Efficiency and Budget Balance

Swaprava Nath(✉) and Tuomas Sandholm

Carnegie Mellon University, Pittsburgh, USA
{swapravn,sandholm}@cs.cmu.edu

Abstract. We study *efficiency* and *budget balance* for designing mechanisms in general quasi-linear domains. Green and Laffont [13] proved that one cannot generically achieve both. We consider strategyproof budget-balanced mechanisms that are approximately efficient. For deterministic mechanisms, we show that a strategyproof and budget-balanced mechanism must have a *sink* agent whose valuation function is ignored in selecting an alternative, and she is compensated with the payments made by the other agents. We assume the valuations of the agents come from a bounded open interval. This result strengthens Green and Laffont's impossibility result by showing that even in a restricted domain of valuations, there does not exist a mechanism that is strategyproof, budget balanced, and takes every agent's valuation into consideration—a corollary of which is that it cannot be efficient. Using this result, we find a tight lower bound on the inefficiencies of strategyproof, budget-balanced mechanisms in this domain. The bound shows that the inefficiency asymptotically disappears when the number of agents is large—a result close in spirit to Green and Laffont [13, Theorem 9.4]. However, our results provide worst-case bounds and the best possible rate of convergence. Next, we consider minimizing any convex combination of inefficiency and budget imbalance. We show that if the valuations are unrestricted, no deterministic mechanism can do asymptotically better than minimizing inefficiency alone. Finally, we investigate randomized mechanisms and provide improved lower bounds on expected inefficiency. We give a tight lower bound for an interesting class of strategyproof, budget-balanced, randomized mechanisms. We also use an optimization-based approach—in the spirit of *automated mechanism design*—to provide a lower bound on the minimum achievable inefficiency of any randomized mechanism. Experiments with real data from two applications show that the inefficiency for a simple randomized mechanism is 5–100 times smaller than the worst case. This relative difference increases with the number of agents.

1 Introduction

Consider a group a friends deciding which movie to watch together. The movie can be watched in someone's home by renting it or at any of a number of

We are grateful to Ioannis Caragiannis, Vincent Conitzer, Debasis Mishra, Hervé Moulin, Ariel Procaccia, and Arunava Sen for useful discussions. Nath is funded by the Fulbright-Nehru postdoctoral fellowship. Sandholm is funded by the National Science Foundation under grants 1320620, 1546752, and 1617590, and by the ARO under award W911NF-16-1-0061.

Y. Cai and A. Vetta (Eds.): WINE 2016, LNCS 10123, pp. 369–383, 2016.
DOI: 10.1007/978-3-662-54110-4_26

movie theaters. Each of these choices incurs a cost. Since individual preferences are different and sometimes conflicting, the final choice may not make everybody maximally satisfied. This may cause some of the agents to misreport their preferences or drop out of the plan. To alleviate this problem, one can think of monetary transfers so that the friends who get their more-preferred choice pay more than the friends that get their less-preferred choice. Desirable properties of such a choice and payment rule are that (1) the total side payments (transfers among the friends) should sum to zero, so there is no surplus or deficit, and (2) the choice is efficient, that is, the movie that is selected maximizes the sum of all the friends' valuations. Since the valuations are private information of the friends, an efficient decision requires the valuations to be revealed truthfully. This simple example is representative of many joint decision-making problems that often involve monetary transfers. Consider, for example, a group of firms sharing time on a jointly-owned supercomputer, city dwellers deciding on the location and choice of a public project (e.g., stadium, subway, or library), mobile service providers dividing spectrum among themselves, or a student body deciding which musician or art performer to invite to entertain at their annual function. These problems all call for efficient joint decision making and involve—or could involve depending on the application—monetary transfers.

This is a ubiquitous problem in practice and a classic problem in the academic literature. We study the standard model of this problem where the agents' utilities are *quasi-linear*: each agent's utility is her valuation for the selected alternative (e.g., the choice of movie) minus the money she has to pay. A classic goal is to select an *efficient* alternative, that is, the one that maximizes the sum of the agents' valuations (also known as *social welfare*). We will study the problem of designing strategyproof mechanisms, that is, mechanisms where each agent is best off revealing the truth regardless of what other agents reveal.

Even though there are mechanisms that select efficient alternatives in a truthful manner (e.g., the Vickrey-Clarke-Groves (VCG) mechanism [5,14,34]), the transfers by the individuals do not sum to zero (in public goods settings, the VCG mechanisms leads to too much money being collected from the agents). The execution of such a mechanism needs an external mediator who consumes the surplus (or may need to pay the deficit), to keep the mechanism truthful and efficient—a phenomenon known as 'money burning' in literature. In our movie selection example, this implies that we need a third party who will collect the additional money paid by the individuals, which is highly impractical in many settings. This has attracted significant criticism of the VCG mechanism [30]. Ideally, one would like to design strategyproof mechanisms that are efficient and *budget balanced*, that is, they do not have any surplus or deficit. Green and Laffont [13] proved a seminal impossibility for this setting: in the general quasi-linear domain, strategyproof, efficient mechanisms cannot be budget balanced.

In this paper, we primarily focus on the problem of minimizing inefficiency subject to budget balance in the general setting of quasi-linear utilities. This is because, in the applications of interest to this paper (e.g., movie selection), budget balance is more critical than efficiency. However, we show that for a large set of agents, the per-agent inefficiency vanishes. We also show that for deterministic settings, optimizing the sum (or any convex combination) of efficiency and budget balance—which seems to be the most sensible objective—does not provide any asymptotic benefit over maximizing efficiency subject to budget balance.

1.1 Contributions of this Paper

In this paper, we assume that the agents' valuations are picked from a bounded open interval. In Sect. 3, we characterize the structure of truthful, budget balanced, *deterministic* mechanisms in this restricted domain, and show that any such mechanism must have a *sink* agent,[1] whose reported valuation function does not impact the choice of alternative and she gets the payments made by the other agents (Theorem 1). This result strengthens the Green and Laffont impossibility by showing that even in a restricted domain of bounded valuations, there does not exist a mechanism that is strategyproof, budget balanced, and takes every agent's valuation into consideration—a corollary of which is that it cannot be efficient. With the help of this characterization, we find the optimal deterministic mechanism that minimizes the inefficiency. This provides a tight lower bound on the inefficiency of deterministic, strategyproof, budget-balanced mechanisms. By inefficiency of a mechanism in this paper, we mean the worst-case inefficiency over all valuation profiles. We provide a precise rate of decay ($\frac{1}{n}$) of the inefficiency with the increase in the number of agents (Theorem 2). This implies that the inefficiency vanishes for large number of agents. To contrast this mechanism with the class of mechanisms that minimize budget imbalance subject to efficiency, we considered the joint minimization problem of a convex combination of *inefficiency* and budget imbalance, and observed that it does *not* provide any asymptotic benefit over the previous problem. Due to limited space, we discuss this only in the full version of this paper [28].

We investigate the advantages of randomized mechanisms in Sect. 4. We first consider the class of *generalized sink* mechanisms. These mechanisms have, for every possible valuation profile, a probability distribution over the agents that determines each agent's chance of becoming the sink. This class of mechanisms is budget balanced by design. We show examples where mechanisms from this class are not strategyproof (Algorithm 2), and then isolate an interesting subclass whose mechanisms are strategyproof, the *modified irrelevant sink mechanisms* (Algorithm 3). We show that no mechanism from this class can perform better than the deterministic mechanisms if the number of alternatives is greater than the number of agents (Theorem 3). Since inefficiency (weakly) increases with the

[1] Mechanisms using this idea have been presented with different names in the literature. The original paper by Green and Laffont [13] refers to this kind of agents as a *sample* of the population. Later Gary-Bobo and Jaaidane [11] formalized the randomized version of this mechanism which is known as *polling* mechanism. Faltings [9] refers to this as an *excluded coalition* (when there are multiple such agents) and Moulin [25] mentions this as *residual claimants*. However, we use the term 'sink' for brevity and convenience, and our paper considers a different setup and optimization objective.

number of alternatives (Theorem 4), we consider the extreme case of two alternatives and compare the performances of different mechanisms. We show that a naïve uniform random sink mechanism and the modified irrelevant sink mechanism (Algorithm 3) perform equally well (Theorems 5 and 6) and reduce the inefficiency by a constant factor of 2 from that of the deterministic mechanisms. However, the optimal, strategyproof, budget-balanced, randomized mechanism performs better than these mechanisms. Since the structure of strategyproof randomized mechanisms for general quasi-linear utilities is unknown,[2] we take an optimization-based approach to find the best mechanism for the special case of two agents. This approach is known in the literature as *automated mechanism design* [6]. For an overview, see [32]. We discretize the range of the valuations into finite levels and show that when the number of levels increases—thereby making the lower bound tighter to the actual open-interval problem—the improvement factor reduces to less than 5 (Fig. 1). This is a significant improvement over the class of randomized sink mechanisms, which only improve over the best deterministic mechanism by a factor of 2.

We present experiments using real data from two applications. They show that in practice the inefficiency is significantly smaller and has a faster rate of decay than the worst case bounds (Sect. 5). We conclude the paper in Sect. 6 and present future research directions. Owing to the page limitation, the complete details of the results and the proofs are available in the full version of this paper [28].

1.2 Relationships to Prior Literature

The Green-Laffont impossibility result motivated the research direction of designing efficient mechanisms that are minimally budget imbalanced. The approach is to redistribute the surplus money in a way that satisfies truthfulness of the mechanism [3,4]. The *worst case optimal* and *optimal in expectation* guarantees have been given for this class of mechanisms in restricted settings [16,17,25]. The performance of this class of *redistribution* mechanisms has been evaluated in interesting special domains such as allocating single or multiple (identical or heterogeneous) objects [15]. Also, mechanisms have been developed and analyzed that are budget balanced (or no deficit) and minimize the inefficiency in special settings [18,22,24]. Characterization of strategyproof budget-balanced mechanisms in the setting of cost-sharing is explored by Moulin and Shenker [26] and its quantitative guarantees are presented by Roughgarden and Sundararajan [31]. If the distribution of the agents' valuations is known and we assume common knowledge among the agents over those priors, the strategyproofness requirement can be weakened to Bayesian incentive compatibility. In that weaker framework, mechanisms can extract full expected efficiency and achieve budget balance [1,7]. However, those mechanisms use knowledge of the priors. Therefore, in the general quasi-linear setting, for mechanisms without priors, it is an important open

[2] For randomized mechanisms, results involving special domains are known, e.g., facility location [10,29,33], auctions [8], kidney exchange [2], and most of these mechanisms aim for specific objectives.

question to characterize the class of strategyproof budget-balanced mechanisms, to find such mechanisms that minimize inefficiency, and to find strategyproof mechanisms that minimize the sum (or other convex combination) of inefficiency and budget imbalance. This paper addresses this important research gap in the general quasi-linear setting, for both deterministic and randomized settings. Our approach is also prior-free—the strategyproofness guarantees consider the worst-case scenarios. We show that the answers are asymptotically positive: even in such a general setup, the Green-Laffont impossibility is not too restrictive when the number of agents is large, and our mechanisms seem to work well on real-world datasets.

2 Model and Definitions

We denote the set of agents by $N = \{1, 2, \ldots, n\}$ and the set of alternatives by $A = \{a_1, a_2, \ldots, a_m\}$. We assume that each agent's valuation is drawn from an open interval $(-\frac{M}{2}, \frac{M}{2}) \subset \mathbb{R}$, that is, the valuation of agent i is a mapping $v_i : A \rightarrow (-\frac{M}{2}, \frac{M}{2}), \forall i \in N$ and is a private information. Denote the set of all such valuations of agent i as V_i and the set of valuation profiles by $V = \times_{i \in N} V_i$.

A *mechanism* is a tuple of two functions $\langle f, \mathbf{p} \rangle$, where f is called the *social choice function* (SCF) that selects the *allocation* and $\mathbf{p} = (p_1, p_2, \ldots, p_n)$ is the vector of *payments*, $p_i : V \rightarrow \mathbb{R}, \forall i \in N$. The utility of agent i for an alternative a and valuation profile $v \equiv (v_i, v_{-i})$ is given by the *quasi-linear* function: $v_i(a) - p_i(v_i, v_{-i})$. For *deterministic* mechanisms, $f : V \rightarrow A$ is a deterministic mapping, while for *randomized* mechanisms, the allocation function f is a lottery over the alternatives, that is, $f : V \rightarrow \Delta A$. With a slight abuse of notation, we denote $v_i(f(v_i, v_{-i})) \equiv \mathbb{E}_{a \sim f(v_i, v_{-i})} v_i(a) = f(v_i, v_{-i}) \cdot v_i$ to be the expected valuation of agent i for a randomized mechanism. The following definitions are standard in the mechanism design literature.

Definition 1 (Strategyproofness). *A mechanism $\langle f, \mathbf{p} \rangle$ is* strategyproof *if for all $v \equiv (v_i, v_{-i}) \in V$,*

$$v_i(f(v_i, v_{-i})) - p_i(v_i, v_{-i}) \geq v_i(f(v_i', v_{-i})) - p_i(v_i', v_{-i}), \quad \forall\ v_i' \in V_i, i \in N.$$

Definition 2 (Efficiency). *An allocation f is* efficient *if it maximizes social welfare, that is, $f(v) \in \text{argmax}_{a \in A} \sum_{i \in N} v_i(a),\ \forall v \in V$.*

Definition 3 (Budget Balance). *A payment function $p_i : V \rightarrow \mathbb{R}, i \in N$ is budget balanced if $\sum_{i \in N} p_i(v) = 0,\ \forall v \in V$.*

In addition, in parts of this paper we will consider mechanisms that are oblivious to the alternatives—a property known as *neutrality*. To define this, we consider a permutation $\pi : A \rightarrow A$ of the alternatives. Therefore, π over a randomized mechanism and over a valuation profile will imply that the probability masses and the valuations of the agents are permuted over the alternatives according to π, respectively.[3]

[3] We have overloaded the notation of π following the convention in social choice literature (see, e.g., Myerson [27]). The notation $\pi(v)$ denotes the valuation profile where the alternatives are permutated according to π.

Definition 4 (Neutrality). *A mechanism* $\langle f, \mathbf{p} \rangle$ *is* neutral *if for every permutation of the alternatives* π *(where* $\pi(v) \neq v$*) we have*

$$\pi(f(v)) = f(\pi(v)) \quad \text{and} \quad p_i(\pi(v)) = p_i(v), \quad \forall v \in V, \forall i \in N.$$

Note that efficient social choice functions are neutral and the Green-Laffont result implicitly assumes this property.

The most important class of allocation functions in the context of deterministic mechanisms are *affine maximizers*, defined as follows.

Definition 5 (Affine Maximizers). *An allocation function* f *is an* affine maximizer *if there exist real numbers* $w_i \geq 0, i \in N$*, not all zeros, and a function* $\kappa : A \to \mathbb{R}$ *such that* $f(v) \in \text{argmax}_{a \in A} \left(\sum_{i \in N} w_i v_i(a) + \kappa(a)\right)$.

As we will explain in the body of this paper, we will focus on *neutral* affine maximizers [23], where the function κ is zero.

$$f(v) \in \underset{a \in A}{\text{argmax}} \sum_{i \in N} w_i v_i(a) \qquad \textbf{neutral affine maximizer} \qquad (1)$$

The following property of the mechanism ensures that two different payment functions of an agent, say i, that implement the same social choice function differ from each other by a function that does not depend on the valuation of agent i.[4]

Definition 6 (Revenue Equivalence). *An allocation* f *satisfies* revenue equivalence *if for any two payment rules* p *and* p' *that make* f *strategyproof, there exist functions* $h_i : V_{-i} \to \mathbb{R}$*, such that*

$$p_i(v_i, v_{-i}) = p'_i(v_i, v_{-i}) + h_i(v_{-i}), \;\; \forall v_i \in V_i, \forall v_{-i} \in V_{-i}, \forall i \in N.$$

The metrics of inefficiency we consider in this paper are defined as follows.

Definition 7 (Sample Inefficiency). *The* sample inefficiency *for a deterministic mechanism* $\langle f, \mathbf{p} \rangle$ *is:*

$$r_n^M(f) := \frac{1}{nM} \sup_{v \in V} \left[\max_{a \in A} \sum_{i \in N} v_i(a) - \sum_{i \in N} v_i(f(v)) \right]. \qquad (2)$$

The metric is adapted to expected sample inefficiency *for randomized mechanisms:*

$$r_n^M(f) := \frac{1}{nM} \sup_{v \in V} \left\{ \mathbb{E}_{f(v)} \left[\max_{a \in A} \sum_{i \in N} v_i(a) - \sum_{i \in N} v_i(f(v)) \right] \right\}. \qquad (3)$$

The majority of this paper is devoted to finding strategyproof and budget balanced mechanisms $\langle f, \mathbf{p} \rangle$ that minimize the sample inefficiency.

[4] This definition is a generalization of auction revenue equivalence and is commonly used in the social choice literature (see, e.g., Heydenreich et al. [21]).

A different, but commonly used, metric of inefficiency in the literature is the worst-case ratio of the social welfare of the mechanism and the maximum social welfare: $\inf_{v \in V} \frac{\sum_{i \in N} v_i(f(v))}{\max_{a \in A} \sum_{i \in N} v_i(a)}$. A conclusion similar to what we prove in this paper: "*inefficiency vanishes when* $n \to \infty$", holds in that metric as well, but unlike our metric, that metric would require an additional assumption that the valuations are positive, which is not always the case in a quasi-linear domain.

We are now ready to start presenting our results. We begin with deterministic mechanisms that are strategyproof and budget balanced.

3 Deterministic, Strategyproof, Budget-Balanced Mechanisms

Before presenting the main result of this section, we formally define a class of mechanisms we call *sink* mechanisms. A sink mechanism has one or more *sink* agents, given by the set $S \subset N$, picked a priori, whose valuations are not used when computing the allocation (i.e., $f(v) = f(v_{-S})$) and the sink agents do not pay anything and together they receive the payments made by the other agents. The advantage of a sink mechanism is that it is strategyproof if it is strategyproof for the agents other than the sink agents and the surplus is divided among the sink agents in some reasonable manner, and sink mechanisms are budget balanced by design. An example of a sink mechanism is where $S = \{i_s\}$ (only one sink agent) and $f(v_{-i_s})$ chooses an alternative that would be efficient had agent i_s not exist, that is, $f(v_{-i_s}) = \text{argmax}_{a \in A} \sum_{i \in N \setminus \{i_s\}} v_i(a)$. The Clarke [5] payment rule can be applied here to make the mechanism strategyproof for the rest of the agents—that is, for agents other than i_s, $p_i(v_{-i_s}) = \max_{a \in A} \sum_{j \in N \setminus \{i_s, i\}} v_j(a) - \sum_{j \in N \setminus \{i_s, i\}} v_j(f(v_{-i_s})), \ \forall i \in N \setminus \{i_s\}$. Paying agent i_s the 'leftover' money (that is, $p_{i_s}(v_{-i_s}) = -\sum_{j \in N \setminus \{i_s\}} p_j(v_{-i_s})$) makes the mechanism budget balanced. Our first result establishes that the existence of a sink agent is not only sufficient but also *necessary* for deterministic mechanisms.

Theorem 1. *Any deterministic, strategyproof, budget-balanced, neutral mechanism* $\langle f, \mathbf{p} \rangle$ *in the domain* V *has at least one sink agent.*[5]

All proofs are provided in the full version of this paper [28]. This proof involves two steps. First, we leverage the fact that a mechanism that satisfies the stated axioms must necessarily be a neutral affine maximizer (Eq. 1) and has a specific structure for payments. The characterization of the payment structure comes

[5] Green and Laffont's impossibility result holds for efficient mechanisms, and all efficient mechanisms are neutral. However, there could be instances where multiple alternatives are efficient, i.e., there is a tie. The neutrality of an efficient rule is up to tie-breaking, and Green-Laffont applies no matter how the tie is broken. Similarly, our result also holds irrespective of how the tie is broken. Therefore, this theorem covers and generalizes that result since having at least one sink agent implies that the outcome cannot be efficient.

from the revenue equivalence result. The second part of the proof shows that for such payment functions, it is impossible to have no sink agents (identified as agents that have zero weights, $w_i = 0$, in the affine maximizer). This is shown in a contrapositive manner—assuming that there is no sink agent, we construct valuation profiles that lead to a contradiction to budget balance.

The next goal is to find the mechanism in this class that gives the *lowest* sample inefficiency (Eq.2). In the proof of the next theorem (presented in [28]) we show that this is achieved when there is exactly one sink and the neutral affine maximizer weights are equal for all other agents. This, in turn, yields the following lower bound on inefficiency.

Theorem 2. *For every deterministic, strategyproof, budget-balanced, neutral mechanism* $\langle f, \mathbf{p} \rangle$ *over* V*,* $r_n^M(f) \geq \frac{1}{n}$*. This bound is tight.*

4 Randomized, Strategyproof, Budget-Balanced Mechanisms

In Sect. 3, we saw that the best sample inefficiency achieved by a deterministic budget balanced mechanism is $\frac{1}{n}$. In this section, we discuss how the inefficiency can be reduced by considering randomized mechanisms. An intuitive approach is to consider a mechanism where each agent is picked as a sink with probability $\frac{1}{n}$.

Definition 8 (Naïve Randomized Sink). *A* naïve randomized sink *(NRS) mechanism picks every agent as a sink w.p.* $\frac{1}{n}$ *and takes the efficient allocation without that agent. The payments of the non-sink agents are VCG payments without the sink. The surplus is transferred to the sink.*

Clearly, this mechanism is strategyproof, budget balanced, and neutral by design. One can anticipate that this may not yield the best achievable inefficiency bound. Unlike deterministic mechanisms, very little is known about the structure of randomized strategyproof mechanisms in the general quasi-linear setting. Furthermore, we consider mechanisms that are budget-balanced in addition. Hence, even though we can obtain an upper bound on the expected sample inefficiency ($r_n^M(f)$) by considering specific mechanisms like the NRS mechanism described above, the problem of providing a lower bound (i.e., no randomized mechanism can achieve a smaller $r_n^M(f)$ than a given number), seems elusive in the general quasi-linear setting.

Therefore, in the following two subsections, we consider two approaches, respectively. First, we show lower bounds in a special class of strategyproof, budget-balanced, randomized mechanisms. Second, we provide a lower bound of the optimal, strategyproof, budget-balanced, randomized mechanism for two agents and two alternatives, using a discrete relaxation of the original problem (in the spirit of *automated mechanism design* [6,32]). However, the problems of finding a mechanism that matches this lower bound and extending the lower bound to any number of agents and alternatives are left as future work.

4.1 Generalized Sink Mechanisms

In the first approach, we consider a broad class of randomized, budget-balanced mechanisms, which we coin *generalized sink mechanisms*. In this class, the probability of an agent i to become a sink is dependent on the valuation profile $v \in V$, and we consider mechanisms with only *one* sink, i.e., if the probability vector returned by a generalized sink mechanism is $g(v)$, then w.p. $g_i(v)$, agent i is treated as the *only* sink agent.[6] Clearly, the naïve randomized sink mechanism belongs to this class. Once agent i is picked as a sink, the alternative chosen is the *efficient* one *without* agent i. All agents $j \neq i$ are charged a Clarke tax payment in the world without i, and the surplus amount of money is transferred to the sink agent i. Algorithm 1 shows the steps of a generic mechanism in this class.

ALGORITHM 1. Generalized Sink Mechanisms, $\mathcal{G}$

Input: a valuation profile $v \in V$
A generic mechanism in this class is characterized by a probability distribution over the agents N (which may depend on the valuation profile), $g : V \to \Delta N$
The mechanism randomly picks one agent i in N with probability $g_i(v)$
Treat agent i as the sink

Clearly, not every mechanism in this class is strategyproof. The crucial aspect is how the probabilities of choosing the sink are decided. If the probability $g_i(v)$ depends on the valuation of agent i, that is, v_i, then there is a chance for agent i to misreport v_i to have higher (or lower) probability of being a sink (being a sink could be beneficial since she gets all the surplus). For example, the *irrelevant sink* mechanism given in Algorithm 2 is *not* strategyproof.

ALGORITHM 2. Irrelevant Sink Mechanism (not strategyproof)

Input: a valuation profile $v \in V$
for agent i in N **do**
    Define: $a^*(v_{-i}) = \text{argmax}_{a \in A} \sum_{j \neq i} v_j(a)$
    if $\sum_{j \neq i} v_j(a^*(v_{-i})) - \sum_{j \neq i} v_j(a) > M$ for all $a \in A \setminus \{a^*(v_{-i})\}$ **then**
        Call i an irrelevant agent
if irrelevant agent is found **then**
    Arbitrarily pick one of them as a sink with probability 1
else
    Pick an agent i with probability $\frac{1}{n}$ and treat as sink

In the full version of this paper [28], we provide an counterexample to strategyproofness of this mechanism. However, a small modification of the previous mechanism leads to a strategyproof generalized sink mechanism. This shows

[6] One can think of a more general class of sink mechanisms where multiple agents are treated as sink agents simultaneously. However, it is easy to see—by a similar argument to that in the context of deterministic mechanisms—that using multiple sinks cannot decrease inefficiency.

that the class of generalized sink mechanisms is indeed richer than the constant probability sink mechanisms. In the modified version, we pick a default sink with a certain probability, which will be the sink if there exists no irrelevant agent among the rest of the agents. The change here is that when an agent is picked as a default sink, her valuation has no effect in deciding the sink. See Algorithm 3.

ALGORITHM 3. Modified Irrelevant Sink Mechanism (strategyproof)

```
Input: a valuation profile v ∈ V
Pick agent i as a default sink with probability p_i
for agent j in N \ {i} do
    if irrelevant agent(s) found within N \ {i} then
        Arbitrarily pick one of them as a sink
        Irrelevant agent is found
if no irrelevant agent is found within N \ {i} then
    Treat agent i as sink
```

It is easy to verify that this mechanism is strategyproof. Interestingly, no generalized sink mechanism can improve the expected sample inefficiency over deterministic mechanisms if there are more alternatives than agents ($m > n$).

Theorem 3 (Generalized Sink for $m > n$). *If $m > n$, every generalized sink mechanism has expected sample inefficiency $\geq \frac{1}{n}$.*

The proof is critically dependent on $m > n$. However, we can hope for a smaller inefficiency if the number of alternatives is small. We state this intuition formally as follows.

Theorem 4 (Increasing Inefficiency with m). *For every mechanism f and for a fixed number of agents n, the expected sample inefficiency is non-decreasing in m, i.e., $r^M_{n,m_1}(f) \geq r^M_{n,m_2}(f), \forall m_1 > m_2$.*[7]

Theorems 3 and 4 suggest that in order to minimize inefficiency, one must have a small number of alternatives. So from now on, we consider the extreme case with $m = 2$, where we investigate the advantages of randomization.

For two alternatives, the following theorem shows that the naïve randomized sink (NRS) mechanism reduces the inefficiency by a factor of two.

Theorem 5 (Naïve Randomized Sink). *For $m = 2$, the expected sample inefficiency of the NRS mechanism is $\frac{1}{n^2}\left\lceil \frac{n}{2} \right\rceil \sim \frac{1}{2n}$.*

Even though the modified irrelevant sink (MIS) mechanism (Algorithm 3) is more sophisticated than NRS, it turns out that both of them have the same inefficiency on every valuation profile. Both mechanisms choose a single agent as a sink. The default sink for MIS is chosen uniformly at random, identical to the choice of the sink for NRS. If there does not exist an irrelevant sink in the rest of the agents, the inefficiency remains the same as that for the default sink,

[7] We overload the notation for the expected sample inefficiency r_n with $r_{n,m}$ to make the number of alternatives explicit.

which is identical to the inefficiency of NRS for that choice of sink. But even if an irrelevant sink exists, by the construction of the irrelevant sink, the resulting alternative is the efficient alternative for the agents except the default sink. This outcome would have resulted even if the default sink were chosen as the sink. Therefore, the inefficiencies in MIS and NRS mechanisms are the same. Hence, we get the following theorem.

Theorem 6 (Modified Irrelevant Sink). *For $m = 2$, the expected sample inefficiency of the MIS mechanism (Algorithm 3) is at least $\frac{1}{n^2}\left\lceil\frac{n}{2}\right\rceil \sim \frac{1}{2n}$.*

4.2 Unrestricted Randomized Mechanisms

We now move on to study optimal randomized mechanisms without restricting attention necessarily to generalized sink mechanisms. For a fixed number of agents, minimizing the expected sample inefficiency is equivalent to minimizing the expected *absolute* inefficiency given by $nr_n^M(f)$. Finding a mechanism that achieves the minimum absolute inefficiency can be posed as the following optimization problem.

$$
\begin{aligned}
\min_{f,\mathbf{p}} \quad & \sup_{v\in V}\left[\max_{a\in A}\sum_{i\in N} v_i(a) - \sum_{i\in N} v_i(f(v))\right] \\
\text{s.t.} \quad & v_i(f(v_i, v_{-i})) - p_i(v_i, v_{-i}) \\
& \qquad \geq v_i(f(v_i', v_{-i})) - p_i(v_i', v_{-i}), \ \forall v_i, v_i', v_{-i}, \forall i \in N \\
& \sum_{a\in A} f_a(v) = 1, \ \forall v \in V, \\
& \sum_{i\in N} p_i(v) = 0, \ \forall v \in V, \\
& f_a(v) \geq 0, \ \forall v \in V, a \in A.
\end{aligned} \tag{4}
$$

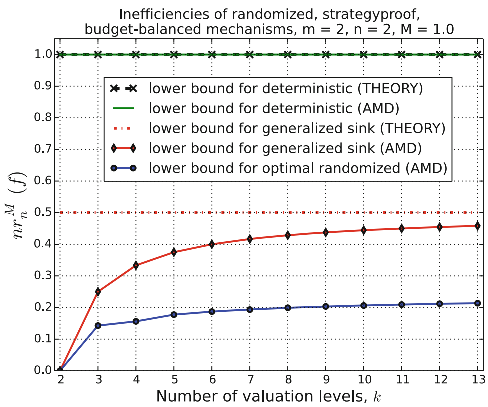

Fig. 1. Lower bound for the discrete relaxation of the inefficiency minimization LP.

The objective function denotes the absolute inefficiency. The first set of inequalities in the constraints denote the strategyproofness requirement, where the term $v_i(f(v)) = v_i \cdot f(v)$ denotes the expected valuation of agent i due to the randomized mechanism f. The second and last set of inequalities ensure that the $f_a(v)$'s are valid probability distributions, and the third set of inequalities ensure that the budget is balanced. The optimization is over the social choice functions f and the payments $\mathbf{p}$, where the f variables are non-negative but the p variables are unrestricted. Clearly, this is a linear program (LP), which has an uncountable number of constraints (because the equalities and inequalities have to be satisfied at all $v \in V$, which are the profiles of valuation functions mapping alternatives to an open interval). We address this optimization problem using finite constrained optimization techniques by discretizing the valuation levels. We assume that each agent's valuations are uniformly discretized with k levels in $[-M/2, M/2]$, which makes the set of valuation profiles V finite. The optimal value of such a discretized relaxation of the constraints provides a lower bound on the optimal value of the original problem. This is because the discretized relaxation of the valuations only increases the feasible set since some of the constraints are removed, that is, more f's and p's satisfy the constraints, allowing a potentially lower value to be achieved for the minimization objective.

We conducted a form of automated mechanism design [6,32] by solving this LP using Gurobi [19] for increasing values of k. We apply the same optimization-based approach for generalized sink and the deterministic cases as well, even though for these cases we have theoretical bounds. The solid lines in Fig. 1 show the optimization-based results (denoted as AMD) and the dotted lines show the theoretical bounds. Note that for deterministic case, the theoretical and optimization-based approaches overlap since the inefficiency is unity even with two valuation levels. The convergence of the optimization-based approach for generalized sink mechanism shows the efficacy of the approach and helps to predict the convergence point for the optimal randomized mechanism. One can see that the lower bound is greater than 0.2 for the optimal mechanism, but it seems to converge to a value much lower than 0.5.

5 Experiments with Real Data

In this section we investigate the average and worst-sink performances of the the naïve randomized sink (NRS) (Definition 8) mechanism on real datasets of user preferences. Going back to the example of movie selection by a group of friends (Sect. 1), we consider several sizes of the group. A small group consists of tens of friends, while if the decision involves screening a movie at a school auditorium, the group size could easily be in the hundreds. This is why we consider group sizes spanning from 10 to 210 in steps of 50.

A similar situation occurs when a group of people decides which comedian/musician to invite in a social gathering, where they need to pay the cost of bringing the performer. Keeping these motivating situations in mind, we used two datasets that closely represent the scenarios discussed. We used the MovieLens 20M dataset [20] and the Jester dataset [12] to compare the average and worst-case performances of NRS. The first dataset contains preferences for movies,

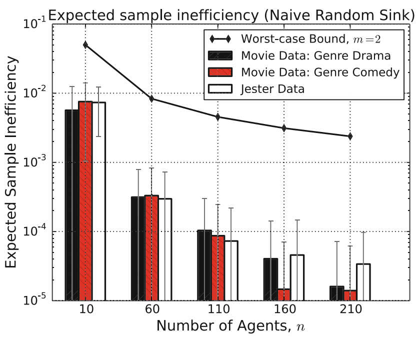

Fig. 2. Naïve random sink mechanism

while the second contains preferences for online jokes. The MovieLens 20M dataset (`ml-20m`) describes users' ratings between 1 and 5 stars from MovieLens, a movie recommendation service. It contains 20,000,263 ratings across 27,278 movies. These data were created from the ratings of 138,493 users between January 09, 1995 and March 31, 2015. For our experiment, we sampled the preferences of a specific number of users (shown as agents on the x-axis of Figs. 2 and 3) 100 times uniformly at random from the whole set of users that rated a particular genre of movies, and computed the sample inefficiency on this sampled set and plotted it along with the standard deviation.

The Jester dataset (`jester-data-1`) used in our experiment contains data from 24,983 users who have rated 36 or more jokes, a matrix with dimensions 24983 X 100, and is obtained from Jester, an online joke recommendation system.[8]

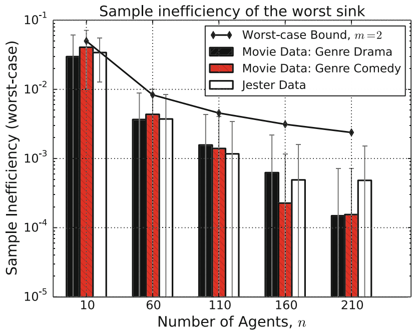

Fig. 3. Worst-sink behavior

Figure 2 shows that the real preferences of users yield much lower expected sample inefficiencies for the naïve randomized sink (NRS) mechanism than the theoretical worst-case guarantee. The improvement ranges from roughly a factor of 5 (for a group size of 10) to almost 100 (for a group size of 210). This also indicates that the rate of decay of the inefficiency with the size of the group is faster than the theoretical guarantee. The bars in Figs. 2 and 3 show the *average* (w.r.t. the randomly selected users) expected sample inefficiency (Eq. 3) and the inefficiency of the worst sink of the NRS mechanism respectively with the standard deviations around them.

By the arguments preceding Theorem 6 and since MIS (Algorithm 3) also picks exactly one sink, it is easy to see that the average inefficiency and inefficiency of the worst sink of MIS will be same as NRS.

6 Conclusions and Future Research

In this paper, we considered the classic question of the interplay between efficiency and budget balance, properties that are incompatible with strategyproofness due to the Green-Laffont impossibility result, in the general quasi-linear setting. We proved the limits of possibility in the context of deterministic

[8] In both datasets there are missing values because a user has typically not rated all movies/jokes. Before our experiment, we filled the missing values with a random realization of ratings drawn from the empirical distribution for that alternative (movie or joke). The empirical distribution of an alternative is created from the histogram of the available ratings of the users. We cleaned the dataset by keeping only those alternatives that have at least 10 or more available ratings and filled the rest using their empirical distributions.

mechanisms for both efficiency and budget balance. For randomized mechanisms, we identified a class of mechanisms that perform better than deterministic ones. We used an optimization-based scheme to find the optimal randomized mechanism. Experiments with real datasets showed that the values (rate of decay) of inefficiency are significantly smaller (faster) than those of the theoretical worst case. Future research includes studying the structure of the optimal randomized mechanisms that achieve the (theoretical) improved efficiency. Future work also includes investigating the rate of improvement of the optimal bound for a general number of agents.

References

1. Arrow, K.: The property rights doctrine and demand revelation under incomplete information. In: Economics and Human Welfare. New York Academic Press (1979)
2. Ashlagi, I., Fischer, F., Kash, I.A., Procaccia, A.D.: Mix and match: a strategyproof mechanism for multi-hospital kidney exchange. Games Econ. Behav. **91**, 284–296 (2015)
3. Bailey, M.J.: The demand revealing process: to distribute the surplus. Public Choice **91**(2), 107–126 (1997)
4. Cavallo, R.: Optimal decision-making with minimal waste: strategyproof redistribution of VCG payments. In: Proceedings of the Conference on Autonomous Agents and Multiagent Systems (AAMAS), pp. 882–889 (2006)
5. Clarke, E.: Multipart pricing of public goods. Public Choice **8**, 19–33 (1971)
6. Conitzer, V., Sandholm, T.: Complexity of mechanism design. In: Proceedings of the Conference on Uncertainty in Artificial Intelligence (UAI), pp. 103–110. Morgan Kaufmann Publishers Inc. (2002)
7. d'Aspremont, C., Gérard-Varet, L.A.: Incentives and incomplete information. J. Public Econ. **11**(1), 25–45 (1979)
8. Dobzinski, S., Nisan, N., Schapira, M.: Truthful randomized mechanisms for combinatorial auctions. In: Proceedings of the Annual ACM Symposium on Theory of Computing, pp. 644–652. ACM (2006)
9. Faltings, B.: A budget-balanced, incentive-compatible scheme for social choice. In: Faratin, P., Rodríguez-Aguilar, J.A. (eds.) AMEC 2004. LNCS (LNAI), vol. 3435, pp. 30–43. Springer, Heidelberg (2006). doi:10.1007/11575726_3
10. Feldman, M., Wilf, Y.: Randomized strategyproof mechanisms for facility location, the mini-sum-of-squares objective arXiv:1108.1762 (2011)
11. Robert, R.J., Jaaidane, T.: Polling mechanisms and the demand revelation problem. J. Public Econ. **76**(2), 203–238 (2000)
12. Goldberg, K., Roeder, T., Gupta, D., Perkins, C.: Eigentaste: a constant time collaborative filtering algorithm. Inf. Retr. **4**(2), 133–151 (2001)
13. Green, J.R., Laffont, J.-J.: Incentives in Public Decision Making. North-Holland, Amsterdam (1979)
14. Groves, T.: Incentives in teams. Econometrica **41**(4), 617–631 (1973)
15. Gujar, S.P., Narahari, Y.: Redistribution mechanisms for assignment of heterogeneous objects. J. Artif. Intell. Res. **41**, 131–154 (2011)
16. Guo, M., Conitzer, V.: Optimal-in-expectation redistribution mechanisms. In: Proceedings of the Conference on Autonomous Agents and Multiagent Systems (AAMAS), pp. 1047–1054 (2008)

17. Guo, M., Conitzer, V.: Worst-case optimal redistribution of VCG payments in multi-unit auctions. Games Econ. Behav. **67**(1), 69–98 (2009)
18. Guo, M., Conitzer, V.: Better redistribution with inefficient allocation in multi-unit auctions. Artif. Intell. **216**, 287–308 (2014)
19. Gurobi. Gurobi optimizer reference manual (2015)
20. Harper, F.M., Konstan, J.A.: The movielens datasets: history and context. ACM Trans. Interact. Intell. Syst. (TiiS) **5**(4), 19 (2016)
21. Heydenreich, B., Müller, R., Uetz, M., Vohra, R.V.: Characterization of revenue equivalence. Econometrica **77**(1), 307–316 (2009)
22. Massó, J., Nicolò, A., Sen, A., Sharma, T., Ülkü, L.: On cost sharing in the provision of a binary and excludable public good. J. Econ. Theory **155**, 30–49 (2015)
23. Mishra, D., Sen, A.: Roberts' theorem with neutrality: a social welfare ordering approach. Games Econ. Behav. **75**(1), 283–298 (2012)
24. Mishra, D., Sharma, T.: On the optimality of the green-laffont mechanism. Preliminary Draft (2016)
25. Moulin, H.: Almost budget-balanced VCG mechanisms to assign multiple objects. J. Econ. Theory **144**(1), 96–119 (2009)
26. Moulin, H., Shenker, S.: Strategyproof sharing of submodular costs: budget balance versus efficiency. Econ. Theor. **18**(3), 511–533 (2001)
27. Myerson, R.B.: Fundamentals of social choice theory. Q. J. Polit. Sci. **8**(3), 305–337 (2013)
28. Nath, S., Sandholm, T.: Efficiency and budget balance in general quasi-linear domains. arXiv:1610.01443 [cs.GT] (2016)
29. Procaccia, A.D., Tennenholtz, M.: Approximate mechanism design without money. In: Proceedings of the 10th ACM Conference on Electronic Commerce (EC), pp. 177–186. ACM (2009)
30. Rothkopf, M.H.: Thirteen reasons why the Vickrey-Clarke-Groves process is not practical. Oper. Res. **55**(2), 191–197 (2007)
31. Roughgarden, T., Sundararajan, M.: Quantifying inefficiency in cost-sharing mechanisms. J. ACM (JACM) **56**(4), 23 (2009)
32. Sandholm, T.: Automated mechanism design: a new application area for search algorithms. In: Rossi, F. (ed.) CP 2003. LNCS, vol. 2833, pp. 19–36. Springer, Heidelberg (2003). doi:10.1007/978-3-540-45193-8_2
33. Thang, N.K.: On randomized strategy-proof mechanisms without payment for facility location games. In: Proceedings of the Web and Internet Economics (WINE), Stanford, USA, pp. 13–16 (2010)
34. Vickrey, W.: Counterspeculation, auctions, and competitive sealed tenders. J. Financ. **16**(1), 8–37 (1961)