Efficient Money Burning in General Domains*

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

[1] National Technical University of Athens, 157 80 Athens, Greece
[2] Massachusetts Institute of Technology, Cambridge, MA 02139
fotakis@cs.ntua.gr, dtsip@corelab.ntua.gr, tzamos@mit.edu, mzampet@mit.edu

Abstract. We study mechanism design where the payments charged to the agents are not in the form of monetary transfers, but are effectively burned. In this setting, the objective is to maximize *social utility*, i.e., the social welfare minus the payments charged. We consider a general setting with m discrete outcomes and n multidimensional agents. We present two essentially orthogonal randomized truthful mechanisms that extract an $O(\log m)$ fraction of the maximum welfare as social utility. Moreover, the first mechanism achieves a $O(\log m)$-approximation for the social welfare, 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 "smoothed" space, thus achieving a nice and smooth trade-off between welfare approximation and the payments charged.

1 Introduction

The extensive use of monetary transfers in the Algorithmic Game Theory is due to the fact that so little can be implemented truthfully in their absence (see e.g., [14]). On the other hand, if monetary transfers are available (and acceptable for the particular application), the famous Vickrey-Clarke-Groves (VCG) mechanism (see e.g., [14]) succeeds in truthfully maximizing the *social welfare*, i.e., the total value generated for the agents, albeit with possible very large monetary transfers from the agents to the center. This is acceptable as long as the payments generate revenue for the center (e.g., the government for public good allocation or the auctioneer for allocation of private goods), since the funds are not lost, but are transferred to the center. Then, the funds could be redistributed among the agents (see e.g., [9, 10]) or invested in favor of the society.

However, there are settings where 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 "computational" challenges (e.g., captcha), waiting times (e.g., waiting lists in hospitals [2] or in popular events or places), or reduction in service quality (see also [11, 4] for more examples). In these settings, the natural objective

* This research was 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. Additional support by NSF Award CCF-1101491 and an Award for Graduate Students in TCS by the Simons Foundation.

is to maximize the net gain of the society, measured by the social welfare minus the payments charged and usually referred to as the *social utility* (or the *social surplus*).

In the AGT community, the general idea of money burning and social utility maximization by truthful mechanisms was first considered by Hartline and Roughgarden [11]. 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 (wrt. the social utility) 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 social utility extracted by an optimal Bayesian mechanism with 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 social utility of a truthful mechanism to the maximum social welfare, then the best possible approximation guarantee for k-unit auctions is $\Theta(1 + \log \frac{n}{k})$, where n is the number of agents.

Contribution. In this work, we consider social utility maximization by truthful mechanisms in a general mechanism design setting with m discrete possible outcomes and multidimensional agents with positive valuations. Due to the fact that social utility maximization is closely related to revenue maximization, coming up with a characterization of the optimal (wrt. the social utility) truthful Bayesian mechanism, as in [11], is a daunting task and far beyond the scope of this work. Instead, we evaluate the performance of our mechanisms by comparing their social utility to the maximum social welfare (achievable by an optimal algorithm that does not need to be truthful). In fact, we seek for mechanisms that achieve nontrivial approximation guarantees wrt. both social utility and social welfare. Our main contribution is two randomized truthful mechanisms, based on essentially orthogonal approaches, that approximate social utility within a best possible factor of $O(\log m)$, thus extending the last result of [11] to our general mechanism design setting.

Probably the simplest candidate mechanisms for utility maximization are the random allocation, where each outcome is implemented with probability $1/m$, and the VCG mechanism. Clearly, the approximation ratio of random allocation for both the social utility and the social welfare is m, while VCG cannot approximate within a ratio of m even for the natural case of uniform i.i.d. bidders. A natural way to approximate social utility is through a careful tradeoff between VCG, which optimizes welfare but may result in poor utility due to high payments, and random allocation on appropriate sets of outcomes, which is truthful without payments and thus, translates all welfare into utility.

Exploiting this intuition and building on the mechanism of [11, Theorem 5.2], we present a randomized truthful mechanism that approximates both the social utility and the social welfare within a factor of $O(\log m)$. The idea is to select a random integer j from 0 to $\log m$, and then, select a random outcome i among the best (in total value) 2^j outcomes, and apply VCG payments. The key step in establishing the approximation guarantee is to show that in terms of utility maximization, the worst-case instances correspond to single item auctions. Then, the upper bound of [11, Theorem 5.2] carries over to our more general setting. Moreover, since the single item auction is a special case of our setting, the lower bound of [11, Proposition 5.1] implies that our approximation ratio is aymptotically tight.

Our second mechanism optimizes the social welfare (using VCG) over a carefully defined subspace of the unit simplex with all probability distributions over the outcomes. Intuitively, if we optimized over the unit simplex, we would have optimal welfare but probably poor utility, due to 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 unit simplex, thus achieving a significant reduction in the payments if the best outcomes are close in total value. Due to this fact, this mechanism is *partial*, in the sense that with probability $1-\epsilon$ it may not implement any outcome (see [5] for another use of partial allocation to induce truthfulness). For any $\epsilon > 0$, the approximation ratio is $1+\epsilon$ for the social welfare and $O(\epsilon^{-1}\log m)$ for the social utility. Hence, this mechanism achieves a best possible approximation ratio for the social utility and a constant approximation for the social welfare, thus significantly improving on our first mechanism. The main idea behind this mechanism is to "smoothen" the solution space so that we achieve a smooth tradeoff between welfare approximation and the payments charged, where for mechanisms close to the optimal, payments are reduced significantly faster that social welfare. On the technical side, this mechanism bears a resemblance to proper scoring rules in [8]. We believe that such mechanisms, which are based on carefully chosen "smoothed" subspaces and provide smooth tradeoffs between approximation 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 $\mathbf{NP}$-hard problems. In certain domains, we can combine our mechanisms with existing Maximal-in-Range mechanisms so that everything runs in polynomial time (e.g., for subadditive Combinatorial Public Projects, we can use the Maximal-in-Range mechanism of [15, Sec. 3.2] and obtain a randomized polynomial-time truthful mechanism that with $O(\min\{k,\sqrt{u}\})$-approximation for the social welfare and $O(\min\{k,\sqrt{u}\}\log u)$-approximation for the social utility, where u is the number of items and k is the size of the project).

Related Work. There is much work on (mostly polynomial-time) truthful mechanisms with monetary transfers that seek to maximize (exactly or approximately) the social welfare. In this general agenda, our work is closest in spirit to mechanisms with frugal payments (see e.g., [1, 6]). In addition to [11], Chakravarty and Kaplan [4] characterized the Bayesian mechanism of maximum social utility in multi-unit (unit demand) auctions. More recently, Braverman et al. [2] considered utility optimization in health care service allocation, but they focused on the complexity of computing efficient equilibrium allocations, instead of approximate truthful mechanisms.

An orthogonal direction is that of revenue redistribution (see e.g, [3, 9, 10] 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 social utility optimization relaxing the requirement for social welfare maximization (see e.g., [10]). 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., [12]). This is infeasible our setting, where money-burning payment schemes (e.g., computational challenges or waiting time) make redistribution infeasible.

2 Preliminaries and Notation

For any integer m, $[m] \equiv \{1, \ldots, m\}$. We denote the j-th coordinate of a vector $\boldsymbol{x}$ by x_j. For a vector $\boldsymbol{x} = (x_1, \ldots, x_m)$ and $i \in [m]$, $\boldsymbol{x_{-i}}$ is $\boldsymbol{x}$ without coordinate i. For a vector $\boldsymbol{x} \in \mathbb{R}^m$ and some $\ell \geq 0$, $\boldsymbol{x}^\ell = (x_1^\ell, \ldots, x_m^\ell)$ is the coordinate-wise power of $\boldsymbol{x}$ and $\|\boldsymbol{x}\|_\ell = (\sum_{j=1}^m x_j^\ell)^{1/\ell}$ is the ℓ-norm of $\boldsymbol{x}$. For convenience, we let $\|\boldsymbol{x}\|_1 = |\boldsymbol{x}|$. Moreover, $\|\boldsymbol{x}\|_\infty = \max_{j \in [m]}\{x_j\}$ is the infinity norm of $\boldsymbol{x}$.

The Setting. There is a finite set of possible outcomes O and we denote $|O| = m$. We consider a set of n strategic agents, each with a *private, non-negative value* for each outcome. For agent i, we denote his *valuation* as a vector $\boldsymbol{x_i} \in \mathbb{R}_+^m$, that is, agent i receives value x_{ij} for outcome j. We call the vector of all valuations $\boldsymbol{x} = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n)$ a *valuation profile*. For a valuation profile $\boldsymbol{x}$, $\boldsymbol{w}(\boldsymbol{x}) = \boldsymbol{x_1} + \ldots + \boldsymbol{x_n}$ is the vector of *weights* for the outcomes. We will write $\boldsymbol{w}$ instead of $\boldsymbol{w}(\boldsymbol{x})$ and $\boldsymbol{w_{-i}}$ instead of $\boldsymbol{w}(\boldsymbol{x_{-i}})$ when $\boldsymbol{x}$ is clear from the context.

Allocation Rules and Mechanisms. For a finite set S, $\Delta(S)$ denotes the unit simplex over S. A (randomized) allocation rule is a function $f : (\mathbb{R}_+^m)^n \to \Delta(O)$, mapping valuation profiles to probability distributions over outcomes. Then $f_j(\boldsymbol{x})$ is the probability of outcome j on valuation profile $\boldsymbol{x}$. It follows that the expected value of agent i is $\boldsymbol{x_i} \cdot f(\boldsymbol{x})$. We consider allocation rules that are *strongly anonymous*, in the sense that $f(\boldsymbol{x})$ depends only on $\boldsymbol{w}(\boldsymbol{x})$, and we therefore write the allocation rule only in terms of the weight vector.

A *payment rule* is a function $p : (\mathbb{R}_+^m)^n \to \mathbb{R}^n$ mapping valuation profiles to *payment vectors*. A mechanism is a pair $\mathcal{M} = (f, p)$ that given some valuation profile $\boldsymbol{x}$ outputs the probability distribution $f(\boldsymbol{x})$ and charges agent i the amount $p_i(\boldsymbol{x})$. We focus on symmetric payment rules, and we therefore represent the amount charged to agent i as $p(\boldsymbol{x_{-i}}, \boldsymbol{x_i})$. The expected *utility* of agent i on valuation profile $\boldsymbol{x}$ under mechanism $\mathcal{M} = (f, p)$ is

$$\boldsymbol{x_i} \cdot f(\boldsymbol{x}) - p(\boldsymbol{x_{-i}}, \boldsymbol{x_i})$$

and is the amount he aims to maximize.

We require that our mechanisms are truthful and individually rational in expectation. A mechanism (f, p) is *truthful* (in expectation) if for any agent i, valuation profile $\boldsymbol{x}$ and valuation $\boldsymbol{x_i'}$,

$$\boldsymbol{x_i} \cdot f(\boldsymbol{x}) - p(\boldsymbol{x_{-i}}, \boldsymbol{x_i}) \geq \boldsymbol{x_i} \cdot f(\boldsymbol{x_{-i}}, \boldsymbol{x_i'}) - p(\boldsymbol{x_{-i}}, \boldsymbol{x_i'})$$

and *individually rational (IR)* if for any agent i and valuation profile $\boldsymbol{x}$,

$$\boldsymbol{x_i} \cdot f(\boldsymbol{x}) - p(\boldsymbol{x_{-i}}, \boldsymbol{x_i}) \geq 0$$

Objectives and Approximation. Let some mechanism $\mathcal{M} = (f, p)$ and valuation profile $\boldsymbol{x}$. We denote the total payments of $\mathcal{M}$ on input $\boldsymbol{x}$ by $P[\boldsymbol{x}] = \sum_i p(\boldsymbol{x_{-i}}, \boldsymbol{x_i})$. The quantities we are interested in maximizing are the social welfare and the social utility. The *social welfare* of $\mathcal{M}$ on $\boldsymbol{x}$ is $SW[\boldsymbol{x}] = \sum_i \boldsymbol{x_i} \cdot f(\boldsymbol{x}) = \boldsymbol{w} \cdot f(\boldsymbol{x})$ and the *social utility* of $\mathcal{M}$ on $\boldsymbol{x}$ is $U[\boldsymbol{x}] = SW[\boldsymbol{x}] - P[\boldsymbol{x}]$. The maximum possible social utility and social welfare of the mechanism (ignoring truthfulness constraints) on input $\boldsymbol{x}$ is $\|\boldsymbol{w}\|_\infty$. We say that mechanism $\mathcal{M}$, ρ-approximates social welfare (resp. social utility) if for any input $\boldsymbol{x}$, $SW[\boldsymbol{x}] \geq \frac{1}{\rho}\|\boldsymbol{w}(\boldsymbol{x})\|_\infty$ (resp. $U[\boldsymbol{x}] \geq \frac{1}{\rho}\|\boldsymbol{w}(\boldsymbol{x})\|_\infty$). For a mechanism $\mathcal{M}$ that ρ_1-approximates social welfare and ρ_2-approximates social utility, we say that it approximates *social efficiency* within (ρ_1, ρ_2).

Implementable Rules. For every set $S \subseteq R_+^m$, the mechanism $\mathcal{M} = (f, p)$ such that $f(\boldsymbol{x}) = \arg\max_{\boldsymbol{s} \in S} \boldsymbol{s} \cdot \boldsymbol{w}$ and $p(\boldsymbol{x_{-i}}, \boldsymbol{x_i}) = \boldsymbol{w_{-i}} \cdot f(\boldsymbol{x_{-i}}) - \boldsymbol{w_{-i}} \cdot f(\boldsymbol{x})$ is truthful and individually rational. This follows directly from the analysis of the VCG mechanism [14]. We refer to such mechanisms as *Maximal in Distributional Range* (MIDR) and to the corresponding payment rule as the *VCG payment scheme*.

3 Best-Possible Guarantees for Social Utility

In contrast to social welfare maximization, where monetary transfers can be used freely to truthfully elicit the agents' preferences, in the case of social utility maximization, the transfers needed for the implementation of some mechanisms may be a significant part of the social welfare, thus prohibiting any non-trivial approximation guarantees.

Since the model we consider is so rich, the single item auction is a special case of it, when we restrict the domain to m outcomes and m agents, where agent i has a value $v_i \geq 0$ for outcome i and zero for the rest. By proving lower bounds to the approximation of social utility maximization in this special case, we get the same lower bounds for the general model. Our main tool here is Myerson's characterization of the revenue of any truthful auction in the single parameter environment.

Theorem 1 (Myerson [13]). *For any truthful mechanism $\mathcal{M} = (f, p)$ and valuation profile $\boldsymbol{x}$, where agent i has some value $v_i \geq 0$ only for outcome i and v_i is drawn independently from distribution $\mathcal{F}$ with cumulative distribution function $F_{\mathcal{F}}(v)$ and probability density function $f_{\mathcal{F}}(v)$, $\mathbb{E}[P(\boldsymbol{x})] = \mathbb{E}[\boldsymbol{\phi} \cdot f(\boldsymbol{x})]$, where $\phi_i = v_i - \frac{1-F_{\mathcal{F}}(v_i)}{f_{\mathcal{F}}(v_i)}$.*

Theorem (1) completely deftermines the expected amount of payments for any truthful allocation rule. This in turn determines the expected utility in terms of the allocation rule. By plugging in an appropriate distribution we can come up with lower bounds to the social utility of truthful mechanisms.

Corollary 1. *The Vickrey Auction when bidders are drawn i.i.d. from the uniform distribution $\mathcal{U}(0, 1)$, cannot approximate social utility within a factor better than m.*

This shows that the VCG mechanism for the natural case of uniform i.i.d. bidders performs no better that a random allocation. By aiming to maximize the social welfare, it has to charge every bidder his critical price which results to a high amount of payments, negating the welfare it produces. We therefore need to come up with mechanisms that instead of maximizing social welfare, employ suboptimal allocations to reduce payments, while preserving some amount of welfare. Our goal is to achieve the best possible worst-case guarantee for social utility maximization. A lower bound on the best approximation ratio in our setting can be obtained from [11, Proposition 5.1], which we prove here for completeness.

Corollary 2 ([11]). *No truthful mechanism can approximate social utility within a factor of $o(\log m)$.*

Proof. If agents are drawn from the exponential distribution, that is $f_{\mathcal{E}}(x) = e^{-x}$, $F_{\mathcal{E}}(x) = 1 - e^{-x}$, then $\phi_i = v_i - 1$ and by applying Theorem (1) we get that

$$\mathbb{E}[P(\boldsymbol{x})] = \mathbb{E}[\boldsymbol{w} \cdot f(\boldsymbol{x}) - |f(\boldsymbol{x})|] = \mathbb{E}[SW[\boldsymbol{x}] - |f(\boldsymbol{x})|]$$

and by linearity of expectation $\mathbb{E}[U[\boldsymbol{x}]] \leq 1$ It is straightforward to show that the expectation of the maximum of m i.i.d. exponential random variables equals H_m where H_m the m-th harmonic number. Then

$$\mathbb{E}[U(\boldsymbol{x})] \leq \mathbb{E}\left[\frac{\|\boldsymbol{w}\|_\infty}{H_m}\right]$$

and by the probabilistic method we get that there is some profile $\boldsymbol{x}$ for which the approximation ratio is logarithmic. □

We will now describe a mechanism that matches this lower bound in the general domain.

Definition 1. *For some $k \in [m]$, the Top_k allocation rule on input $\boldsymbol{x}$, orders outcomes in decreasing weight order, $w_1 \geq \ldots \geq w_m$ (breaking ties arbitrarily) and assigns probability $\frac{1}{k}$ to the first k. Formally, $\mathrm{Top}_k(\boldsymbol{x}) = \arg\max_{\boldsymbol{s} \in S_k} \boldsymbol{s} \cdot \boldsymbol{w}$, where S_k is the set of vectors in $\mathbb{R}_+^m$ with exactly k coordinates equal to $\frac{1}{k}$ and $m - k$ equal to 0.*

Since Top_k are welfare maximizers, they can be turned into truthful and IR mechanisms with the VCG payment scheme. We denote mechanisms of this family by $\mathcal{M}_k = (\mathrm{Top}_k, p_k)$. Each of these mechanism achieves different approximation guarantees with respect to social welfare and social utility in different settings, with respect to k. Thus by randomizing over them we can provide worst-case guarantees. In Mechanism 1 we achieve such an optimal social utility approximation guarantee by randomizing over exponentially increasing values of k. For simplicity we assume that m is a power of 2.

Mechanism 1 A $\log m$-approximate mechanism for Social Utility

Choose j uniformly at random from $\{0, 1, 2, \ldots, \log m\}$
Let $k \leftarrow 2^j$
Output the probability distribution $\mathrm{Top}_k(\boldsymbol{x})$ over outcomes
Charge agent i the amount $\boldsymbol{w_{-i}} \cdot \mathrm{Top}_k(\boldsymbol{x_{-i}}) - \boldsymbol{w_{-i}} \cdot \mathrm{Top}_k(\boldsymbol{x})$

The complete mechanism is randomization over $\mathcal{M}_k$ for some k independent of the input. As a result, Mechanism 1 is truthful and IR as a whole. In order to quantify the efficiency of the mechanism in terms of utility maximization, we first show that the worst-case instances are those of the single item auction, that is for each outcome i there is exactly one single-minded agent with valuation v_i for it (a bidder i is called single-minded if his utility is $v_i \geq 0$ for some outcome $j \in [m]$ and zero for the rest).

Lemma 1. *For any valuation profile $\boldsymbol{x} = (\boldsymbol{x_1}, \ldots, \boldsymbol{x_n})$, the utility of Mechanism 1 on $\boldsymbol{x}$ is higher than the utility on the valuation profile $\boldsymbol{y} = (\boldsymbol{y_1}, \ldots, \boldsymbol{y_m})$, where $\boldsymbol{y_i}$ is a single-minded agent for outcome i ($y_{ij} = 0$ for any $i \neq j$).*

Proof (sketch). Since the complete mechanism is a randomization over mechanisms $\mathcal{M}_k$ is suffices to show this property for each $\mathcal{M}_k$ separately. We prove the claim in two 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 total utility. This holds since the “competition” between agents is increased, and as a result, so do the payments, thus decreasing the total utility (the social welfare remains unaltered since the mechanism depends only on the weight of each outcome). By induction we transform any input to one with single-minded agents without increasing the utility.
- Then we show that if there are multiple single-minded bidders for the same outcome, joining their values into a single agent can only decrease the total utility. The reason for this is that the value the agents must “prove” (in the form of payments) to the mechanism is initially split amongst 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 input with single-minded agents to an input with one single-minded agent per outcome.

The technical details can be found in the full version of the paper. □

Theorem 2. *Mechanism 1 is a $(O(\log m), O(\log m))$ approximation to the social efficiency.*

Proof. By Proposition (1) and the analysis of [11], Mechanism 1, is a $O(\log m)$-approximation to social utility (and therefore social welfare). For the instance $\boldsymbol{x} = (\boldsymbol{x_1}, 0, \ldots, 0)$, where i is a single minded agent with value v_1 for outcome 1, the approximation ratio of $\log m$ is tight for both the welfare and the utility. □

4 Optimizing Social Utility Without Sacrificing Social Welfare

The mechanism of Section 3, approximates utility within an optimal logarithmic factor. However, it also approximates Social Welfare within the same logarithmic factor. The impossibility of Corollary 2 implies that no mechanism can do better than $(O(1), O(\log m))$-approximate social efficiency. So the question of simultaneously optimizing social welfare stands. We answer this question affirmatively by presenting a mechanism that optimizes welfare on a smooth probability space.

Theorem 3. *For any $\epsilon > 0$, there is a mechanism $\mathcal{M}$ that $\left((1+\epsilon), \frac{(1+\epsilon)^2}{\epsilon}\ln m\right)$-approximates social efficiency.*

Remark 1. We can $(O(1), O(\log m))$-approximate social efficiency simply by randomizing, with constant probability, between the VCG mechanism and Mechanism 1. However, the mechanism of Theorem 3 follows from a more general approach that yields a smooth mechanism and may be of independent interest.

4.1 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 optimizes the expected welfare 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 corners, 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 surfaces:

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

For any $k \geq 1$ or for $k = \infty$, we define the mechanism $f_k(\boldsymbol{x}) = \arg\max_{\boldsymbol{s} \in S_k} \boldsymbol{s} \cdot \boldsymbol{w}(\boldsymbol{x})$.

The reason VCG is not working for utility maximization is that if the weight vector for e.g. 2 outcomes is $(1, 1+\epsilon)$, the mechanism will output the second outcome instead of a mixture of both. Such a mechanism requires a high amount of payments in order to truthfully distinguish between the outcomes, leading to minimal utility. In contrast, the mechanism with allocation f_k outputs a "smooth max" over outcomes leading to a reduced amount of payments (Figure 1).

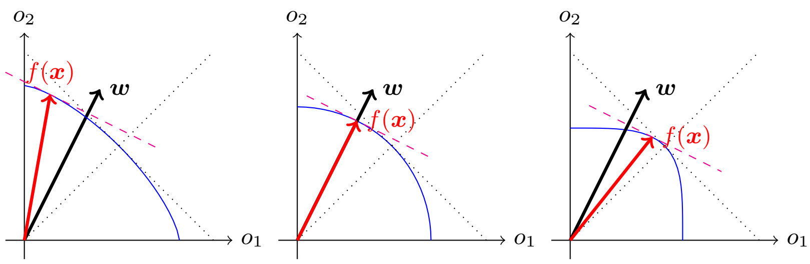

Fig. 1. Optimizing on the curved surfaces for $m = 2$

Lemma 2. *The closed form of the mechanism f_k is*

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

Proof. The outcome of the mechanism is the vector $\boldsymbol{s}$ the optimizes $\boldsymbol{w} \cdot \boldsymbol{s}$ subject to $\|\boldsymbol{s}\|_k \leq m^{-\frac{k-1}{k}}$. By the Minkowski inequality, Equation (1) defines a strictly convex space. Therefore the optimal point will lie on the boundary of the space, at the extremal point in the direction of $\boldsymbol{w}$. The boundary is defined by

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

and since we seek the extremal 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 $\boldsymbol{w}$, that is there is some c such the

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

Moreover $\boldsymbol{s_*}$ needs to be to be on the surface, and therefore

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

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

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

$$f_\ell(\boldsymbol{x}) = \frac{1}{m^{1/\ell}} \frac{\boldsymbol{w}^{\ell-1}}{\|\boldsymbol{w}^{\ell-1}\|_{\frac{\ell}{\ell-1}}} \tag{2}$$

The reader is invited to verify that the mechanism exhibits a smooth transition between the VCG mechanism (for $\ell \to \infty$) and the uniform mechanism (for $\ell = 1$). Moreover, the mechanism is partial in the sense that for $\ell \in (1, \infty)$, $|f_\ell(\boldsymbol{x})| < 1$ and there is a positive probability that f_ℓ does not implement any outcome.

4.2 Social Welfare Guarantees

Lemma 3. *For any $\ell \geq 1$, the mechanism of Equation (2) approximates the social welfare within $m^{1/\ell}$.*

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

$$\frac{\|\boldsymbol{a}^{\boldsymbol{\ell}}\|_1}{\|\boldsymbol{a}^{\ell-1}\|_{\frac{\ell}{\ell-1}}} = \|\boldsymbol{a}\|_\ell \tag{3}$$

The approximation ratio follows from

$$\frac{\boldsymbol{w} \cdot f(\boldsymbol{x})}{\|\boldsymbol{w}\|_\infty} = \frac{1}{m^{1/\ell}} \frac{\boldsymbol{w} \cdot \boldsymbol{w}^{\ell-1}}{\|\boldsymbol{w}^{\ell-1}\|_{\frac{\ell}{\ell-1}} \|\boldsymbol{w}\|_\infty} \overset{\text{Equation (3)}}{=} m^{1/\ell} \frac{\|\boldsymbol{w}\|_\ell}{\|\boldsymbol{w}\|_\infty} \geq \frac{1}{m^{1/\ell}}$$

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

4.3 Bounds to the Revenue of the Mechanism

We will now study the amount of payments charged by the mechanism. The payments of player i are computed as follows

$$\begin{aligned} p(\boldsymbol{x_{-i}}, \boldsymbol{x_i}) &= \boldsymbol{w_{-i}} \cdot f(\boldsymbol{x_{-i}}) - \boldsymbol{w_{-i}} \cdot f(\boldsymbol{x}) \\ &= \frac{1}{m^{1/\ell}} \left(\|\boldsymbol{w_{-i}}\|_\ell - \|\boldsymbol{w}\|_\ell + \frac{\boldsymbol{x_i} \cdot \boldsymbol{w}^{\ell-1}}{\|\boldsymbol{w}\|_\ell^{\ell-1}} \right) \end{aligned}$$

Therefore we can now bound the total amount of payments

Lemma 4. *For any integer $\ell \geq 1$, the mechanism of Equation (2) charges the set of agents at most*

$$P[\boldsymbol{x}] \leq \frac{1}{m^{1/\ell}} \left(1 - \frac{1}{\ell}\right) \|\boldsymbol{w}(\boldsymbol{x})\|_\ell \tag{4}$$

Proof. By summing up the individual payments.

$$\begin{aligned} \sum_{i=1}^{n} p(\boldsymbol{x_{-i}}, \boldsymbol{x_i}) &= \frac{1}{m^{1/\ell}} \left(\frac{(\sum_i \boldsymbol{x_i}) \cdot \boldsymbol{w}}{\|\boldsymbol{w}\|_\ell^{l-1}} - \sum_i (\|\boldsymbol{w}\|_\ell - \|\boldsymbol{w_{-i}}\|_\ell) \right) \\ &= \frac{1}{m^{1/\ell}} \left(\|\boldsymbol{w}\|_\ell - \sum_i (\|\boldsymbol{w}\|_\ell - \|\boldsymbol{w_{-i}}\|_\ell) \right) \end{aligned}$$

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 bound as follows

$$\begin{aligned} \|\boldsymbol{w}\|_\ell^\ell - \|\boldsymbol{w_{-i}}\|_\ell^\ell &= (\|\boldsymbol{w}\|_\ell - \|\boldsymbol{w_{-i}}\|_\ell) \cdot \sum_{k=0}^{l-1} (\|\boldsymbol{w}\|_\ell^{l-1-k} \|\boldsymbol{w_{-i}}\|_\ell^k) \\ &\leq (\|\boldsymbol{w}\|_\ell - \|\boldsymbol{w_{-i}}\|_\ell) \cdot \ell \|\boldsymbol{w}\|_\ell^{\ell-1} \end{aligned}$$

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

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

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

$$\boldsymbol{w}^{\ell}[j] - \boldsymbol{w_{-i}}^{\ell}[j] \geq (\boldsymbol{w}[j] - \boldsymbol{w_{-i}}[j])\boldsymbol{w}^{\ell-1}[j] = \boldsymbol{x_i}[j]\boldsymbol{w}^{\ell-1}[j]$$

and summing over i gives us

$$\sum_i (\boldsymbol{w}^{\ell}[j] - \boldsymbol{w_{-i}}^{\ell}[j]) \geq \sum_i \boldsymbol{x_i}[j]\boldsymbol{w}^{\ell-1}[j] = \boldsymbol{w}[j]\boldsymbol{w}^{\ell-1}[j] = \boldsymbol{w}^{\ell}[j]$$

concluding the proof. □

4.4 Maximizing Utility

The utility of the mechanism is therefore

$$U[\boldsymbol{x}] = \boldsymbol{w} \cdot f(\boldsymbol{x}) - P[\boldsymbol{x}] \geq \frac{\|\boldsymbol{w}\|_{\ell}}{\ell m^{1/\ell}} \geq \frac{\|\boldsymbol{w}\|_{\infty}}{\ell m^{1/\ell}}$$

We summarize our results in the following theorem.

Theorem 4. *For every integer $\ell \geq 1$, there is a truthful mechanism that $(m^{1/\ell}, \ell m^{1/\ell})$-approximates social efficiency.*

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

Corollary 3. *There is a truthful mechanism that $(e, e \ln m)$-approximates social efficiency.*

Alternatively by setting $\ell = \frac{\ln m}{\ln(1+\epsilon)}$ we get the following.

Corollary 4. *There is a truthful mechanism, that for any $\epsilon > 0$, $\left(1+\epsilon, \frac{(1+\epsilon)^2}{\epsilon} \ln m\right)$-approximates social efficiency.*

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. These mechanism are known as *Maximal in Range (MIR)*, and are tailored to obtain suboptimal welfare guarantees in polynomial time for NP-hard problems. Our mechanisms can be run on these modified outcome spaces with no modification preserving welfare guarantees and providing social utility logarithmic to the number of outcomes.

Corollary 5. *Let some MIR mechanism with outcome space S, that a-approximates social welfare. Then, it can be modified to $((1+\epsilon)a)$-approximate social welfare and $\left(\frac{(1+\epsilon)^2}{\epsilon} a \ln |S|\right)$-approximate social utility.*

Remark 2. We have shown that the mechanism is IR in expectation, however there are examples where players net negative utility for certain random outcomes. Nonetheless the mechanism can be modified to be universally IR. Consider some agent i. Let $P_i = \boldsymbol{w} \cdot f(\boldsymbol{x_{-i}}) - \boldsymbol{w} \cdot f(\boldsymbol{x})$ denote the payments that induce truthfulness. If outcome j is realized we charge this agent $p_{ij} = \frac{P_i}{\boldsymbol{x_i} \cdot f(\boldsymbol{x})} x_{ij}$. It is easy to verify that the expected payments are unaltered so truthfulness is preserved. Moreover, the mechanism is now universally IR. A similar technique can be found in [7].

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)*, pages 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)*, pages 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*, page 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)*, pages 251–268, 2013.
6. E. Elkind, A. Sahai, and K. Steiglitz. Frugality in Path Auctions. In *Proc. of the 15th ACM-SIAM Symposium on Discrete Algorithms(SODA '04)*, pages 701–709, 2004.
7. Peter Eso and Gabor Futo. Auction design with a risk averse seller. *Economics Letters*, 65(1):71–74, October 1999.
8. T. Gneiting and A.E. Rafterys. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
9. 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.
10. M. Guo and V. Conitzer. Better redistribution with inefficient allocation in multi-unit auctions. *Artificial Intelligence*, 216:287–308, 2014.
11. 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)*, pages 75–84, 2008.
12. 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)*, pages 1201–1208, 2014.
13. R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
14. N. Nisan. Introduction to Mechanism Design (for Computer Scientists). *Algorithmic Game Theory*, 9:209–242, 2007.
15. M. Schapira and Y. Singer. Inapproximability of combinatorial public projects. In *Proc. of the 4th Workshop on Internet and Network Economics(WINE '08)*, volume 5385 of *LNCS*, pages 351–361. Springer-Verlag, 2008.