



## King's Research Portal

DOI:

[10.1007/s10458-019-09435-9](https://doi.org/10.1007/s10458-019-09435-9)

*Document Version*

Peer reviewed version

[Link to publication record in King's Research Portal](#)

*Citation for published version (APA):*

Serafino, P., Ventre, C., & Vidali, A. (2020). Truthfulness on a budget: trading money for approximation through monitoring. *Autonomous Agents and Multi-Agent Systems*, 34(1), [5]. <https://doi.org/10.1007/s10458-019-09435-9>

### **Citing this paper**

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

### **General rights**

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

### **Take down policy**

If you believe that this document breaches copyright please contact [librarypure@kcl.ac.uk](mailto:librarypure@kcl.ac.uk) providing details, and we will remove access to the work immediately and investigate your claim.

# Truthfulness on a Budget: Trading Money for Approximation through Monitoring

Paolo Serafino · Carmine Ventre ·  
Angelina Vidali

the date of receipt and acceptance should be inserted later

**Abstract** Albeit a pervasive desideratum when computing in the presence of selfish agents, truthfulness typically imposes severe limitations to what can be implemented. The price of these limitations is typically paid either *economically*, in terms of the financial resources needed to enforce truthfulness, or *algorithmically*, in terms of restricting the set of implementable objective functions, which often leads to renouncing optimality and resorting to approximate allocations. In this paper, with regards to utilitarian problems, we ask two fundamental questions: (i) what is the minimum sufficient budget needed by optimal truthful mechanisms, and (ii) whether it is possible to sacrifice optimality in order to achieve truthfulness with a lower budget. To answer these questions, we connect two streams of work on mechanism design and look at monitoring – a paradigm wherein agents’ actual costs are bound to their declarations. In this setting, we prove that the social cost is always a sufficient budget, even for collusion-resistant mechanisms, and, under mild conditions, also a necessary budget for a large class of utilitarian problems that encompass set system problems. Furthermore, for two well-studied problems outside of this class, namely facility location and obnoxious facility location, we draw a novel picture about the relationship between (additive) approximation and frugality. While for optimal mechanisms we prove that the social cost is always a sufficient and necessary budget for both problems, for approximate mechanisms we do have a dichotomy: for the facility location problem (i.e., agents want to be close to the facilities) we show that “good” approximations still need a budget equal to the social cost; on the contrary, for the obnoxious facility location problem (i.e. agents want to be as far away from the facilities as possible) we show that it is possible to trade approximation for frugality, thus obtaining truthfulness with a lower budget.

---

P. Serafino  
Gran Sasso Science Institute, Italy. E-mail: paolo.serafino@gssi.it

C. Ventre  
King’s College London, United Kingdom. E-mail: carmine.ventre@kcl.ac.uk

A. Vidali  
University of Athens, Greece. E-mail: angvid@gmail.com

**Keywords** First keyword · Second keyword · More

## 1 Introduction

*Truthful* (also termed *strategyproof*) mechanisms represent the principal tool to compute in the presence of selfish agents. They guarantee that the objectives of the designer (e.g., minimize a certain global function) are aligned with the agents' individual objectives. This compatibility to the incentives, however, often requires extravagantly large monetary transfers from the mechanism to the agents. Consider, e.g., the problem of buying a path on a given graph whose  $n$  edges are owned by selfish agents. This is an instance of a *utilitarian* problem – a problem wherein each agent incurs a cost to implement the mechanism's outcome and the designer wants to minimize the sum of individual agents' costs. In this case, while truthfulness can be easily guaranteed by using Vickrey-Clarke-Groves (VCG), the amount of money needed to achieve incentive-compatibility can be as large as  $n$  times the cost of the second cheapest path. The truthfulness guarantee then loses some of its appeal, as it cannot be ensured when the budget of the designer is not big enough [39].

This motivates the definition of *budget-feasible mechanisms* [47], wherein, together with truthfulness, we require that the total amount of payments is within a given budget. The literature on this notion (cf., e.g., [16,7,13,30,2,25,3,11]) mainly looks at procurement auctions, in which sellers (bidders) have private costs to produce items, and the buyer (auctioneer) aims at maximizing a social valuation function on subsets of items, under the budget constraint on the total payment. Alas, this setting is quite different from the path auction setting (and, more generally, from any utilitarian problem) in that the designer's objective does not depend on the agents' declarations. The literature on *frugality* of truthful mechanisms for utilitarian problems, on the other hand, proves that the overpayment of VCG is not accidental: any truthful mechanism for path auctions *must* overpay by as much [19]. Further research proves that similar results hold true for other utilitarian problems, including minimum spanning trees, set systems and matroids [48,28]. Therefore, generally speaking, the budget might be as high as the total payments of VCG, making truthfulness too expensive to afford for a mechanism designer on a budget.

A related strand of work, which considers the weaker solution concept of Nash equilibria instead of dominant strategies, studies *subsidies* to selfish agents as a means of enforcing efficient equilibria in non-cooperative games. For instance, subsidies in price competition games among single product vendors are studied in [10]. In [8] a cost sharing system is studied where users purchase public services equally sharing their cost, and a central authority enforces “good” purchasing choices by reinvesting tax revenue in the form of subsidies to the public services. In [5] Augustine et al. study network design games where a set of agents share the cost of building edges on a network, and subsidies (partially) cover the cost of some of the edges included in the equilibrium spanning tree. Augustine et al. prove that to enforce a particular spanning tree as a Nash equilibrium it is sufficient and necessary to subsidize a constant fraction of the cost of the minimum spanning tree.

In this work we pursue two main desiderata: firstly, we want to make truthfulness less expensive to obtain, and, secondly, we want to characterize the payments – and hence the budget needed – for *any* utilitarian problem. As from the discussion above, making truthfulness less expensive requires us to depart from classical mechanism design; it turns out that our choice of alternative mechanism design paradigm goes some way towards the latter objective (i.e., obtaining a characterization of frugal payments). In detail, we look at *mechanisms with monitoring*, studied in [32,33,23].<sup>1</sup> Mechanisms with monitoring assume that the designer is able to monitor the agents’ costs so to guarantee that the agents who have exaggerated theirs end up paying the reported (higher) cost. Intuitively, it is assumed that the designer can ‘watch’ the agents at “work” and force them to spend more than their real cost, if they overbid. The designer, however, neither punishes agents who underbid their cost nor imposes any extra fines to overbidding agents. For example, let us consider again the path auction setting where each edge weight represents its latency time. If an agent overbids her cost, then the designer can detect that, as the observed latency time will be lower than the reported one (unless the agent waits and “simulates” the reported cost – note that in doing so the agent will effectively pay her reported cost). In this scenario, it is possible to design a mechanism with monitoring as the designer can easily enforce this higher cost by, e.g., charging the difference to the agents or keeping them idle for the time difference.

### 1.1 Our contribution

We begin by observing that, irrespectively of the algorithm of interest, paying an agent her declared cost is sufficient to guarantee *collusion-resistant* mechanisms with monitoring, i.e., the strongest form of incentive-compatibility, resistant to collusive behavior and side payments. This result has three important implications. Firstly, for mechanisms with monitoring, the authors of [33] prove that any algorithm admits payments leading to truthful mechanisms with monitoring when agents bid from finite domains (i.e., the possible cost functions they can declare to the mechanism belong to a finite set). Our result extends theirs in that: (i) we do not make any assumption on the bidding domain and (ii) we significantly strengthen the incentive-compatibility guarantee to obtain resistance to collusive behavior and side payments. The second consequence concerns the related research agenda in mechanisms with verification [39,41]. The notion defined in [39] is more restrictive than monitoring since underbidding is punished. We show that the compensation and bonus mechanism in [39] is unnecessarily complex: simple compensation without punishments guarantees both truthfulness and collusion-resistance. Lastly, we show that the social cost (i.e., sum of the agents’ costs) is a sufficient budget for *any* (not necessarily utilitarian) problem.

---

<sup>1</sup> The terminology is introduced in [33] while [32,23] study the same model under the misnomer of mechanisms with verification. In the verification model of [39] it is effectively assumed that the designer can monitor overbidding *and* punish underbidding; in [41] (and related literature) instead the designer only punishes underbidding but does not monitor overbidding – the difference between the models is furthermore studied therein. We here only monitor overbidding as in [32,33,23].

*Discounting the payments for a cheaper budget.* We subsequently investigate the extent to which the social cost is a necessary budget. To do so, we relax voluntary participation and study the feasibility of diminishing our payment function via discounts (for the designer) while maintaining incentive-compatibility and non-negative payments (clearly, *we do not want to charge agents*, for otherwise any budget would be feasible). There are many real-life scenarios where voluntary participation is not satisfied and agents are somehow ‘forced’ to play (e.g., abiding by a law that increases one’s own cost).

We assume that the discounts are bid-independent functions; e.g., the discount applied to the payment to player  $i$  does not depend on her bid  $b_i$ . While in the classical setting this assumption is without loss of generality, for mechanisms with monitoring it excludes some possible payment schemes. For example, the payment function in [33] can be seen as a first-price payment discounted by a bid-dependent function. However, we make two important observations on more general bid-dependent discounts. Firstly, we show that bid-dependent discounts *do restrict* the class of algorithms that can be used by incentive-compatible mechanisms. Having bid-independent discounts is then necessary to keep the focus on algorithms as general as possible. Furthermore, studying budget feasibility in general would require an algorithmic characterization of truthful mechanisms with monitoring in order to relate the quality of the algorithmic solution to the sum of the payments needed. Understanding how truthfulness depends upon the individual bid-dependent discounts is an open problem that requires the development of novel techniques, as all the known ones only deal with bid-independent payments. An initial study for a specific discount is done in [21]. We however give some indication that bid-dependent discounts are not much cheaper than bid-independent discounts, and prove how a particular subclass of the former is actually not cheaper at all.

*Can we trade money for approximation?* Under the assumption of bid-independent discounts, we first show the necessity of the social cost as a budget for the large class of *set system problems* (which include, essentially, all problems on graphs such as path auctions, shortest path trees, minimum spanning trees, etc.). We then focus on two well-known variants of the facility location problem, namely the  $K$ -facility location and the obnoxious 1-facility location. We want to determine the location of some facilities ( $K$  desirable facilities and one obnoxious facility, respectively) on the real line on input the locations of  $n$  agents, with the objective of minimizing the social cost. In the  $K$ -facility location problem, facilities provide some desirable service to the agents, who then want to be as close as possible to them: the social cost is defined as the sum of the individual agents’ distances between their location and the nearest facility. In the obnoxious 1-facility location problem, the facility causes some discomfort to the agents (e.g., it is a landfill site), hence they want to be as far away from it as possible: the social cost is defined as the sum of individual agents’ costs that are non-increasing functions with respect to the distance between their location and that of the facility.

The motivation to study the budget for this paradigmatic problem in algorithmic mechanism design is twofold: (i) it has a combinatorial structure significantly different from the aforementioned set system problems that requires an ad hoc analysis; and (ii) the state of the art on the problem suggests that money ought to be used to obtain both truthfulness and good approximations [22, 20].

For  $K$ -facility location, we provide a number of results proving the necessity of the social cost as a budget, depending on the algorithm of interest. Arguably, the most interesting contribution relates (additive) approximation guarantee with frugality and answers the question of whether truthful implementations of approximate solutions might be cheaper than that of the optimum. We prove that there is no such tradeoff as long as “good” approximations are concerned. For obnoxious 1-facility location, we show that the social cost is a necessary budget for optimal mechanisms, but, contrarily to the facility location problem, we show that money and (additive) approximation guarantee can be traded for one another: a designer can here in fact know how good a solution they can truthfully implement given their available budget, or, alternatively, how much they can save on their budget if they are willing to accept a certain degree of approximation.

*Roadmap.* The remainder of the paper is organized as follows. In Section 2 we discuss some related work. In Section 3 we give some formal definitions. Section 4 shows that the social cost is a sufficient budget for collusion-resistant mechanisms with monitoring (Corollary 1) and identifies some properties for this to be also a necessary budget for any algorithm (Theorem 2). Subsequent sections study these properties for a number of optimization problems studied in the literature, namely: set systems (Section 5),  $K$ -facility location (Section 6) and obnoxious facility location (Section 7).

## 2 Related Work

We here discuss some previous works related to facility location, obnoxious facility location and set system.

### 2.1 Set systems

Talwar [48] analyzes the overpayment the VCG mechanism incurs to enforce truthfulness in  $k$ -set cover problems, where, given a universe  $E$  and a collection of sets  $\{S_1, \dots, S_m\}$ , each having a cost  $c(S_i)$  and  $|S_i| \leq k$ , the problem consists in finding a minimum cost cover whose union equals  $E$ . Talwar shows that the payment of the VCG mechanism is at most  $k \cdot c(OPT')$ , where  $OPT'$  is the cost of the best cover *disjoint* from the optimum cover  $OPT$ . For the same  $k$ -set cover problem as [48], Calinescu [9] proposes two *approximate* truthful mechanisms whose payments are bounded by  $(k-1) \cdot c(OPT) + k \cdot c(OPT')$ . Furthermore, he shows that the same two approximation algorithms have payments bounded by  $k \cdot (c(OPT) + c(OPT'))$  when applied to more general set systems, which include  $k$ -polymatroid cover (a problem related to Steiner Tree computations). The total payments of these algorithms are bounded by  $q \cdot k^2$  times the total payment of the VCG mechanism, where  $q$  is such that an element in a  $k$ -set cover instance appears in at most  $q$  sets.

To address those scenarios where the concept of frugality ratio given in [48] is not well defined, Karlin et al. [28] define a new notion of frugality ratio for procurement auctions. In particular, they compare the payments of a truthful auction to the minimum payments of a *first-price auction* (i.e., each bidder is paid their bid) at a *Nash equilibrium* of the game (since first-price auctions are

not truthful and hence have no dominant strategies). Under this new definition of frugality ratio, they generalize results from [48] and prove that VCG mechanisms have frugality ratio 1 for auctions on set systems if and only if the feasible sets of the set system are the basis of a *monopoly-free matroid*. Furthermore, they develop truthful mechanisms with constant frugality ratio for  $r$ -out-of- $k$  set systems<sup>2</sup> and path auction problems.

Elkind et al. [18] propose a novel truthful polynomial-time auction for the vertex cover problem and bound its frugality ratio. In particular, they show both (i) that the solution quality of a truthful auction is within a *constant factor of the optimum*, and (ii) that its frugality ratio is within a *constant factor* of the best possible worst-case bound. Notably, this is the first auction for this problem to have these properties. Furthermore, they show how *any* truthful auction can be transformed into a frugal one while preserving the approximation ratio. They also consider two natural modifications of the definition of frugality ratio given in [28], and analyze the properties of the resulting payment bounds.

Kempe et al. [29] study truthful frugal mechanisms for vertex cover auctions,  $k$ -flow auctions (i.e., the auctioneer wants to purchase  $k$  edge-disjoint  $s$ - $t$  paths), and cut auctions (i.e., the auctioneer wants to purchase an  $s$ - $t$  cut) and present truthful mechanisms with *constant* frugality ratios for all three set systems problems, i.e., the maximum overpayment of the mechanism is within a constant factor of the maximum overpayment of any truthful mechanism, for every set system in the class. In particular, for vertex cover they propose a mechanism based on scaling each bid by a multiplier derived from the dominant eigenvector of a certain matrix, whereas the mechanism for  $k$ -flows prunes the graph to be minimally  $(k + 1)$ -connected, and then applies the vertex cover mechanism. The mechanism for cut auctions contracts the graph until all  $s$ - $t$  paths have length exactly 2, and then applies the vertex cover mechanism.

Chen et al. [12] devise a uniform scheme for designing frugal truthful mechanisms for general set systems based on scaling the agents bids using the eigenvector of a matrix that encodes the interdependencies between the agents. They show how the  $r$ -out-of- $k$ -system mechanism and the mechanism for buying a path in a graph presented in [28] can be viewed as instantiations of this scheme. Furthermore, they apply their scheme to vertex cover systems and  $k$ -path systems, bounding the frugality of the mechanism in terms of the largest eigenvalue of the respective interdependency matrix.

Hajiaghayi et al. [26] introduce a new benchmark for measuring the frugality ratio of an auction and show two attractive properties lacking in the previous proposals, namely: *monotonicity* and *smoothness*. Based on this benchmark, they provide positive results for vertex cover and knapsack auctions. In particular, for vertex cover auctions they prove that the mechanism proposed in [29] is optimal up to a constant factor, whereas for knapsack auctions they give a truthful mechanism with an almost-tight bounded frugality ratio.

---

<sup>2</sup> A set system  $(E, \mathcal{F})$  is an  $r$ -out-of- $k$ -system if there exists a partition of  $E$  into  $k$  disjoint sets  $S_1, \dots, S_k$ , such that every set  $F \in \mathcal{F}$  contains exactly  $r$  out of these  $k$  sets of elements.

## 2.2 Facility Location

The facility location problem has proved to be a fertile research problem in many research communities. In his classical paper ([38]) Moulin characterizes the class of generalized median voter schemes as the only deterministic truthful mechanisms for *single-peaked* agents on the line. Schummer and Vohra [43] extend the result of Moulin to the more general setting where *continuous graphs* are considered, characterizing truthful mechanisms on *continuous lines and trees*. They show that on circular graphs every truthful mechanism must be dictatorial. Procaccia and Tennenholtz [42] initiated the field of approximate mechanism design without money by suggesting the idea of leveraging *approximation* to obtain strategyproofness in all those scenarios where (i) the optimal allocation is not truthful and (ii) it is not possible to resort to monetary transfers. For the 2-facility location problem, they propose an  $(n - 2)$ -approximate truthful mechanism, where  $n$  is the number of agents. Furthermore, they provide a lower bound of  $3/2$  on the approximation ratio of any truthful algorithm for the facility location problem on the line and conjecture a lower bound of  $\Omega(n)$ . The latter conjecture was later proven by Fotakis et al. [22] by providing a characterization of deterministic truthful mechanisms with *bounded approximation ratio* for the 2-facility location problem on the line. Lu et al. [36], improve several bounds studied in [42]. Particularly, as regards deterministic algorithms they prove a better (with respect to [42]) lower bound of  $2 - \mathcal{O}(\frac{1}{n})$ . Furthermore, they prove a 1.045 lower bound for randomized mechanisms for the 2-facility location problem on the line and present a randomized  $n/2$ -approximate mechanism. In [35], Lu et al. focus on general metric spaces for 2-facility games. They prove an  $\Omega(n)$  lower bound for the approximation of deterministic strategyproof mechanisms and propose a randomized algorithm that achieves *constant* approximation ratio. Alon et al. [1] derive a linear (in the number of agents) lower bound for truthful mechanisms on continuous cycles and a constant approximation bound for randomized mechanisms in the same settings. Dokow et al [17] study facility location problems on *discrete* lines and cycles. They prove that, in this setting, truthful mechanisms are *nearly-dictatorial* in that all agents can effect the outcome to a certain extent. Contrarily to the case of continuous cycles studied in [43], for small discrete graphs Dokow et al. prove that there are anonymous truthful mechanisms. Furthermore, they prove a linear lower bound in the number of agents for the approximation ratio of truthful mechanisms on discrete cycles. Facility location for discrete graphs and heterogeneous facilities (i.e., serving different purposes) has been studied [44–46] also in presence of heterogeneous preferences [49, 4, 34] where agents can either like or dislike the facilities.

Ferraioli et al. [20] focus on identifying a set of verification assumption that yield good approximation guarantees for (homogenous)  $K$ -facility location problems on the line. In particular they identify sets of cost-based and topological verification assumptions that guarantee truthfulness of the optimal mechanisms for both the social cost (for a generic  $K$  and max cost (for  $K = 2$ ) objective functions, showing that these sets of assumptions are minimal in the sense that relaxing any assumption in the set breaks truthfulness. Furthermore, they conjecture that these verification assumptions are necessary not only to achieve truthfulness with optimal algorithms, but also for any constant-approximation (or even sublinear-approximation) algorithm.

### 2.3 Obnoxious Facility Location

Cheng et al. [15] are the first to study the obnoxious facility location problem from a mechanism design without money perspective. They prove a 3-approximation group-strategyproof deterministic mechanism, which is the best possible if the only two candidates are the endpoints of the path, and a group-strategyproof randomized mechanism with tight approximation ratio of  $\frac{3}{2}$ . For circles and trees they propose 3-approximation group-strategyproof deterministic mechanisms and, for unrestricted networks, a 4-approximation group-strategyproof deterministic mechanism and a 2-approximation group-strategyproof randomized mechanism. Ibara and Nagamochi [27] prove that if there are more than two candidates (i.e. locations output by the mechanism for some reported locations), there are no truthful mechanisms without money in the linear metric. Furthermore, they characterize group-strategyproofness for mechanisms with exactly two candidates in the general metric and show that there always exists a 4-approximation group-strategyproof mechanism in any metric.

In [48], Talwar gives an upper bound for the total payments in a facility location game, but his setting is quite different to the one we explore in this paper, as agents control the facilities and connection costs are publicly known.

### 3 Preliminaries

We have a set of outcomes  $\mathcal{O}$  and  $n$  selfish agents. Each selfish agent has a *cost function*, also called *type*  $t_i : \mathcal{O} \rightarrow \mathbb{R}_{\geq 0}$ . For  $x \in \mathcal{O}$ ,  $t_i(x)$  is the cost paid by agent  $i$  to implement  $x$ . The type  $t_i$  is *private knowledge* of agent  $i$ . The set of all admissible cost functions  $t_i$  is called the *domain* of agent  $i$  and is denoted  $D_i$ . Assuming that each agent has (truthfully or otherwise) reported, or *bid*, a cost function  $b_i \in D_i$ , a *mechanism* determines an allocation  $x \in \mathcal{O}$ . Furthermore, depending the different  $b_i$ , it will determine payments  $p = (p_1, \dots, p_n)$  to be given to the agents. In summary, by letting  $D = \times_{i=1}^n D_i$ , a *mechanism* is a pair  $(f, p)$  that, on input the agents' costs functions, uses algorithm  $f : D \rightarrow \mathcal{O}$  to return an outcome in  $\mathcal{O}$  and function  $p : D \rightarrow \mathbb{R}_{\geq 0}^n$  to determine payments to be awarded to each agent  $i$ .

For mechanism  $(f, p)$  and a bid vector  $\mathbf{b}$ , we let  $u_i^{(f,p)}(\mathbf{b})$  denote the *utility* of agent  $i$  for the output computed by  $(f, p)$  on input  $\mathbf{b}$ . To ease the notation we drop  $(f, p)$  when this is clear from the context. Since the types  $t_i$  are private knowledge of the agents, they might find it profitable to bid  $b_i \neq t_i$ . We are interested in mechanisms for which *truth-telling is a dominant strategy* for each agent.

**Definition 1** A mechanism  $(f, p)$  is *truthful* if for any  $i$ , any bids  $\mathbf{b}_{-i}$  of the agents other than  $i$ , and any  $b_i \in D_i$ ,  $u_i(t_i, \mathbf{b}_{-i}) \geq u_i(b_i, \mathbf{b}_{-i})$ .

*Collusion resistance* is a stronger requirement that demands truth-telling be a dominant strategy for coalitions of agents that can exchange side payments among themselves.

**Definition 2** A mechanism  $(f, p)$  is *collusion-resistant* if for any  $C \subseteq [n]$  any bids  $\mathbf{b}_{-C} \in \times_{i \notin C} D_i$  and any  $\mathbf{b}_C \in \times_{i \in C} D_i$ ,  $\sum_{i \in C} u_i(\mathbf{t}_C, \mathbf{b}_{-C}) \geq \sum_{i \in C} u_i(\mathbf{b}_C, \mathbf{b}_{-C})$ ,  $\mathbf{t}_C$  denoting the vector  $(t_i)_{i \in C}$ .

A mechanism  $(f, p)$  satisfies *Voluntary Participation* (VP) if  $u_i(t_i, \mathbf{b}_{-i}) \geq 0$  for all agents  $i$ ,  $t_i \in D_i$  and  $\mathbf{b}_{-i} \in \times_{j \neq i} D_j$ . We also say that a mechanism is *budget-feasible* (with respect to budget  $B$ ) if, for any  $\mathbf{b} \in D$ ,  $\sum_{i=1}^n p_i(\mathbf{b}) \leq B$ . Observe that budget-feasibility makes sense when payments are non-negative (as from our model above) as any mechanism charging agents would be budget-feasible (with respect to budget 0). This property on the sign of the payments is sometimes known in literature as *No Positive Transfer* (NPT).

In the literature,  $u_i$  is usually defined as a linear combination of the payment and the agent’s *true* cost. In mechanisms with monitoring, this quasi-linear definition is maintained but costs paid by the agents for the allocated resource are more strictly tied to their declarations. Intuitively, it is assumed that the designer can ‘watch’ the agents at “work” and force them to spend more than their real cost, if they overbid. Note that the designer neither punishes agents who underbid their cost nor imposes any extra fines to overbidding agents. In this sense, such a notion is much less restrictive than the concept of mechanisms with verification studied, e.g., in [39, 41]. In a *mechanism with monitoring*  $(f, p)$ , we have

$$u_i^{(f,p)}(b_i, \mathbf{b}_{-i}) := p_i(\mathbf{b}) - \max\{t_i(f(\mathbf{b})), b_i(f(\mathbf{b}))\}.$$

It is important to stress that the designer does not observe the agents’ private information but only checks that agents are not more economical (e.g., faster) than declared. Agents can pretend to have a higher cost (e.g., processing time) at the expense of being “busy” that long. Agents can still underbid and at execution time have a higher cost (e.g., saying to have underestimated their cost or work). There are many applicative scenarios in which this can be implemented. In the context of facility location, one might think of facilities as meeting/conference places, and of costs as travel expenses to reach the facilities. When agents’ costs are reimbursed, receipts must be presented. For the agents to be consistent with overbidding they need to pay the exaggerated (reported) cost. Receipts are then a “monitoring” tool.

For an instance  $\mathbf{b}$  of a *utilitarian cost-minimization problem*  $\Pi$ , we focus on algorithms  $f^*$  optimizing the *social cost*, i.e.,  $f^*(\mathbf{b}) \in \arg \min_{x \in \mathcal{O}} \text{cost}(x, \mathbf{b})$ , where  $\text{cost}(x, \mathbf{b}) = \sum_{i=1}^n b_i(x)$ . We say that algorithm  $f$  for a utilitarian cost-minimization problem  $\Pi$  is (*additively*)  $\alpha$ -*approximate* if, for every  $\mathbf{b}$ , it returns a solution whose cost is at most  $\text{cost}(f^*(\mathbf{b}), \mathbf{b}) + \alpha$ .

We argue that adopting additive approximation instead of the classical notion of multiplicative approximation is better suited to our goals. Indeed, our aim here is not merely to optimize the approximation, but rather that of identifying a tool that allows us to relate the budget with the quality of solutions. As we will prove later on, in our setting the additive approximation term is more appropriate to serve this purpose, as it immediately tells us the amount of extra payments (prior to the discounts, if any) needed for implementation. Interestingly, [24] also argues the merits of additive approximation for facility location.

#### 4 Bounding the budget

For an algorithm  $f$ , define  $p_i(\mathbf{b}) = b_i(f(\mathbf{b}))$  and call  $(f, p)$  a *first-price mechanism*.

**Theorem 1** *For any  $f$ , the first-price mechanism  $(f, p)$  is a collusion-resistant mechanism with monitoring.*

*Proof.* Fix  $C$  and  $\mathbf{b}_{-C}$ . Let  $\mathbf{t}$  be a shorthand for the bid vector  $(\mathbf{t}_C, \mathbf{b}_{-C})$  and, for a given  $\mathbf{b}_C$ , let  $\mathbf{b}$  be a shorthand for  $(\mathbf{b}_C, \mathbf{b}_{-C})$ . We observe that

$$\sum_{i \in C} u_i^{(f,p)}(\mathbf{t}) = \sum_{i \in C} p_i(\mathbf{t}) - t_i(f(\mathbf{t})) = \sum_{i \in C} t_i(f(\mathbf{t})) - t_i(f(\mathbf{t})) = 0;$$

while

$$\begin{aligned} \sum_{i \in C} u_i^{(f,p)}(\mathbf{b}) &= \sum_{i \in C} p_i(\mathbf{b}) - \max\{t_i(f(\mathbf{b})), b_i(f(\mathbf{b}))\} \\ &= \sum_{i \in C} b_i(f(\mathbf{b})) - \max\{t_i(f(\mathbf{b})), b_i(f(\mathbf{b}))\} \leq 0. \end{aligned}$$

This concludes the proof.  $\square$

We can then conclude that the social cost is a sufficient budget.

**Corollary 1** *Let  $\Pi$  be a (utilitarian) cost-minimization problem. There exists a collusion-resistant budget-feasible first-price mechanism  $(f, p)$  with monitoring and budget  $B$  if for any instance  $\mathbf{b}$  of  $\Pi$ ,  $\text{cost}(f(\mathbf{b}), \mathbf{b}) \leq B$ .*

#### 4.1 When is the social cost necessary?

The minimum budget needed for a budget-feasible first-price mechanism using  $f$  for problem  $\Pi$  is  $B^* = \max_{\mathbf{b} \in D} \text{cost}(f(\mathbf{b}), \mathbf{b})$ . Is it possible to maintain incentive-compatibility with a budget smaller than  $B^*$ ? A first observation is that, if we want our mechanisms to satisfy voluntary participation, then  $B^*$  is the smallest budget possible. However, in some applications, it is reasonable to adopt mechanisms that do not satisfy this property and therefore it makes sense to consider the necessity of  $B^*$  for those mechanisms. For instance, drawing a parallel with the setting where subsidies are employed to enforce efficient equilibria, [8, 5, 10], relaxing voluntary participation is equivalent to providing partial subsidies that do not cover the whole cost of the equilibrium solution. We pursue this research direction in the remainder of the paper.

##### 4.1.1 Payment structure

It is easy to see that the incentive-compatibility of first-price mechanisms is preserved if

$$p_i(\mathbf{b}) = b_i(f(\mathbf{b})) + h_i \quad \text{for collusion-resistance} \quad (1)$$

$$p_i(\mathbf{b}) = b_i(f(\mathbf{b})) + h_i(\mathbf{b}_{-i}) \quad \text{for truthfulness} \quad (2)$$

This is an if and only if for classical mechanisms (see, for instance, [40, Theorem 9.37]) in that payments of truthful mechanisms are unique up to bidder-dependent discounts that depend only on the declarations of the other bidders. We call the mechanisms that use the payment functions above (*bid-independent*)

*discounted first-price mechanisms.* Since we want both to save on the payments of a first-price mechanism and to have non-negative payments, we have that for all  $i$  and  $\mathbf{b}$ :

$$-b_i(f(\mathbf{b})) \leq h_i, h_i(\mathbf{b}_{-i}) \leq 0. \quad (3)$$

We can then prove the following result.

**Theorem 2** (i) *If for all  $i$  there exists a bid vector  $\mathbf{b}$  such that  $b_i(f(\mathbf{b})) = 0$  then there is no collusion-resistant discounted first-price budget-feasible mechanism with monitoring using  $f$  and whose budget is smaller than  $B^*$ .*

(ii) *If for all  $i$  and  $\mathbf{b}_{-i}$  there exists  $b_i$  such that  $b_i(f(\mathbf{b})) = 0$  then there is no truthful discounted first-price budget-feasible mechanism with monitoring using  $f$  and whose budget is smaller than  $B^*$ .*

*Proof.* (i) Let us assume for the sake of contradiction that there exists a collusion-resistant discounted first-price budget-feasible mechanism with monitoring  $f$  such that for all  $i$  there exists  $\mathbf{b}$  such that  $b_i(f(\mathbf{b})) = 0$ , and still  $\max_{\mathbf{b} \in D} \sum_{i=1}^n (p_i(\mathbf{b}) + h_i) < B^*$ . Let  $\mathbf{b}^* \in \operatorname{argmax}_{\mathbf{b} \in D} \operatorname{cost}(f(\mathbf{b}), \mathbf{b})$ . Since  $\sum_{i=1}^n (p_i(\mathbf{b}^*) + h_i) < B^*$ , there must exist an  $\hat{i}$  such that  $h_{\hat{i}} < 0$ . Let  $\bar{\mathbf{b}}$  be the vector such that  $\bar{b}_{\hat{i}}(f(\bar{\mathbf{b}})) = 0$ . We have that  $h_{\hat{i}} < -\bar{b}_{\hat{i}}(f(\bar{\mathbf{b}}))$ , which violates (3).

(ii) Let us assume by contradiction that there exists a truthful discounted first-price budget-feasible mechanism with monitoring  $f$  such that for all  $i$  and for all  $\mathbf{b}_{-i}$  there exists  $b_i$  such that  $b_i(f(\mathbf{b})) = 0$ , and still  $\max_{\mathbf{b} \in D} \sum_{i=1}^n (p_i(\mathbf{b}) + h_i(\mathbf{b}_{-i})) < B^*$ . Let  $\mathbf{b}^* \in \operatorname{argmax}_{\mathbf{b} \in D} \operatorname{cost}(f(\mathbf{b}), \mathbf{b})$ . Since we have assumed that  $\sum_{i=1}^n (p_i(\mathbf{b}^*) + h_i(\mathbf{b}_{-i}^*)) < B^*$  holds, there must exist an  $\hat{i}$  such that  $h_{\hat{i}}(\mathbf{b}_{-i}^*) < 0$ . Let  $\bar{b}_{\hat{i}}$  be such that  $\bar{b}_{\hat{i}}(f(\bar{b}_{\hat{i}}, \mathbf{b}_{-i}^*)) = 0$ . We have that  $h_{\hat{i}}(\mathbf{b}_{-i}^*) < -\bar{b}_{\hat{i}}(f(\bar{b}_{\hat{i}}, \mathbf{b}_{-i}^*))$ , which violates (3).  $\square$

The theorem above gives us a condition for the necessity of  $B^*$  for discounted first-price mechanisms. Whenever there is no  $b_i$  and  $\mathbf{b}_{-i}$  such that  $b_i(f(\mathbf{b})) = 0$ , we do have a discount in terms of payments. It is not hard to see that to maximize this discount, and reduce the budget of a truthful discounted first-price mechanism as much as possible, we simply need to set

$$h_i(\mathbf{b}_{-i}) = \inf_{b \in D_i} b(f(b, \mathbf{b}_{-i})). \quad (4)$$

#### 4.1.2 The limits of more general discounts

We justify the focus of Theorem 2 on payment functions (1) and (2) by showing the limitations of more general payment functions, such as,

$$p_i(\mathbf{b}) = b_i(f(\mathbf{b})) + h_i(\mathbf{b})$$

that we call *fully discounted* with respect to  $(h_i)_i$ , and the very special case

$$p_i(\mathbf{b}) = b_i(f(\mathbf{b})) + h_i(b_i)$$

---

<sup>3</sup> One might wonder whether relaxing the upper bound for some agents, and allowing a surplus for them, would allow to save more on the remaining agents. Our results of necessity cover also this case since they prove that no *single* agent can have a negative discount.

termed *single-bid discounted*. Similar to the case of discounted first price mechanisms, we observe that the following inequality must hold for all  $i$  and all  $\mathbf{b}$ :

$$-b_i(f(\mathbf{b})) \leq h_i(b_i), h_i(\mathbf{b}) \leq 0. \quad (5)$$

*General discounts (might) destroy truthfulness*

Theorem 2 is as general as it can be without restricting the class of algorithms it applies to. We show, in fact, that fully discounted payments restrict the class of algorithms that can be used by truthful mechanisms.

An algorithm  $f$  is said *monotone* with respect to  $h_i$  if  $b_i(f(\mathbf{b})) > b'_i(f(\mathbf{b}))$  implies  $h_i(b_i, \mathbf{b}_{-i}) > h_i(b'_i, \mathbf{b}_{-i})$  for some  $\mathbf{b}_{-i}$  and  $b'_i \neq b_i$ . It is then not too hard to prove:

**Proposition 1** *An algorithm  $f$  monotone with respect to  $h_i$  is not truthful with monitoring when coupled with a payment function that is fully discounted with respect to  $(h_j)_j$ .*

*Proof.* Let  $\mathbf{b}_{-i}$  be the vector for which  $f$  is monotone with respect to  $h_i$ . Consider a  $b_i \in D_i$  and define agent  $i$ 's true type to be  $t_i \neq b_i, t_i \in D_i$  such that  $b_i(f(\mathbf{b})) \geq t_i(f(\mathbf{b}))$  (such a  $t_i$  always exists as long as the domain is sufficiently rich). Then, agent  $i$  is better off declaring  $b_i$  when  $f$  is coupled with a fully discounted payment function. In fact,  $u_i(t_i, \mathbf{b}_{-i}) = h_i(t_i, \mathbf{b}_{-i})$  and  $u_i(b_i, \mathbf{b}_{-i}) = h_i(b_i, \mathbf{b}_{-i})$ . By monotonicity of  $f$ , we have  $h_i(b_i, \mathbf{b}_{-i}) > h_i(t_i, \mathbf{b}_{-i})$ .  $\square$

*Money is needed*

The observation here directly follows from the results in [33]. The authors prove that mechanisms with monitoring without money have very limited power in the context of truthful allocation of RAM to selfish jobs. Specifically there is no truthful mechanism with monitoring without money with bounded approximation ratio. This implies that, in general, money needs to be used (i.e.,  $h_i(\mathbf{b}) > -b_i(f(\mathbf{b}))$ ) and then savings, if any, are somewhat limited. Furthermore, this suggests that the trade-off between approximation and frugality is strongly unbalanced.

*Single-bid discounts are as bad as bid-independent discounts*

The following result, similar in nature to Theorem 2, gives us a condition for the necessity of  $B^*$  for single-bid discounted mechanisms.

**Theorem 3** *If for all  $i$  and for all  $b_i$  there exists  $\mathbf{b}_{-i}$  such that  $b_i(f(\mathbf{b})) = 0$ , then there is no truthful single-bid discounted first price mechanism with monitoring using  $f$  and whose budget is smaller than  $B^*$ .*

*Proof.* Let us assume by contradiction that there exists a truthful single-bid discounted budget-feasible mechanism with monitoring  $f$  such that for all  $i$  and for all  $b_i$  there exists  $\mathbf{b}_{-i}$  such that  $b_i(f(\mathbf{b})) = 0$ , and still  $\max_{\mathbf{b} \in D} \sum_{i=1}^n (p_i(\mathbf{b}) + h_i(b_i)) < B^*$ . Let  $\mathbf{b}^* \in \operatorname{argmax}_{\mathbf{b} \in D} \operatorname{cost}(f(\mathbf{b}), \mathbf{b})$ . Since  $\sum_{i=1}^n (p_i(\mathbf{b}^*) + h_i(b_i)) < B^*$ , there must exist an  $\hat{i}$  such that  $h_{\hat{i}}(b_{\hat{i}}^*) < 0$ . Let  $\bar{\mathbf{b}}_{-\hat{i}}$  be such that  $b_{\hat{i}}^*(f(b_{\hat{i}}^*, \bar{\mathbf{b}}_{-\hat{i}})) = 0$ . We have that  $h_{\hat{i}}(b_{\hat{i}}^*) < -b_{\hat{i}}^*(f(b_{\hat{i}}^*, \bar{\mathbf{b}}_{-\hat{i}}))$ , which violates (5).  $\square$

In the following sections we will also briefly discuss the application of single-bid discounts to set systems, facility location and obnoxious facility location. This serves the purpose of justifying our emphasis on bid-independent<sup>4</sup> discounts. Indeed, we show that more general discount functions such as single-bid discounts do not exhibit superior performances in terms of budget savings, whereas fully discounted payment functions (as we have already shown) may even impose restrictions on the set of possible algorithms.

## 5 Set systems

In a *set system*  $(E, \mathcal{F})$  we are given a set  $E$  of elements and a family  $\mathcal{F} \subseteq 2^E$  of feasible subsets of  $E$ . A set system is upward closed if for every  $S \in \mathcal{F}$  and every superset  $T$ , with  $S \subset T \subseteq E$ , we have  $T \in \mathcal{F}$ . Each element  $e \in E$  is associated with a non-negative cost function  $c : E \rightarrow \mathbb{R}_{\geq 0}$ ; for  $S \subseteq E$ , we let  $c(S) = \sum_{e \in S} c(e)$ . The goal is to find a feasible subset of minimum total cost. For these problems, we can then assume without loss of generality that  $(E, \mathcal{F})$  is upward closed and therefore that it is fully defined by the minimal sets (bases) in  $\mathcal{F}$ , defined as  $\mathcal{B} = \{S \in \mathcal{F} : \text{no proper subset of } S \text{ is in } \mathcal{F}\}$ . As noted above, several problems such as minimum spanning tree, shortest path (tree), etc. can be cast in this framework. To map this setting with the general mechanism design setting above, we have that  $\mathcal{O} = \mathcal{B}$  and the set of agents is  $E$  – in this context, for bidder  $i$ , we let  $t_i = t(i)$  denote  $i$ 's true type and  $b_i = b(i)$  a generic bid in  $D_i$ . In particular, we here have

$$t_i(S) = \begin{cases} t_i & \text{if } i \in S \\ 0 & \text{otherwise} \end{cases}$$

i.e., an element that does not belong to a set has no cost to “implement” that set. For example, an agent controlling an edge that does not belong to the minimum spanning tree computed by the mechanism has no cost, since she will not be part of the output. The extent to which monitoring can be implemented for this class of problems is discussed above in the introduction.

**Theorem 4** *Let  $f^*$  be an optimal algorithm for a set system. For all  $i$  and  $\mathbf{b}_{-i}$  there exists  $b_i$  such that  $b_i(f^*(\mathbf{b})) = 0$ .*

*Proof.* Fix  $i$  and  $\mathbf{b}_{-i}$ . Consider  $f^*(\mathbf{b}_{-i}) \in \mathcal{O}$  and let  $C_{-i} = c(f^*(\mathbf{b}_{-i}))$ . Let  $S$  be the set in  $\mathcal{O}$  including  $i$  with minimum cost and let  $C_{i=0} = \sum_{e \in S, e \neq i} b_e$ . If  $C_{-i} < C_{i=0}$  then no matter  $b_i$ ,  $i \notin f^*(\mathbf{b})$  and  $b_i(f^*(\mathbf{b})) = 0$ . If  $C_{-i} \geq C_{i=0}$  then by setting  $b_i = 0$ , no matter whether  $i$  is in  $f^*(\mathbf{b})$  or not, we have that  $b_i(f^*(\mathbf{b})) = 0$ .  $\square$

By Theorem 2, we have the following result.

**Corollary 2** *There exists no optimal truthful discounted first-price budget-feasible mechanism with monitoring for set systems using a budget that is smaller than  $\max_{\mathbf{b} \in D} \text{cost}(f(\mathbf{b}), \mathbf{b})$ .*

<sup>4</sup> For the sake of brevity, unless otherwise stated, in the remainder we will always refer to bid-independent discounted first-price mechanisms simply as discounted first-price mechanisms, whereas we will explicitly write single-bid discounted first-price mechanisms to refer to payment functions that employ single-bid discounts.

We now investigate single-bid discounts and prove that they perform as bad as bid-independent discounts.

**Proposition 2** *There exists no optimal truthful single-bid discounted first-price budget-feasible mechanism with monitoring for set systems using a budget that is smaller than  $B^*$ .*

*Proof.* By Theorem 3, to prove the claim we need to show that for all  $i$  and for all  $b_i$  there exists  $\mathbf{b}_{-i}$  such that  $b_i(f^*(\mathbf{b})) = 0$ . Let  $b_i = \delta > 0$  (for otherwise  $b_i(f^*(\mathbf{b})) = 0$  regardless of  $f^*(\mathbf{b})$ ). Without loss of generality, we can assume that there is a feasible solution  $S$  such that  $i \notin S$  (the set system is *monopoly-free*), for otherwise we can exclude  $i$  from the instance since no matter what she declares she is going to be selected (note that incentive-compatibility and the budget are not affected in our setting with monitoring). Let  $b_j = \frac{\delta - \varepsilon}{|S|}$  for each  $j \neq i$ , for some  $\delta > \varepsilon > 0$  arbitrarily small. But then we have that  $i \notin f^*(\mathbf{b})$  since  $\delta - \varepsilon = c(S) < \delta \leq c(S')$  for any feasible  $S'$  such that  $i \in S'$ .  $\square$

## 6 $K$ -facility location

In  $K$ -facility location (or simply facility location) the set of feasible solutions is comprised of all the  $K$ -tuples of possible allocations of the facilities, whilst the domain of each agent is the real line. For a given algorithm  $f$  and  $t_i \in D_i$ ,  $t_i(f(t_i, \mathbf{b}_{-i})) = |t_i - f_{t_i}(t_i, \mathbf{b}_{-i})|$ , where  $f_{t_i}(t_i, \mathbf{b}_{-i})$  denotes the location of the facility output by  $f(t_i, \mathbf{b}_{-i})$  closer to location  $t_i$ . In other words,  $t_i(f(t_i, \mathbf{b}_{-i}))$  denotes the distance between  $t_i$  and the location of  $f_{t_i}(t_i, \mathbf{b}_{-i})$ .

The following theorem proves that VCG mechanisms charge the agents rather than paying them in this setting.

**Theorem 5** *VCG payments for facility location are always non-positive.*

*Proof.* The VCG payment for agent  $i$  is defined as  $p_i(\mathbf{b}) = \text{cost}(f^*(\mathbf{b}_{-i}), \mathbf{b}_{-i}) - \text{cost}(f^*(\mathbf{b}), \mathbf{b}_{-i})$ , where  $f^*$  is an optimal algorithm. The thesis follows by observing that  $\text{cost}(f^*(\mathbf{b}_{-i}), \mathbf{b}_{-i}) \leq \text{cost}(f^*(\mathbf{b}), \mathbf{b}_{-i})$  by optimality.  $\square$

The fact that the VCG payments are always non-positive makes VCG a non-viable option for many real life incarnations of facility location, and motivates the study of monitoring in this setting. To give an example of such a scenario, consider a crowdsourcing data mining problem where we want to cluster/classify<sup>5</sup> a set of agents for market-segmentation purposes. The agents might lie in order to obtain a classification/clustering that more closely matches their own true labels/cluster membership, thereby creating a bias in the data. In this case we cannot assume that the agents will pay to take part to our market analysis study. For another example, consider a scenario similar to the one in [8] where the payments are interpreted as subsidies taken by the general taxation and paid by the government in order to support the adoption of environmentally sustainable public transportation systems (the facilities to be used by the agents) in a given urban area. Clearly, the subsidies here are meant to be a positive payment from the mechanism to the agents and not

<sup>5</sup> See [37] for the similarities among facility location and clustering/classification with strategic data sources.

vice versa. Finally, let us note that we can implement monitoring in this context whenever evidences of the cost can be provided (and cannot be counterfeited) – e.g., distance is measured in miles/time and/or receipts are needed for cost reimbursement, as from the discussion above.

### 6.1 Optimal mechanisms

We next prove that the social cost is a necessary budget for optimal mechanisms.

**Lemma 1** *For all  $i$ , and  $\mathbf{b}_{-i}$ , if  $b_i \in f^*(\mathbf{b}_{-i})$  then*

$$\text{cost}(f^*(b_i, \mathbf{b}_{-i}), (b_i, \mathbf{b}_{-i})) = \text{cost}(f^*(\mathbf{b}_{-i}), \mathbf{b}_{-i})$$

*Proof.* Let us first assume that:

$$\text{cost}(f^*(b_i, \mathbf{b}_{-i}), (b_i, \mathbf{b}_{-i})) < \text{cost}(f^*(\mathbf{b}_{-i}), \mathbf{b}_{-i}). \quad (6)$$

Then, the following holds:

$$\text{cost}(f^*(b_i, \mathbf{b}_{-i}), \mathbf{b}_{-i}) \leq \text{cost}(f^*(b_i, \mathbf{b}_{-i}), (b_i, \mathbf{b}_{-i}))$$

which combined with (6), yields:

$$\text{cost}(f^*(b_i, \mathbf{b}_{-i}), \mathbf{b}_{-i}) < \text{cost}(f^*(\mathbf{b}_{-i}), \mathbf{b}_{-i})$$

thus contradicting the optimality of  $f^*$ . We observe that

$$\text{cost}(f^*(b_i, \mathbf{b}_{-i}), (b_i, \mathbf{b}_{-i})) > \text{cost}(f^*(\mathbf{b}_{-i}), \mathbf{b}_{-i})$$

does not hold either, as otherwise  $f^*(\mathbf{b}_{-i})$  would be a better solution for  $(b_i, \mathbf{b}_{-i})$  than  $f^*(b_i, \mathbf{b}_{-i})$ . Indeed,  $\text{cost}(f^*(\mathbf{b}_{-i}), (b_i, \mathbf{b}_{-i})) = \text{cost}(f^*(\mathbf{b}_{-i}), \mathbf{b}_{-i})$  holds, since  $b_i(f^*(\mathbf{b}_{-i})) = 0$ .  $\square$

**Theorem 6** *Let  $f^*$  be an optimal algorithm for  $K$ -facility location. Then for all  $i$  and  $\mathbf{b}_{-i}$  there exists  $b_i$  such that  $b_i(f^*(\mathbf{b})) = 0$ .*

*Proof.* Fix  $i$  and  $\mathbf{b}_{-i}$  and let  $\mathbf{y} = \{y_1, \dots, y_K\} = f^*(\mathbf{b}_{-i})$  be the optimal allocation for bid vector  $\mathbf{b}_{-i}$ . We set  $b_i = y_j$  for a generic  $j \in \{1, \dots, K\}$ . We need to prove that  $b_i \in f^*(b_i, \mathbf{b}_{-i})$ , from which the theorem follows, since in this case  $b_i(f^*(b_i, \mathbf{b}_{-i})) = 0$ . For the sake of contradiction, let us suppose that  $b_i \notin f^*(b_i, \mathbf{b}_{-i})$ , which implies that  $b_i(f^*(b_i, \mathbf{b}_{-i})) > 0$ . Then  $f^*(b_i, \mathbf{b}_{-i})$  is a better allocation for  $\mathbf{b}_{-i}$  than  $f^*(\mathbf{b}_{-i})$ , which contradicts the optimality of  $f^*$ . Indeed, the following holds:  $\text{cost}(f^*(b_i, \mathbf{b}_{-i}), \mathbf{b}_{-i}) < \text{cost}(f^*(b_i, \mathbf{b}_{-i}), (b_i, \mathbf{b}_{-i})) = \text{cost}(f^*(\mathbf{b}_{-i}), \mathbf{b}_{-i})$  where the inequality holds because  $b_i(f^*(\mathbf{b})) > 0$  by hypothesis, and the equality follows from Lemma 1.<sup>6</sup>  $\square$

By Theorem 2 we have the following result.

**Corollary 3** *There is no optimal truthful discounted first-price budget-feasible mechanism with monitoring for facility location whose budget is smaller than  $B^*$ .*

<sup>6</sup> A simpler proof works when agents have unrestricted domains (i.e., simply set  $b_i$  big enough with respect to  $\mathbf{b}_{-i}$  so to force  $f^*$  to locate one facility on  $b_i$ ). Our argument applies also to more restricted settings as considered in, e.g., [6].

## 6.2 Collusion-resistant mechanisms

An algorithm  $f$  is *simple*, if  $f(\mathbf{b}) \subseteq \mathbf{b}$ , i.e. the facilities are allocated only at agents' locations. Observe that the focus on simple algorithms is without loss of generality from the approximation guarantee point of view. Moreover, as already discussed in [20], this is the most natural class of (deterministic) algorithms for the problem. An algorithm is *anonymous* if its output does not change if agents swap declarations. Formally, let  $\pi : [n] \rightarrow [n]$  be a permutation of the agents and let  $\mathbf{b}^\pi$  be the declaration vector  $\mathbf{b}$  where entries are permuted according to  $\pi$ ;  $f$  is anonymous if  $f(\mathbf{b}) = f(\mathbf{b}^\pi)$ .

**Theorem 7** *Let  $f$  be a simple and anonymous algorithm for facility location. Then for all  $i$  there exists  $\mathbf{b}$  such that  $b_i(f(\mathbf{b})) = 0$ .*

*Proof.* In order to prove the claim, we need to prove that  $\forall i, \exists \mathbf{b}_{-i} \exists b_i$  such that  $b_i(f(\mathbf{b})) = 0$ , for all simple anonymous mechanisms  $f$ . Since  $f$  is simple, for all  $\mathbf{b}$  there exists an agent  $j$  such that  $b_j(f(\mathbf{b})) = 0$ . For all  $i \neq j$  we construct  $\mathbf{b}^\pi$  from  $\mathbf{b}$  where we swap the declarations of agent  $i$  and agent  $j$ :  $b_i^\pi = b_j$ ,  $b_j^\pi = b_i$  and  $b_\ell^\pi = b_\ell \forall \ell \notin \{i, j\}$ . Since  $f$  is anonymous,  $f(\mathbf{b}) = f(\mathbf{b}^\pi)$  and  $b_i^\pi(f(\mathbf{b}^\pi)) = b_j(f(\mathbf{b})) = 0$ .  $\square$

By Theorem 2 we have the following result.

**Corollary 4** *There is no collusion-resistant discounted first-price budget-feasible mechanism with monitoring for facility location that utilizes a simple and anonymous algorithm  $f$  whose budget is smaller than  $B^*$*

## 6.3 Trading approximation for frugality

Given an optimal algorithm  $f^*$ , we define  $f_\varepsilon^*$  as the algorithm that shifts<sup>7</sup> by  $\varepsilon$  the location of the facilities returned by  $f^*$ . Formally, let  $f_\varepsilon^*(\mathbf{b}) = (F_j + \varepsilon)_{j=1, \dots, K}$ , where  $(F_1, \dots, F_K)$  denotes the output of  $f^*(\mathbf{b})$ . We ask whether, by moving the allocation from  $f^*$  to  $f_\varepsilon^*$ , the cost paid in terms of approximation guarantee can be compensated by a lower budget.

It is well known that the optimal solution for an instance  $\mathbf{b}$  partitions the agents into  $\{S_1, \dots, S_K\}$  so that all  $i \in S_j$  are served by  $F_j$ , defined as a median of  $S_j$ . Our analysis assumes that  $\varepsilon$  is “small enough”; this intuitively means that the optimal partitions of  $f^*(\mathbf{b})$  do not change when agents' locations and the locations of the facilities shift by  $\varepsilon$ . Formally, let  $\mathbf{b}_{\ell, k}^\varepsilon = (b_k + \varepsilon, \mathbf{b}_{-\ell})$  for some  $k \in \{1, \dots, n\}$ . In words, in  $\mathbf{b}_{\ell, k}^\varepsilon$  we set the location of agent  $\ell$  to  $b_k + \varepsilon$  (i.e.,  $\varepsilon$  to the right of the location of some other agent  $k$ ). Let  $\{S_1^\varepsilon, \dots, S_K^\varepsilon\}$  and  $\{S_1^{\varepsilon, \ell, k}, \dots, S_K^{\varepsilon, \ell, k}\}$  be the optimal partitions with respect to the allocations  $f^*(\mathbf{b}_{\ell, k}^\varepsilon)$  and  $f_\varepsilon^*(\mathbf{b}_{\ell, k}^\varepsilon)$ ,

<sup>7</sup> We justify our focus on  $f_\varepsilon^*$  by highlighting how useful a tool it is to study the tradeoff between frugality and approximation in this setting, as it allows us to both: (i) easily predict the location of the facility; and (ii) compute the approximation ratio achievable given the particular configuration of the instance at hand (see Theorem 8). If we were to focus, for instance, on the class of mechanisms having a given approximation guarantee, we would certainly retain (ii) but we would have to restrict ourselves to specific algorithms within this class to also retain (i).

respectively. We say that  $\varepsilon$  is “small enough” if, for all  $\ell, k \in \{1, \dots, n\}$ , the following conditions hold: (i)  $S_j = S_j^\varepsilon$  for all  $j \in \{1, \dots, K\}$  and (ii)  $S_j = S_j^{\varepsilon, \ell, k}$  for all  $j \in \{1, \dots, K\}$ .

The following partition of the agents of a given instance is useful.

**Definition 3** Let  $(F_1, \dots, F_k)$  denote the output of  $f^*$  on input  $\mathbf{b}$ . For all  $1 \leq j \leq K$ , we partition the agents in the following sets:

- $\Delta_1^j = \{i | b_i \in (\max \Delta_3^{j-1}, F_j]\}$ ,  $|\Delta_1^j| = \delta_1^j$ ;
- $\Delta_2^j = \{i | b_i \in (F_j, F_j + \varepsilon]\}$ ,  $|\Delta_2^j| = \delta_2^j$ ;
- $\Delta_3^j = \{i | b_i \in (F_j + \varepsilon, \min \Delta_1^{j+1})\}$ ,  $|\Delta_3^j| = \delta_3^j$ ,

where, for notational convenience, we set  $\max_{\ell \in \Delta_3^0} b_\ell = -\infty$  and  $\min_{\ell \in \Delta_1^{K+1}} b_\ell = \infty$ . We let  $\Delta_i = \cup_j \Delta_i^j$  and  $\delta_i = \sum_j \delta_i^j$ .

Informally,  $\Delta_1^j$  denotes all the agents that are located to the left of  $F_j$  but to the right of the rightmost agent in  $\Delta_3^{j-1}$ ,  $\Delta_2^j$  denotes the agents that are located between  $F_j$  and  $F_j + \varepsilon$  (inclusive) and  $\Delta_3^j$  denotes the agents that are located to the right of  $F_j + \varepsilon$  but to the left of the leftmost agent in  $\Delta_1^{j+1}$ .

We now prove the approximation guarantee of algorithm  $f_\varepsilon^*$ .

**Theorem 8** *If  $\varepsilon$  is small enough, then  $f_\varepsilon^*(\mathbf{b})$  is  $((\delta_1 + \delta_2)\varepsilon)$ -approximate. More specifically, we have  $(\delta_1 + \delta_2)\varepsilon > \text{cost}(f_\varepsilon^*(\mathbf{b}), \mathbf{b}) - \text{cost}(f^*(\mathbf{b}), \mathbf{b}) \geq (\delta_1 - \delta_2 - \delta_3)\varepsilon$ .*

*Proof.* Fix  $\mathbf{b}$  and  $j \in \{1, \dots, K\}$ . We first observe that, since  $\varepsilon$  is small enough for the allocation  $f_\varepsilon^*(\mathbf{b})$ , the agents in  $\Delta_1^j$  have to pay the same cost they paid for  $f^*(\mathbf{b})$  plus  $\varepsilon$ . For the lower bound to  $\text{cost}(f_\varepsilon^*(\mathbf{b}), \mathbf{b}) - \text{cost}(f^*(\mathbf{b}), \mathbf{b})$ , we similarly observe that for the solution  $f_\varepsilon^*(\mathbf{b})$ , all the agents in  $\Delta_2^j \cup \Delta_3^j$  can save (at most)  $\varepsilon$  on their cost to implement  $f^*(\mathbf{b})$ . (That is, the agents in  $\Delta_3^j$  save exactly  $\varepsilon$  whilst those in  $\Delta_2^j$  can save less than that.) For the upper bound, instead, we note that agents in  $\Delta_2^j$  might need to pay at most  $\varepsilon$  more in  $f_\varepsilon^*(\mathbf{b})$  than in  $f^*(\mathbf{b})$  (when their location is very close to  $F_j$ ); in the worst case, the agents in  $\Delta_3^j$  will pay exactly as much in  $f_\varepsilon^*(\mathbf{b})$  as in  $f^*(\mathbf{b})$  (when their location is very close to  $F_j + \varepsilon$ ). By summing up on all  $j$ , we prove the claim.  $\square$

We next show that, in this context, there cannot be a tradeoff between approximation quality and frugality of mechanisms. We begin by establishing a useful property of certain optimal algorithms. We let  $f^*$  be the optimal algorithm for facility location that uses a *fixed tie-breaking rule*, i.e., for every pair of bid vectors  $\mathbf{b}, \mathbf{b}'$ , if both  $\text{cost}(f^*(\mathbf{b}), \mathbf{b}) = \text{cost}(f^*(\mathbf{b}'), \mathbf{b})$  and  $\text{cost}(f^*(\mathbf{b}'), \mathbf{b}') = \text{cost}(f^*(\mathbf{b}), \mathbf{b}')$  hold, then  $f^*(\mathbf{b}) = f^*(\mathbf{b}')$ . It is easy to check that an optimal algorithm with fixed tie-breaking always exists. For example, the rule that chooses the lexicographically minimal allocation among all optimal allocations satisfies the property above.

**Lemma 2** *Let  $f^*$  be an optimal algorithm for facility location that uses a fixed tie-breaking rule. Then for all  $i$ , all  $\mathbf{b}_{-i}$  and all  $b_i$  such that  $b_i \in f^*(\mathbf{b}_{-i})$ , the following holds:  $f^*(b_i, \mathbf{b}_{-i}) = f^*(\mathbf{b}_{-i})$ .*

*Proof.* Let us fix  $i$  and  $\mathbf{b}_{-i}$  and consider the output of  $f^*(\mathbf{b}_{-i})$ . We claim the following hold:

$$\text{cost}(f^*(\mathbf{b}_{-i}), \mathbf{b}_{-i}) = \text{cost}(f^*(\mathbf{b}), \mathbf{b}_{-i}) \quad (7)$$

$$\text{cost}(f^*(\mathbf{b}_{-i}), \mathbf{b}) = \text{cost}(f^*(\mathbf{b}), \mathbf{b}), \quad (8)$$

by which, we get  $f^*(\mathbf{b}_{-i}) = f^*(\mathbf{b})$  because of the fixed-tie breaking rule. Indeed, for (7) we have that

$$\text{cost}(f^*(\mathbf{b}_{-i}), \mathbf{b}_{-i}) \leq \text{cost}(f^*(\mathbf{b}), \mathbf{b}_{-i}) \leq \text{cost}(f^*(\mathbf{b}), \mathbf{b})$$

where the first inequality follows from optimality of  $f^*(\mathbf{b}_{-i})$  and the second inequality follows from non-negativity of costs.

Equation (7) then follows since  $\text{cost}(f^*(\mathbf{b}_{-i}), \mathbf{b}_{-i}) = \text{cost}(f^*(\mathbf{b}), \mathbf{b})$  by Lemma 1. Equation (8) holds by Lemma 1, because  $\text{cost}(f^*(\mathbf{b}_{-i}), \mathbf{b}_{-i}) = \text{cost}(f^*(\mathbf{b}), \mathbf{b})$  and  $b_i(f^*(\mathbf{b}_{-i})) = 0$  by definition.  $\square$

We are now ready to state and prove the following.

**Theorem 9** *Let  $f^*$  be an optimal algorithm for  $K$ -facility location that returns the lexicographically minimal (or maximal) allocation among all optimal allocations. For  $\varepsilon$  small enough, let  $f_\varepsilon^*$  be defined as above upon  $f^*$ . If  $n > 2K$ , then for all  $i$  and  $\mathbf{b}_{-i}$  there exists  $b_i$  such that  $b_i(f_\varepsilon^*(\mathbf{b})) = 0$ .*

*Proof.* We start from a vector  $\mathbf{b}$  and use  $f^*(\mathbf{b})$  to prove that for all agents  $i$  there exists a  $b'_i \in D_i$  such that  $b'_i(f_\varepsilon^*(b'_i, \mathbf{b}_{-i})) = 0$ .

Fix  $i$  and let  $S$  be the member of the partition of agents induced by  $f^*(\mathbf{b})$  to which  $i$  belongs. We let  $F$  denote the median of  $S$  and, with a slight abuse of notation,  $\Delta_1$ ,  $\Delta_2$  and  $\Delta_3$  be the partitions of the agents in  $S$  as from Definition 3.

We begin by considering the case that  $|S| > 1$  is odd. In this case, the optimal location is  $F = b_k$  for some agent  $k$ . If  $i \in \Delta_2 \cup \Delta_3$ , then both  $F$  and  $F + \varepsilon$  do not change as long as agent  $i$  declares a location that does not change the optimal partition and is to the right of  $F$ . Hence, since  $\varepsilon$  is small enough, the claim is proved for  $b'_i = F + \varepsilon$ . Let us now consider the case when  $i \in \Delta_1$ . Let  $k$  denote the median agent of  $S$  (breaking ties by lexicographic order if there is more than one agent at the median location) and let  $\ell$  be the leftmost agent such that  $F \leq b_\ell$  and  $\ell \neq k$ . If  $i$  declares  $b'_i = b_\ell + \varepsilon$ , then since  $\varepsilon$  is small enough, we have  $F = b_\ell$  and the claim is proved.

We now consider the case that  $|S| > 2$  is even. In this case,  $S$  has two medians; we denote their locations as  $a$  and  $b$ , with  $a \leq b$ , respectively. Note that, by hypothesis,  $f^*$  returns either  $a$  or  $b$  (in general, all the locations in  $[a, b]$  are optimal for  $S$ ). An argument similar to the case when  $|S| > 1$  and  $|S|$  is odd holds in this case, and we omit it for the sake of brevity. Finally, we now deal with  $|S| \leq 2$ . In this case, we look at the biggest partition  $S_j$  output of  $f^*(\mathbf{b}_{-i})$  and set  $b'_i = F_j$ . As  $n - 1 \geq 2K$  then  $|S_j| \geq \lceil \frac{2K-1}{K} \rceil \geq 2$ . Moreover, since  $f^*$  uses a fixed tie-breaking rule (as noted above, returning lexicographically minimal/maximal solutions is a particular fixed-tie breaking rule) we can conclude by Lemma 2 that  $f^*(\mathbf{b}_{-i}) = f^*(b'_i, \mathbf{b}_{-i})$ . Hence, the partition  $S_j$  will have now one more agent than before. But then putting things together we note that  $i$  belongs to a set of the partition induced by  $f^*(b'_i, \mathbf{b}_{-i})$  with at least 3 agents. We can then apply recursively on  $(b'_i, \mathbf{b}_{-i})$  the arguments above.  $\square$

By Theorem 2 we have the following result.

**Corollary 5** *There is no truthful discounted first-price budget-feasible mechanism with monitoring for facility location, that uses  $f_\varepsilon^*$  defined upon an optimal algorithm that returns lexicographically minimal (or maximal) optimal allocation whose budget is smaller than  $\max_{\mathbf{b} \in D} \text{cost}(f_\varepsilon^*(\mathbf{b}), \mathbf{b})$ .*

#### 6.4 Single-bid Discounts

We now investigate single-bid discounts and prove that they perform as bad as bid-independent discounts.

**Proposition 3** *There is no truthful single-bid discounted first price mechanism for the facility location problem that uses a budget lower than  $B^*$  and has bounded multiplicative approximation ratio.*

*Proof.* Any algorithm  $f$  with bounded multiplicative approximation ratio satisfies Theorem 3, i.e., for all  $i$  and for all  $b_i$  there exists  $\mathbf{b}_{-i}$  such that  $b_i(f(\mathbf{b})) = 0$ . Indeed, given  $i$  and  $b_i$ , we construct an instance where all agents are located at  $b_i$ , i.e.,  $b_j = b_i, \forall j$ . Note that the optimum algorithm allocates at least one facility at  $b_i$ , hence  $b_i(f(\mathbf{b})) = 0$ . Furthermore, since the optimal allocation has a cost of 0, all allocations such that  $b_i(f(\mathbf{b})) > 0$  (i.e., allocate no facilities at  $b_i$ ) have unbounded approximation ratio.  $\square$

**Proposition 4** *There is no truthful single-bid discounted first-price budget-feasible mechanism with monitoring for facility location that uses  $f_\varepsilon^*$  whose budget is smaller than  $\max_{\mathbf{b} \in D} \text{cost}(f_\varepsilon^*(\mathbf{b}), \mathbf{b})$ .*

*Proof.* By Theorem 3, we need to prove that for all  $i$  and for all  $b_i$  there exists a  $\mathbf{b}_{-i}$  such that  $b_i(f_\varepsilon^*(\mathbf{b})) = 0$ . We assume that  $n \geq 2K + 1$ . Given  $b_i$ , construct  $\mathbf{b}_{-i}$  as follows. We define  $K$  groups of  $\lfloor \frac{n-1}{K} \rfloor$  bidders, such that each group is located at  $b_i - j \cdot \varepsilon$ , for  $j = 1, \dots, K$ . Each of the remaining  $q = n - 1 - K \cdot \lfloor \frac{n-1}{K} \rfloor$  agents (if any) is located to  $b_i - j \cdot \varepsilon$ , for  $j = 1, \dots, q$ . Now it is easy to check that the optimal algorithm locates the  $K$  facilities at  $b_i - j \cdot \varepsilon$  for  $j = 1, \dots, K$ , while  $f_\varepsilon^*$  will locate one facility at  $b_i$ , hence  $b_i(f_\varepsilon^*(\mathbf{b})) = 0$ .  $\square$

### 7 Obnoxious Facility Location

In the obnoxious facility location problem we have to locate one obnoxious facility on a fixed interval in such a way as to minimize the social cost incurred by the agents. The set of feasible solutions is the set of points in the interval, namely  $\mathcal{O} = \{x \in R \mid 0 \leq x \leq \ell\}$ , where  $\ell$  denotes the length of the interval.<sup>8</sup> For a given algorithm  $f$  and  $t_i \in D_i$ ,  $t_i(f(t_i, \mathbf{b}_{-i})) = \ell - |t_i - f(t_i, \mathbf{b}_{-i})|$ . It is well known from [14] that an optimal algorithm allocates the facility either at 0 or at  $\ell$ . In particular, if  $\sum_{i=1}^n t_i \geq \sum_{i=1}^n (\ell - t_i)$  then the facility is allocated at 0, otherwise it is allocated at  $\ell$ .

<sup>8</sup> We note that unless the length of the interval is specified, the problem is not well-defined, as a solution with lower social cost can always be obtained by moving the facility farther away.

Just like for facility location, VCG payments would charge the agents (it is easy to check that Theorem 5 holds for obnoxious facility location as well), and an analog argument suggests that VCG payments are unsuitable in some scenarios where we cannot assume that agents will have to pay to participate in the game.

We then study the budget of truthful discounted first-price mechanisms with monitoring and show that the social cost is a necessary budget for the optimal mechanism.

**Theorem 10** *Let  $f^*$  be an optimal algorithm for the obnoxious facility location problem. For all  $i$  and for all  $\mathbf{b}_{-i}$  there exists  $b_i$  such that  $b_i(f^*(\mathbf{b})) = 0$ .*

*Proof.* Let us consider a generic  $\mathbf{b}_{-i}$  and let us assume first that  $\sum_{j \neq i} b_j > \sum_{j \neq i} (\ell - b_j)$ . If this is the case, then  $f^*(\mathbf{b}_{-i}) = 0$ . If we choose  $b_i = \ell$ , it is easy to see that  $\sum_{j=1}^n b_j > \sum_{j=1}^n (\ell - b_j)$  still holds and hence  $f^*(\mathbf{b}) = 0$ . In this case,  $b_i(f(\mathbf{b})) = 0$ . A similar argument holds if  $\sum_{j \neq i} b_j < \sum_{j \neq i} (\ell - b_j)$  and  $f^*(\mathbf{b}_{-i}) = \ell$ . The remaining case to consider is  $\sum_{j \neq i} b_j = \sum_{j \neq i} (\ell - b_j)$ . This can only occur when  $\mathbf{b}_{-i}$  is symmetric with respect to  $\frac{\ell}{2}$ . However, irrespective of the output of  $f^*(\mathbf{b}_{-i})$ , by setting  $b_i = \ell$ , we have  $\sum_{j=1}^n b_j > \sum_{j=1}^n (\ell - b_j)$  and hence  $f^*(\mathbf{b}) = 0$ , implying  $b_i(f(\mathbf{b})) = 0$ .  $\square$

By Theorem 2 we have the following result.

**Corollary 6** *There is no truthful discounted first-price budget-feasible mechanism with monitoring for obnoxious facility location, that uses an optimal algorithm  $f^*$ , whose budget is smaller than  $B^*$*

### 7.1 Trading approximation for frugality

Given an optimal algorithm  $f^*$ , we define  $f_\varepsilon^*$  as the algorithm that shifts by  $\varepsilon$  the location of the facility returned by  $f^*$ . In particular, if  $f^*(\mathbf{b}) = 0$  then  $f_\varepsilon^*(\mathbf{b}) = \varepsilon$ , whereas if  $f^*(\mathbf{b}) = \ell$  then  $f_\varepsilon^*(\mathbf{b}) = \ell - \varepsilon$ . Unlike the case of facility location, the shift of the facility operated by  $f_\varepsilon^*$  is not unidirectional: this is to prevent the facility from being allocated outside of the interval. We next prove that for  $f_\varepsilon^*$  savings are possible.

**Theorem 11** *For all agents  $i$  and for all  $\mathbf{b}_{-i}$  there is no  $b_i$  such that  $b_i(f_\varepsilon^*(\mathbf{b})) = 0$ .*

*Proof.* It is easy to see that  $b_i(f_\varepsilon^*(\mathbf{b})) = 0$  only if  $|b_i - f_\varepsilon^*(\mathbf{b})| = \ell$ , which is the case only when the facility is located at one extreme of the interval and  $b_i$  is at the other extreme. For  $\varepsilon > 0$ , this case can never occur, as  $f_\varepsilon^*(\mathbf{b}) \notin \{0, \ell\}$ .  $\square$

**Definition 4** Let  $f^* = f^*(\mathbf{b})$  and  $f_\varepsilon^* = f_\varepsilon^*(\mathbf{b})$ . We define the following partition of agents:  $\Delta^+ = \{i \mid b_i(f_\varepsilon^*) > b_i(f^*)\}$ ,  $\delta^+ = |\Delta^+|$ ;  $\Delta^- = \{i \mid b_i(f_\varepsilon^*) < b_i(f^*)\}$ ,  $\delta^- = |\Delta^-|$ ; and  $\Delta^\pm = \{i \mid b_i(f_\varepsilon^*) = b_i(f^*)\}$ ,  $\delta^\pm = |\Delta^\pm|$ .

Put in words,  $\Delta^+$ ,  $\Delta^-$  and  $\Delta^\pm$  denote the set of agents whose cost with respect to the allocation  $f_\varepsilon^*$  has, respectively, increased, decreased or stayed the same when compared with the optimal allocation  $f^*$ . Of course,  $\delta^+ + \delta^- + \delta^\pm = n$  holds.

Lemma 3 proves that  $f_\varepsilon^*$  is  $((\delta^+ - \delta^-) \cdot \varepsilon)$ -approximate; however, as proved in Theorem 12, this loss in approximation can be recovered by the budget savings possible.

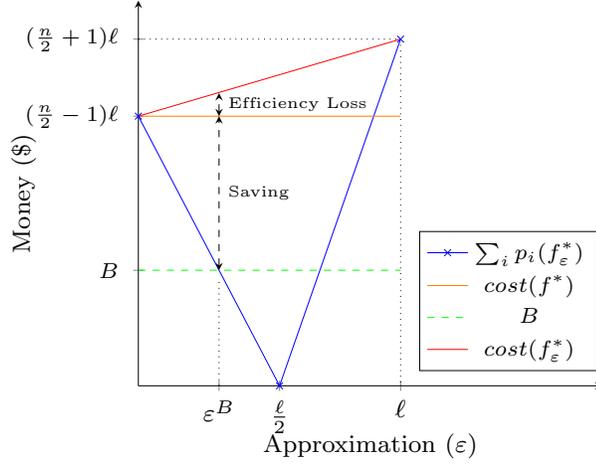


Fig. 1 Approximation vs Frugality for Obnoxious Facility

**Lemma 3** Let  $f^*$  be an optimal allocation algorithm that always locates the facility at one extreme of the interval<sup>9</sup> and let  $f_\varepsilon^*$  be the shifted allocation function defined upon  $f^*$ . It holds:  $cost(f_\varepsilon^*(\mathbf{b}), \mathbf{b}) = cost(f^*(\mathbf{b}), \mathbf{b}) + (\delta^+ - \delta^-) \cdot \varepsilon$ .

*Proof.* By definition, the only agents whose cost varies when the facility is allocated according to  $f_\varepsilon^*$  instead of  $f^*$  are those in  $\Delta^+$  and  $\Delta^-$ . In particular, their cost changes of exactly  $\varepsilon$  with respect to the cost they incur when the facility is allocated optimally. The following equality follows from the simple observation above:  $cost(f_\varepsilon^*(\mathbf{b}), \mathbf{b}) = cost(f^*(\mathbf{b}), \mathbf{b}) + (\delta^+ - \delta^-) \cdot \varepsilon$ .  $\square$

**Theorem 12** For  $0 < \varepsilon < \frac{\ell}{2}$ , algorithm  $f_\varepsilon^*$  achieves a budget saving of  $\varepsilon(n + \delta^- - \delta^+)$  with respect to  $f^*$ . For  $\frac{\ell}{2} < \varepsilon < \ell$  algorithm  $f_\varepsilon^*$  achieves a budget saving of  $n\ell - (n + \delta^- - \delta^+)\varepsilon$  with respect to  $f^*$ .

*Proof.* By Theorem 10, the budget needed to implement  $f^*$  is  $\sum_i b_i(f^*)$ . By Theorem 11 and Lemma 3, the budget needed to implement  $f_\varepsilon^*$  is  $cost(f^*(\mathbf{b}), \mathbf{b}) + (\delta^+ - \delta^-)\varepsilon + n \cdot h_i(\mathbf{b}_{-i})$  where  $h_i(\mathbf{b}_{-i})$  is defined in (4).

For  $0 < \varepsilon \leq \frac{\ell}{2}$ ,  $h_i(\mathbf{b}_{-i}) = -\varepsilon$ . This yields a budget saving of  $(n + \delta^- - \delta^+)\varepsilon$ . We observe that the saving is always non-negative (e.g., the budget required to implement  $f_\varepsilon^*$  is never higher than the one required to implement  $f^*$ ) since  $n \geq \delta^+$ .

For  $\frac{\ell}{2} < \varepsilon < \ell$ ,  $h_i(\mathbf{b}_{-i}) = \varepsilon - \ell$ . This yields a budget saving of  $n\ell - (n + \delta^- - \delta^+)\varepsilon$ . We observe that the saving in this case can be negative (e.g., the budget required to implement  $f_\varepsilon^*$  can be higher than the one required to implement  $f^*$ ) depending on the value of  $\varepsilon$  and on the structural configuration of the instance at hand.  $\square$

Theorem 12 provides a way to compute the budget needed to implement a solution with a desired level of approximation (with respect to the optimal solution) or, put otherwise, to estimate in advance the quality of service (in terms of

<sup>9</sup> Note that we are assuming a fixed tie-breaking rule when there are multiple optimal allocations (i.e.,  $n$  is even and  $n/2$  agents are at 0 and  $n/2$  agents at  $\ell$ ). We assume that the algorithm will allocate consistently the facility at one extreme of the interval.

approximation) we can afford given a set budget. To illustrate this, Figure 1 refers to an instance  $\mathbf{b}$  of obnoxious facility location where  $\frac{n}{2} + 1$  agents reside at 0 and  $\frac{n}{2} - 1$  agents reside at  $\ell$ . The horizontal and vertical axes represent, respectively, the shift  $\varepsilon$  controlling the approximation, and the monetary cost of implementing the solution. The optimal solution for this instance is to allocate the facility at  $\ell$ , at a cost  $(\frac{n}{2} - 1)\ell$ , which is above the budget  $B$  and cannot be implemented. To find out the level of efficiency  $\varepsilon^B$  we can afford, we just need to find the leftmost intersection between the budget line  $B$  and the cost of implementing  $f_\varepsilon^*$ , denoted  $\sum_i p_i(f_\varepsilon^*)$ . In this case,  $\sum_i p_i(f_\varepsilon^*)$  is constituted of two segments, defined as:

$$\sum_i p_i(f_\varepsilon^*) = \begin{cases} (\frac{n}{2} - 1)\ell + (\delta^+ - \delta^- - n)\varepsilon & \text{if } 0 < \varepsilon \leq \frac{\ell}{2} \\ -\frac{n+2}{2}\ell + (\delta^+ - \delta^- + n)\varepsilon & \text{if } \frac{\ell}{2} < \varepsilon < \ell \end{cases}$$

In the general case,  $\sum_i p_i(f_\varepsilon^*)$  is piece-wise linear, where the various segments originate both from  $\varepsilon$  ranging over  $(0, \ell)$  as the location of the facility is shifted, and from the change of  $\delta^+$ ,  $\delta^-$  and  $\delta^-$ , that, in turn modifies the slope of the line segments. In this particular case, the sets  $\Delta^+$ ,  $\Delta^-$  and  $\Delta^-$  do not change since the agents reside at 0 and  $\ell$ . We note that, consistently with Figure 1, for  $\varepsilon \in (0, \frac{\ell}{2}]$  the slope of the curve  $(\delta^+ - \delta^- - n)$  is always negative, whereas for  $\varepsilon \in (\frac{\ell}{2}, \ell)$  the slope of the curve  $(\delta^+ - \delta^- + n)$  is always positive. The efficiency loss we incur in implementing  $f_{\varepsilon^B}^*(\mathbf{b})$  can be computed analytically as  $cost(f_{\varepsilon^B}^*(\mathbf{b}), \mathbf{b}) - cost(f^*(\mathbf{b}), \mathbf{b})$ , where  $cost(f_{\varepsilon^B}^*(\mathbf{b}), \mathbf{b})$  can be computed using Lemma 3.

## 7.2 Single-bid Discounts

Single-bid discounts can potentially lead to savings just as bid-independent discounts. Indeed, for the optimal algorithm, note that for any  $b_i \in (0, \ell)$  we cannot define a  $\mathbf{b}_{-i}$  for which agent  $i$  has cost 0. In particular, this means that the discount for player  $i$  ( $h_i(b_i) = \inf_{\mathbf{b}_{-i}} b_i(f(\mathbf{b}))$ ) is  $\min\{b_i, \ell - b_i\}$ . If  $b_i$  is on the left (right) of  $\ell/2$  then the worst  $\mathbf{b}_{-i}$  we can define has all the agents located at 0 ( $\ell$ ). This in particular means that there are instances that can be implemented for free (all agents on  $\ell/2$ ) and instances for which we need the optimal social cost as budget ( $b_i \in \{0, \ell\}$  for all  $i$ ).

## 8 Conclusions

We study budget-feasible mechanisms with monitoring for utilitarian cost-minimization problems. Our results suggest the amount of financial resources a designer needs to acquire when facing these kinds of problems. We show that the social cost is a sufficient and, in many cases, necessary budget to implement incentive-compatible mechanisms. We also study a case wherein the social cost is not necessary to enforce truthfulness, and more frugal mechanisms are possible, thus allowing a certain degree of freedom in terms of resources accumulation (paid for with a loss in approximation guarantee). In such a case, plots like the one in Figure 1 can be rather informative for decision makers facing budgeting issues. Our results can also be read in terms of frugality: while the overpayment implied by truthfulness is linear in the cost of the second-best solution, we do prove that there is no

overpayment at all when monitoring can be implemented. In particular, we find the necessity (in many cases) of having a budget  $C$  to implement a solution of cost  $C$  to be an element that introduces some fairness in the dichotomy between the designer and the agents, as there is no side with an advantage. This is in sharp contrast with classical mechanisms wherein agents have the upper hand over the designer.

Our research agenda is very much linked to the concept of *extremal transfers* studied in the literature on truthful mechanisms [31], i.e., the payments that charge the maximum/minimum amount to the agents. We here care only about the minimum and show that for a number of problems/algorithms the extremal transfers cannot be lower than the social cost. Clearly, our results are a first attempt at characterizing extremal transfers for incentive compatibility with monitoring and show an alternative technique to the graph-theoretic reasoning used for classical truthfulness.

There remain many open questions to address in future research efforts. Firstly, it would be interesting to study other problems (e.g., machine scheduling) where our results can have interesting applications. It would also be important to know what payment functions, other than discounts of first prices, exist and possibly characterize them. As discussed above, this would restrict the class of algorithms that can be used in truthful mechanisms, thus raising the question of establishing a trade-off between incentive-compatible allocation rules and corresponding payment functions. More generally, our research suggests that even though the vast majority of the mechanism design literature focuses more on the allocation rules, and less on the payments, there are many unanswered questions about the computation of truthful payments and the ability to find the “best” payments for a given algorithm/problem.

## References

1. N. Alon, M. Feldman, A. D. Procaccia, and M. Tennenholtz. Strategyproof approximation of the minimax on networks. *Mathematics of Operations Research*, 3:513–526, 2010.
2. G. Amanatidis, G. Birmpas, and E. Markakis. On budget-feasible mechanism design for symmetric submodular objectives. In *Web and Internet Economics - 13th International Conference, WINE 2017, Bangalore, India, December 17-20, 2017, Proceedings*, pages 1–15, 2017.
3. G. Amanatidis, P. Kleer, and G. Schäfer. Budget-feasible mechanism design for non-monotone submodular objectives: Offline and online. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019.*, pages 901–919, 2019.
4. E. Anastasiadis and A. Deligkas. Heterogeneous facility location games. In *AAMAS*, pages 623–631, 2018.
5. J. Augustine, I. Caragiannis, A. Fanelli, and C. Kalaitzis. Enforcing efficient equilibria in network design games via subsidies. *Algorithmica*, 72(1):44–82, May 2015.
6. M. Babaioff, M. Feldman, and M. Tennenholtz. Mechanism design with strategic mediators. *ACM Trans. Economics and Comput.*, 4(2):7, 2016.
7. X. Bei, N. Chen, N. Gravin, and P. Lu. Worst-case mechanism design via bayesian analysis. *SIAM J. Comput.*, 46(4):1428–1448, 2017.
8. N. Buchbinder, L. Lewin-Eytan, J. Naor, and A. Orda. Non-cooperative cost sharing games via subsidies. In B. Monien and U.-P. Schroeder, editors, *Algorithmic Game Theory*, pages 337–349, Berlin, Heidelberg, 2008. Springer Berlin Heidelberg.
9. G. Calinescu. Bounding the payment of approximate truthful mechanisms. *Theoretical Computer Science*, 562:419 – 435, 2015.

10. I. Caragiannis, X. Chatzigeorgiou, P. Kanellopoulos, G. A. Krimpas, N. Protopapas, and A. A. Voudouris. Efficiency and complexity of price competition among single-product vendors. *Artif. Intell.*, 248:9–25, 2017.
11. I. Caragiannis and A. A. Voudouris. The efficiency of resource allocation mechanisms for budget-constrained users. In *Proceedings of the 2018 ACM Conference on Economics and Computation, Ithaca, NY, USA, June 18-22, 2018*, pages 681–698, 2018.
12. N. Chen, E. Elkind, N. Gravin, and F. Petrov. Frugal mechanism design via spectral techniques. In *Proceedings of the 2010 IEEE 51st Annual Symposium on Foundations of Computer Science, FOCS '10*, pages 755–764, Washington, DC, USA, 2010. IEEE Computer Society.
13. Y. Cheng, N. Gravin, K. Munagala, and K. Wang. A simple mechanism for a budget-constrained buyer. In *Web and Internet Economics - 14th International Conference, WINE 2018, Oxford, UK, December 15-17, 2018, Proceedings*, pages 96–110, 2018.
14. Y. Cheng, W. Yu, and G. Zhang. *Mechanisms for Obnoxious Facility Game on a Path*, pages 262–271. Springer Berlin Heidelberg, Berlin, Heidelberg, 2011.
15. Y. Cheng, W. Yu, and G. Zhang. Strategy-proof approximation mechanisms for an obnoxious facility game on networks. *Theoretical Computer Science*, 497:154 – 163, 2013. Combinatorial Algorithms and Applications.
16. S. Dobzinski, C. H. Papadimitriou, and Y. Singer. Mechanisms for complement-free procurement. In *Proceedings of the 12th ACM Conference on Electronic Commerce, EC '11*, pages 273–282, New York, NY, USA, 2011. ACM.
17. E. Dokow, M. Feldman, R. Meir, and I. Nehama. Mechanism design on discrete lines and cycles. In *ACM EC*, pages 423–440, 2012.
18. E. Elkind, L. A. Goldberg, and P. W. Goldberg. Frugality ratios and improved truthful mechanisms for vertex cover. In *Proceedings of the 8th ACM Conference on Electronic Commerce, EC '07*, pages 336–345, New York, NY, USA, 2007. ACM.
19. E. Elkind, A. Sahai, and K. Steiglitz. Frugality in path auctions. In *SODA*, pages 701–709, 2004.
20. D. Ferraioli, P. Serafino, and C. Ventre. What to verify for optimal truthful mechanisms without money. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, Singapore, May 9-13, 2016*, pages 68–76, 2016.
21. D. Fotakis, P. Krysta, and C. Ventre. Equal-cost mechanism design with monitoring. Submitted.
22. D. Fotakis and C. Tzamos. On the power of deterministic mechanisms for facility location games. In *Automata, Languages, and Programming - 40th International Colloquium, ICALP 2013, Riga, Latvia, July 8-12, 2013, Proceedings, Part I*, pages 449–460, 2013.
23. Y. Giannakopoulos, E. Koutsoupias, and M. Kyropoulou. The anarchy of scheduling without money. In *SAGT*, pages 302–314, 2016.
24. I. Golomb and C. Tzamos. Truthful facility location with additive errors. *CoRR*, abs/1701.00529, 2017.
25. N. Gravin, Y. Jin, P. Lu, and C. Zhang. Optimal budget-feasible mechanisms for additive valuations. In *Proceedings of the 2019 ACM Conference on Economics and Computation, EC 2019, Phoenix, AZ, USA, June 24-28, 2019.*, pages 887–900, 2019.
26. M. Hajiaghayi, M. R. Khani, and S. Seddighin. Frugal auction design for set systems: Vertex cover and knapsack. In *Proceedings of the 2018 ACM Conference on Economics and Computation, EC '18*, pages 645–662, New York, NY, USA, 2018. ACM.
27. K. Ibara and H. Nagamochi. Characterizing mechanisms in obnoxious facility game. In *Combinatorial Optimization and Applications - 6th International Conference, COCOA 2012, Banff, AB, Canada, August 5-9, 2012. Proceedings*, pages 301–311, 2012.
28. A. R. Karlin, D. Kempe, and T. Tamir. Beyond VCG: frugality of truthful mechanisms. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005), 23-25 October 2005, Pittsburgh, PA, USA, Proceedings*, pages 615–626, 2005.
29. D. Kempe, M. Salek, and C. Moore. Frugal and truthful auctions for vertex covers, flows and cuts. In *Proceedings of the 2010 IEEE 51st Annual Symposium on Foundations of Computer Science, FOCS '10*, pages 745–754, Washington, DC, USA, 2010. IEEE Computer Society.
30. P. J. Khalilabadi and É. Tardos. Simple and efficient budget feasible mechanisms for monotone submodular valuations. In *Web and Internet Economics - 14th International Conference, WINE 2018, Oxford, UK, December 15-17, 2018, Proceedings*, pages 246–263, 2018.

31. N. Kos and M. Messner. Extremal incentive compatible transfers. *Journal of Economic Theory*, 148:134–164, 2012.
32. E. Koutsoupias. Scheduling without payments. In *Proc. of SAGT*, volume 6982 of *LNCS*, pages 143–153, 2011.
33. A. Kovács, U. Meyer, and C. Ventre. Mechanisms with monitoring for truthful RAM allocation. In *WINE*, volume 9470 of *Lecture Notes in Computer Science*, pages 398–412. Springer, 2015.
34. M. Kyropoulou, C. Ventre, and X. Zhang. Mechanism design for constrained heterogeneous facility location. In *SAGT 2019*, 2019. In press.
35. P. Lu, X. Sun, Y. Wang, and Z. A. Zhu. Asymptotically optimal strategy-proof mechanisms for two-facility games. In *ACM EC*, pages 315–324, 2010.
36. P. Lu, Y. Wang, and Y. Zhou. Tighter bounds for facility games. In *WINE*, pages 137–148, 2009.
37. R. Meir, A. D. Procaccia, and J. S. Rosenschein. Algorithms for strategyproof classification. *Artif. Intell.*, 186:123–156, July 2012.
38. H. Moulin. On strategy-proofness and single-peakedness. *Public Choice*, 35:437–455, 1980.
39. N. Nisan and A. Ronen. Algorithmic Mechanism Design. *Games and Economic Behavior*, 35:166–196, 2001.
40. N. Nisan, T. Roughgarden, E. Tardos, and V. V. (Eds.). *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.
41. P. Penna and C. Ventre. Optimal collusion-resistant mechanisms with verification. *Games and Economic Behavior*, 86:491–509, 2014.
42. A. D. Procaccia and M. Tennenholtz. Approximate mechanism design without money. *ACM Trans. Economics and Comput.*, 1(4):18, 2013.
43. J. Schummer and R. V. Vohra. Strategy-proof location on a network. *Journal of Economic Theory*, 104:405–428, 2002.
44. P. Serafino and C. Ventre. Heterogeneous facility location without money on the line. In *ECAI - Including PAIS*, pages 807–812, 2014.
45. P. Serafino and C. Ventre. Truthful mechanisms without money for non-utilitarian heterogeneous facility location. In *AAAI*, pages 1029–1035, 2015.
46. P. Serafino and C. Ventre. Heterogeneous facility location without money. *Theor. Comput. Sci.*, 636:27–46, 2016.
47. Y. Singer. Budget feasible mechanisms. In *FOCS*, pages 765–774, 2010.
48. K. Talwar. The price of truth: Frugality in truthful mechanisms. In *STACS 2003, 20th Annual Symposium on Theoretical Aspects of Computer Science, Berlin, Germany, February 27 - March 1, 2003, Proceedings*, pages 608–619, 2003.
49. S. Zou and M. Li. Facility location games with dual preference. In *AAMAS*, pages 615–623, 2015.