

Truthfulness on a Budget: Trading Money for Approximation through Monitoring

Paper #XXX

ABSTRACT

Albeit a pervasive desideratum when computing a novel paradigm wherein agents' declarations can be partially checked against their actual costs, 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 a 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 facility location, a well-studied problem outside of this class, we draw a novel picture about the relationship between approximation and frugality.

1 INTRODUCTION

Truthful mechanisms represent the principal tool to compute in 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 suffers 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 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 large enough [21].

This motivates the definition of *budget-feasible mechanisms* [24], 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., [3, 6, 8]) 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, 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 [9]. Further research proves that similar results hold true for other utilitarian problems, including minimum spanning trees, set systems and matroids [15, 25]. 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 equilibrium 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 [5]. In [4] 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 [1] 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.

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 [13, 17, 18].¹ Mechanisms with monitoring assume that the designer is able to monitor the agents' costs so to guarantee that the agents who have exaggerated their 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

Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019), N. Agmon, M. E. Taylor, E. Elkind, M. Veloso (eds.), May 2019, Montreal, Canada

© 2019 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.
<https://doi.org/doi>

¹The terminology is introduced in [18] while [13, 17] study the same model under the misnomer of mechanisms with verification. In the verification model of [20] it is effectively assumed that the designer can monitor overbidding *and* punish underbidding; in [23] (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 [13, 17, 18].

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. This result has three important implications. Firstly, for mechanisms with monitoring, the authors of [18] prove that any algorithm admits payments leading to truthful mechanisms with monitoring when agents bid from finite domains. 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 [21, 23]. The notion defined in [21] is more restrictive than monitoring since underbidding is punished. We show that the compensation and bonus mechanism in [21] 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.

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 a law increasing 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 w.l.o.g., for mechanisms with monitoring it excludes some possible payment schemes. For example, the payment function in [18] 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 [11]. 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 in 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 for truthfulness and good approximations [10, 12].

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 approximation guarantee with frugality and answers the question of whether the truthful implementation 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 facility location, we show that money and 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 give some formal definitions. Section 3 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 3.2). Subsequent sections study these properties for a number of optimization problems studied in the literature: set systems (Section 4), K -facility location (Section 5) and obnoxious facility location (Section 6).

2 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 reported or *bid* a (true or false) cost function $b_i \in D_i$, a *mechanism* determines an allocation $x \in \mathcal{O}$. Furthermore, depending on the b_i 's, it will determine payments to be given to the agents. Payments will be denoted by $p = (p_1, \dots, p_n)$. In summary, by letting $D = \times_{i=1}^n D_i$, a *mechanism* is a pair (f, p) that on input the agents' costs uses algorithm $f : D \rightarrow \mathcal{O}$ to return a feasible solution 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. A mechanism (f, p) is *truthful* if for any i , any bids \mathbf{b}_{-i} of the agents other than i , and any $b_i \in D_i$, $u_i(t_i, \mathbf{b}_{-i}) \geq u_i(b_i, \mathbf{b}_{-i})$. A stronger requirement demands truth-telling be a dominant strategy for coalitions of agents. A mechanism (f, p) is *collusion-resistant* if for any $C \subseteq [n]$ any bids $\mathbf{b}_{-C} \times_{i \notin C} D_i$ 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* (w.r.t. 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 any mechanism charging agents would be budget-feasible.

Commonly, u_i is 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. 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 refer to the interactive CDR Review Guidance for full instructions. 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* Π , 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 Π is α -approximate² if, for every \mathbf{b} , it returns a solution whose cost is at most $\text{cost}(f^*(\mathbf{b}), \mathbf{b}) + \alpha$.

3 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 3.1. *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

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

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

COROLLARY 1. *Let Π 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 Π ,*

$$\text{cost}(f(\mathbf{b}), \mathbf{b}) \leq B.$$

3.1 When is the social cost necessary?

The minimum budget needed for a budget-feasible first-price mechanism using f for problem Π is $B^* = \max_{\mathbf{b} \text{ instance of } \Pi} \text{cost}(f(\mathbf{b}), \mathbf{b})$. 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, [1, 4, 5], relaxing individual rationality is equivalent to provide partial subsidies that do not cover the whole cost of the equilibrium solution. We pursue this research direction in the remainder of the paper.

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

$$\begin{aligned} p_i(\mathbf{b}) &= b_i(f(\mathbf{b})) + h_i && \text{for collusion-resistance} \\ p_i(\mathbf{b}) &= b_i(f(\mathbf{b})) + h_i(\mathbf{b}_{-i}) && \text{for truthfulness} \end{aligned}$$

This is an if and only if for classical mechanisms, see, e.g., [22, Theorem 9.37], in that payments of truthful mechanisms are unique up to bidder-dependent discounts dependent on the declarations of the others. We call the mechanisms that use the payment functions

²Here we depart from the classical definition of multiplicative approximation. Our aim is not to optimize the approximation but that of identifying a tool that allows to relate budget with the quality of solutions. In this context, the additive term immediately tells us the amount of extra payments (prior to the discounts, if any) needed for implementation. Interestingly, [14] argues the merits of additive approximation for facility location.

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

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

We then have proved:

THEOREM 3.2. *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^* .*

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^ .*

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

$$h_i(\mathbf{b}_{-i}) = \inf_{b \in D_i} b(f(\mathbf{b}, \mathbf{b}_{-i})). \quad (1)$$

3.1.2 The limits of general discounts. We justify our focus on Theorem 3.2 by showing the limitations of more general payments, such as,

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

that we call *fully discounted* w.r.t. $(h_i)_i$ and, the very special case,

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

termed *single-bid discounted*.

General discounts (might) destroy truthfulness. Theorem 3.2 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* w.r.t. 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 w.r.t. h_i is not truthful with monitoring when coupled with a payment function that is fully discounted w.r.t. $(h_j)_j$ for some $h_k, k \neq i$.*

PROOF. Let \mathbf{b}_{-i} be the vector for which f is monotone w.r.t. 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 $h_i(\mathbf{b}) > h_i(t_i, \mathbf{b}_{-i})$. \square

Money is needed. The observation here directly follows from the results in [18]. 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

³One might wonder whether relaxing the upper bound for some agent, 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.

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. We show in appendix that B^* is a necessary budget for set systems and K -facility location, even for single-bid discounts.

4 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 $O = \mathcal{B}$ 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 4.1. *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 $f^*(\mathbf{b}_{-i}) \in O$ and let $C_{-i} = c(f^*(\mathbf{b}_{-i}))$. Let S be the set in 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 then $b_i(f^*(\mathbf{b})) = 0$. If, instead, $C_{-i} \geq C_{i=0}$ then by setting $b_i = 0$, no matter whether i is in $f^*(\mathbf{b})$ or not, we have $b_i(f^*(\mathbf{b})) = 0$. \square

5 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 5.1. *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/classify⁴ 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 [4] 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 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 discussion above.

5.1 Optimal mechanisms

We next prove that 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 $\text{cost}(f^*(b_i, \mathbf{b}_{-i}), (b_i, \mathbf{b}_{-i})) \leq \text{cost}(f^*(\mathbf{b}_{-i}), \mathbf{b}_{-i})$.*

PROOF. Let us first assume that:

$$\text{cost}(f^*(b_i, \mathbf{b}_{-i}), (b_i, \mathbf{b}_{-i})) < \text{cost}(f^*(\mathbf{b}_{-i}), \mathbf{b}_{-i}) \quad (2)$$

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 (2), 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 $b_i \in f^*(\mathbf{b}_{-i}) = 0$. \square

THEOREM 5.2. *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 1.⁵ \square

By Theorem 3.2 we have the following result.

COROLLARY 2. *There is no truthful discounted first-price budget-feasible mechanism with monitoring for facility location, that uses an optimal algorithm f^* whose budget is smaller than $\max_{\mathbf{b} \in D} \text{cost}(f^*(\mathbf{b}), \mathbf{b})$.*

5.2 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 w.l.o.g. from the approximation guarantee point of view. Moreover, as already discussed in [10], 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}^π be the declaration vector \mathbf{b} where entries are permuted according to π ; f is anonymous if $f(\mathbf{b}) = f(\mathbf{b}^\pi)$.

THEOREM 5.3. *Let f be a simple and anonymous algorithm for facility location. Then for all i there $\exists \mathbf{b}$ s.t. $b_i(f(\mathbf{b})) = 0$.*

PROOF. Since f is simple, for all \mathbf{b} there exists an agent j s.t. $b_j(f(\mathbf{b})) = 0$. For all $i \neq j$ we construct \mathbf{b}^π 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 $b_i(f(\mathbf{b}^\pi)) = b_j(f(\mathbf{b})) = 0$. \square

By Theorem 3.2 we have the following result.

COROLLARY 3. *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 $\max_{\mathbf{b} \in D} \text{cost}(f(\mathbf{b}), \mathbf{b})$.*

5.3 Trading approximation for frugality

Given an optimal algorithm f^* , we define f_ε^* as the algorithm that shifts by ε 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_ε^* , 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 . Let $S_1^\varepsilon, \dots, S_K^\varepsilon$ be similarly defined w.r.t. the allocation $(F_j + \varepsilon)_{j=1, \dots, K}$, that is, all $i \in S_j$ are served by $F_j + \varepsilon$. Our analysis assumes that ε 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 ε . Formally, let us rewrite \mathbf{b} as $(\mathbf{b}_1, \dots, \mathbf{b}_K)$, where $\mathbf{b}_j = (b^1, \dots, b^{|S_j|})$ is the vector of all the locations of the agents connecting to F_j . Now, let \mathbf{b}^ε be the

⁴See [19] for the similarities among facility location and clustering/classification with strategic data sources.

⁵A simpler proof works when agents have unrestricted domains (i.e., simply set b_j 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., [2].

vector $(\mathbf{b}_1, \dots, \mathbf{b}_j^\varepsilon, \dots, \mathbf{b}_K)$, where $\mathbf{b}_j^\varepsilon = (b^1, \dots, b^\ell + \varepsilon, \dots, b^{|S_j|})$ for some $1 \leq j \leq K$ and $\ell \in S_j$. In words, in \mathbf{b}^ε we shift the location of one agent of ε to the right. We say that ε is ‘‘small enough’’ if $f^*(\mathbf{b}) = f^*(\mathbf{b}^\varepsilon)$ and $S_j = S_j^\varepsilon$ for all $j \in \{1, \dots, K\}$ and for all $\ell \in \{1, \dots, |S_j|\}$. The following partition of the agents of a given instance is useful.

DEFINITION 1. 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: $\Delta_1^j = \{i | b_i \in (\max \Delta_3^{j-1}, F_j]\}$, $|\Delta_1^j| = \delta_1^j$; $\Delta_2^j = \{i | b_i \in (F_j, F_j + \varepsilon]\}$, $|\Delta_2^j| = \delta_2^j$; and $\Delta_3^j = \{i | b_i \in (F_j + \varepsilon, \min \Delta_1^{j+1})\}$, $|\Delta_3^j| = \delta_3^j$, where $\Delta_3^0 = -\infty$ and $\Delta_1^{K+1} = \infty$. We let $\Delta_i = \cup_j \Delta_i^j$ and $\delta_i = \sum_j \delta_i^j$.

We next prove the approximation guarantee of our algorithm f_ε^* .

THEOREM 5.4. If ε 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 $j \in \{1, \dots, K\}$. Since ε is small enough, for the solution $f_\varepsilon^*(\mathbf{b})$, the agents in Δ_1^j have to pay the same cost they paid for $f^*(\mathbf{b})$ plus ε . 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) ε on their cost to implement $f^*(\mathbf{b})$. (That is, the agents in Δ_3^j save exactly ε whilst those in Δ_2^j can save less than that.) For the upper bound, instead, we note that agents in Δ_2^j might need to pay at most ε 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 Δ_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 and \mathbf{b}_{-i} , $f^*(b_i, \mathbf{b}_{-i}) = f^*(\mathbf{b}_{-i})$ for all b_i such that $b_i \in f^*(\mathbf{b}_{-i})$.

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

$$\text{cost}(f^*(\mathbf{b}_{-i}), \mathbf{b}_{-i}) = \text{cost}(f^*(\mathbf{b}), \mathbf{b}_{-i}) \quad (3)$$

$$\text{cost}(f^*(\mathbf{b}_{-i}), \mathbf{b}) = \text{cost}(f^*(\mathbf{b}), \mathbf{b}). \quad (4)$$

by which, by the fixed-tie breaking rule, we get $f^*(\mathbf{b}_{-i}) = f^*(\mathbf{b})$. Indeed, for (3) 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 one follows from non-negativity of costs. Equation (3) then follows since $\text{cost}(f^*(\mathbf{b}_{-i}), \mathbf{b}_{-i}) = \text{cost}(f^*(\mathbf{b}), \mathbf{b})$ by Lemma

1. Equation 4 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 5.5. Let f^* be an optimal algorithm for K -facility location that returns the lexicographically minimal (or maximal) allocation among all optimal allocations. For ε small enough, let f_ε^* 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. The proof starts from a vector \mathbf{b} and uses $f^*(\mathbf{b})$ 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, Δ_1, Δ_2 and Δ_3 be the partitions of the agents in S as from Definition 1.

We begin by considering the case that $|S| > 1$ is odd. If $i \in \Delta_2 \cup \Delta_3$, then the location of F (and then $F + \varepsilon$) does 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 ε 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 ℓ be the leftmost agent such that $F \leq b_\ell$ and $\ell \neq k$. If i declares $b'_i = b_\ell + \varepsilon$, then since ε 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 3.2 we have the following result.

COROLLARY 4. There is no truthful discounted first-price budget-feasible mechanism with monitoring for facility location, that uses f_ε^* 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 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 $O = \{x \in R \mid 0 \leq x \leq \ell\}$,

where ℓ denotes the length of the interval.⁶ 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 [7] that an optimal algorithm allocates the facility either at 0 or at ℓ . 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 ℓ .

Just like for facility location, VCG payments would charge the agents (Theorem 5.1 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 mechanisms with monitoring and show that the social cost is a necessary budget for the optimal mechanism.

THEOREM 6.1. *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 3.2 we have the following result.

COROLLARY 5. *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 $\max_{\mathbf{b} \in D} \text{cost}(f^*(\mathbf{b}), \mathbf{b})$.*

6.1 Trading approximation for frugality

Given an optimal algorithm f^* , we define f_ε^* as the algorithm that shifts by ε 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_ε^* is not unidirectional: this is to prevent the facility from being allocated outside of the interval. We next prove that for f_ε^* savings are possible.

THEOREM 6.2. *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 2. *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 $\Delta^\pm = \{i \mid b_i(f_\varepsilon^*) = b_i(f^*)\}$, $\delta^\pm = |\Delta^\pm|$.*

⁶We 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.

Put in words, Δ^+ , Δ^- and Δ^\pm denote the set of agents whose cost with respect to the allocation f_ε^* 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_ε^* is $((\delta^+ - \delta^-) \cdot \varepsilon)$ -approximate; however, as proved in Theorem 6.3, this loss in approximation can be recovered by the budget savings possible.

LEMMA 3. *Let f^* be an optimal allocation algorithm that always locates the facility at one extreme of the interval⁷ and let f_ε^* be the shifted allocation function defined upon f^* . It holds: $\text{cost}(f_\varepsilon^*(\mathbf{b}), \mathbf{b}) = \text{cost}(f^*(\mathbf{b}), \mathbf{b}) + (\delta^+ - \delta^-) \cdot \varepsilon$.*

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

THEOREM 6.3. *For $0 < \varepsilon < \frac{\ell}{2}$, algorithm f_ε^* achieves a budget saving of $\varepsilon(n + \delta^- - \delta^+)$ with respect to f^* . For $\frac{\ell}{2} < \varepsilon < \ell$ algorithm f_ε^* achieves a budget saving of $n\varepsilon - (n + \delta^- - \delta^+)\varepsilon$ with respect to f^* .*

PROOF. By Theorem 6.1, the budget needed to implement f^* is $\sum_i b_i(f^*)$. By Theorem 6.2 and Lemma 3, the budget needed to implement f_ε^* is $\text{cost}(f_\varepsilon^*(\mathbf{b}), \mathbf{b}) + (\delta^+ - \delta^-)\varepsilon + n \cdot h_i(\mathbf{b}_{-i})$ where $h_i(\mathbf{b}_{-i})$ is defined in (1).

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_ε^* 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 $n\varepsilon - (n + \delta^- - \delta^+)\varepsilon$. We observe that the saving in this case can be negative (e.g., the budget required to implement f_ε^* can be higher than the one required to implement f^*) depending on the value of ε and on the structural configuration of the instance at hand. \square

Theorem 6.3 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 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 ℓ . The horizontal and vertical axes represent, respectively, the shift ε controlling the approximation, and the monetary cost of implementing the solution. The optimal solution for this instance is to allocate the facility at ℓ , 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 ε^B we can afford, we just need to find the leftmost intersection between the budget line B and the cost of implementing f_ε^* , denoted $\sum_i p_i(f_\varepsilon^*)$. In this case, $\sum_i p_i(f_\varepsilon^*)$ is

⁷Note 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 ℓ). We assume that the algorithm will allocate consistently the facility at one extreme of the interval.




Figure 1: Approximation vs Frugality for Obnoxious Facility Location

constituted of two segments, defined as:

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

In the general case, $\sum_i p_i(f_\varepsilon^*)$ is piece-wise linear, where the various segments originate both from ε ranging over $(0, \ell)$ as the location of the facility is shifted, and from the change of δ^+ , δ^- and δ^- , that in turn modifies the slope of the line segments. In this particular case, the sets Δ^+ , Δ^- and Δ^- do not change since the agents reside at 0 and ℓ . 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 $\text{cost}(f_{\varepsilon^B}^*(\mathbf{b}), \mathbf{b}) - \text{cost}(f^*(\mathbf{b}), \mathbf{b})$, where $\text{cost}(f_{\varepsilon^B}^*(\mathbf{b}), \mathbf{b})$ can be computed using Lemma 3.

7 CONCLUSIONS

We study budget-feasible mechanisms with monitoring for utilitarian cost-minimization problems. Our results suggest the kind 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 for the implementation of 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 [16], 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 interesting 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 there are many interesting, yet unanswered, questions about the computation of truthful payments and the ability to find the best payments for a given algorithm/problem.

REFERENCES

- [1] J. Augustine, I. Caragiannis, A. Fanelli, and C. Kalaitzis. 2015. Enforcing Efficient Equilibria in Network Design Games via Subsidies. *Algorithmica* 72, 1 (May 2015), 44–82. <https://doi.org/10.1007/s00453-013-9845-5>
- [2] M. Babaioff, M. Feldman, and M. Tennenholtz. 2016. Mechanism Design with Strategic Mediators. *ACM Trans. Economics and Comput.* 4, 2 (2016), 7.
- [3] X. Bei, N. Chen, N. Gravin, and P. Lu. 2012. Budget Feasible Mechanism Design: From Prior-free to Bayesian. In *Proceedings of the Forty-fourth Annual ACM Symposium on Theory of Computing (STOC '12)*. ACM, New York, NY, USA, 449–458. <https://doi.org/10.1145/2213977.2214020>
- [4] N. Buchbinder, L. Lewin-Eytan, J. Naor, and A. Orda. 2008. Non-cooperative Cost Sharing Games Via Subsidies. In *Algorithmic Game Theory*, Burkhard Monien and Ulf-Peter Schroeder (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 337–349.
- [5] I. Caragiannis, X. Chatzigeorgiou, P. Kanellopoulos, G. A. Krimpas, N. Protopapas, and A. A. Voudouris. 2017. Efficiency and complexity of price competition among single-product vendors. *Artif. Intell.* 248 (2017), 9–25. <https://doi.org/10.1016/j.artint.2017.03.001>
- [6] N. Chen, N. Gravin, and P. Lu. 2011. Mechanism Design without Money via Stable Matching. *CoRR* abs/1104.2872 (2011).
- [7] Y. Cheng, W. Yu, and G. Zhang. 2011. *Mechanisms for Obnoxious Facility Game on a Path*. Springer Berlin Heidelberg, Berlin, Heidelberg, 262–271. https://doi.org/10.1007/978-3-642-22616-8_21
- [8] S. Dobzinski, C. H. Papadimitriou, and Y. Singer. 2011. Mechanisms for Complement-free Procurement. In *Proceedings of the 12th ACM Conference on Electronic Commerce (EC '11)*. ACM, New York, NY, USA, 273–282. <https://doi.org/10.1145/1993574.1993615>
- [9] E. Elkind, A. Sahai, and K. Steiglitz. 2004. Frugality in path auctions. In *SODA*. 701–709.
- [10] D. Ferraioli, P. Serafino, and C. Ventre. 2016. 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*. 68–76.
- [11] D. Fotakis, P. Krysta, and C. Ventre. [n. d.]. Equal-Cost Mechanism Design with Monitoring. ([n. d.]). Submitted.
- [12] D. Fotakis and C. Tzamos. 2013. 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*. 449–460.
- [13] Y. Giannakopoulos, E. Koutsoupias, and M. Kyropoulou. 2016. The Anarchy of Scheduling Without Money. In *SAGT*. 302–314.

- [14] I. Golomb and C. Tzamos. 2017. Truthful Facility Location with Additive Errors. *CoRR* abs/1701.00529 (2017). <http://arxiv.org/abs/1701.00529>
- [15] A. R. Karlin, D. Kempe, and T. Tamir. 2005. 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*. 615–626.
- [16] N. Kos and M. Messner. 2012. Extremal incentive compatible transfers. *Journal of Economic Theory* 148 (2012), 134–164.
- [17] E. Koutsoupias. 2011. Scheduling without Payments. In *Proc. of SAGT (LNCS)*, Vol. 6982. 143–153.
- [18] A. Kovács, U. Meyer, and C. Ventre. 2015. Mechanisms with Monitoring for Truthful RAM Allocation. In *WINE (Lecture Notes in Computer Science)*, Vol. 9470. Springer, 398–412.
- [19] R. Meir, A. D. Procaccia, and J. S. Rosenschein. 2012. Algorithms for Strategyproof Classification. *Artif. Intell.* 186 (July 2012), 123–156. <https://doi.org/10.1016/j.artint.2012.03.008>
- [20] N. Nisan and A. Ronen. 1999. Algorithmic Mechanism Design (Extended Abstract). In *STOC*. 129–140.
- [21] N. Nisan and A. Ronen. 2001. Algorithmic Mechanism Design. *Games and Economic Behavior* 35 (2001), 166–196.
- [22] N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani (Eds.). 2007. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA.
- [23] P. Penna and C. Ventre. 2014. Optimal Collusion-Resistant Mechanisms with Verification. *Games and Economic Behavior* 86 (2014), 491–509.
- [24] Y. Singer. 2010. Budget Feasible Mechanisms. In *FOCS*. 765–774.
- [25] K. Talwar. 2003. 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*. 608–619.