$u_{i_1}(\boldsymbol{v}_{-i}, \boldsymbol{y}_1) = \begin{cases} |\boldsymbol{y}_1|/k & \text{if } \boldsymbol{w}_j(\boldsymbol{v}_{-i}) \geq w_{crit}, \\ 0 & \text{if } \boldsymbol{w}_j(\boldsymbol{v}_{-i}, \boldsymbol{y}_1) < w_{crit}, \\ (|\boldsymbol{y}_1| + \boldsymbol{w}_j(\boldsymbol{v}_{-i}) - w_{crit})/k & \text{if } \boldsymbol{w}_j(\boldsymbol{v}_{-i}, \boldsymbol{y}_1) \geq w_{crit} > \boldsymbol{w}_j(\boldsymbol{v}_{-i}). \end{cases} \tag{3}$$

Since $f_j(\boldsymbol{v}_{-i}, \boldsymbol{y}_1) = 1/k$ if $\boldsymbol{w}_j(\boldsymbol{v}_{-i}, \boldsymbol{y}_1) \geq w_{crit}$, and $f_j(\boldsymbol{v}_{-i}, \boldsymbol{y}_1) = 0$ otherwise, we can rewrite (3) as

$$u_{i_1}(\boldsymbol{v}_{-i}, \boldsymbol{y}_1) = f_j(\boldsymbol{v}_{-i}, \boldsymbol{y}_1)(|\boldsymbol{y}_1| + \min\{\boldsymbol{w}_j(\boldsymbol{v}_{-i}) - w_{crit}, 0\})\,. \tag{4}$$

Let now w'_{crit} denote the minimum weight of an outcome with positive probability on valuation profile $(\boldsymbol{v}_{-i}, \boldsymbol{y}_2)$. Repeating the analysis above with the valuation profile $(\boldsymbol{v}_{-i}, \boldsymbol{y}_2)$ in the place of the valuation profile $\boldsymbol{v}_{-i}$, we show that

$$u_{i_1}(\boldsymbol{v}_{-i}, \boldsymbol{y}_1, \boldsymbol{y}_2) = f_j(\boldsymbol{v}_{-i}, \boldsymbol{y}_1, \boldsymbol{y}_2)(|\boldsymbol{y}_1| + \min\{\boldsymbol{w}_j(\boldsymbol{v}_{-i}, \boldsymbol{y}_2) - w'_{crit}, 0\})\,. \tag{5}$$

By construction $\boldsymbol{w}_j(\boldsymbol{y}_2) = 0$, implying that $\boldsymbol{w}_j(\boldsymbol{v}_{-i}, \boldsymbol{y}_2) = \boldsymbol{w}_j(\boldsymbol{v}_{-i})$. Also, we have that $w'_{crit} \geq w_{crit}$, since valuation $\boldsymbol{y}_2$ is non-negative and the mechanism is monotone. Therefore,

$$\min\{\boldsymbol{w}_j(\boldsymbol{v}_{-i}, \boldsymbol{y}_2) - w'_{crit}, 0\} \leq \min\{\boldsymbol{w}_j(\boldsymbol{v}_{-i}) - w_{crit}, 0\}\,.$$

Moreover, since $\boldsymbol{w}_j(\boldsymbol{v}_{-i}, \boldsymbol{y}_1) = \boldsymbol{w}_j(\boldsymbol{v}_{-i}, \boldsymbol{y}_2, \boldsymbol{y}_1)$ and the total weight of any outcome in $\boldsymbol{w}(\boldsymbol{v}_{-i})$ does not exceed the total weight of the corresponding outcome in $\boldsymbol{w}(\boldsymbol{v}_{-i}, \boldsymbol{y}_2)$, we obtain that

$$f_j(\boldsymbol{v}_{-i}, \boldsymbol{y}_1, \boldsymbol{y}_2) \leq f_j(\boldsymbol{v}_{-i}, \boldsymbol{y}_1)\,.$$

Multiplying these inequalities (since f is non-negative), and using (4) and (5), we get that

$$u_{i_1}(\boldsymbol{v}_{-i}, \boldsymbol{y}_1, \boldsymbol{y}_2) \leq u_{i_1}(\boldsymbol{v}_{-i}, \boldsymbol{y}_1).$$

Hence, we have showed that by splitting the valuation $\boldsymbol{v}_i$ of agent i into an agent i_2 with valuation $\boldsymbol{y}_2$ and a single-minded agent i_1 with valuation $\boldsymbol{y}_1$, the residual surplus of the mechanism can only decrease. Applying the claim inductively, we can replace every agent with a set of single-minded agents, without increasing the residual surplus of the mechanism.

We now proceed to show that by joining the single-minded agents of an outcome can only decrease the residual surplus. The technical details are quite similar to the first part of the proof. Let $\boldsymbol{v}$ some valuation profile containing only single-minded agents. For some outcome j, assume there exist two agents with positive value for it, $\boldsymbol{v}_a$ and $\boldsymbol{v}_b$. We combine their valuations into a new valuation $\boldsymbol{v}_c = \boldsymbol{v}_a + \boldsymbol{v}_b$. Denoting by $\boldsymbol{v}_{-ab}$ the profile without them, we need to show that

$$u_c(\boldsymbol{v}_{-ab}, \boldsymbol{v}_c) \leq u_a(\boldsymbol{v}) + u_b(\boldsymbol{v}) \,,$$

which, by calculations similar to those after (2), is equivalent to

$$u_a(\boldsymbol{v}_{-ab}, \boldsymbol{v}_a) \leq u_a(\boldsymbol{v}_{-ab}, \boldsymbol{v}_a, \boldsymbol{v}_b) \,.$$

Thus, it suffices to prove that the utility of a single-minded agent increases when there are more single-minded agents bidding on the same outcome. Similarly to the previous part, let w_{crit} denote the mimimum weight of an outcome allocated positive probability on the valuation profile $\boldsymbol{v}_{-ab}$ and w'_{crit} the same quantity on the valuation profile $(\boldsymbol{v}_{-ab}, \boldsymbol{v}_b)$. Then, in parallel to (4),

$$u_a(\boldsymbol{v}_{-ab}, \boldsymbol{v}_a) = f_j(\boldsymbol{v}_{-ab}, \boldsymbol{v}_a)(|\boldsymbol{v}_a| + \min\{\boldsymbol{w}_j(\boldsymbol{v}_{-ab}) - w_{crit}, 0\}) \,,$$

$$u_a(\boldsymbol{v}_{-ab}, \boldsymbol{v}_a, \boldsymbol{v}_b) = f_j(\boldsymbol{v}_{-ab}, \boldsymbol{v}_a, \boldsymbol{v}_b)(|\boldsymbol{v}_a| + \min\{\boldsymbol{w}_j(\boldsymbol{v}_{-ab}, \boldsymbol{v}_b) - w'_{crit}, 0\}) \,.$$

Since $\boldsymbol{w}_i(\boldsymbol{v}_b) = 0$ for $i \neq j$ and $\boldsymbol{w}_j(\boldsymbol{v}_b) \geq 0$,

$$f_j(\boldsymbol{v}_{-ab}, \boldsymbol{v}_a) \leq f_j(\boldsymbol{v}_{-ab}, \boldsymbol{v}_a, \boldsymbol{v}_b) \text{ and } \boldsymbol{w}_j(\boldsymbol{v}_{-ab}) \leq \boldsymbol{w}_j(\boldsymbol{v}_{-ab}, \boldsymbol{v}_b) \,.$$

Multiplying the inequalities (since f is non-negative) and substituting, we conclude that

$$u_a(\boldsymbol{v}_{-ab}, \boldsymbol{v}_a) \leq u_a(\boldsymbol{v}_{-ab}, \boldsymbol{v}_a, \boldsymbol{v}_b) \,.$$

□

Using Lemma 1, we can lower bound the performance of Mechanism 1 in the general case, by its performance in the case of the single-item auction.

Theorem 2. *Mechanism 1 is truthful, individually rational and* $(O(\log m), O(\log m))$*-approximate for the social surplus and the residual surplus.*

Proof. We first observe that the transformation described in the proof of Lemma 1 does not affect the social surplus achieved by Top_k allocation rules. The reason is that Top_k allocations are strongly anonymous, i.e., the resulting probability vector over outcomes only depends on the total value $\boldsymbol{w}$ of the outcomes. It is not hard to verify that the weight vector $\boldsymbol{w}(\boldsymbol{v})$ remains unchanged during the agent splitting and the agent joining operations applied to $\boldsymbol{v}$ in the proof of Lemma 1. Therefore, applying Lemma 1 does not affect the social surplus achieved by Mechanism 1.

We can now use Lemma 1 and the analysis of [13, Theorem 5.2] and show that Mechanism 1 is $(O(\log m), O(\log m))$-approximate for the social surplus and for the residual surplus. For completeness, we provide a formal proof of the approximation ratio below.

Claim (Hartline and Roughgarden [13, Theorem 5.2]). In the special case of single-item auctions, Mechanism 1 is $(2(\log m + 1), 2(\log m + 1))$-approximate for the social surplus and the residual surplus.

Proof (of the Claim). For single-item auctions, the valuation of each agent i consists of a single number $v_i \geq 0$. Suppose that the agents are ordered in nonincreasing order of values, that is $v_1 \geq v_2 \geq \ldots \geq v_n$. The VCG payments of Top_k are identical for the top k agents and equal to v_{k+1}. Therefore, the expected residual surplus of mechanism $\mathcal{M}_k$ is

$$R_k(\boldsymbol{v}) = \sum_{i=1}^{k} (v_i - v_{k+1})/k\,.$$

For $k = 1$, $R_1(\boldsymbol{v}) = v_1 - v_2 \geq (v_1 - v_2)/2$. For $k > 1$, setting $u_{n+1} = 0$, we get that

$$R_k(\boldsymbol{v}) \geq (v_{k/2+1} - v_{k+1})/2\,.$$

Each mechanism $\mathcal{M}_k$ is invoked with probability $\frac{1}{1+\log m}$. So, we obtain that

$$R(\boldsymbol{v}) = \frac{1}{1+\log m} \sum_{k} (v_{k/2+1} - v_{k+1})/2 = \frac{v_1}{2(1+\log m)}\,,$$

where v_1 is the maximum social surplus (and the maximum residual surplus) that a mechanism could achieve. ☐

Lemma 1 shows that the single-item auction is the worst case instance for Mechanism 1 w.r.t. the residual surplus. Thus, the analysis of this case implies an upper bound on the approximation ratio of Mechanism 1. In addition to the fact that the expected social surplus of Mechanism 1 is not affected by Lemma 1 and thus, the approximation ratio for the social surplus is at most $2(\log m + 1)$ by the claim above, we observe that a slightly stronger approximation ratio for the social surplus follows directly from the fact that the outcome with the maximum total value is chosen with probability $1/(1 + \log m)$. ☐

The analysis above is essentially tight, as we can see by the case where $\boldsymbol{v} = (\boldsymbol{v_1}, 0, \ldots, 0)$ and each agent $i \in [n]$ is a single-minded agent with value v_1/n for outcome 1. Then, the resulting social surplus and residual surplus are

$$\frac{1}{1+\log m}\left(\sum_{i=0}^{\log_2 m} \frac{\|\boldsymbol{w}\|_\infty}{2^i}\right) \leq \frac{2\|\boldsymbol{w}\|_\infty}{1+\log m}\,.$$

Hence, the approximation ratio of $O(\log m)$ is tight for both the social surplus and the residual surplus.

5 Optimizing Residual Surplus without Sacrificing Social Surplus

The mechanism presented in Section 4 approximates residual surplus within a logarithmic factor, which is best possible. However, it does so by wasting a large portion of the optimal social surplus, since the approximation ratio for the social surplus is also logarithmic. The impossibility result of Proposition 2 does not imply that a better approximation ratio for social surplus is impossible. In this section, we present an MIDR mechanism that approximates both social and residual surplus within optimal factors simultaneously.

Theorem 3. *For any $\varepsilon > 0$, there is a truthful and individually rational mechanism $\mathcal{M}$ that is $(1+\varepsilon, \frac{(1+\varepsilon)^2}{\varepsilon}\ln m)$-approximate for the social surplus and the residual surplus.*

Remark 1. We should highlight that one can obtain the guarantees of Theorem 3 simply by randomizing, with some constant probability, between the VCG mechanism and Mechanism 1. Nevertheless, the mechanism of Theorem 3 follows from a more principled approach, which yields smooth allocation rules (in the analytical sense, i.e., the probability distribution of the mechanism is a smooth function of any component of the input), and may be of independent interest.

5.1 The Mechanism

Similarly to the previous mechanism, we need a careful tradeoff between the VCG mechanism and suboptimal allocations close to the uniform mechanism. We note that the VCG mechanism achieves an optimal social surplus by selecting the best outcome in the unit simplex $\Delta(O)$. Here, we optimize on a surface that is close to the unit simplex, but slightly curved towards the pure outcomes, in order to reduce the payments when the best outcomes are close in weight. To this end, we define a mechanism by optimizing on the following family of convex subsets of $\Delta(O)$:

$$S_k = \left\{ \boldsymbol{s} \in \mathbb{R}_+^m \mid \|\boldsymbol{s}\|_k \leq \frac{1}{m^{1-1/k}} \right\} . \tag{6}$$

For any $k \geq 1$ or for $k = \infty$, we define the allocation rule $f_k(\boldsymbol{v}) = \arg\max_{\boldsymbol{s}\in S_k} \boldsymbol{s} \cdot \boldsymbol{w}(\boldsymbol{v})$. Since this allocation rule optimizes over the convex set S_k, it can be combined with VCG payments and give an MIDR mechanism that is truthful in expectation and individually rational, as explained in Section 2.4. So, from now on, we do not distinguish between the allocation rule f_k and the corresponding mechanism.

The reason that VCG cannot provide non-trivial approximation guarantees for residual welfare maximization is that if the weight vector for, say, two outcomes is $(1, 1+\varepsilon)$, the mechanism outputs the second outcome instead of a mixture of both. Such a mechanism requires high payments in order to truthfully distinguish between the outcomes, leading to a negligible residual surplus. In contrast, the mechanism with allocation f_k outputs a "smooth max" over outcomes leading to a significantly reduced amount of payments (see also Figure 1).

Lemma 2. *For any $k \geq 1$, the closed form of the allocation rule f_k is*

$$f_k(\boldsymbol{v}) = \frac{1}{m^{1-1/k}} \frac{\boldsymbol{w}(\boldsymbol{v})^{\frac{1}{k-1}}}{||\boldsymbol{w}(\boldsymbol{v})^{\frac{1}{k-1}}||_k} .$$

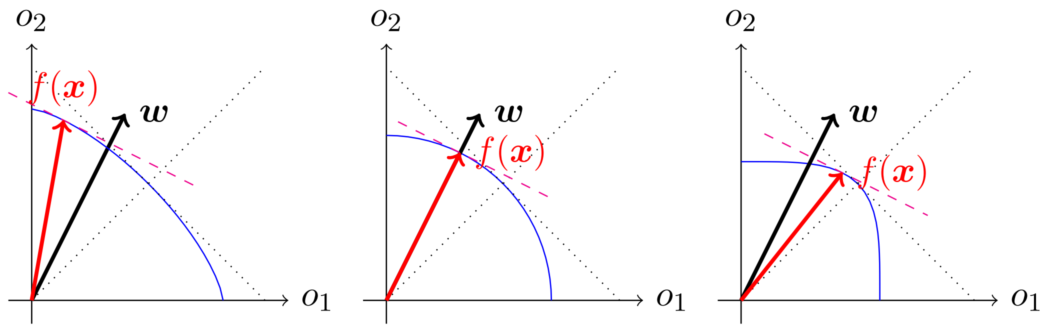

Fig. 1. Optimizing on the convex sets S_k defined in (6) for $m = 2$ outcomes and for $k = 1.4$, $k = 2$ and $k = 4$, respectively.

Proof. The outcome of the allocation rule $f_k(\boldsymbol{v})$ is the vector $\boldsymbol{s}$ that optimizes $\boldsymbol{w} \cdot \boldsymbol{s}$ subject to $\|\boldsymbol{s}\|_k \leq m^{-\frac{k-1}{k}}$. By Minkowski's inequality, (6) defines a strictly convex space. Therefore, the optimal point lies on the boundary of the space S_k, at the extreme point in the direction of $\boldsymbol{w}$. The boundary is defined by

$$\|\boldsymbol{s}\|_k = m^{-\frac{k-1}{k}} \iff \|\boldsymbol{s}\|_k^k = m^{-(k-1)} .$$

Since we seek the extreme point in the direction of $\boldsymbol{w}$, $\boldsymbol{w}$ must be perpendicular to the boundary at the optimal point. Therefore, at the optimal point $\boldsymbol{s}_*$, the gradient of the surface is in the direction of $\boldsymbol{w}$. Namely, there is some c such that

$$\nabla(\|\boldsymbol{s}_*\|_k^k) = c\,\boldsymbol{w} \iff \boldsymbol{s}_* = \left(\frac{c}{k}\right)^{\frac{1}{k-1}} \boldsymbol{w}^{\frac{1}{k-1}} .$$

Moreover, $\boldsymbol{s}_*$ needs to be on the surface of S_k. Thus,

$$\|\boldsymbol{s}_*\|_k^k = \frac{1}{m^{k-1}} \iff \left(\frac{c}{k}\right)^{\frac{1}{k-1}} = \frac{1}{m^{\frac{k-1}{k}} \|\boldsymbol{w}\|_{\frac{k}{k-1}}^{\frac{1}{k-1}}} .$$

Substituting in the equation for $\boldsymbol{s}_*$ concludes the proof. □

We are interested in allocation rules with S_k close to S_1. So, we set $k = \ell/(\ell-1)$ for some integer $\ell \geq 1$. The resulting allocation is

$$f_\ell(\boldsymbol{v}) = \frac{1}{m^{1/\ell}} \frac{\boldsymbol{w}(\boldsymbol{v})^{\ell-1}}{||\boldsymbol{w}(\boldsymbol{v})^{\ell-1}||_{\frac{\ell}{\ell-1}}} . \tag{7}$$

Applying (7) for $\ell \to \infty$, we obtain an allocation with probability that tends to 1 for the outcome with the highest total value. On the other hand, applying (7) for $\ell = 1$, we obtain the random allocation with probability $1/m$ for each outcome. Hence, the allocation rules defined by (7) exhibit a transition between the optimal solution and the random allocation. Moreover, the allocation is partial in the sense that for $\ell \in (1, \infty)$, $|f_\ell(\boldsymbol{v})| < 1$ and there is a positive probability that f_ℓ does not implement any outcome. Intuitively, sampling outcomes according to the ℓ-th power corresponds to a notion of "smooth max" operator. The higher the value of ℓ, the higher the resemblance to the true max, and the higher the sensitivity to input changes. Very roughly speaking, high input sensitivity implies that the mechanism relies on the exact bid values, leading to a high amount of payments, since payments act as "proof of value" for the agents.

5.2 Social Surplus Guarantees

The social surplus of the allocation rule f_ℓ depends on how well its range approximates the unit simplex.

Lemma 3. *For any $\ell \geq 1$, the allocation rule f_ℓ described by (7) has an expected social surplus of*

$$S(\boldsymbol{v}) = \|\boldsymbol{w}\|_\ell / m^{1/\ell}$$

and approximates the maximum social surplus within a factor of $m^{1/\ell}$.

Proof. For any vector $\boldsymbol{a}$,

$$\frac{\|\boldsymbol{a}^\ell\|_1}{\|\boldsymbol{a}^{\ell-1}\|_{\frac{\ell}{\ell-1}}} = \|\boldsymbol{a}\|_\ell \,. \tag{8}$$

The approximation ratio follows from

$$\boldsymbol{w} \cdot f(\boldsymbol{v}) = \frac{1}{m^{1/\ell}} \cdot \frac{\boldsymbol{w} \cdot \boldsymbol{w}^{\ell-1}}{\|\boldsymbol{w}^{\ell-1}\|_{\frac{\ell}{\ell-1}}} \overset{\text{Eq. (8)}}{=} \frac{\|\boldsymbol{w}\|_\ell}{m^{1/\ell}} \geq \frac{\|\boldsymbol{w}\|_\infty}{m^{1/\ell}} \,.$$

The analysis is tight, because when $\boldsymbol{v}$ consists of a single-minded agent with unit value, $\boldsymbol{w} \cdot f(\boldsymbol{v}) = \frac{1}{m^{1/\ell}}$ and $\|\boldsymbol{w}\|_\infty = 1$. □

We highlight that the approximation ratio is exactly the distance of range S_k from the unit simplex at the extreme points of the unit simplex.

5.3 Bounding the Total Payments

We proceed to study the amount of payments charged by the mechanism. The payments of agent i are computed as follows

$$\begin{aligned} p_i(\boldsymbol{v}) &= \boldsymbol{w}_{-i} \cdot f(\boldsymbol{v}_{-i}) - \boldsymbol{w}_{-i} \cdot f(\boldsymbol{v}) \\ &= \frac{1}{m^{1/\ell}} \cdot \left(\|\boldsymbol{w}_{-i}\|_\ell - \|\boldsymbol{w}\|_\ell + \frac{\boldsymbol{v_i} \cdot \boldsymbol{w}^{\ell-1}}{\|\boldsymbol{w}\|_\ell^{\ell-1}} \right) . \end{aligned}$$

Next, summing up the payment charged to each agent i by the mechanism, we bound the total amount of payments.

Lemma 4. *For any integer $\ell \geq 1$, the total amount of payments charged by the mechanism described by (7) to the agents is at most*

$$P(\boldsymbol{v}) \leq \frac{1}{m^{1/\ell}} \cdot \left(1 - \frac{1}{\ell}\right) \cdot \|\boldsymbol{w}(\boldsymbol{v})\|_\ell \tag{9}$$

Proof. Summing up the individual payments, we obtain that

$$\sum_{i=1}^{n} p_i(\boldsymbol{v}) = \frac{1}{m^{1/\ell}} \cdot \left(\frac{(\sum_i \boldsymbol{v_i}) \cdot \boldsymbol{w}}{\|\boldsymbol{w}\|_\ell^{l-1}} - \sum_i (\|\boldsymbol{w}\|_\ell - \|\boldsymbol{w}_{-i}\|_\ell) \right)$$
$$= \frac{1}{m^{1/\ell}} \cdot \left(\|\boldsymbol{w}\|_\ell - \sum_i (\|\boldsymbol{w}\|_\ell - \|\boldsymbol{w}_{-i}\|_\ell) \right) .$$

Therefore it suffices to show that

$$\sum_i (\|\boldsymbol{w}\|_\ell - \|\boldsymbol{w}_{-i}\|_\ell) \geq \frac{\|\boldsymbol{w}\|_\ell}{\ell} .$$

The ℓ-th power difference is bounded as follows

$$\|\boldsymbol{w}\|_\ell^\ell - \|\boldsymbol{w}_{-i}\|_\ell^\ell = (\|\boldsymbol{w}\|_\ell - \|\boldsymbol{w}_{-i}\|_\ell) \cdot \sum_{k=0}^{\ell-1} (\|\boldsymbol{w}\|_\ell^{\ell-1-k} \|\boldsymbol{w}_{-i}\|_\ell^k)$$
$$\leq (\|\boldsymbol{w}\|_\ell - \|\boldsymbol{w}_{-i}\|_\ell) \cdot \ell \cdot \|\boldsymbol{w}\|_\ell^{\ell-1} .$$

For the rest of the proof, for a vector $\boldsymbol{a}$, we denote its j-th coordinate by $\boldsymbol{a}[j]$. Then,

$$\sum_i \frac{\|\boldsymbol{w}\|_\ell^\ell - \|\boldsymbol{w}_{-i}\|_\ell^\ell}{\ell \cdot \|\boldsymbol{w}\|_\ell^{\ell-1}} \geq \frac{\|\boldsymbol{w}\|_\ell}{\ell} \iff \sum_i (\|\boldsymbol{w}\|_\ell^\ell - \|\boldsymbol{w}_{-i}\|_\ell^\ell) \geq \|\boldsymbol{w}\|_\ell^\ell$$
$$\iff \sum_{i=1}^{n} \left(\sum_{j=1}^{m} \boldsymbol{w}^\ell[j] - \sum_{j=1}^{n} \boldsymbol{w}_{-i}^\ell[j] \right) \geq \sum_{j=1}^{m} \boldsymbol{w}^\ell[j]$$
$$\iff \sum_{j=1}^{m} \sum_{i=1}^{n} (\boldsymbol{w}^\ell[j] - \boldsymbol{w}_{-i}^\ell[j]) \geq \sum_{j=1}^{m} \boldsymbol{w}^\ell[j] .$$

We now prove that the inequality holds for each term separately. It holds that

$$\boldsymbol{w}^\ell[j] - \boldsymbol{w}_{-i}^\ell[j] \geq (\boldsymbol{w}[j] - \boldsymbol{w}_{-i}[j]) \cdot \boldsymbol{w}^{\ell-1}[j] = \boldsymbol{v_i}[j] \cdot \boldsymbol{w}^{\ell-1}[j] .$$

Summing over i, we obtain that

$$\sum_i (\boldsymbol{w}^\ell[j] - \boldsymbol{w}_{-i}^\ell[j]) \geq \sum_i \boldsymbol{v_i}[j] \cdot \boldsymbol{w}^{\ell-1}[j] = \boldsymbol{w}[j] \cdot \boldsymbol{w}^{\ell-1}[j] = \boldsymbol{w}^\ell[j] ,$$

which concludes the proof. □

5.4 Obtaining a Bound on Residual Surplus

Lemma 3 quantifies the expected social surplus of the mechanism and Lemma 4 provides an upper bound on the total amount of payments charged to the agents. Combining them, we obtain the following lower bound on the residual surplus of the mechanism:

$$R(\boldsymbol{v}) = \boldsymbol{w} \cdot f(\boldsymbol{v}) - P(\boldsymbol{v}) \geq \frac{\|\boldsymbol{w}\|_\ell}{\ell m^{1/\ell}} \geq \frac{\|\boldsymbol{w}\|_\infty}{\ell m^{1/\ell}} .$$

We summarize our results in the following theorem.

Theorem 4. *For every integer $\ell \geq 1$, there is a truthful and individually rational mechanism that is $(m^{1/\ell}, \ell m^{1/\ell})$-approximate for the social surplus and the residual surplus.*

The optimal point of this tradeoff in terms of residual surplus maximization is when $\ell = \ln m$ (for simplicity, we assume that if ℓ is not an integer, it is rounded to the smallest integer exceeding the given value).

Corollary 1. *There is a truthful and individually rational mechanism that is $(e, e \ln m)$-approximate for the social surplus and the residual surplus.*

Alternatively, setting $\ell = \frac{\ln m}{\ln(1+\varepsilon)}$, we get the following:

Corollary 2. *There is a truthful and individually rational mechanism which for any $\varepsilon > 0$, is $(1+\varepsilon, \frac{(1+\varepsilon)^2}{\varepsilon} \ln m)$-approximate for the social surplus and the residual surplus.*

An interesting property of our mechanism is that the set of outcomes can be a priori restricted to some subset of the original outcome space O. If there is a subset (or range) $O' \subseteq O$ of the set of all outcomes such that optimizing the social surplus over O' provides a good approximation guarantee to the optimal social surplus, we can apply our mechanism on O' (instead of O, without any other changes) and obtain a truthful mechanism with a similar approximation guarantee for the social surplus and a logarithmic approximation guarantee for the residual surplus. Mechanisms based on exact optimization over such a restricted set of outcomes O' are known as *Maximal-in-Range* (MIR) and have been extensively used to obtain truthful VCG-based mechanisms that approximate the social surplus, when computing the optimal social surplus corresponds to an **NP**-hard optimization problem (see e.g., the MIR mechanism of [23, Section 3.2] that is truthful and approximates the social surplus for subadditive Combinatorial Public Projects within a factor of $O(\min\{k, \sqrt{u}\})$, where u is the number of items and k is the size of the project). Hence, we obtain the following

Corollary 3. *Let $O' \subseteq O$ be a subset of outcomes such that optimizing the social surplus over O' results in a ρ-approximation to the social surplus obtained by optimizing over the set of all outcomes O. Then, we can obtain a truthful and individually rational mechanism which for any $\varepsilon > 0$, is $((1+\varepsilon)\rho, \frac{(1+\varepsilon)^2 \rho}{\varepsilon} \ln |O'|)$-approximate for the social surplus and the residual surplus.*

6 Discussion and Directions for Further Research

We believe that the idea of reducing the total amount of payments charged to the agents by optimizing over a (carefully selected) "smoothed" subspace of the simplex of all probability distributions, which results in a partial allocation that does not depend heavily on the bid of any particular agent, may be of independent interest and may have interesting applications to other mechanism design settings where an upper bound on the payments is required. In Section 5, we employ this approach and show how to translate the fact that the resulting partial allocation does not depend heavily on any particular bid into a tradeoff between the resulting social surplus and the total amount of payments required for truthfulness. This directly implies the approximation result for the residual surplus. In another recent work [10], we use the same approach and show how to approximate the optimal social surplus within a factor of $1 + \varepsilon$ by randomized mechanisms without money that achieve truthfulness through selective verification of $O(\log m)$ agents. In retrospect, we think that the use of differential privacy to the design of almost truthful or truthful mechanisms without money [17, 21] and the use of the exponential mechanism to achieve truthful, differentially private and almost optimal smooth allocations [14] can be regarded, at least to some extent, as instances of the same general approach. That is because the exponential mechanism, on which the mechanisms of [17, 21, 14] are based, is again MIDR with the entire unit simplex as range and a smooth objective function equal to a convex combination of the expected social surplus and the entropy of the resulting probability vector. The entropy in the objective function results in a smooth allocation rule, which allows [17] to achieve a bounded deviation from truthfulness, [21] to achieve truthfulness without money and a good additive approximation guarantee, and [14] to achieve a tradeoff between additive approximation and the total amount of payments.

Another direction for further research is to investigate where our smooth tradeoff between the social welfare and the total amount of payments required for truthfulness could have any applications to mechanism design settings with budget-restricted bidders (see e.g. [6] and the references therein). In a complementary direction, we found it quite interesting that quasi-proportional mechanisms, very similar in spirit to our power-proportional mechanisms, together with an all-pay or a winners-pay (non-truthful) payment scheme, were shown in [18] to achieve high revenue at their unique pure Nash equilibrium.

References

1. A. Archer and É. Tardos. Frugal path mechanisms. *ACM Transactions on Algorithms*, 3(1), 2007.
2. M. Braverman, J. Chen, and S. Kannan. Optimal provision-after-wait in healthcare. In *Proc. of the 5th Conference on Innovations in Theoretical Computer Science (ITCS '14)*, pp. 541-542, 2014.
3. R. Cavallo. Efficiency and redistribution in dynamic mechanism design. In *Proc. of the 9th ACM Conference on Electronic Commerce (EC '08)*, pp. 220-229, 2008.
4. S. Chakravarty and T.R. Kaplan. Manna from heaven or forty years in the desert: Optimal allocation without transfer payments. *Social Science Research Network*, `http://dx.doi.org/10.2139/ssrn.939389`, 2006.
5. R. Cole, V. Gkatzelis, and G. Goel. Mechanism design for fair division: allocating divisible items without payments. In *Proc. of the 14th ACM Conference on Electronic Commerce (EC '13)*, pp. 251-268, 2013.
6. S. Dobzinski and R. Paes Leme. Efficiency guarantees in auctions with budgets. In *Proc. of the 41st International Colloquium on Automata, Languages and Programming (ICALP '14)*, LNCS 8572, pp. 392-404, 2014.
7. E. Elkind, A. Sahai, and K. Steiglitz. Frugality in Path Auctions. In *Proc. of the 15th ACM-SIAM Symposium on Discrete Algorithms (SODA '04)*, pp. 701-709, 2004.
8. Peter Eso and Gabor Futo. Auction design with a risk averse seller. *Economics Letters*, 65(1):71-74, October 1999.
9. D. Fotakis, D. Tsipras, C. Tzamos, and E. Zampetakis. Efficient Money Burning in General Domains. In *Proc. of the 8th International Symposium on Algorithmic Game Theory (SAGT '15)*, LNCS 9347, pp. 85-97, 2015.
10. D. Fotakis, C. Tzamos, and E. Zampetakis. Mechanism Design with Selective Verification. In *Proc. of the 17th ACM Conference on Electronic Commerce (EC '16)*, 2016.
11. 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.

12. M. Guo and V. Conitzer. Better redistribution with inefficient allocation in multi-unit auctions. *Artificial Intelligence*, 216:287-308, 2014.
13. J.D. Hartline and T. Roughgarden. Optimal mechanism design and money burning. In *Proc. of the 40th ACM Symposium on Theory of Computing (STOC '08)*, pp. 75-84, 2008.
14. Z. Huang and S. Kannan. The exponential mechanism for social welfare: Private, truthful, and nearly optimal. In *Proc. of the 53rd IEEE Symposium on Foundations of Computer Science (FOCS '12)*, pp. 140-149, 2012.
15. J. Lesca, T. Todo, and M. Yokoo. Coexistence of utilitarian efficiency and false-name-proofness in social choice. In *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS '14)*, pp. 1201-1208, 2014.
16. A.M. Manelli and D.R. Vincent. Multidimensional Mechanism Design: Revenue Maximization and the Multiple-Good Monopoly. *Journal of Economic Theory*, 137(1):153-185, 2007.
17. F. McSherry and K. Talwar. Mechanism design via differential privacy. In *Proc. of the 48th IEEE Symposium on Foundations of Computer Science (FOCS '07)*, pp. 94-103, 2007.
18. V.S. Mirrokni, S. Muthukrishnan, and U. Nadav. Quasi-proportional mechanisms: Prior-free revenue maximization. In *Proc. of the 9th Latin American Symposium on Theoretical Informatics (LATIN '10)*, LNCS 6034, pp. 565-576, 2010.
19. R. Myerson. Optimal Auction Design. *Mathematics of Operations Research*, 6(1):58-73, 1981.
20. N. Nisan. Introduction to Mechanism Design (for Computer Scientists). *Algorithmic Game Theory*, 9:209-242, 2007.
21. K. Nissim, R. Smorodinsky, and M. Tennenholtz. Approximately optimal mechanism design via Differential Privacy. In *Proc. of the 3rd Conference on Innovations in Theoretical Computer Science (ITCS '12)*, pp. 203-213, 2012.
22. J.-C. Rochet and L.A. Stole. The Economics of Multidimensional Screening. *Econometric Society Monographs*, 35:150-197, 2003.
23. M. Schapira and Y. Singer. Inapproximability of Combinatorial Public Projects. In *Proc. of the 4th Workshop on Internet and Network Economics (WINE '08)*, LNCS 5385, pp. 351-361, 2008.