Efficiency and Budget Balance *

Swaprava Nath[1] and Tuomas Sandholm[1]

[1]Computer Science Department, Carnegie Mellon University

March 9, 2016

Abstract

We study *efficiency* and *budget balance* in mechanism design in the quasi-linear domain. Green and Laffont (1979) 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 given the payments made by the other agents. We assume the valuations of the agents are drawn 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 (1979, 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 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.

1 Introduction

Mechanism design with monetary transfers forms a cornerstone of multi-agent decision making. It has extremely broad applications ranging from resource allocation in the physical world and computer systems, to building public projects, to trading in electronic markets. We study the problem in the standard framework of *quasi-linear* utilities: each agent's utility is her valuation for the selected alternative (e.g., an allocation of resources) minus the amount she has to pay. For example, in an auction, an agent's utility is her valuation for the items that she receives minus what

*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 and 1546752.

she has to pay. The framework also captures the setting where an agent's valuation may depend on how items that she does not receive get allocated. The classic goal is to select an alternative that maximizes *efficiency*, that is, the sum of the agents' valuations.

In the setting where valuations are private information, a mechanism needs to be designed that incentivizes the agents to reveal their valuations truthfully (by the revelation principle, there is no loss in objective from restricting attention to such direct-revelation mechanisms). 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.

Achieving strategyproofness is not always possible, but with quasilinear utilities one can always achieve strategyproofness using the Vickrey-Clarke-Groves (VCG) mechanism (Vickrey, 1961; Clarke, 1971; Groves, 1973), which selects an efficient alternative. However, the efficiency measured only by agents' valuations ignores the overall sum of utilities, which involves the sum of payments too. If the sum of payments is not zero—we call a mechanism *budget balanced* when this sum is zero—then it induces a solution that may not be 'efficient' from the perspective of maximizing the sum of the agents' *utilities*. In particular, (the Clarke tax version of) the VCG mechanism has a positive outflow of money from the agents in aggregate, and this excess money needs to be burned in order to ensure strategyproofness. This has attracted significant criticism of the VCG mechanism (Rothkopf, 2007). Similarly, some mechanisms may require an external benefactor to subsidize the mechanism because there is an inflow of money to the agents in the mechanism. There too, that inflow should be subtracted from the valuations-based efficiency to evaluate the true utility-based efficiency.

Ideally, one would design strategyproof mechanisms that are efficient and budget balanced—that is, they do not require money burning or an external benefactor. Green and Laffont (1979) proved that in the general quasi-linear domain strategyproof efficient mechanisms cannot be budget balanced. This 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 (Bailey, 1997; Cavallo, 2006). 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 (Gujar and Narahari, 2011). The *worst case optimal* and *optimal in expectation* guarantees are given for this class of mechanisms by Moulin (2009); Guo and Conitzer (2008, 2009) in restricted settings. On the other hand, mechanisms have been developed and analyzed that are budget balanced (or no deficit) and minimize the inefficiency in special settings (Massó et al., 2015; Guo and Conitzer, 2014; Mishra and Sharma, 2016). Characterization of strategyproof budget-balanced mechanisms in a cost-sharing setting is explored by Moulin and Shenker (2001).

If the distribution of the valuations of the agents 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 d'Aspremont and Gérard-Varet (1979); Arrow (1979). However, those mechanisms use knowledge of the priors. Therefore, in the general quasi-linear setting, for mechanisms without priors, it is an important open question to characterize the class of strategyproof budget-balanced mechanisms, to find such mechanisms that minimize inefficiency, and to find strategyproof mechanisms that minimize a convex combination of inefficiency and budget imbalance.

1.1 Contributions of this paper

In this paper, we consider the problem of minimizing inefficiency subject to budget balance, and minimizing inefficiency and budget imbalance jointly in the general setting of quasi-linear utilities where agents' valuations are drawn from a bounded open interval. In Section 3, we characterize the structure of truthful, budget balanced, *deterministic* mechanisms in this restricted domain, and show that it must have a *sink* agent[1] whose valuations do 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 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 the deterministic, strategyproof, budget-balanced mechanisms. By inefficiency of a mechanism in this paper, we mean the largest inefficiency of the mechanism 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).

To contrast this mechanism with the class of mechanisms that minimize budget imbalance subject to efficiency, in Section 4 we consider the joint objective of *efficiency-budget spillover*, which is a convex combination of inefficiency and budget imbalance. We prove that no deterministic, strategyproof mechanism can reduce this spillover at a rate faster than $\frac{1}{n}$ (Theorem 3). In other words, minimizing the joint objective does not give any asymptotic advantage over the solution of minimizing inefficiency with the constraint that the mechanism is budget balanced!

We investigate the advantages of randomization in Section 5. We first consider the class of *generalized sink* mechanisms. These mechanisms draw a probability distribution over the agents for every valuation profile which determines their chance of becoming a 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 that is strategyproof (Algorithm 3). We show that no mechanism from this class can perform better than the deterministic mechanisms if the number of alternatives is more than the number of agents (Theorem 4). Since, for a fixed number of agents, increase in the number of alternatives does not decrease the inefficiency (Theorem 5), 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) both perform equally well (Theorems 6 and 7) and reduces 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

[1]Mechanisms using this idea have been presented with different names in the literature. The original paper by Green and Laffont (1979) refers to this kind of agents as a *sample* of the population. Later Gary-Bobo and Jaaidane (2000) formalized the randomized version of this mechanism which is known as *polling* mechanism. Faltings (2004) refers to this as an *excluded coalition* (when there are multiple such agents) and Moulin (2009) 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.

[2]For randomized mechanisms, results involving special domains are known, e.g., facility location (Thang, 2010; Procaccia and Tennenholtz, 2009; Feldman and Wilf, 2011), auctions (Dobzinski et al., 2006), kidney exchange (Ashlagi et al., 2013), and most of these mechanisms aim for specific objectives.

of two agents[3]. We prove that for a discrete valuation space with 3 levels, the optimal inefficiency is reduced by a factor of 7 (Theorem 9). However, when the number of levels increases—thereby making the lower bound tighter to the actual problem of valuations being drawn from an open interval—the improvement factor reduces to less than 5 (Figure 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. In Section 6 we present conclusions and future research directions.

2 Model and definitions

The set of agents is denoted 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 $v_i : A \to (-\frac{M}{2}, \frac{M}{2}), \forall i \in N$ 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 $V = \prod_{i \in N} V_i$.

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*, $p_i : V \to \mathbb{R}, \forall i \in N$. 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, $f : V \to A$ is a deterministic mapping, while for *randomized* mechanisms, the allocation function f is a lottery over the alternatives, that is, $f : V \to \Delta A$. 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, $f(v) \in \operatorname{argmax}_{a \in A} \sum_{i \in N} v_i(a),\ \forall v \in V$.*

Definition 3 **(Budget Balance)** *A payment function $p_i : V \to \mathbb{R}, i \in N$ 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 $\pi : A \to A$ 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.

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*

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

Later in the paper, we will also use another form of agent symmetry called *anonymity*, which ensures that the SCF is insensitive to the agent identities.

[3]This approach is known in the literature as *automated mechanism design* (Conitzer and Sandholm, 2002). For an overview, see Sandholm (2003).

Definition 5 **(Anonymity)** *A mechanism* $\langle f, \mathbf{p}\rangle$ *is* anonymous *if for every permutation of the agents* λ *we have*

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

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

Definition 6 **(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, in parts of the paper we will focus on *neutral* affine maximizers Mishra and Sen (2012), where the function κ is zero.

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

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.

Definition 7 **(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 8 **(Sample Inefficiency)** *The* sample inefficiency *for a deterministic mechanism* $\langle f, \mathbf{p}\rangle$ *is:*

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

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

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

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

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: $\inf_{v\in V} \frac{\sum_{i\in N} v_i(f(v))}{\max_{a\in A} \sum_{i\in N} v_i(a)}$. Our conclusions hold in that metric as well, but unlike our metric, that metric would require an additional assumption that the valuations are non-negative, which is not always the case in a quasi-linear domain.

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

3 Deterministic, strategyproof, budget-balanced mechanisms

In this section we study deterministic budget-balanced mechanisms. We provide characterization results and a tight lower bound on the sample inefficiency.

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 (it does not matter how those payments are divided among the sink 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 sink mechanisms are obviously 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 if agent i_s did 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 tax payment rule Clarke (1971) 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, and neutral mechanism $\langle f, \mathbf{p} \rangle$ in the domain V has at least one sink agent.*

Proof: Consider the class of deterministic, strategyproof, and neutral mechanisms. Mishra and Sen (2012) have shown that in the domain V, an allocation that satisfies the properties above must be a neutral affine maximizer (Definition 6), that is, there exists $w_i \geq 0, \forall i \in N$, not all zero, such that,

$$f(v) \in \underset{a \in A}{\text{argmax}} \sum_{i \in N} w_i v_i(a). \tag{3}$$

Additionally, the result by Rockafellar (1997) and Krishna and Maenner (2001) states that for any convex type space, if the valuations are linear in type, then a strategyproof allocation satisfies revenue equivalence (Definition 7). In our setting, the types of the agents are their valuations, which implies, trivially, that the valuations are linear in type. Also, they are drawn from the interval $(-\frac{M}{2}, \frac{M}{2})$, which is convex. So, revenue equivalence holds for the allocations in our setting. The following payment implements the affine maximizer allocation f given by Equation (3):

$$p_i(v_i, v_{-i}) = \begin{cases} \dfrac{1}{w_i} \left(\displaystyle\sum_{j \neq i} w_j v_j(f(v)) \right), & w_i > 0 \\ 0, & w_i = 0 \end{cases} \tag{4}$$

for all $i \in N$. Since revenue equivalence holds in this setting, we conclude that any payment $\hat{p}_i, i \in N$ that makes $\langle f, \hat{\mathbf{p}} \rangle$ strategyproof, will be different from the above mentioned payments $\mathbf{p}$ by an additive factor $h_i(v_{-i})$ for each agent i in every valuation profile.

Now, we turn to proving the result of the theorem. We have the functional form of deterministic, strategyproof, neutral mechanisms given by Equation (3). If, on this class of mechanisms, we show that one cannot have weights $w_i > 0$ for all $i \in N$ while imposing budget balance, then we are done. This is because, if there exists one agent $i \in N$, for which $w_i = 0$, that agent is a *sink* agent as her valuations are never used by the social choice function and she is charged no payment. By revenue equivalence, any other payment that can implement the same allocation f is $h_i(v_{-i})$. Putting this in the budget balance equation, we get $h_i(v_{-i}) = -\sum_{j \in N \setminus \{i\}} p_j(v)$, that is, she receives the payments made by the other agents. Thus agent i is a sink agent. Hence, the proof is completed by proving the following claim.

LEMMA **1 (Existence of $w_i = 0$ Agent)** *A budget balanced mechanism $\langle f, \mathbf{p} \rangle$, where f is a neutral affine maximizer on the domain V, must have at least one agent i that has $w_i = 0$.*

Proof: Suppose for contradiction that $w_i > 0, \forall i \in N$. Since f is a neutral affine maximizer (Equation (3)) and revenue equivalence holds in V (Equation (4)), we know that the payments are of the form $p_i(v_i, v_{-i}) = h_i(v_{-i}) + \frac{1}{w_i}\left(\sum_{j \neq i} w_j v_j(f(v))\right), \forall v \in V, \forall i \in N$.

Additionally, since the mechanism $\langle f, \mathbf{p} \rangle$ is also budget balanced, we have

$$
\begin{aligned}
& \sum_{i=1}^{n} \left(h_i(v_{-i}) + \frac{1}{w_i} \left(\sum_{j \neq i} w_j v_j(f(v)) \right) \right) = 0, \ \forall v \in V \\
\Rightarrow \quad & \sum_{i=1}^{n} h_i(v_{-i}) + \sum_{i=1}^{n} \left(\sum_{j \neq i} \frac{1}{w_j} \right) w_i v_i(f(v)) = 0, \ \forall v \in V. \qquad (5)
\end{aligned}
$$

For an easier exposition, we first explain the proof technique for $n = 2$. Later the same proof is generalized to any number of agents.

By assumption, $w_1, w_2 > 0$. Pick two valuation profiles (v_1^+, v_2) and (v_1^-, v_2) such that the affine maximizer alternative in the first is a_1 while that in the second is a_2, that is,

$$
w_1 v_1^+(a_1) + w_2 v_2(a_1) > w_1 v_1^+(a_2) + w_2 v_2(a_2) \qquad (6)
$$

$$
w_1 v_1^-(a_1) + w_2 v_2(a_1) < w_1 v_1^-(a_2) + w_2 v_2(a_2) \qquad (7)
$$

This can be done by choosing $v_1^+(a_2) = v_1^-(a_2) = v_1(a_2)$ (say) small and v_2 to be small enough for both alternatives, so that the valuation of agent 1 for a_1 determines the resulting alternative of f. Therefore, the RHS of the inequalities above are the same. Since the inequality of Equation (7) is strict, let the difference of the RHS and LHS be $\delta > 0$. The allocations at these two profiles are: $f(v_1^+, v_2) = a_1$ and $f(v_1^-, v_2) = a_2$. Since the payments satisfy revenue equivalence and budget balance, Equation (5) holds, which gives

$$
\frac{1}{w_1} w_2 v_2(a_1) + h_1(v_2) + \frac{1}{w_2} w_1 v_1^+(a_1) + h_2(v_1^+) = 0
$$

$$
\frac{1}{w_1} w_2 v_2(a_2) + h_1(v_2) + \frac{1}{w_2} w_1 v_1^-(a_2) + h_2(v_1^-) = 0.
$$

Subtracting the first equation from the second and rearranging, we get

$$
\frac{1}{w_1} w_2 (v_2(a_1) - v_2(a_2)) = \frac{1}{w_2} w_1 (v_1^-(a_2) - v_1^+(a_1)) + h_2(v_1^-) - h_2(v_1^+). \qquad (8)
$$

Note that the RHS is independent of v_2. Therefore, if $v_2(a_1)$ is increased by a small amount ($< \delta/w_2$), both the inequalities given by Equations (6) and (7) still hold, but Equation (8) fails to hold, which is a contradiction.

The general proof of this lemma extends this idea to any number of agents $n \geq 2$. We prove this for a set of alternatives $A = \{0, 1\}$. Consider this setting as that of a public project. In alternative 0, the project is not undertaken, yielding every agent a value of zero, and when 1 is chosen—i.e., the project is undertaken—the valuation of each agent is denoted by a single real number. This assumption helps us reduce the notational complexity. The proof, however, is completely general for any number of alternatives.

Let the agents be numbered in decreasing order of their weights WLOG, that is, $w_i \geq w_{i+1}, i = 1, 2, \dots, n-1$. We consider the following valuation profile: $(v_1 + \delta, v_2 + \delta, \dots, v_{n-1} + \delta, v_n), \delta > 0$ such that

$$-\delta \sum_{i=1}^{n-1} w_i < \sum_{i=1}^{n} w_i v_i < -\delta \sum_{i=1}^{n-2} w_i \tag{9}$$

The above inequalities imply that the affine maximizer alternative given by Equation (3) for the profile mentioned above is 1. However, if any agent i's, $i = 1, 2, \dots, n-1$, valuation changes from $v_i + \delta$ to v_i, the alternative changes to 0. We use a generic notation v^k to denote this profile, where k denotes the agent(s) whose valuation(s) is(are) v_k while all other agents $j \neq k$ have valuations $v_j + \delta$. Hence, v^n is the profile mentioned before: $(v_1 + \delta, v_2 + \delta, \dots, v_{n-1} + \delta, v_n)$ and $v^{n-1,n}$ is the profile: $(v_1 + \delta, v_2 + \delta, \dots, v_{n-1}, v_n)$, for example.

Since, $f(v^n) = 1$, from Equation (5) we have,

$$\left(\sum_{i=1}^{n-1} h_i(v_{-i}^n) + h_n(v_{-n}^n) \right) + \left(\sum_{i=1}^{n} \left(\sum_{j \neq i} \frac{1}{w_j} \right) w_i v_i + \sum_{i=1}^{n-1} \left(\sum_{j \neq i} \frac{1}{w_j} \right) w_i \delta \right) = 0. \tag{10}$$

The idea of the proof is to make a series of substitutions in the first parentheses of the expression above, leaving the terms in the other parentheses unchanged. Note that, the expression in the second parentheses depends on v_n, while the expression $h_n(v_{-n}^n)$ does not. The substitutions sequentially eliminate the dependency on v_n from all the terms in the first parentheses, similar to what we did in the two agent case before. This leads to a contradiction, since v_n can be perturbed to be arbitrarily small so that it continues to satisfy the inequalities of Equation (9), our only assumed condition, but violates the equality in Equation (10).

The substitutions will involve the term $\sum_{i=1}^{n-1} h_i(v_{-i}^n)$ in the first parentheses of Equation (10). Consider the profiles $v^{j,n}, j = 1, \dots, n-1$. In each of these profiles, $f(v^{j,n}) = 0$ (due to the choice of v^n in Equation (9)). Hence,

$$\sum_{i=1}^{n-1} h_i(v_{-i}^{j,n}) + h_n(v_{-n}^{j,n}) = 0, \ \forall j \in \{1, \dots, n-1\}. \tag{11}$$

Note that $v_{-i}^{i,n} = v_{-i}^n$. Hence, we can substitute terms from Equation (11) to the terms in the first parentheses of Equation (10) to get,

$$\left(- \sum_{i=1}^{n-1} \sum_{j \neq \{i,n\}} h_j(v_{-j}^{i,n}) - \sum_{j \neq n} h_n(v_{-n}^{j,n}) + h_n(v_{-n}^n) \right) + \left(\sum_{i=1}^{n} \left(\sum_{j \neq i} \frac{1}{w_j} \right) w_i v_i + \sum_{i=1}^{n-1} \left(\sum_{j \neq i} \frac{1}{w_j} \right) w_i \delta \right) = 0. \tag{12}$$

We continue replacing the terms $h_j(v_{-j}^{i,n})$ in the first summation of the first parentheses above. All other terms in that parentheses are h_n functions and, therefore, are independent of v_n. For every $i \neq n$, consider the valuation profiles $v^{j,i,n}, j \neq i, n$. By Equation (9), $f(v^{j,i,n}) = 0$, so we get an equality similar to Equation (11): $\sum_{k=1}^{n-1} h_k(v_{-k}^{j,i,n}) + h_n(v_{-n}^{j,i,n}) = 0, \ \forall j \neq i, n$. Also, $v_{-j}^{j,i,n} = v_{-j}^{i,n}$. So, we follow the same procedure to replace the terms $h_j(v_{-j}^{i,n})$ in Equation (12) to yield a similar equality where more terms that were dependent on v_n are now replaced with h_n functions, which are independent of v_n. Since the number of agents is finite, this process will stop after a finite number of iterations, reducing the terms in the first parentheses to only consisting of h_n functions. This construction shows that a small perturbation of v_n, which keeps Equation (9) unaffected, will violate the equality obtained through the iterative procedure described above. This completes the proof of the lemma. ■

Lemma 1 shows that there exists an agent with weight zero, which is a sink agent, and hence the proof of Theorem 1 is complete. ■

Theorem 1 states that a deterministic, strategyproof, budget-balanced, neutral mechanism must necessarily be a neutral affine maximizer (Equation (3)) that has at least one sink agent. Our next goal is to find the mechanism in this class that gives the *lowest* sample inefficiency (Equation (1)). We show that it is achieved when there is exactly one sink and weights are equal for all agents other than the sink.

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.*

Proof: From Theorem 1, we know that any f that satisfies the properties mentioned in the statement of the current theorem must be a neutral affine maximizer with at least one agent i^* that has $w_{i^*} = 0$. We now show that the minimum sample inefficiency $r_n^M(f)$ is achieved when there is exactly *one* such agent i^* and the weights of the other agents $i \in N \setminus \{i^*\}$ are equal. This immediately proves the theorem since the ensuing mechanism will have sample inefficiency $\frac{1}{n}$. This mechanism picks the welfare maximizing allocation not considering the sink agent i^*, that is, $f(v) \in \text{argmax}_{a \in A} \sum_{j \in N \setminus \{i^*\}} v_j(a)$. Denoting $a^*(v)$ to be the efficient allocation, we can write the expression:

$$
\begin{aligned}
&\max_{a \in A} \sum_{i \in N} v_i(a) - \sum_{i \in N} v_i(f(v)) = \sum_{i \in N} v_i(a^*(v)) - \sum_{i \in N} v_i(f(v)) \\
&= v_{i^*}(a^*(v)) - v_{i^*}(f(v)) + \left[\sum_{j \in N \setminus \{i^*\}} v_j(a^*(v)) - \sum_{j \in N \setminus \{i^*\}} v_j(f(v)) \right] < \left(\frac{M}{2} - \left(-\frac{M}{2} \right) \right) + 0 = M
\end{aligned} \tag{13}
$$

The first part of the above inequality comes from the fact that the valuations are drawn from $(-\frac{M}{2}, \frac{M}{2})$ so the difference in valuation can at most be M. The second part of the inequality is due to $f(v)$ being the welfare maximizing allocation when excluding agent i^*.

It is easy to verify that this inequality is tight at the following valuation profile: $v_{i^*}(a) = \frac{M}{2} - \delta, v_{i^*}(z) = -\frac{M}{2} + \gamma, \ \forall z \neq a$, and $v_j(b) = -\frac{M}{2} + \epsilon, v_j(z) = -\frac{M}{2} + \frac{\epsilon}{2}, \ \forall z \neq b, \forall j \neq i^*$, where $\delta, \gamma, \epsilon > 0$ are arbitrarily small. The alternatives are: $a^*(v) = a, f(v) = b$. Clearly, this satisfies the above inequality and by taking $\delta, \gamma, \epsilon \to 0$, we get that the supremum of the difference term approaches M, and hence the sample inefficiency becomes $\frac{1}{n}$.

This example also shows that having more than one sink agent will make the sample inefficiency worse—that is, larger. This is because we can replicate the valuation of i^* for every other sink and the inequality above will be tightly upper bounded at $2M$ for 2 sinks, $3M$ for 3 sinks, etc. Consequently, the sample inefficiency increases to $\frac{2}{n}$, $\frac{3}{n}$ etc.

We, therefore, need to prove the following lemma to complete the proof of the theorem.

LEMMA **2 (Lower Bound)** *In the class of neutral affine maximizers given by Equation (3) having a sink agent i^* (i.e., $w_{i^*} = 0$), the lowest sample inefficiency is achieved when $w_i = w$ for all $i \in N \setminus \{i^*\}$.*

Proof: Suppose not, that is, $\exists j, j' \in N \setminus \{i^*\}$ such that, WLOG, $w_j > w_{j'}$. Consider the following valuation profile:

$$\begin{array}{rclrcl} v_{i^*}(a) & = & -\frac{M}{2} + \gamma, & v_{i^*}(b) & = & \frac{M}{2} - \delta, \\ v_j(a) & = & \frac{M}{2} - \delta, & v_j(b) & = & \frac{M}{2} - \frac{w_{j'}}{w_j}M - \epsilon, \\ v_{j'}(a) & = & -\frac{M}{2} + \gamma, & v_{j'}(b) & = & \frac{M}{2} - \delta, \end{array} \quad \begin{array}{rcl} v_i(a) & = & v_i(b), \ \forall i \in N \setminus \{i^*, j, j'\}, \\ v_i(z) & = & -\frac{M}{2} + \frac{\gamma}{2}, \ \forall z \in A \setminus \{a, b\}, \forall i \in N. \end{array}$$

The constants $\delta, \gamma, \epsilon > 0$ are arbitrarily small. It is easy to verify that on this profile, by choosing the constants δ, γ, ϵ appropriately small, the affine maximizer will return a. But the efficient alternative is $a^*(v) = b$. Consider the term $\sum_{i \in N} v_i(a^*(v)) - \sum_{i \in N} v_i(f(v))$. To this term, agent i^* contributes inefficiency M, agent j' contribute M, and agent j contributes $\left(-\frac{w_{j'}}{w_j}M\right)$, taking the limiting values of $\delta, \gamma, \epsilon \to 0$. Therefore, the inefficiency term on this profile equals $M + \left(1 - \frac{w_{j'}}{w_j}\right) M > M$, the maximum inefficiency when the weights are equal except i^* (by the arguments just before Lemma 2). Hence, this mechanism cannot achieve the lowest sample inefficiency, which is a contradiction. ■

Lemma 2 and the arguments before it complete the proof of the theorem. ■

4 Jointly minimizing budget imbalance and inefficiency

In the previous section, we considered strategyproof, budget-balanced mechanisms that are minimally inefficient. We achieved a sample inefficiency lower bound of $\frac{1}{n}$. *Could one do better by, instead of requiring budget balance and minimizing inefficiency, relaxing budget balance by allowing money burning, and then minimizing the inefficiency from the allocation plus the inefficiency caused by money burning (or required subsidy from outside the mechanism)?*

In this section, we consider the joint problem of minimizing inefficiency and budget imbalance. We consider a convex combination of these two quantities since in the quasi-linear domain both of them contribute additively in the agents' utilities and social welfare. Therefore, the metric to minimize in this context is the *efficiency-budget spillover* defined as follows.

$$\rho_n(f, \mathbf{p}) := \lim_{M \to \infty} \frac{1}{nM} \sup_{v \in V} \left[\lambda \cdot T_1^n(f, v) + (1 - \lambda) \cdot T_2^n(\mathbf{p}, v)\right], \tag{14}$$

Where $T_1^n(f, v) = \left(\max_{a \in A} \sum_{i \in N} v_i(a) - \sum_{i \in N} v_i(f(v))\right)$ and $T_2^n(\mathbf{p}, v) = \left|\sum_{i \in N} p_i(v)\right|$. For $\lambda = 1$, that is, when budget imbalance is not a concern, one can use the VCG mechanism to get

$\rho_n(f, \mathbf{p}) = 0$. Similarly, for $\lambda = 0$, a sink mechanism will give $\rho_n(f, \mathbf{p}) = 0$. So, the interesting cases are when $\lambda \in (0, 1)$, and for this we have a solution that decays as $1/n$. In this section, we will assume that $\lambda, 0 < \lambda < 1$ is exogenous. Our goal is to find a strategyproof and neutral mechanism $\langle f, \mathbf{p} \rangle$ that minimizes the objective ρ_n. We have shown in Section 3 that $T_1^n(f, v)$ can at most be a constant when $T_2^n(\mathbf{p}, v)$ is zero for every v. Hence, for any improvement in the efficiency-budget spillover metric, that is, for $\rho_n(f, \mathbf{p}) = o(r_n(f))$, it is necessary that the term $\sup_{v \in V} [\lambda \cdot T_1^n(f, v) + (1 - \lambda) \cdot T_2^n(\mathbf{p}, v)]$ be $o(1)$. Since both $T_1^n(f, v)$ and $T_2^n(\mathbf{p}, v)$ are non-negative, it is necessary that the factor $T_2^n(\mathbf{p}, v) = o(1)$ for every $v \in V$. Our next result shows that it is impossible to have $T_2^n(\mathbf{p}, v) = o(1)$, $\forall v \in V \Leftrightarrow \lim_{n \to \infty} \sup_{v \in V} T_2^n(\mathbf{p}, v) = 0$. Hence, for deterministic mechanisms, the bound on inefficiency *with no budget imbalance* (presented in Section 3) is *asymptotically* optimal for this joint optimization problem as well!

THEOREM **3 (Unimprovability)** *For every deterministic, strategyproof, and neutral mechanism $\langle f, \mathbf{p} \rangle$ over V and for every $\lambda \in (0, 1)$, $\rho_n(f, \mathbf{p}) = \Omega\left(\frac{1}{n}\right)$. This bound is tight. For $\lambda = 0$, a sink mechanism, and for $\lambda = 1$, the VCG mechanism, achieves zero spillover.*

We defer the proof of this theorem to the Appendix. This result shows that as the number of agents grows, an inefficiency-minimizing budget-balanced mechanism is optimal in the class of deterministic mechanisms. It does not, however, claim optimality for every n. This leaves open the question of finding a solution to this joint problem that yields a lower spillover $\rho_n(f, \mathbf{p})$ for a given finite n than the degenerate problem of minimizing inefficiency subject to budget-balance.

5 Randomized, strategyproof, budget-balanced mechanisms

We saw that the best sample inefficiency achieved by a deterministic mechanism is $\frac{1}{n}$. In this section, we discuss how the inefficiency can be reduced by considering randomized mechanisms.

A naïve approach is to consider a *randomized sink* mechanism, where each agent is picked as a sink with probability $\frac{1}{n}$. 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 require budget balance. Hence, even though we can obtain an upper bound on the expected sample inefficiency ($r_n^M(f)$) by considering specific mechanisms like the naïve randomized sink mechanism described above, the problem of providing a lower bound (i.e., no randomized mechanism can achieve a lower $r_n^M(f)$ 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 analytically find the lower bound of the optimal, strategyproof, budget-balanced, randomized mechanism for two agents and two alternatives. We also use *automated mechanism design* Conitzer and Sandholm (2002); Sandholm (2003) to supplement the analysis. (The problems of finding a mechanism matching this lower bound and extending the lower bound to any number of agents and alternatives are left as future work.)

5.1 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 [4]. This makes these mechanisms different from the naïve sink mechanism. 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 generic steps of mechanisms in this class.

Algorithm 1 Generalized Sink Mechanisms, $\mathcal{G}$

Input: a valuation profile $v \in V$
Output: a probability distribution g over the agents N
for agent i in N picked with probability $g_i(v)$ **do**
  consider the valuations v_{-i} and agents $j \in N \setminus \{i\}$
  resulting alternative: $a^*(v_{-i}) = \text{argmax}_{a \in A} \sum_{j \in N \setminus \{i\}} v_j(a)$
  agent $j \in N \setminus \{i\}$ pays $p_j(v) = \max_{a \in A} \sum_{k \neq j,i} v_j(a) - \sum_{k \neq j,i} v_j(a^*(v_{-i}))$: Clarke tax without i
  agent i (the sink) receives $\sum_{j \in N \setminus \{i\}} p_j(v)$
end for

It is clear that not every mechanism in this class is strategyproof. The crucial thing is how the probabilities of choosing the sinks 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. 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
  end if
end for
if irrelevant agent is found **then**
  treat that agent as sink with probability 1
else
  pick an agent i with probability $\frac{1}{n}$ and treat as sink
end if

The operation of this intuitive mechanism lies in the fact that if an agent i's maximum valuation sweep ($-M/2$ to $M/2$) cannot change the alternative, this *irrelevant* agent can be selected as a sink, which yields the efficient alternative. However, when there is no such irrelevant agent, the decision of choosing every agent equi-probably leads to a chance of manipulation. An agent whose true valuation report does not lead her to become a sink can misreport

[4]One can think of a more general class of sink mechanisms where multiple agents are treated as sink agents simultaneously. But it is easy to see (by a similar argument to that in the context of deterministic mechanisms, just before Lemma 2) that multiple sinks will only increase the inefficiency.

a valuation so that there is no irrelevant agent, thereby increasing her own probability of being selected as a sink. In particular, consider two valuation profiles with three agents (numbered $1, 2, 3$) and three alternatives (a, b, c), and $M = 1$. The agents' valuations in the first profile are $v_1 = (0.5, 0, -0.5), v_2 = (-0.5, 0, 0.5), v_3 = (0, -0.5, 0.5)$, and in the second profile they are $v_1' = (0.5, 0, -0.5), v_2' = (-0.5, 0, 0.5), v_3' = (-0.5, 0, 0.5)$. The mechanism returns agent 1 as the irrelevant agent in the first profile and therefore picks alternative c with probability 1. There is no irrelevant agent in the second profile and hence each agent is picked as a sink with uniform probability, leading to the probability vector $(2/3, 0, 1/3)$ for the alternatives a, b, c. But agent 3 strictly gains by moving from the first profile to the second. [5]

However, a small modification of the previous mechanism leads to a strategyproof generalized sink mechanism. This shows 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 uniform 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
for agent i in N picked with probability p_i do
  for agent j in N \ {i} do
    if agent j is an irrelevant agent within N \ {i} then
      treat agent j as sink
      irrelevant agent is found
    end if
  end for
  if irrelevant agent is not found within N \ {i} then
    treat agent i as sink
  end if
end for
```

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

Proof: Assume $m = n + 1$. The proof generalizes to any $m > n$. Consider the valuation profile $v_i = (v_i(a_1), \ldots, v_i(a_n), v_i(a_{n+1}))$ where $v_i(a_i) = -M/2 + \epsilon/2, v_i(a_{n+1}) = M/2 - \epsilon$ and $v_i(a_j) = M/2 - \epsilon/2, \forall j \neq i, n+1$, and $\epsilon > 0$ is arbitrarily small, $i \in N$. This profile is possible to construct since $m > n$. Clearly, the efficient alternative is a_{n+1}, but if any agent i is picked as a sink, the alternative changes to a_i, which has inefficiency of $M - \epsilon$. Therefore the expected sample inefficiency for any generalized sink mechanism is $\frac{1}{nM}(M - \epsilon)$. Taking $\epsilon \to 0$ proves the theorem. ■

[5]One can also verify that the weak monotonicity condition, which is a necessary condition for strategyproofness, is violated for agent 3 between these two profiles.

This profile hinges on $m > n$, and is not possible to reproduce if $m \leq n$. We can hope for a smaller inefficiency if the number of alternatives is small. We state this intuition formally as follows.

THEOREM **5 (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., $r^M_{n,m_1}(f) \geq r^M_{n,m_2}(f), \forall m_1 > m_2$.*[6]

Proof: Suppose, for m_2 alternatives the valuation profile v^* yields the worst inefficiency $r^M_{n,m_2}(f)$. Clearly, we can append the other alternatives when we increase the number of alternatives to m_1 with values arbitrarily close to $-M/2$ so that they never change the optimal alternative. Hence the inefficiency cannot decrease. ■

Theorems 4 and 5 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 randomization scheme reduces the inefficiency by a factor of two.

THEOREM **6 (Naïve Randomized Sink)** *For $m = 2$, the expected sample inefficiency of the naïve randomized sink mechanism (NRS, every agent is selected as a sink with probability $\frac{1}{n}$) is $r^M_n(f_{NRS}) = \frac{1}{n^2}\left\lceil \frac{n}{2} \right\rceil \sim \frac{1}{2n}$.*

Proof: Consider an arbitrary agent i. If agent i is chosen as a sink, the maximum *absolute* inefficiency that it can produce is M (by the same argument as in Equation (13), and we refer to the unnormalized difference term $\sum_{i \in N} v_i(a^*(v)) - \sum_{i \in N} v_i(f(v))$ as the absolute inefficiency). Say the efficient alternative is a. This inefficiency is achieved when the sum of valuations of the agents other than i at the other alternative b is just higher than that of those agents at a, i.e., $\sum_{j \neq i} v_j(b) = \sum_{j \neq i} v_j(a) + \epsilon$, where $\epsilon > 0$ and small, and also the difference in valuations of agent i at these two alternatives is maximum, i.e., $v_i(a) - v_i(b) = M - \delta$, where $\delta > 0$ and small. This implies that, without i, the population is almost equally divided among the alternatives a and b with a marginal bias to b and agent i is 'maximally' in favor of a. The difference $v_i(a) - v_i(b) = M - \delta$ is achieved only when the values are close to $v_i(a) = \frac{M}{2} - \frac{\delta}{2}$ and $v_i(b) = -\frac{M}{2} + \frac{\delta}{2}$, since all valuations must lie within $\left(\frac{M}{2}, -\frac{M}{2}\right)$. Agent i, therefore, is *pivotal* for making the decision in favor of a.

Now, we claim that there cannot be more than $\left\lceil \frac{n}{2} \right\rceil$ such pivotal agents i. Suppose for contradiction that there are $> \left\lceil \frac{n}{2} \right\rceil$ such pivotal agents. We present the argument for $\left\lceil \frac{n}{2} \right\rceil + 1$ for brevity, but it generalizes to the case with more number of pivotal agents. For each of these agents k, $v_k(a)$ is arbitrarily close to $\frac{M}{2}$ and $v_k(b)$ is arbitrarily close to $-\frac{M}{2}$. Therefore, if any of them, say agent k^*, is chosen as a sink, there are $\left\lceil \frac{n}{2} \right\rceil$ other similar agents who still make the decision of the mechanism in favor of a (since the sum valuation for a will be larger than b by an unsurmountable value $\approx \left\lceil \frac{n}{2} \right\rceil M$) irrespective of the valuation profiles of the other agents. This implies that k^* is not a pivotal agent, which is a contradiction.

The mechanism chooses each sink agent with probability $\frac{1}{n}$. Therefore, the expected inefficiency can at most be $\frac{1}{n} \cdot \left\lceil \frac{n}{2} \right\rceil \cdot M$. Divide this by nM to get the expected sample inefficiency. It is easy to

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

see that this bound is tight. The valuation profile that achieves this bound has $\lceil \frac{n}{2} \rceil$ agents having valuations $v(a) \approx M/2, v(b) \approx -M/2$ and the rest of the agents have the reverse valuations. ■

Even though the modified irrelevant sink mechanism (Algorithm 3) is sophisticated in its use of the valuation profile, it is easy to check that even that mechanism yields the same inefficiency on the profile illustrated in the proof above. Thus we have the following theorem.

THEOREM **7 (Modified Irrelevant Sink)** *For $m = 2$, the expected sample inefficiency of the modified irrelevant sink mechanism (MIS, Algorithm 3), $r_n^M(f_{MIS}) \geq \frac{1}{n^2} \lceil \frac{n}{2} \rceil \sim \frac{1}{2n}$.*

The above result does not say much about the lowest achievable expected sample inefficiency (even in this special class of generalized sink mechanisms). Therefore to understand the limit of lowest achievable inefficiency for randomized mechanisms in general, we focus on the case of two agents and two alternatives. The following theorem gives a lower bound on the inefficiency for the class of generalized sink mechanisms in that setting. Since we now fix the number of agents in the analysis, minimizing the expected sample inefficiency is equivalent to minimizing the expected *absolute* inefficiency given by $nr_n^M(f)$ which is $\frac{1}{M} \sup_{v \in V} \left\{ \mathbb{E}_{f(v)} \left[\max_{a \in A} \sum_{i \in N} v_i(a) - \sum_{i \in N} v_i(f(v)) \right] \right\}$. We will also WLOG assume $M = 1$. From now on, we let 'inefficiency' mean the expected absolute inefficiency.

THEOREM **8 (Lower Bound of Generalized Sink)** *For $n = m = 2$, the expected absolute inefficiency of every strategyproof generalized sink mechanism is lower bounded by $\frac{1}{2}$.*

Proof: For $n = m = 2$, either the two agents prefer the same candidate or opposing candidates. (Ties are broken in favor of a_1, say.) If they agree, any probability of the generalized sink mechanism yields zero absolute inefficiency, since the efficient alternative will be chosen irrespective of which agent is the sink. Note that the payments are always zero for two-agent generalized sink mechanisms. When the agents oppose, we claim that in every opposing profile, a strategyproof generalized sink mechanism must have the same probability of picking the sinks. Suppose not, that is, for a specific generalized sink mechanism $g : V \to \Delta N$, the probabilities of picking the sinks are different in profiles $v = (v_1, v_2)$ and $v' = (v_1', v_2')$, i.e., $g(v) \neq g(v')$. Consider the transition: $v = (v_1, v_2) \to (v_1', v_2) \to (v_1', v_2') = v'$. The sink-picking probabilities g must have changed in at least one of these two transitions, that is, either $g(v_1, v_2) \neq g(v_1', v_2)$ or $g(v_1', v_2) \neq g(v_1', v_2')$. But this is a contradiction to strategyproofness since at least one agent will misreport in that profile pair. She will prefer to increase the probability of the other agent becoming sink so that her favorite candidate has higher probability of being selected, which increases her utility since payment is zero. For example, suppose in the first transition, the probability of agent 1 being sink is higher in the profile (v_1, v_2), and consequently, probability of agent 2 being sink is lower. Then agent 1 will misreport her valuation to v_1'. Now, among all fixed probability distributions, $(0.5, 0.5)$ gives the minimum absolute inefficiency which is $\frac{1}{2}$. ■

5.2 Unrestricted randomized mechanisms

We now move on to study optimal randomized mechanisms without restring attention necessarily to generalized sink mechanisms. Finding a mechanism that achieves the minimum absolute inefficiency can be posed as the following optimization problem.

$$
\begin{aligned}
\min_{f,\mathbf{p}} \quad & \sup_{v\in V}\left[\max_{a\in A}\sum_{i\in N} v_i(a) - \sum_{i\in N} v_i(f(v))\right] \\
\text{s.t.} \quad & v_i(f(v_i, v_{-i})) - p_i(v_i, v_{-i}) \geq v_i(f(v_i', v_{-i})) - p_i(v_i', v_{-i}),\ \forall v_i, v_i', v_{-i}, \forall i \in N \\
& \sum_{a\in A} f_a(v) = 1,\ \forall v \in V, \\
& \sum_{i\in N} p_i(v) = 0,\ \forall v \in V, \\
& f_a(v) \geq 0,\ \forall v \in V, a \in A.
\end{aligned}
\tag{15}
$$

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. With two agents 1 and 2, and two alternatives a and b, we can write the optimization problem 15 as a linear program (LP) as follows.

$$
\begin{aligned}
& \min_{f,\mathbf{p}} \quad \ell \\
& \text{s.t.} \\
& [v_1(a)\cdot f_a(v_1,v_2) + v_1(b)\cdot f_b(v_1,v_2) - p_1(v_1,v_2)] - [v_1(a)\cdot f_a(v_1',v_2) \\
& \qquad + v_1(b)\cdot f_b(v_1',v_2) - p_1(v_1',v_2)] \geq 0,\ \forall v_1, v_1', v_2, && \text{Agent 1, SP} \\
& [v_2(a)\cdot f_a(v_1,v_2) + v_2(b)\cdot f_b(v_1,v_2) - p_2(v_1,v_2)] - [v_2(a)\cdot f_a(v_1,v_2') \\
& \qquad + v_2(b)\cdot f_b(v_1,v_2') - p_2(v_1,v_2')] \geq 0,\ \forall v_1, v_2, v_2', && \text{Agent 2, SP} \\
& f_a(v) + f_b(v) = 1,\ \forall v \in V, && \text{SCF} \\
& p_1(v) + p_2(v) = 0,\ \forall v \in V, && \text{Budget Balance} \\
& \ell + (v_1(a)+v_2(a))\cdot f_a(v) + (v_1(b)+v_2(b))\cdot f_b(v) \\
& \qquad \geq \max_{x\in\{a,b\}} (v_1(x)+v_2(x)),\ \forall v \in V, && \text{Max Inefficiency} \\
& f_a(v), f_b(v) \geq 0,\ \forall v \in V, a \in A
\end{aligned}
\tag{16}
$$

The objective function is the maximum inefficiency that we want to minimize, and the constraints are the requirements that our mechanism needs to satisfy. The formulation considers all possible valuations, which makes the number of constraints uncountable. Therefore, to solve this optimization problem with finite constrained optimization techniques, we need to discretize the valuation levels. We assume that each agent's valuations are uniformly discretized with k levels in $(-M/2, M/2)$ (with the lowest and highest values arbitrarily close to the boundary), which makes the set of valuation profiles V a discrete finite set. However, note that the optimal value of such a discretized relaxation of the constraints provides a lower bound on the optimal value of the original

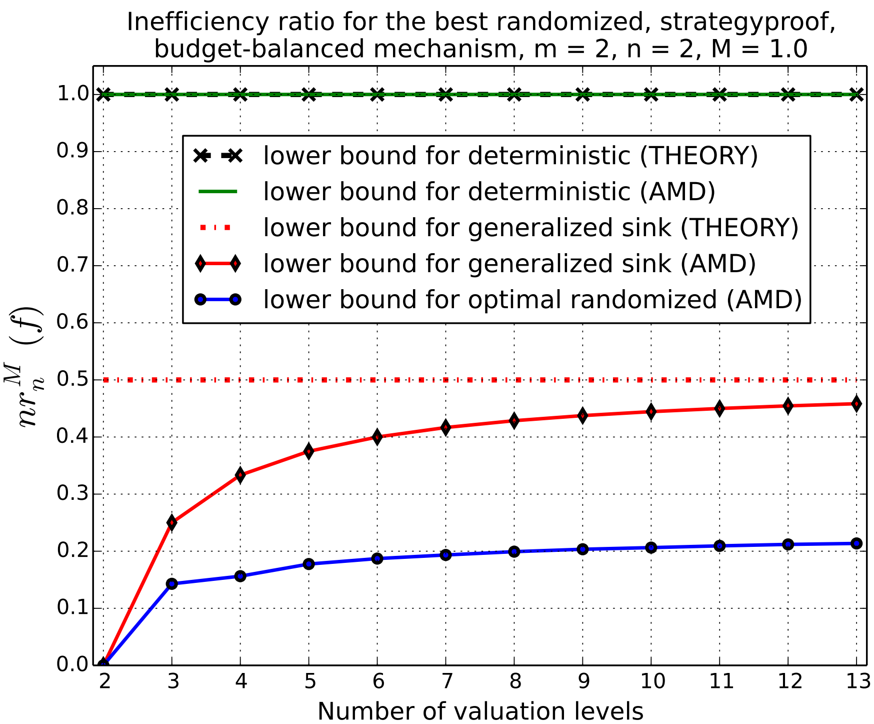

Figure 1: Lower bound for the discrete relaxation of the inefficiency minimization LP.

problem. This is because the discretized relaxation of the valuations only increases the feasible set as the number of equalities and inequalities is reduced, that is, more f's and p's satisfy the constraints. When the valuations are richer, there will be fewer feasible solutions to this LP and therefore the optimal value will increase. Hence, by solving the above LP, we are providing a lower bound on the actual inefficiency. We now prove a lower bound when the number of discretized levels is three.

Theorem **9** **(Lower Bound of Inefficiency for Randomized Mechanisms)** *For $n = m = 2$, and for $k = 3$ discrete levels of valuations, the absolute inefficiency is lower bounded by $\frac{1}{7} = 0.142857$.*

The proof is deferred to the Appendix.

In the appendix we also provide structural insights on the problem, such as showing that anonymity and neutrality can be imposed on the mechanism without loss. In fact, our proof of this fact applies to any number of agents and alternatives.

The proof technique can be extended to a larger number of discrete levels to obtain a tighter lower bound on the actual inefficiency. We conducted a form of automated mechanism design (Conitzer and Sandholm, 2002; Sandholm, 2003) by solving this LP using Gurobi (2015) 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 Figure 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.

In summary, for two agents and two alternatives, we found that the best expected absolute inefficiency achievable by a deterministic mechanism is 1.0, while for the generalized randomized sink mechanisms the absolute inefficiency improves to 0.5. We showed a lower bound over 0.2 on strategyproof, budget-balanced randomized mechanisms. So, randomization reduces expected inefficiency. The structure of the optimal randomized mechanism is an open problem.

6 Conclusions and future research

We provided several new results on the classic question of the interplay between efficiency and budget balance in the quasi-linear setting. We studied strategyproof mechanisms both in the deterministic and randomized framework.

We proved characterization results, and a tight lower bound for inefficiency, for deterministic budget-balanced mechanisms. We also proved that minimizing inefficiency and budget imbalance together does not provide any asymptotic advantage in the deterministic paradigm over requiring budget balance and minimizing inefficiency!

We provided results that show that randomization helps—particularly when the number of alternatives is small compared to the number of agents. In the case of two agents and two alternatives, generalized sink mechanisms reduce inefficiency by a factor $\frac{1}{2}$, and optimal randomized mechanisms offer further reductions as we showed analytically and using automated mechanism design.

Future research includes studying the structure of the optimal randomized mechanism that achieves this improved efficiency. Future work also includes investigating the rate of improvement of the optimal bound for a general number of agents.

References

Arrow, K. (1979): *The property rights doctrine and demand revelation under incomplete information*, Economics and Human Welfare. New York Academic Press.

Ashlagi, I., F. Fischer, I. A. Kash, and A. D. Procaccia (2013): "Mix and match: A strategyproof mechanism for multi-hospital kidney exchange," *Games and Economic Behavior*, 1–13.

Bailey, M. J. (1997): "The demand revealing process: to distribute the surplus," *Public Choice*, 91, 107–126.

Cavallo, R. (2006): "Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments," in *Proceedings of the fifth international joint conference on autonomous agents and multiagent systems (AAMAS)*, 882–889.

Clarke, E. (1971): "Multipart Pricing of Public Goods," *Public Choice*, 19–33.

Conitzer, V. and T. Sandholm (2002): "Complexity of mechanism design," in *Proceedings of the Eighteenth conference on Uncertainty in Artificial Intelligence*, Morgan Kaufmann Publishers Inc., 103–110.

d'Aspremont, C. and L. A. Gérard-Varet (1979): "Incentives and Incomplete Information," *Journal of Public Economics*, 11(1), 25–45.

Dobzinski, S., N. Nisan, and M. Schapira (2006): "Truthful randomized mechanisms for combinatorial auctions," in *Proceedings of the thirty-eighth annual ACM Symposium on Theory of Computing*, ACM, 644–652.

Faltings, B. (2004): "A budget-balanced, incentive-compatible scheme for social choice," in *Agent-Mediated Electronic Commerce VI. Theories for and Engineering of Distributed Mechanisms and Systems*, Springer, 30–43.

Feldman, M. and Y. Wilf (2011): "Randomized strategyproof mechanisms for Facility Location and the mini-sum-of-squares objective," *arXiv preprint arXiv:1108.1762*.

Gary-Bobo, R. J. and T. Jaaidane (2000): "Polling mechanisms and the demand revelation problem," *Journal of Public Economics*, 76, 203–238.

Green, J. R. and J.-J. Laffont (1979): *Incentives in public decision making*, North-Holland.

Groves, T. (1973): "Incentives in Teams," *Econometrica*, 41, 617–31.

Gujar, S. P. and Y. Narahari (2011): "Redistribution mechanisms for assignment of heterogeneous objects," *Journal of Artificial Intelligence Research*, 131–154.

Guo, M. and V. Conitzer (2008): "Optimal-in-expectation redistribution mechanisms," in *Proceedings of the 7th international joint conference on autonomous agents and multiagent systems (AAMAS)*, 1047–1054.

——— (2009): "Worst-case optimal redistribution of VCG payments in multi-unit auctions," *Games and Economic Behavior*, 67, 69–98.

——— (2014): "Better redistribution with inefficient allocation in multi-unit auctions," *Artificial Intelligence*, 216, 287–308.

Gurobi (2015): "Gurobi Optimizer Reference Manual," .

Krishna, V. and E. Maenner (2001): "Convex potentials with an application to mechanism design," *Econometrica*, 69, 1113–1119.

Massó, J., A. Nicolò, A. Sen, T. Sharma, and L. Ülkü (2015): "On cost sharing in the provision of a binary and excludable public good," *Journal of Economic Theory*, 155, 30–49.

Mishra, D. and A. Sen (2012): "Roberts' theorem with neutrality: A social welfare ordering approach," *Games and Economic Behavior*, 75, 283–298.

Mishra, D. and T. Sharma (2016): "On the Optimality of the Green-Laffont Mechanism," Preliminary Draft.

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

Moulin, H. and S. Shenker (2001): "Strategyproof sharing of submodular costs: budget balance versus efficiency," *Economic Theory*, 18, 511–533.

Procaccia, A. D. and M. Tennenholtz (2009): "Approximate mechanism design without money," in *Proceedings of the 10th ACM conference on Electronic Commerce (EC)*, ACM, 177–186.

Rockafellar, R. T. (1997): *Convex analysis*, vol. 28, Princeton university press.

Rothkopf, M. H. (2007): "Thirteen reasons why the Vickrey-Clarke-Groves process is not practical," *Operations Research*, 55, 191–197.

Sandholm, T. (2003): "Automated mechanism design: A new application area for search algorithms," in *Principles and Practice of Constraint Programming–CP 2003*, Springer, 19–36.

Thang, N. K. (2010): "On randomized strategy-proof mechanisms without payment for facility location games," in *Proceedings Workshop on Internet and Network Economics (WINE). Stanford, USA*, 13–16.

Vickrey, W. (1961): "Counterspeculation, auctions, and competitive sealed tenders," *The Journal of Finance*, 16, 8–37.

APPENDIX

Proof of Theorem 3

Proof: Suppose, there exists a deterministic, strategyproof, and neutral mechanism $\langle f, \mathbf{p} \rangle$ that also satisfies $\lim_{n\to\infty} \sup_{v \in V} T_2^n(\mathbf{p}, v) = 0$. It implies that, at the limit, the mechanism has no budget imbalance, i.e., $\lim_{n\to\infty} \sup_{v\in V} |\sum_{i=1}^n p_i(v)| = 0$. From the arguments in Theorem 1 (Equations (3) and (4)), we know that f is a neutral affine maximizer and payments are of the form $p_i(v_i, v_{-i}) = h_i(v_{-i}) + \frac{1}{w_i}\left(\sum_{j\neq i} w_j v_j(f(v))\right), \forall w_i > 0$. We already have the sink mechanism where at least one $w_i = 0$ and the above sum can be made smallest (exactly zero) for every profile $v \in V$. However, that yields a constant upper bound for the term $\lambda \cdot T_1^n(f, v) + (1-\lambda)\cdot T_2^n(\mathbf{p}, v)$. Hence, we need to consider the case $w_i > 0, \forall i$, which implies that

$$\lim_{n\to\infty} \sup_{v\in V} \left| \sum_{i=1}^n \left(h_i(v_{-i}) + \frac{1}{w_i} \left(\sum_{j\neq i} w_j v_j(f(v)) \right) \right) \right| = 0.$$

This implies that for every $\epsilon > 0$, there exists $N_\epsilon \in \mathbb{Z}_{\geq 0}$ such that for all $n \geq N_\epsilon$,

$$\left| \sum_{i=1}^n \left(h_i(v_{-i}) + \frac{1}{w_i} \left(\sum_{j\neq i} w_j v_j(f(v)) \right) \right) \right| < \epsilon, \ \forall v \in V. \tag{17}$$

We show that this identity leads to a contradiction for an appropriately chosen v. Note that this immediately proves the theorem, because if there does not exist any mechanism $\langle f, \mathbf{p} \rangle$ that satisfies the properties mentioned in the theorem statement and makes $\lim_{n\to\infty} \sup_{v\in V} T_2^n(\mathbf{p}, v) = 0$, then the best possible lower bound is a constant, i.e., $\sup_{v\in V} T_2^n(\mathbf{p}, v) = \Omega(1)$. Therefore, the best lower bound for the spillover factor $\rho_n(f, \mathbf{p})$ is $\Omega\left(\frac{1}{n}\right)$ and this is achievable by the sink mechanism.

We prove this for a set of alternatives $A = \{0, 1\}$ for similar reasons mentioned in Lemma 1. As an illustration of the general proof, let us consider the same argument when $N_\epsilon = 2$. Let the valuations are $(v_1 + \delta, v_2)$ for alternative 1 and zero otherwise ($\delta > 0$). Also assume that the numbers are such that

$$w_1(v_1 + \delta) + w_2 v_2 > 0, \text{ and } w_1 v_1 + w_2 v_2 < 0.$$

That means, the affine maximizer results in 1 at profile $(v_1 + \delta, v_2)$ and 0 at (v_1, v_2). The above inequalities can be written concisely as

$$-w_1\delta < w_1 v_1 + w_2 v_2 < 0. \tag{18}$$

Now by the convergence relation of Equation (17), we have

$$\begin{aligned}
&\left| h_1(v_2) + h_2(v_1 + \delta) + \frac{w_2}{w_1} v_2 + \frac{w_1}{w_2}(v_1 + \delta) \right| < \epsilon \\
&|h_1(v_2) + h_2(v_1)| < \epsilon
\end{aligned}$$

These inequalities imply[7]

$$\begin{aligned}
&\left| h_2(v_1 + \delta) + \frac{w_2}{w_1} v_2 + \frac{w_1}{w_2}(v_1 + \delta) - h_2(v_1) \right| < 2\epsilon \\
\Rightarrow \quad &\left| \frac{w_2}{w_1} v_2 - \left(h_2(v_1) - h_2(v_1 + \delta) - \frac{w_1}{w_2}(v_1 + \delta) \right) \right| < 2\epsilon
\end{aligned} \tag{19}$$

But this inequality is violated by choosing a large enough δ and large negative v_2 in Equation (18). This is possible to pick since the valuations are picked from $(-M/2, M/2)$ and M is large by definition of ρ_n (Equation (14)). Also note that, the term within parentheses in Equation (19) is independent of v_2, hence changes in the v_2 will not affect them. Our only assumed relation is Equation (18), and a suitable choice satisfying it violates Equation (19).

The general proof of this theorem extends this idea for any $N_\epsilon = n \geq 2$. Let the agents are numbered in the decreasing order of their weights WLOG, i.e., $w_i \geq w_{i+1}, i = 1, 2, \ldots, n-1$. We consider the valuation profile $(v_1 + \delta, v_2 + \delta, \ldots, v_{n-1} + \delta, v_n), \delta > 0$ such that

$$-\delta \sum_{i=1}^{n-1} w_i < \sum_{i=1}^{n} w_i v_i < -\delta \sum_{i=1}^{n-2} w_i \tag{20}$$

The above inequalities imply that the affine maximizer alternative for the profile mentioned above is 1. However, if any agent i's, $i = 1, 2, \ldots, n-1$, valuation changes to v_i from $v_i + \delta$, the alternative changes to 0. We use a generic notation v^k to denote this profile, where k denotes the agent(s) whose valuation(s) is(are) v_k while all other agents $j \neq k$ have valuations $v_j + \delta$. Hence, v^n is the profile mentioned before: $(v_1 + \delta, v_2 + \delta, \ldots, v_{n-1} + \delta, v_n)$ and $v^{n-1,n}$ is the profile: $(v_1 + \delta, v_2 + \delta, \ldots, v_{n-1}, v_n)$, for example. Note that, the following term in Equation (17) can be reorganized as

$$\sum_{i=1}^{n} \frac{1}{w_i} \left(\sum_{j \neq i} w_j v_j(f(v)) \right) = \sum_{i=1}^{n} \left(\sum_{j \neq i} \frac{1}{w_j} \right) w_i v_i(f(v)).$$

[7] If $|x + z| < \epsilon$ and $|y + z| < \epsilon$, then $|x - y| = |x + z - (y + z)| \leq |x + z| + |y + z| < 2\epsilon$.

Since, $f(v^n) = 1$, from Equation (17) we have

$$\left| \left(\sum_{i=1}^{n-1} h_i(v_{-i}^n) + h_n(v_{-n}^n) \right) + \left(\sum_{i=1}^{n} \left(\sum_{j \neq i} \frac{1}{w_j} \right) w_i v_i + \sum_{i=1}^{n-1} \left(\sum_{j \neq i} \frac{1}{w_j} \right) w_i \delta \right) \right| < \epsilon. \tag{21}$$

The idea of the proof is to make a series of substitutions in the first parentheses of the expression above, leaving the terms in the other parentheses unchanged. Note that, the expression in the second parentheses depends on v_n, while the expression $h_n(v_{-n}^n)$ does not. The substitutions sequentially eliminate the dependency on v_n from all the terms in the first parentheses, similar to what we did in the two agent case before. This will also increase the RHS of the inequality in Equation (21), but it will be a finite constant factor of ϵ. This leads to a contradiction, since v_n can be chosen arbitrarily large negative by choosing a large positive δ, and still continues to satisfy Equation (20) but violates Equation (21).

The substitutions will involve the term $\sum_{i=1}^{n-1} h_i(v_{-i}^n)$ in the first parentheses of Equation (21). Consider the profiles $v^{j,n}, j = 1, \ldots, n-1$. In each of these profiles, $f(v^{j,n}) = 0$ (due to the choice of v^n in Equation (20)). Hence,

$$\left| \sum_{i=1}^{n-1} h_i(v_{-i}^{j,n}) + h_n(v_{-n}^{j,n}) \right| < \epsilon, \ \forall j \in \{1, \ldots, n-1\}. \tag{22}$$

Note that $v_{-i}^{i,n} = v_{-i}^n$. Hence, we can substitute terms from Equation (22) to the terms in the first parentheses of Equation (21) to get

$$\begin{aligned} \Bigg| \Bigg(-\sum_{i=1}^{n-1} \sum_{j \neq \{i,n\}} h_j(v_{-j}^{i,n}) - \sum_{j \neq n} h_n(v_{-n}^{j,n}) + h_n(v_{-n}^n) \Bigg) + \\ \left(\sum_{i=1}^{n} \left(\sum_{j \neq i} \frac{1}{w_j} \right) w_i v_i + \sum_{i=1}^{n-1} \left(\sum_{j \neq i} \frac{1}{w_j} \right) w_i \delta \right) \Bigg| < n\epsilon. \end{aligned} \tag{23}$$

We now replace the terms $h_j(v_{-j}^{i,n})$ in the first summation of the first parentheses above. All other terms in that parentheses are h_n functions and, therefore, are independent of v_n. For every $i \neq n$, consider the valuation profiles $v^{j,i,n}, j \neq i, n$. By Equation (20), $f(v^{j,i,n}) = 0$, hence, we get an inequality similar to Equation (22):

$$\left| \sum_{k=1}^{n-1} h_k(v_{-k}^{j,i,n}) + h_n(v_{-n}^{j,i,n}) \right| < \epsilon, \ \forall j \neq i, n. \tag{24}$$

Also, note that, $v_{-j}^{j,i,n} = v_{-j}^{i,n}$. So, we follow the same procedure to replace the terms $h_j(v_{-j}^{i,n})$ in Equation (23) to yield a similar inequality where the RHS is replaced by a larger term. Since the number of agents is finite, this process will stop after a finite number of iterations, reducing the terms in the first parentheses only consisting of h_n functions, which are independent of v_n, and the RHS of the inequality being a finite factor $K(n)\epsilon$ (say). This construction shows that the choice of a suitably large δ and negative v_n, which keeps Equation (20) unaffected, can violate the inequality obtained through the iterative procedure described above. Hence the claim is established. ■

Proof of Theorem 9

Proof: For $k = 3$, each agent has $3^2 = 9$ valuations, since the number of alternatives is 2, and therefore, the number of valuation profiles is 81. The optimization variables are

$$\mathbf{x} := (f_a(v^0), f_b(v^0), \dots, f_a(v^{80}), f_b(v^{80}), p_1(v^0), p_2(v^0), \dots, p_1(v^{80}), p_2(v^{80}), \ell)^\top.$$

Here the 81 valuation profiles are indexed from 0 to 80 and are denoted by the superscripts. Hence there are $81 \times 4 + 1 = 325$ variables to the discretized relaxation of the primal problem of Equation (16). However, we can significantly reduce the number of variables using the symmetry of the LP. The symmetry that we consider are *anonymity*, i.e., the SCF alternative is invariant to the permutation of the agents, and the payments are permutated according to the permutation of the agents (Definition 5), and *neutrality*, i.e., the relabeling of the alternatives changes the alternative according to the same relabeling (Definition 4).

LEMMA **3** *For every strategyproof, budget-balanced, randomized mechanism that achieves the minimum absolute inefficiency, there exists an* anonymous, neutral*, strategyproof, budget-balanced, randomized mechanism that achieves the same absolute inefficiency.*

Proof: We prove this for two agents and two alternatives. The same argument generalizes to any number of agents and alternatives. Consider an optimal solution of the optimization problem of Equation (16). This yields a solution $\mathbf{x}^*$ (say). Suppose, we relabel the agents 1 and 2 by swapping their identities, which changes the valuation profiles accordingly. For example, now the payment $p_1(v_1, v_2)$ is swapped with $p_2(v_2, v_1)$. We keep the SCF alternatives identical, i.e., $f_a(v_1, v_2) = f_a(v_2, v_1)$. Now consider the resulting vector of variables $\mathbf{x}^*_{\text{AGENT-SWAP}}$. Note that, this permutation of the variables reorders the set of constraints in Equation (16). The SP constraints of agent 1 now becomes the SP constraints of agent 2 and vice-versa. SCF constraints remain identical, budget balance constraints are reordered but same, and the max-inefficiency constraints are also reordered. Hence $\mathbf{x}^*_{\text{AGENT-SWAP}}$ is a feasible solution of the LP (Equation (16)) and since $\mathbf{x}^*_{\text{AGENT-SWAP}}$ and $\mathbf{x}^*$ has the same value for ℓ, $\mathbf{x}^*_{\text{AGENT-SWAP}}$ is an optimal solution of the LP (Equation (16)).

Similarly, we swap the alternatives a and b and the valuations accordingly to obtain a different reordered vector $\mathbf{x}^*_{\text{ALT-SWAP}}$. This relabeling of the alternatives again reorders all the constraints in a different way than the earlier case, with ℓ remaining same in both these cases. In a similar way as before, we argue that $\mathbf{x}^*_{\text{ALT-SWAP}}$ is an optimal solution of the LP (Equation (16)).

Now, we swap both the alternatives and agents to obtain $\mathbf{x}^*_{\text{AGENT-ALT-SWAP}}$ which reorders the constraints in a two-fold manner, but the last variable of this vector remains ℓ as before, and therefore it is also an optimal solution of the LP (Equation (16)).

Now, we have 4 optimal solutions given the original optimal solution $\mathbf{x}^*$, which are complementary to each other in terms of agents and alternatives, but all of them are strategyproof, budget-balanced, randomized mechanisms (since they are feasible solutions of Equation (16)). Consider the average of all these solutions:

$$\mathbf{x}^{A,N} = \frac{1}{4}(\mathbf{x}^* + \mathbf{x}^*_{\text{AGENT-SWAP}} + \mathbf{x}^*_{\text{ALT-SWAP}} + \mathbf{x}^*_{\text{AGENT-ALT-SWAP}}).$$

By construction, $\mathbf{x}^{A,N}$ is anonymous and neutral, but this is also another optimal solution of the LP (Equation (16)) (since the set of constraints is convex). Hence, we have proved the lemma for two agents and two alternatives. For n agents and m alternatives, we consider all $n!$ and $m!$ possible permutations of the agents and alternatives respectively and take the mean of them to obtain our resulting optimal solution that is both anonymous and neutral. ■

Hence, it is WLOG to consider neutral and anonymous mechanisms to solve the optimization problem of Equation (16). This reduces the number of variables in the primal problem, since for valuations that are either agent permutated or alternative permutated or both permutated version of a valuation profile we have already considered, we can replace their constraints with the already considered variables. We write the coefficient matrix of the constraint set of the LP in Equation (16) denoted by A as follows.

$$\begin{array}{c}
\begin{array}{ccccccccccccc}
f_a(v^0) & f_b(v^0) & \dots & \dots & f_a(v^{80}) & f_b(v^{80}) & p_1(v^0) & p_2(v^0) & \dots & \dots & p_1(v^{80}) & p_2(v^{80}) & \ell
\end{array} \\
\left(\begin{array}{ccccccccccccc}
v_1^0(a) & v_1^0(b) & -v_1^8(a) & -v_1^8(b) & 0 & 0 & -1 & 0 & 1 & 0 & 0 & 0 & 0 \\
v_2^0(a) & v_2^0(b) & -v_2^1(a) & -v_2^1(b) & 0 & 0 & 0 & -1 & 0 & 1 & 0 & 0 & 0 \\
0 & 0 & \dots & \dots & 0 & 0 & 0 & 0 & \dots & \dots & 0 & 0 & 0 \\
0 & 0 & -v_2^{79}(a) & -v_2^{79}(b) & v_2^{80}(a) & v_2^{80}(b) & 0 & 0 & 0 & 1 & 0 & -1 & 0 \\
0 & 0 & \dots & \dots & 0 & 0 & 0 & 0 & \dots & \dots & 0 & 0 & 0 \\
\hdashline
1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & \dots & \dots & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
\hdashline
0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \dots & \dots & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 \\
\hdashline
w^0(a) & w^0(b) & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\
0 & 0 & \dots & \dots & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\
0 & 0 & 0 & 0 & w^{80}(a) & w^{80}(b) & 0 & 0 & 0 & 0 & 0 & 0 & 1
\end{array}\right)
\end{array} \tag{25}$$

Where $w^p(x) = v_1^p(x) + v_2^p(x), x \in \{a, b\}$, and p denotes the profile index. The header of the matrix shows the primal variables. The sections showed in dotted lines corresponds to the strategyproofness, valid SCF, budget balance, and maximum inefficiency constraints respectively. The RHS of the constrained inequalities of the LP is a vector $\mathbf{b}$ that looks as follows:

$$\mathbf{b} := (0, \dots, 0, \mathbf{1}_{|V|}, \mathbf{0}_{|V|}, \max_x w^0(x), \dots, \max_x w^{80}(x))^\top.$$

Denoting the cost vector of the LP as, $\mathbf{c} := (\mathbf{0}_{4|V|}, 1)^\top$, we can represent the LP of Equation (16) in the standard form:

$$\textbf{primal} \quad \begin{array}{cc} \min\limits_{\mathbf{x}} & \mathbf{c}^\top \mathbf{x} \\ \text{s.t.} & A\mathbf{x} \geq \mathbf{b} \end{array} \qquad\qquad \textbf{dual} \quad \begin{array}{cc} \max\limits_{\mathbf{y}} & \mathbf{b}^\top \mathbf{y} \\ \text{s.t.} & \mathbf{y}^\top A \leq \mathbf{c}^\top \end{array} \tag{26}$$

Our goal is to provide a lower bound of the optimal value of the primal. Hence, we consider its dual, and provide a feasible solution. By weak duality lemma, the value of the dual objective at that feasible point will be a lower bound of the primal. The dual variables represented by $\mathbf{y}$ consists of $(\lambda, \gamma, \mu, \delta)$. The λ variables refer to the dual variables corresponding to the strategyproofness constraints, and we denote the dual variable that represent the strategyproofness of agent i between the profiles v^k and v^l by λ_{i,v^k,v^l}. By this representation, we consider only such pairs of profiles v^k and v^l where only agent i's valuation changes. The γ variables are the dual variables corresponding to the constraint that the SCF must add to unity, and we denote the dual variable corresponding to value profile v as γ_v. Dual variables μ and δ corresponds to the budget balance and the maximum inefficiency constraints. Since SCF and budget balance constraints are equalities, γ and μ are unrestricted, while λ and δ are non-negative. Additionally, in the primal problem the payment variable p_i's were unrestricted, hence in the dual the corresponding constraints are equalities.

We now provide a dual feasible solution, which is represented with respect to the reduced set of dual variables. Using symmetry according to Lemma 3, we reduce the number of valuation profiles. We number the profiles from 0 to 80 in the following way: for the valuation $(-0.5, -0.5)$ of agent 1, all possible valuations of agent 2 from $(-0.5, -0.5)$ to $(0.5, 0.5)$ (9 profiles) are listed, and then the valuation of agent 1 is moved to $(-0.5, 0)$. Due to symmetry, setting a primal variable $f_a(v)$ to a certain value also fixes 3 other variables that are agent-swapped or alternative-swapped or both-swapped versions of this variable. Denote the reduced set of valuation profiles by V_R. This also reduces the dual variables γ, μ, δ from 81 to 27 independent variables. However, for the λ variables we need to list all of them since they correspond to constraints that involve two valuation profiles. Consider the following set of dual variables (numbers of v and v' correspond to the valuation profile numbers in the listing discussed above):

i	v	v'	λ
1	11	2	4/7
1	12	21	4/7
1	30	66	1/14
1	52	16	2/7
1	57	30	1/7
1	60	33	3/14
1	68	32	2/7
1	78	15	4/7
2	12	16	4/7
2	14	10	1/14
2	18	26	4/7
2	20	19	3/14

v	γ
2	2/7
6	3/28
7	1/7
8	2/7
10	−1/28
11	2/7
14	−3/14
19	−4/7
24	−1/7

v	μ
6	−1/14
7	1/7
8	4/7
11	−4/7
12	−4/7
14	3/14
24	2/7

v	δ
2	2/7
6	1/7
11	4/7

All other entries of the variables that are not shown in the list above are zero. Note that for only the λ variables, the valuation profiles listed go beyond the index 27, but for all other dual variables they are represented by the 27 independent variables listed in V_R.

We claim that this is a feasible solution of the dual. The proof requires an exhaustive verification for each of the inequalities in the constraint set of the dual. However, we provide a few cases to give an insight how this example is picked. Consider, the variables $\lambda(1, 52, 16) = 2/7$ and $\lambda(1, 68, 32) = 2/7$. Note that, $v^{24} = ((0, 0.5), (0.5, 0)) = v^{52}$ and $v^{68} = ((0.5, 0), (0, 0.5))$ is an alternative swapped version of v^{24}. Also, none of the other variables involve any agent or alternative or both swap of this profile in the example we gave. Therefore, now we need to concentrate on the column $f_b(v^{24})$ in the matrix of Equation (25). Note that the matrix of Equation (25) is also reduced on the column and the rows. On the column, each of the f and p columns are reduced to $|V_R|$, and on the rows, only the strategyproofness constraints retain the original number, but the

SCF, budget balance and maximum inefficiency constraints reduce to $|V_R|$ in size. Carrying out the product with the terms we get $0.5 \times 2/7 + 0 \times 2/7 = 1/7$. While inspecting other variables, we find $\gamma(v^{24}) = -1/7$. Hence the sum of the products on the column $f_b(v^{24})$ gives $1/7 - 1/7 = 0$ which satisfies the inequality. This is not an isolated case, in all the columns $f_x(v)$ (the numbers of such variables are reduced because of symmetry), the examples are chosen such that the sum of the non-zero products in the SP constraints section and one non-zero product in the SCF constraints section add up to a non-positive number (for example, repeat the same argument for v^{78} and v^{19}).

Similarly, consider the column $p_2(v^{12})$: the variable $\lambda(2, 12, 16) = 4/7$ gets multiplied with -1 in this column since the constraint for agent 2 in the profile v^{12} gives a -1 coefficient for $p_2(v^{12})$. However, the variable $\mu(v^{12}) = -4/7$ which is multiplied with 1 in this column, and we can inspect that no other product is non-zero on this column. Hence the sum of the products is $-8/7$ which is non-positive, and satisfies the dual constraint.

The easiest thing to verify is the last column, where the sum of the δ_v for the reduced set of v's add to unity $(2/7 + 1/7 + 4/7)$. Therefore, the example provided is a dual feasible solution. We compute the objective value of the solution:

$$\begin{aligned}
&\sum_{v \in V_R} \gamma_v + \sum_{v \in V_R} \delta_v \max_{x \in \{a,b\}} w^v(x) \\
&= \frac{2}{7} + \frac{3}{28} + \frac{1}{7} + \frac{2}{7} - \frac{1}{28} + \frac{2}{7} - \frac{3}{14} - \frac{4}{7} - \frac{1}{7} + 0.5 \times \frac{4}{7} \\
&= \frac{1}{7}
\end{aligned}$$

This completes the proof of the theorem. ■