Budget Feasible Mechanisms for Experimental Design

Thibaut Horel
Technicolor
thibaut.horel@ens.fr

Stratis Ioannidis
Technicolor
stratis.ioannidis
@technicolor.com

S. Muthukrishnan
Rutgers University, Microsoft
Research
muthu@cs.rutgers.edu

ABSTRACT

In the classical *experimental design* setting, an experimenter `E` with a budget B has access to a population of n potential experiment subjects $i \in \{1, \dots, n\}$, each associated with a vector of features $x_i \in \mathbb{R}^d$ as well as a cost $c_i > 0$. Conducting an experiment with subject i reveals an unknown value $y_i \in \mathbb{R}$ to `E`. `E` typically assumes some hypothetical relationship between x_i's and y_i's, *e.g.*, $y_i \approx \beta^T x_i$, and estimates β from experiments. `E`'s goal is to select which experiments to conduct, subject to her budget constraint.

We initiate the study of mechanisms for experimental design. In this setting, subjects are *strategic* and may lie about their costs. In particular, we formulate the *Experimental Design Problem* (EDP) as finding a set S of subjects that maximize $V(S) = \log\det(I_d + \sum_{i\in S} x_i x_i^T)$ under the constraint $\sum_{i\in S} c_i \leq B$; our objective function corresponds to the information gain in β when it is learned through linear regression methods, and is related to the so-called D-optimality criterion. We present the first known deterministic, polynomial time, truthful, budget feasible mechanism for EDP. Our mechanism yields a constant factor (≈ 19.68) approximation, and we show that no truthful, budget-feasible algorithms are possible within a factor 2 approximation. Our approach here generally applies to a wider class of learning problems and obtains polynomial time universally truthful (*i.e.*, randomized) budget feasible mechanism, also within a constant factor approximation.

1. INTRODUCTION

In the classic setting of experimental design [?, ?], an *experimenter* `E` has access to a population of n potential experiment subjects. Each subject $i \in \{1, \dots, n\}$ is associated with a set of parameters (or features) $x_i \in \mathbb{R}^d$, known to the experimenter. `E` wishes to perform an experiment that measures a certain inherent property of the subjects: the outcome for a subject i is denoted y_i, which is unknown to `E` before the experiment is performed. Typically, `E` has a hypothesis of the relationship between x_i's and y_i's, such as, say linear, *i.e.*, $y_i \approx \beta^T x_i$; conducting the experiments and obtaining the measurements y_i lets `E` estimate β.

The above experimental design scenario has many applications, from medical testing to marketing research and others. There is a rich literature about estimation procedures, as well as for means of quantifying the quality of the produced estimate [?]. There is also an extensive theory on how to select subjects if `E` can conduct only a limited number of experiments, so the estimation process returns β that approximates the true parameter of the underlying population [?, ?, ?, ?].

We depart from this classical setup by viewing experimental design in a strategic setting, and by studying mechanism design issues. In our setting, experiments cannot be manipulated and hence measurements are considered reliable. However, there is a cost c_i associated with experimenting on subject i which varies from subject to subject. This may be viewed as the cost subject i incurs when tested and for which she needs to be reimbursed; or, it might be viewed as the incentive for i to participate in the experiment; or, it might be the inherent value of the data. This economic aspect has always been inherent in experimental design: experimenters often work within strict budgets and design creative incentives. However, we are not aware of principled study of this setting from a strategic point of view. When subjects are strategic, they may have an incentive to misreport their cost and the choice of experiments and payments need to be more sophisticated.

Our contributions are as follows.

- We formulate the problem of experimental design subject to a given budget, in the presence of strategic agents who may lie about their costs. In particular, we focus on linear regression. This is naturally viewed as a budget feasible mechanism design problem, in which the objective function is related to the covariance of the x_i's. In particular, we formulate the *Experimental Design Problem* (EDP) as follows: the experimenter `E` wishes to find set S of subjects to maximize
$$V(S) = \log\det\left(I_d + \sum_{i\in S} x_i x_i^T\right) \qquad (1)$$
subject to a budget constraint $\sum_{i\in S} c_i \leq B$, where B is `E`'s budget. The objective function, which is the key, is obtained by optimizing the information gain in β when it is learned through linear regression methods, and is related to the so-called D-optimality criterion.
- The above objective is submodular. There are several recent results in budget feasible mechanisms [?, ?, ?, ?, ?], and some apply to the submodular optimization in EDP. There is a randomized, 7.91-approximate polynomial time mechanism for maximizing a general submodular function that is universally truthful, *i.e.*,

it is sampled from a distribution among truthful mechanisms. Also, there is a 8.34-approximate exponential time deterministic mechanism. There are however no known deterministic, truthful, polynomial time mechanisms for general submodular functions. For specific combinatorial problems such as KNAPSACK or COVERAGE, there exist deterministic, truthful, polynomial constant-approximation algorithms [?, ?, ?], but they do not work for the linear-algebraic objective function in EDP.

We present the first known, polynomial time truthful mechanism for EDP. Our mechanism is a constant factor (≈ 12.98) approximation for EDP. In contrast to this, we show that no truthful, budget-feasible algorithms are possible for EDP within a factor 2 approximation. From a technical perspective, we present a convex relaxation of (1), and show that it is within a constant factor from the so-called multi-linear relaxation of (1). This allows us to adopt the approach followed by prior work in budget feasible mechanisms by Chen *et al.* [?] and Singer [?].

- Our approach to mechanisms for experimental design — by optimizing the information gain in parameters like β which are estimated through the data analysis process — is general. We give examples of this approach beyond linear regression to a general class that includes logistic regression and learning binary functions, and show that the corresponding budgeted mechanism design problem is also expressed through a submodular optimization. Hence, prior work [?, ?] immediately applies, and gives randomized, universally truthful, polynomial time, constant factor approximation mechanisms for problems in this class. Getting deterministic, truthful, polynomial time mechanisms with a constant approximation factor for this class or specific problems in it, like we did for EDP, remains an open problem.

In what follows, we describe related work in Section 2. We briefly review experimental design and budget feasible mechanisms in Section 3 and define EDP formally. In Section 4 we present our mechanism for EDP and prove our main results. The present applications of our general framework are presented in Section 5.

2. RELATED WORK

Budget feasible mechanism design was originally proposed by Singer [?]. Singer considers the problem of maximizing an arbitrary submodular function subject to a budget constraint in the *value query* model, *i.e.* assuming an oracle providing the value of the submodular objective on any given set. Singer shows that there exists a randomized, 112-approximation mechanism for submodular maximization that is *universally truthful* (*i.e.*, it is a randomized mechanism sampled from a distribution over truthful mechanisms). Chen *et al.* [?] improve this result by providing a 7.91-approximate mechanism, and show a corresponding lower bound of 2 among universally truthful mechanisms for submodular maximization.

In contrast to the above results, no deterministic, truthful, constant approximation mechanism that runs in polynomial time is presently known for submodular maximization. However, assuming access to an oracle providing the optimum in the full-information setup, Chen *et al.*, provide a truthful, 8.34-approximate mechanism; in cases for which the full information problem is NP-hard, as the one we consider here, this mechanism is not poly-time, unless P=NP. Chen *et al.* also prove a $1+\sqrt{2}$ lower bound for truthful mechanisms, improving upon an earlier bound of 2 by Singer [?].

Improved bounds, as well as deterministic polynomial mechanisms, are known for specific submodular objectives. For symmetric submodular functions, a truthful mechanism with approximation ratio 2 is known, and this ratio is tight [?]. Singer also provides a 7.32-approximate truthful mechanism for the budget feasible version of MATCHING, and a corresponding lower bound of 2 [?]. Improving an earlier result by Singer, Chen *et al.* [?], give a truthful, $2+\sqrt{2}$-approximate mechanism for KNAPSACK, and a lower bound of $1+\sqrt{2}$. Finally, a truthful, 31-approximate mechanism is also known for the budgeted version of COVERAGE [?, ?].

Beyond submodular objectives, it is known that no truthful mechanism with approximation ratio smaller than $n^{1/2-\epsilon}$ exists for maximizing fractionally subadditive functions (a class that includes submodular functions) assuming access to a value query oracle [?]. Assuming access to a stronger oracle (the *demand* oracle), there exists a truthful, $O(\log^3 n)$-approximate mechanism [?] as well as a universally truthful, $O(\frac{\log n}{\log\log n})$-approximate mechanism for subadditive maximization [?]. Moreover, in a Bayesian setup, assuming a prior distribution among the agent's costs, there exists a truthful mechanism with a 768/512-approximation ratio [?].

A series of recent papers [?, ?, ?, ?] consider the related problem of retrieving data from an *unverified* database: the auctioneer cannot verify the data reported by individuals and therefore must incentivize them to report truthfully. McSherry and Talwar [?] argue that *differentially private* mechanisms offer a form of *approximate truthfulness*: if users have a utility that depends on their privacy, reporting their data untruthfully can only increase their utility by a small amount. Xiao [?], improving upon earlier work by Nissim *et al.* [?], constructs mechanisms that simultaneously achieve exact truthfulness as well as differential privacy. Eliciting private data through a *survey* [?], whereby individuals first decide whether to participate in the survey and then report their data, also falls under the unverified database setting [?]. In the *verified* database setting, Ghosh and Roth [?] and Dandekar *et al.* [?] consider budgeted auctions where users have a utility again captured by differential privacy. Our work departs from the above setups in that utilities do not involve privacy, whose effects are assumed to be internalized in the costs reported by the users; crucially, we also assume that experiments are tamper-proof, and individuals can misreport their costs but not their values.

Our work is closest to the survey setup of Roth and Schoenebeck [?], who also consider how to sample individuals with different features who report a hidden value at a certain cost. The authors assume a prior on the joint distribution between costs and features, and wish to obtain an unbiased estimate of the expectation of the hidden value under the constraints of truthfulness, budget feasibility and individual rationality. Our work departs by learning a more general statistic (a linear model) than data means. We note that, as in [?], costs and features can be arbitrarily correlated (our results are prior-free).

3. PRELIMINARIES

3.1 Experimental Design

The theory of experimental design [?, ?] studies how an experimenter `E` should select the parameters of a set of experiments she is about to conduct. In general, the optimality of a particular design depends on the purpose of the experiment, *i.e.*, the quantity `E` is trying to learn or the hypothesis she is trying to validate. Due to their ubiquity in statistical analysis, a large literature on the subject focuses on learning *linear models*, where `E` wishes to fit a linear map to the data she has collected.

More precisely, putting cost considerations aside, suppose that `E` wishes to conduct k among n possible experiments. Each experiment $i \in \mathcal{N} \equiv \{1, \ldots, n\}$ is associated with a set of parameters (or features) $x_i \in \mathbb{R}^d$, normalized so that $\|x_i\|_2 \leq 1$. Denote by $S \subseteq \mathcal{N}$, where $|S| = k$, the set of experiments selected; upon its execution, experiment $i \in S$ reveals an output variable (the "measurement") y_i, related to the experiment features x_i through a linear function, *i.e.*,

$$\forall i \in \mathcal{N}, \quad y_i = \beta^T x_i + \varepsilon_i \tag{2}$$

where β is a vector in $\mathbb{R}^d$, commonly referred to as the *model*, and ε_i (the *measurement noise*) are independent, normally distributed random variables with mean 0 and variance σ^2.

The purpose of these experiments is to allow `E` to estimate the model β. In particular, under (2), the maximum likelihood estimator of β is the *least squares* estimator: for $X_S = [x_i]_{i \in S} \in \mathbb{R}^{|S| \times d}$ the matrix of experiment features and $y_S = [y_i]_{i \in S} \in \mathbb{R}^{|S|}$ the observed measurements,

$$\begin{aligned}\hat{\beta} &= \max_{\beta \in \mathbb{R}^d} \Pr(y_S; \beta) = \arg\min_{\beta \in \mathbb{R}^d} \sum_{i \in S} (\beta^T x_i - y_i)^2 \\ &= (X_S^T X_S)^{-1} X_S^T y_S \end{aligned} \tag{3}$$

Note that the estimator $\hat{\beta}$ is a linear map of y_S; as y_S is a multidimensional normal r.v., so is $\hat{\beta}$ (the randomness coming from the noise terms ε_i). In particular, $\hat{\beta}$ has mean β (*i.e.*, it is an *unbiased estimator*) and covariance $(X_S^T X_S)^{-1}$.

Let $V : 2^{\mathcal{N}} \to \mathbb{R}$ be a *value function*, quantifying how informative a set of experiments S is in estimating β. The standard optimal experimental design problem amounts to finding a set S that maximizes $V(S)$ subject to the constraint $|S| \leq k$.

A variety of different value functions are used in experimental design [?]; almost all make use of the covariance $(X_S^T X_S)^{-1}$ of the estimator $\hat{\beta}$. A value function preferred because of its relationship to entropy is the *D-optimality criterion*:

$$V(S) = \frac{1}{2} \log \det X_S^T X_S \tag{4}$$

As $\hat{\beta}$ is a multidimensional normal random variable, the D-optimality criterion is equal (up to a constant) to the negative of the entropy of $\hat{\beta}$. Hence, selecting a set of experiments S that maximizes $V(S)$ is equivalent to finding the set of experiments that minimizes the uncertainty on β, as captured by the entropy of its estimator.

Value function (4) is undefined when $\text{rank}(X_S^T X_S) < d$; in this case, we take $V(S) = -\infty$ (so that V takes values in the extended reals). Note that (4) is a submodular set function, *i.e.*, $V(S) + V(T) \geq V(S \cup T) + V(S \cap T)$ for all $S, T \subseteq \mathcal{N}$; it is also monotone, *i.e.*, $V(S) \leq V(T)$ for all $S \subset T$.

3.2 Budget Feasible Reverse Auctions

A *budget feasible reverse auction* [?] comprises a set of items $\mathcal{N} = \{1, \ldots, n\}$ as well as a single buyer. Each item has a cost $c_i \in \mathbb{R}_+$. Moreover, the buyer has a positive value function $V : 2^{\mathcal{N}} \to \mathbb{R}_+$, as well as a budget $B \in \mathbb{R}_+$. In the *full information case*, the costs c_i are common knowledge; the objective of the buyer in this context is to select a set S maximizing the value $V(S)$ subject to the constraint $\sum_{i \in S} c_i \leq B$. We write:

$$OPT = \max_{S \subseteq \mathcal{N}} \Big\{ V(S) \;\Big|\; \sum_{i \in S} c_i \leq B \Big\} \tag{5}$$

for the optimal value achievable in the full-information case.

In the *strategic case*, each items in $\mathcal{N}$ is held by a different strategic agent, whose cost is *a priori* private. A *mechanism* $\mathcal{M} = (f, p)$ comprises (a) an *allocation function* $f : \mathbb{R}_+^n \to 2^{\mathcal{N}}$ and (b) a *payment function* $p : \mathbb{R}_+^n \to \mathbb{R}_+^n$. Given the vector or costs $c = [c_i]_{i \in \mathcal{N}}$, the allocation function f determines the set in $\mathcal{N}$ of items to be purchased, while the payment function returns a vector of payments $[p_i]_{i \in \mathcal{N}}$. Let $s_i(c) = \mathbb{1}_{i \in f_i(c)}$ be the binary indicator of $i \in f(c)$. As in [?, ?], we study mechanisms that are normalized ($i \notin f(c)$ implies $p_i(c) = 0$), individually rational ($p_i(c) \geq c_i \cdot s_i(c)$) and have no positive transfers ($p_i(c) \geq 0$).

In addition to the above, mechanism design in budget feasible reverse auctions seeks mechanisms that have the following properties [?, ?]:

- *Truthfulness.* An agent has no incentive to missreport her true cost. Formally, let c_{-i} be a vector of costs of all agents except i. A mechanism is truthful iff for every $i \in \mathcal{N}$ and every two cost vectors (c_i, c_{-i}) and (c_i', c_{-i}) $p_i(c_i, c_{-i}) - s_i(c_i, c_{-i}) \cdot c_i \geq p_i(c_i', c_{-i}) - s(c_i', c_{-i}) \cdot c_i$.

- *Budget Feasibility.* The sum of the payments should not exceed the budget constraint: $\sum_{i \in \mathcal{N}} p_i \leq B$.

- *Approximation ratio.* The value of the allocated set should not be too far from the optimum value of the full information case (5). Formally, there must exist some $\alpha \geq 1$ such that:

  $$OPT \leq \alpha V(S).$$

  The approximation ratio captures the *price of truthfulness*, *i.e.*, the relative value loss incurred by adding the truthfulness constraint.

- *Computational efficiency.* The allocation and payment function should be computable in polynomial time in the number of agents n.

As noted in [?, ?], budget feasible reverse auctions are *single parameter* auctions: each agent has only one private value. In this case, Myerson's Theorem [?] gives a characterization of truthful mechanisms.

Lemma 1 (Myerson [?]). *A normalized mechanism $\mathcal{M} = (f, p)$ for a single parameter auction is truthful iff: (a) f is* monotone, i.e.*, for any agent i and $c_i' \leq c_i$, for any fixed costs c_{-i} of agents in $\mathcal{N} \setminus \{i\}$, $i \in f(c_i, c_{-i})$ implies $i \in f(c_i', c_{-i})$, and (b) agents are paid* threshold payments, i.e.*, for all $i \in f(c)$, $p_i(c) = \inf\{c_i' : i \in f(c_i', c_{-i})\}$.*

Myerson's Theorem allows us to focus on designing a monotone allocation function. Then, the mechanism will be truthful as long as we give each agent her threshold payment—the caveat being that the latter need to sum to a value below B.

3.3 Budget Feasible Experimental Design

We approach the problem of optimal experimental design from the perspective of a budget feasible reverse auction, as defined above. In particular, we assume the experimenter `E` has a budget $B \in \mathbb{R}_+$ and plays the role of the buyer. Each experiment $i \in \mathcal{N}$ corresponds to a strategic agent, whose cost c_i is private. In order to obtain the measurement y_i, the experimenter needs to pay agent i a price that exceeds her cost.

For example, each i may correspond to a human subject; the feature vector x_i may correspond to a normalized vector of her age, weight, gender, income, *etc.*, and the measurement y_i may capture some biometric information (*e.g.*, her red cell blood count, a genetic marker, etc.). The cost c_i is the amount the subject deems sufficient to incentivize her participation in the study. Note that, in this setup, the feature vectors x_i are public information that the experimenter can consult prior the experiment design. Moreover, though a subject may lie about her true cost c_i, she cannot lie about x_i (*i.e.*, all features are verifiable upon collection) or y_i (*i.e.*, she cannot falsify her measurement).

Ideally, motivated by the D-optimality criterion, we would like to design a mechanism that maximizes (4) within a good approximation ratio. In what follows, we consider a slightly more general objective as follows:

$$\textsc{ExperimentalDesignProblem (EDP)}$$

$$\text{Maximize} \quad V(S) = \log\det(I_d + X_S^T X_S) \tag{6a}$$

$$\text{subject to} \quad \sum_{i \in S} c_i \leq B \tag{6b}$$

where $I_d \in \mathbb{R}^{d \times d}$ is the identity matrix. We present our results with this version of the objective function because it is simple and captures the versions we will need later (including an arbitrary matrix in place of I_d or a zero matrix which will correspond to (4)).

EDP *and* the corresponding problem with objective (4) are NP-hard; to see this, note that Knapsack reduces to EDP with dimension $d = 1$ by mapping the weight of each item w_i to an experiment with $x_i = w_i$. Moreover, (6a) is submodular, monotone and satisfies $V(\emptyset) = 0$.

4. MECHANISM FOR EDP

In this section we present a mechanism for EDP.

Prior approach to submodular optimization problems. Previous works [?, ?, ?, ?, ?] rely on a greedy algorithm. In this algorithm, elements are added to the solution set according to the following greedy selection rule. Assume that $S \subseteq \mathcal{N}$ is the solution set constructed thus far; then, the next element i to be included in S is the one with the highest *marginal-value-per-cost*:

$$i = \arg\max_{j \in \mathcal{N} \setminus S} \frac{V(S \cup \{i\}) - V(S)}{c_i} \tag{7}$$

This process terminates when no more items can be added to S using (7) under budget B. This is the generalization of the *value-per-cost* ratio used in the greedy approximation algorithm for Knapsack. However, in contrast to Knapsack, for general submodular functions, the marginal value of an element in (7) depends on the set to which the element is added. Similarly, the value of an element depends on the set in which it is considered. Unfortunately, even for the full information case, the greedy algorithm gives an unbounded approximation ratio. Instead,

Lemma 2. *[?] Let S_G be the set computed by the greedy algorithm and let $i^* = \arg\max_{i \in \mathcal{N}} V(i)$. We have:*

$$OPT \leq \frac{e}{e-1}\big(3V(S_G) + 2V(i^*)\big).$$

Thus, taking the maximum between $V(S_G)$ and $V(i^*)$ yields an approximation ratio of $\frac{5e}{e-1}$. However, this approach breaks incentive compatibility and therefore cannot be directly applied to the strategic case [?].

For the strategic case,

- When the underlying full information problem (5) can be solved in polynomial time, Chen *et al.* [?] prove that allocating to i^* when $V(i^*) \geq C \cdot OPT_{-i^*}$ (for some constant C) and to S_G otherwise yields a 8.34-approximation mechanism. However, this is not a polytime mechanism when the underlying problem is NP hard (unless P=NP), as is the case for EDP.
- For NP-hard problems, consider the optimal value of a *fractional relaxation* of the function V over the set $\mathcal{N}$. A function $R : [0,1]^n \to \mathbb{R}_+$ defined on the hypercube $[0,1]^n$ is a fractional relaxation of V over the set $\mathcal{N}$ if $R(\mathbb{1}_S) = V(S)$ for all $S \subseteq \mathcal{N}$, where $\mathbb{1}_S$ denotes the

indicator vector of S. The optimization program (5) extends naturally to such relaxations:

$$OPT' = \arg\max_{\lambda\in[0,1]^n}\left\{R(\lambda) \mid \sum_{i=1}^{n}\lambda_i c_i \leq B\right\} \tag{8}$$

Substituting OPT'_{-i^*} for OPT_{-i^*} like before works for specific problems like KNAPSACK [?] and COVERAGE [?]. For other instances of submodular function, this overall technology has to be applied and extended.

Our approach. We build on [?, ?]. We introduce a relaxation specifically tailored to the value function of EDP. $P_{\mathcal{N}}^{\lambda}(S)$ is the probability of choosing the set S if we select each element i in $\mathcal{N}$ independently with probability λ_i:

$$P_{\mathcal{N}}^{\lambda}(S) = \prod_{i\in S}\lambda_i \prod_{i\in\mathcal{N}\setminus S}(1-\lambda_i)$$

Consider the general *multi-linear* extension:

$$F(\lambda) = \mathbb{E}_{S\sim P_{\mathcal{N}}^{\lambda}}\left[V(S)\right] = \sum_{S\subseteq\mathcal{N}} P_{\mathcal{N}}^{\lambda}(S)V(S) \tag{9}$$

For EDP the multi-linear extension can be written:

$$F(\lambda) = \mathbb{E}_{S\sim P_{\mathcal{N}}^{\lambda}}\Big[\log\det\big(I_d + \sum_{i\in S} x_i x_i^T\big)\Big].$$

(8) is not a convex optimization problem, and is not easy to solve directly. We consider an additional relaxation L that follows naturally by swapping the expectation and the $\log\det$ in the above formula:

$$\begin{aligned} L(\lambda) &\equiv \log\det\Big(\mathbb{E}_{S\sim P_{\mathcal{N}}^{\lambda}}\big[I_d + \sum_{i\in S} x_i x_i^T\big]\Big) \\ &= \log\det\left(I_d + \sum_{i\in\mathcal{N}} \lambda_i x_i x_i^T\right) \end{aligned} \tag{10}$$

This function is well-known to be concave and even self-concordant (see *e.g.*, [?]). In this case, the analysis of Newton's method for self-concordant functions in [?], shows that it is possible to find the maximum of L to any precision ε in a number of iterations $O(\log\log\varepsilon^{-1})$. The main challenge will be to prove that OPT' in (8), for the relaxation $R = L$, is close to $V(S_G)$, which we will address later and is the technical bulk of the paper.

The resulting mechanism for EDP is composed of

- the allocation function presented in Algorithm 1, and
- the payment function which pays each allocated agent i her threshold payment as described in Myerson's Theorem. In the case where $\{i^*\}$ is the allocated set, her threshold payment is B (she would be have been dropped on line 1 of Algorithm 1 had she reported a higher cost). In the case where S_G is the allocated set, threshold payments' characterization from [?] gives a formula to compute these payments.

We can now state our main result:

Algorithm 1 Mechanism for EDP

```
1: $\mathcal{N} \leftarrow \mathcal{N} \setminus \{i \in \mathcal{N} : c_i > B\}$
2: $i^* \leftarrow \arg\max_{j\in\mathcal{N}} V(j)$
3: $\xi \leftarrow \arg\max_{\lambda\in[0,1]^n}\{L(\lambda) \mid \lambda_{i^*} = 0, \sum_{i\in\mathcal{N}\setminus\{i^*\}} c_i\lambda_i \leq B\}$

4: if $L(\xi) < C \cdot V(i^*)$ then
5:     return $\{i^*\}$
6: else
7:     $i \leftarrow \arg\max_{1\leq j\leq n} \frac{V(j)}{c_j}$
8:     $S_G \leftarrow \emptyset$
9:     while $c_i \leq \frac{B}{2}\frac{V(S_G\cup\{i\})-V(S_G)}{V(S_G\cup\{i\})}$ do
10:        $S_G \leftarrow S_G \cup \{i\}$
11:        $i \leftarrow \arg\max_{j\in\mathcal{N}\setminus S_G} \frac{V(S_G\cup\{j\})-V(S_G)}{c_j}$
12:    end while
13:    return $S_G$
14: end if
```

THEOREM 1. *The allocation described in Algorithm 1, along with threshold payments, is truthful, individually rational and budget feasible. Furthermore, for any $\varepsilon > 0$, the mechanism runs in time $O\left(poly(n, d, \log\log\varepsilon^{-1})\right)$ and returns a set S^* such that:*

$$\begin{aligned} OPT &\leq \frac{10e - 3 + \sqrt{64e^2 - 24e + 9}}{2(e-1)} V(S^*) + \varepsilon \\ &\simeq 12.98V(S^*) + \varepsilon \end{aligned}$$

In addition, we prove the following simple lower bound.

THEOREM 2. *There is no 2-approximate, truthful, budget feasible, individually rational mechanism for EDP.*

PROOF. Suppose, for contradiction, that such a mechanism exists. Consider two experiments with dimension $d = 2$, such that $x_1 = e_1 = [1, 0]$, $x_2 = e_2 = [0, 1]$ and $c_1 = c_2 = B/2 + \epsilon$. Then, one of the two experiments, say, x_1, must be in the set selected by the mechanism, otherwise the ratio is unbounded, a contradiction. If x_1 lowers its value to $B/2 - \epsilon$, by monotonicity it remains in the solution; by threshold payment, it is paid at least $B/2 + \epsilon$. So x_2 is not included in the solution by budget feasibility and individual rationality: hence, the selected set attains a value $\log 2$, while the optimal value is $2\log 2$. □

4.1 Proof of Theorem 1

We now present the proof of Theorem 1. Truthfulness and individual rationality follows from monotonicity and threshold payments. Monotonicity and budget feasibility follow the same steps as the analysis of Chen *et al.* [?]; for the sake of completeness, we restate their proof in the Appendix. Our proof of the approximation ratio uses a bound on our concave relaxation L (Lemma 4). This is our main technical contribution; the proof of this lemma can be found in Section 4.2.

LEMMA 3. *For any $\varepsilon > 0$, the complexity of the mechanism is $O(poly(n, d, \log\log\varepsilon^{-1}))$.*

PROOF. The value function V in (6a) can be computed in time $O(\text{poly}(n, d))$ and the mechanism only involves a linear number of queries to the function V. The function $\log\det$ is concave and self-concordant (see [?]), so for any ε, its maximum can be found to a precision ε in $O(\log\log\varepsilon^{-1})$ of iterations of Newton's method. Each iteration can be done in time $O(\text{poly}(n, d))$. Thus, line 3 of Algorithm 1 can be computed in time $O(\text{poly}(n, d, \log\log\varepsilon^{-1}))$. Hence the allocation function's complexity is as stated. $\square$

Finally, we prove the approximation ratio of the mechanism. We use the following lemma which establishes that OPT', the optimal value (8) of the fractional relaxation L under the budget constraints is not too far from OPT.

LEMMA 4. $OPT' \leq 2OPT + 2\max_{i\in\mathcal{N}} V(i)$

The proof of Lemma 4 is our main technical contribution, and can be found in Section 4.2.

Finishing Proof of Theorem 1 . Note that the lower bound $OPT' \geq OPT$ also holds trivially, as L is a fractional relaxation of V over $\mathcal{N}$. Using Lemma 4 we can complete the proof of Theorem 1 by showing that, for any $\varepsilon > 0$, if OPT'_{-i}, the optimal value of L when i^* is excluded from $\mathcal{N}$, has been computed to a precision ε, then the set S^* allocated by the mechanism is such that:

$$OPT \leq \frac{10e-3+\sqrt{64e^2-24e+9}}{2(e-1)}V(S^*)+\varepsilon \tag{11}$$

To see this, let OPT'_{-i^*} be the true maximum value of L subject to $\lambda_{i^*} = 0$, $\sum_{i\in\mathcal{N}\setminus i^*} c_i \leq B$. Assume that on line 3 of algorithm 1, a quantity $\tilde{L}$ such that $\tilde{L}-\varepsilon \leq OPT'_{-i^*} \leq \tilde{L}+\varepsilon$ has been computed (Lemma 3 states that this is computed in time within our complexity guarantee). If the condition on line 3 of the algorithm holds, then:

$$V(i^*) \geq \frac{1}{C}OPT'_{-i^*} - \frac{\varepsilon}{C} \geq \frac{1}{C}OPT_{-i^*} - \frac{\varepsilon}{C}$$

as L is a fractional relaxation of V. Also,

$$OPT \leq OPT_{-i^*} + V(i^*)$$

Hence:

$$OPT \leq (1+C)V(i^*) + \varepsilon \tag{12}$$

Note that $OPT'_{-i^*} \leq OPT'$. If the condition does not hold, from Lemmas 4 and 2:

$$\begin{aligned} V(i^*) &\leq \frac{1}{C}OPT'_{-i^*} + \frac{\varepsilon}{C} \leq \frac{1}{C}\big(2OPT + 2V(i^*)\big) + \frac{\varepsilon}{C} \\ &\leq \frac{1}{C}\left(\frac{2e}{e-1}\big(3V(S_G) + 2V(i^*)\big) + 2V(i^*)\right) + \frac{\varepsilon}{C} \end{aligned}$$

Thus, if C is such that $C(e-1) - 6e + 2 > 0$,

$$V(i^*) \leq \frac{6e}{C(e-1)-6e+2}V(S_G) + \frac{(e-1)\varepsilon}{C(e-1)-6e+2}$$

Finally, using again Lemma 2, we get:

$$\begin{aligned} OPT(V,\mathcal{N},B) \leq \frac{3e}{e-1}&\left(1+\frac{4e}{C(e-1)-6e+2}\right)V(S_G) \\ &+ \frac{2e\varepsilon}{C(e-1)-6e+2} \end{aligned} \tag{13}$$

To minimize the coefficients of V_{i^*} and $V(S_G)$ in (12) and (13) respectively, we wish to chose for $C = C^*$ such that:

$$C^* = \arg\min_C \max\left(1+C, \frac{3e}{e-1}\left(1+\frac{4e}{C(e-1)-6e+2}\right)\right)$$

This equation has two solutions. Only one of those is such that $C(e-1)-6e+2 \geq 0$. This solution is:

$$C^* = \frac{8e-1+\sqrt{64e^2-24e+9}}{2(e-1)} \tag{14}$$

For this solution, $\frac{2e\varepsilon}{C^*(e-1)-6e+2} \leq \varepsilon$. Placing the expression of C^* in (12) and (13) gives the approximation ratio in (11), and concludes the proof of Theorem 1. $\square$

4.2 Proof of Lemma 4

We prove that our multi-linear function F has a property which allows to trade one fractional component of the solution for another until one of them becomes integral, without losing any value. This property is referred to in the literature as *cross-convexity* (see, *e.g.*, [?]), or ε-convexity by Ageev and Sviridenko [?]. Formally, if we define:

$$\tilde{F}_\lambda(\varepsilon) \equiv F\big(\lambda + \varepsilon(e_i - e_j)\big)$$

where e_i and e_j are two vectors of the standard basis of $\mathbb{R}^n$, then $\tilde{F}_\lambda$ is convex. Hence its maximum over the interval:

$$I_\lambda = \Big[\max(-\lambda_i, \lambda_j - 1), \min(1-\lambda_i, \lambda_j)\Big]$$

is attained at one of the boundaries of I_λ for which one of the i-th or the j-th component of λ becomes integral.

The lemma below proves that we can similarly trade a fractional component for another until one of them becomes integral *while maintaining the feasibility of the point at which F is evaluated.* Here, by feasibility of a point λ, we mean that it satisfies the budget constraint $\sum_{i=1}^n \lambda_i c_i \leq B$.

LEMMA 5 (ROUNDING). *For any feasible $\lambda \in [0,1]^n$, there exists a feasible $\bar{\lambda} \in [0,1]^n$ such that at most one of its components is fractional and $F_\mathcal{N}(\lambda) \leq F_\mathcal{N}(\bar{\lambda})$.*

PROOF. We give a rounding procedure which, given a feasible λ with at least two fractional components, returns some feasible λ' with one less fractional component such that:

$$F(\lambda) \leq F(\lambda')$$

Applying this procedure recursively yields the lemma's result. Let us consider such a feasible λ. Let i and j be two fractional components of λ and let us define the following function:

$$F_\lambda(\varepsilon) = F(\lambda_\varepsilon) \quad \text{where} \quad \lambda_\varepsilon = \lambda + \varepsilon\left(e_i - \frac{c_i}{c_j}e_j\right)$$

It is easy to see that if λ is feasible, then:

$$\forall\varepsilon \in \Big[\max\Big(-\lambda_i, (\lambda_j - 1)\frac{c_j}{c_i}\Big), \min\Big(1-\lambda_i, \lambda_j\frac{c_j}{c_i}\Big)\Big], \quad \lambda_\varepsilon \text{ is feasible} \tag{15}$$

Furthermore, the function F_λ is convex; indeed:

$$\begin{aligned}F_\lambda(\varepsilon) = \mathbb{E}_{S'\sim P^\lambda_{\mathcal{N}\setminus\{i,j\}}(S')}\Big[&(\lambda_i+\varepsilon)\Big(\lambda_j - \varepsilon\frac{c_i}{c_j}\Big)V(S'\cup\{i,j\})\\ &+ (\lambda_i+\varepsilon)\Big(1-\lambda_j+\varepsilon\frac{c_i}{c_j}\Big)V(S'\cup\{i\})\\ &+ (1-\lambda_i-\varepsilon)\Big(\lambda_j - \varepsilon\frac{c_i}{c_j}\Big)V(S'\cup\{j\})\\ &+ (1-\lambda_i-\varepsilon)\Big(1-\lambda_j+\varepsilon\frac{c_i}{c_j}\Big)V(S')\Big]\end{aligned}$$

Thus, F_λ is a degree 2 polynomial whose dominant coefficient is:

$$\begin{aligned}\frac{c_i}{c_j}\mathbb{E}_{S'\sim P^\lambda_{\mathcal{N}\setminus\{i,j\}}(S')}\Big[&V(S'\cup\{i\}) + V(S'\cup\{i\})\\ &- V(S'\cup\{i,j\}) - V(S')\Big]\end{aligned}$$

which is positive by submodularity of V. Hence, the maximum of F_λ over the interval given in (15) is attained at one of its limit, at which either the i-th or j-th component of λ_ε becomes integral. $\square$

Next, we prove the central result of bounding L appropriately in terms of F.

LEMMA 6. *For all* $\lambda\in[0,1]^n$, $\frac{1}{2}\,L(\lambda) \le F(\lambda) \le L(\lambda)$.

PROOF. The bound $F_{\mathcal{N}}(\lambda) \le L_{\mathcal{N}(\lambda)}$ follows by the concavity of the $\log\det$ function. To show the lower bound, we first prove that $\frac{1}{2}$ is a lower bound of the ratio $\partial_i F(\lambda)/\partial_i L(\lambda)$, where $\partial_i\cdot$ denotes the partial derivative with respect to the i-th variable.

Let us start by computing the derivatives of F and L with respect to the i-th component. Observe that:

$$\partial_i P^\lambda_{\mathcal{N}}(S) = \begin{cases} P^\lambda_{\mathcal{N}\setminus\{i\}}(S\setminus\{i\}) \text{ if } i\in S\\ -P^\lambda_{\mathcal{N}\setminus\{i\}}(S) \text{ if } i\in\mathcal{N}\setminus S\end{cases}$$

Hence:

$$\partial_i F(\lambda) = \sum_{\substack{S\subseteq\mathcal{N}\\ i\in S}} P^\lambda_{\mathcal{N}\setminus\{i\}}(S\setminus\{i\})V(S) - \sum_{\substack{S\subseteq\mathcal{N}\\ i\in\mathcal{N}\setminus S}} P^\lambda_{\mathcal{N}\setminus\{i\}}(S)V(S)$$

Now, using that every S such that $i\in S$ can be uniquely written as $S'\cup\{i\}$, we can write:

$$\partial_i F(\lambda) = \sum_{\substack{S\subseteq\mathcal{N}\\ i\in\mathcal{N}\setminus S}} P^\lambda_{\mathcal{N}\setminus\{i\}}(S)\big(V(S\cup\{i\}) - V(S)\big)$$

The marginal contribution of i to S can be written as

$$\begin{aligned}V(S\cup\{i\}) - V(S) &= \frac{1}{2}\log\det(I_d + X_S^TX_S + x_ix_i^T)\\ &\quad - \frac{1}{2}\log\det(I_d + X_S^TX_S)\\ &= \frac{1}{2}\log\det(I_d + x_ix_i^T(I_d+X_S^TX_S)^{-1})\\ &= \frac{1}{2}\log(1+x_i^TA(S)^{-1}x_i)\end{aligned}$$

where $A(S) = I_d + X_S^TX_S$. Using this,

$$\partial_i F(\lambda) = \frac{1}{2}\sum_{\substack{S\subseteq\mathcal{N}\\ i\in\mathcal{N}\setminus S}} P^\lambda_{\mathcal{N}\setminus\{i\}}(S)\log\Big(1+x_i^TA(S)^{-1}x_i\Big)$$

The computation of the derivative of L uses standard matrix calculus. Writing $\tilde{A}(\lambda) = I_d + \sum_{i\in\mathcal{N}}\lambda_ix_ix_i^T$:

$$\begin{aligned}\det\tilde{A}(\lambda + h\cdot e_i) &= \det\big(\tilde{A}(\lambda) + hx_ix_i^T\big)\\ &= \det\tilde{A}(\lambda)\big(1+hx_i^T\tilde{A}(\lambda)^{-1}x_i\big)\end{aligned}$$

Hence:

$$\log\det\tilde{A}(\lambda+h\cdot e_i) = \log\det\tilde{A}(\lambda) + hx_i^T\tilde{A}(\lambda)^{-1}x_i + o(h)$$

Finally:

$$\partial_i L(\lambda) = \frac{1}{2}x_i^T\tilde{A}(\lambda)^{-1}x_i$$

For two symmetric matrices A and B, we write $A\succ B$ ($A\succeq B$) if $A-B$ is positive definite (positive semi-definite). This order allows us to define the notion of a *decreasing* as well as *convex* matrix function, similarly to their real counterparts. In particular, matrix inversion is decreasing and convex over symmetric positive definite matrices. In particular,

$$\forall S\subseteq\mathcal{N},\quad A(S)^{-1}\succeq A(S\cup\{i\})^{-1}$$

Observe that:

$$\begin{aligned}\forall S\subseteq\mathcal{N}\setminus\{i\},\quad P^\lambda_{\mathcal{N}\setminus\{i\}}(S) \ge P^\lambda_{\mathcal{N}\setminus\{i\}}(S\cup\{i\})\\ \forall S\subseteq\mathcal{N},\quad P^\lambda_{\mathcal{N}\setminus\{i\}}(S)\ge P^\lambda_{\mathcal{N}}(S)\end{aligned}$$

Hence:

$$\begin{aligned}\partial_i F(\lambda) &\ge \frac{1}{4}\sum_{\substack{S\subseteq\mathcal{N}\\ i\in\mathcal{N}\setminus S}} P^\lambda_{\mathcal{N}\setminus\{i\}}(S)\log\Big(1+x_i^TA(S)^{-1}x_i\Big)\\ &\quad+\frac{1}{4}\sum_{\substack{S\subseteq\mathcal{N}\\ i\in\mathcal{N}\setminus S}} P^\lambda_{\mathcal{N}\setminus\{i\}}(S\cup\{i\})\log\Big(1+x_i^TA(S\cup\{i\})^{-1}x_i\Big)\\ &\ge\frac{1}{4}\sum_{S\subseteq\mathcal{N}} P^\lambda_{\mathcal{N}}(S)\log\Big(1+x_i^TA(S)^{-1}x_i\Big)\end{aligned}$$

Using that $A(S)\succeq I_d$ we get that:

$$x_i^TA(S)^{-1}x_i \le \|x_i\|_2^2\le 1$$

Moreover:

$$\forall x\le 1,\ \log(1+x)\ge x$$

Hence:

$$\partial_i F(\lambda) \ge \frac{1}{4}x_i^T\Big(\sum_{S\subseteq\mathcal{N}} P^\lambda_{\mathcal{N}}(S)A(S)^{-1}\Big)x_i$$

Finally, using that the inverse is a matrix convex function over symmetric positive definite matrices:

$$\begin{aligned}\partial_i F(\lambda) &\geq \frac{1}{4} x_i^T \bigg(\sum_{S\subseteq\mathcal{N}} P_{\mathcal{N}}^{\lambda}(S) A(S) \bigg)^{-1} x_i \\ &= \frac{1}{4} x_i^T \tilde{A}(\lambda)^{-1} x_i \\ &= \frac{1}{2}\partial_i L(\lambda)\end{aligned}$$

Having bound the ratio between the partial derivatives, we now bound the ratio $F(\lambda)/L(\lambda)$ from below. Consider the following cases. First, if the minimum of the ratio $F(\lambda)/L(\lambda)$ is attained at a point interior to the hypercube, then it is a critical point, *i.e.*, $\partial_i\big(F(\lambda)/L(\lambda)\big) = 0$ for all $i \in \mathcal{N}$; hence, at such a critical point:

$$\frac{F(\lambda)}{L(\lambda)} = \frac{\partial_i F(\lambda)}{\partial_i L(\lambda)} \geq \frac{1}{2} \tag{16}$$

Second, if the minimum is attained as λ converges to zero in, *e.g.*, the l_2 norm, by the Taylor approximation, one can write:

$$\frac{F(\lambda)}{L(\lambda)} \sim_{\lambda\rightarrow 0} \frac{\sum_{i\in\mathcal{N}} \lambda_i \partial_i F(0)}{\sum_{i\in\mathcal{N}} \lambda_i \partial_i L(0)} \geq \frac{1}{2},$$

i.e., the ratio $\frac{F(\lambda)}{L(\lambda)}$ is necessarily bounded from below by 1/2 for small enough λ. Finally, if the minimum is attained on a face of the hypercube $[0,1]^n$ (a face is defined as a subset of the hypercube where one of the variable is fixed to 0 or 1), without loss of generality, we can assume that the minimum is attained on the face where the n-th variable has been fixed to 0 or 1. Then, either the minimum is attained at a point interior to the face or on a boundary of the face. In the first sub-case, relation (16) still characterizes the minimum for $i < n$. In the second sub-case, by repeating the argument again by induction, we see that all is left to do is to show that the bound holds for the vertices of the cube (the faces of dimension 1). The vertices are exactly the binary points, for which we know that both relaxations are equal to the value function V. Hence, the ratio is equal to 1 on the vertices. □

Proof of Lemma 4. Let us consider a feasible point $\lambda^* \in [0,1]^n$ such that $L(\lambda^*) = OPT'$. By applying Lemma 6 and Lemma 5 we get a feasible point $\bar{\lambda}$ with at most one fractional component such that:

$$L(\lambda^*) \leq 2F(\bar{\lambda}) \tag{17}$$

Let λ_i denote the fractional component of $\bar{\lambda}$ and S denote the set whose indicator vector is $\bar{\lambda} - \lambda_i e_i$. By definition of the multi-linear extension F:

$$F(\bar{\lambda}) = (1-\lambda_i)V(S) + \lambda_i V(S \cup \{i\})$$

Using the submodularity of V:

$$F(\bar{\lambda}) \leq V(S) + V(i)$$

Note that since $\bar{\lambda}$ is feasible, S is also feasible and $V(S) \leq OPT$. Hence:

$$F(\bar{\lambda}) \leq 2OPT + \max_{i\in\mathcal{N}} V(i) \tag{18}$$

Together, (17) and (18) imply the lemma. □

5. EXTENSION TO OTHER PROBLEMS

5.1 Bayesian Experimental Design

We extend Theorem 1 to a more general Bayesian setting, where it is assumed that the experimenter `E` has a *prior* distribution on β: in particular, β has a multivariate normal prior with zero mean and covariance $\sigma^2 R^{-1} \in \mathbb{R}^{d^2}$ (where σ^2 is the noise variance). `E` estimates β through *maximum a posteriori estimation*: *i.e.*, finding the parameter which maximizes the posterior distribution of β given the observations y_S. Under the linearity assumption (2) and the Gaussian prior on β, maximum a posteriori estimation leads to the following maximization [**?**]:

$$\hat{\beta} = \arg\min_{\beta\in\mathbb{R}^d} \sum_i (y_i - \beta^T x_i)^2 + \beta^T R \beta$$

This optimization, commonly known as *ridge regression*, includes an additional penalty term compared to the least squares estimation (3).

Let $H(\beta)$ be the entropy of β under this distribution, and $H(\beta \mid y_S)$ the entropy of β conditioned on the experiment outcomes y_S, for some $S \subseteq \mathcal{N}$. In this setting, a natural objective, originally proposed by Lindley [**?**], is to select a set of experiments S that maximizes her *information gain*:

$$I(\beta; y_S) = H(\beta) - H(\beta \mid y_S).$$

Assuming normal noise variables, the information gain is equal (up to a constant) to the following value function [**?**]:

$$V(S) = \frac{1}{2}\log\det(R + X_S^T X_S) \tag{19}$$

Our objective for EDP clearly follows from (19) by setting $R = I_d$. Hence, the optimization discussed thus far can be interpreted as a maximization of the information gain when the prior distribution has a covariance $\sigma^2 I_d$, and the experimenter is solving a ridge regression problem with penalty term $\|\beta\|_2^2$.

Our results can be extended to the general Bayesian case, by replacing I_d with the positive semidefinite matrix R. First, we re-set the origin of the value function so that $V(\emptyset) = 0$:

$$\begin{aligned}\tilde{V}(S) &= \frac{1}{2}\log\det(R + X_S^T X_S) - \frac{1}{2}\log\det R \qquad (20)\\ &= \frac{1}{2}\log\det(I_d + R^{-1} X_S^T X_S)\end{aligned}$$

Applying the mechanism described in algorithm 1 and adapting the analysis of the approximation ratio, we get the following result which extends Theorem 1:

THEOREM 3. *There exists a truthful, individually rational and budget feasible mechanism for the objective function $\tilde{V}$ given by* (20)*. Furthermore, for any $\varepsilon > 0$, in time $O(poly(n, d, \log\log\varepsilon^{-1}))$, the algorithm computes a set S^* such that:*

$$OPT \leq \frac{5e-1}{e-1}\frac{2\mu}{\log(1+\mu)} V(S^*) + 5.1 + \varepsilon$$

where μ is the smallest eigenvalue of R.

5.2 D-Optimality and Beyond

We now reexamine the classical D-optimality in (4), which is given by objective (19) with R replaced by the zero matrix. Since (4) may take arbitrarily small negative values, to define a meaningful approximation one would consider the (equivalent) maximization of $V(S) = \det X_S^T X_S$. However, the following lower bound implies that such an optimization goal cannot be attained under the constraints of truthfulness, budget feasibility, and individual rationality.

LEMMA 7. *For any $M > 1$, there is no M-approximate, truthful, budget feasible, individually rational mechanism for a budget feasible reverse auction with value function $V(S) = \det X_S^T X_S$. For any $M > 1$, there is no M-approximate, truthful, budget feasible, individually rational mechanism for a budget feasible reverse auction with $V(S) = \det X_S^T X_S$.*

PROOF. Given $M > 1$, consider $n = 4$ experiments of dimension $d = 2$. For e_1, e_2 the standard basis vectors in $\mathbb{R}^2$, let $x_1 = e_1$, $x_2 = e_1$, and $x_3 = \delta e_1$, $x_4 = \delta e_2$, where $0 < \delta < 1/(M-1)$. Moreover, assume that $c_1 = c_2 = 0.5 + \epsilon$, while $c_3 = c_4 = \epsilon$, for some small $\epsilon > 0$. Suppose, for the sake of contradiction, that there exists a mechanism with approximation ratio M. Then, it must include in the solution S at least one of x_1 or x_2: if not, then $V(S) \leq \delta^2$, while $OPT = (1+\delta)\delta$, a contradiction. Suppose thus that the solution contains x_1. By the monotonicity property, if the cost of experiment x_1 reduces to $B/2 - 3\epsilon$, x_1 will still be in the solution. By threshold payments, experiment x_1 receives in this case a payment that is at least $B/2 + \epsilon$. By individual rationality and budget feasibility, x_2 cannot be included in the solution, so $V(S)$ is at most $(1+\delta)\delta$. However, the optimal solution includes all experiments, and yields $OPT = (1+\delta)^2$, a contradiction. $\square$

Beyond D-optimality, several other objectives such as E-optimality (maximizing the smallest eigenvalue of $X_S^T X_S$) or T-optimality (maximizing $\mathrm{trace}(X_S^T X_S)$) are encountered in the literature [?], though they do not relate to entropy as D-optimality. We leave the task of approaching the maximization of such objectives from a strategic point of view as an open problem.

5.3 Beyond Linear Models

Selecting experiments that maximize the information gain in the Bayesian setup leads to a natural generalization to other learning examples beyond linear regression. In particular, in the more general PAC learning setup [?], the features x_i, $i \in \mathcal{N}$ take values in some general set Ω, called the *query space*. Measurements $y_i \in \mathbb{R}$ are given by

$$y_i = h(x_i) + \varepsilon_i \tag{21}$$

where $h \in \mathcal{H}$ for some subset $\mathcal{H}$ of all possible mappings $h : \Omega \to \mathbb{R}$, called the *hypothesis space*, and ε_i are random variables in $\mathbb{R}$, not necessarily identically distributed, that are independent *conditioned on h*. This model is quite broad, and captures many learning tasks, such as support vector machines (SVM) and linear discriminant analysis (LDA); we give a few concrete examples below:

1. **Generalized Linear Regression.** $\Omega = \mathbb{R}^d$, $\mathcal{H}$ is the set of linear maps $\{h(x) = \beta^T x \text{ s.t. } \beta \in \mathbb{R}^d\}$, and ε_i are independent zero-mean normal variables, where $\mathbb{E}[\varepsilon_i^2] = \sigma_i$.
2. **Logistic Regression.** $\Omega = \mathbb{R}^d$, $\mathcal{H}$ is the set of maps $\{h(x) = \frac{e^{\beta^T x}}{1+e^{\beta^T x}} \text{ s.t. } \beta \in \mathbb{R}^d\}$, and ε_i are independent conditioned on h such that
$$\varepsilon_i = \begin{cases} 1 - h(x_i), & \text{w. prob. } h(x_i) \\ -h(x_i), & \text{w. prob. } 1 - h(x_i) \end{cases}$$
3. **Learning Binary Functions with Bernoulli Noise.** $\Omega = \{0,1\}^d$, and $\mathcal{H}$ is some subset of 2^{Ω}, and
$$\varepsilon_i = \begin{cases} 0, & \text{w. prob. } p \\ \bar{h}(x_i) - h(x_i), & \text{w. prob. } 1 - p \end{cases}$$

In this setup, assume that the experimenter has a prior distribution on the hypothesis $h \in \mathcal{H}$. Then, the information gain objective can be written again as the mutual information between β and y_S.

$$V(S) = H(\beta) - H(\beta \mid y_S), \quad S \subseteq \mathcal{N} \tag{22}$$

This is a monotone set function, and it clearly satisfies $V(\emptyset) = 0$. Though, in general, mutual information is not a submodular function, this specific setup leads indeed to a submodular formulation.

LEMMA 8. *The value function given by the information gain* (22) *is submodular.*

PROOF. Using the chain rule for the conditional entropy we get:

$$V(S) = H(y_S) - H(y_S \mid \beta) = H(y_S) - \sum_{i \in S} H(y_i \mid \beta) \tag{23}$$

where the second equality comes from the independence of the y_i's conditioned on β. Recall that the joint entropy of a set of random variables is a submodular function. Thus, the value function is written in (23) as the sum of a submodular function and a modular function. $\square$

This lemma implies that learning an *arbitrary hypothesis, under an arbitrary prior* when noise is conditionally independent leads to a submodular value function. Hence, we can apply the previously known results to get the following corollary:

COROLLARY 1. *For Bayesian experimental design with the objective given by the information gain* (22)*, there exists a randomized, budget feasible, individually rational, and universally truthful mechanism. This mechanism has a complexity $O\big(poly(n,d)\big)$ and an approximation ratio of* 7.91*.*

In cases where maximizing (22) *can be done in polynomial time in the full-information setup, there exists a randomized, budget feasible, individually rational, and truthful mechanism for Bayesian experimental design. This mechanism has a complexity $O\big(poly(n,d)\big)$ and an approximation ratio of* 8.34*.*

Note however that, in many scenarios covered by this model (including the last two examples above), even computing the entropy under a given set might be a hard task—*i.e.*, the value query model may not apply. Hence, identifying learning tasks in the above class for which truthful or universally truthful constant approximation mechanisms exist, or studying these problems in the context of stronger query models such as the demand model [**?**, **?**] remains an interesting open question.

APPENDIX

LEMMA 9. *Our mechanism for EDP is monotone and budget feasible.*

PROOF. Consider an agent i with cost c_i that is selected by the mechanism, and suppose that she reports a cost $c_i' \leq c_i$ while all other costs stay the same. Suppose that when i reports c_i, $L(\xi) \geq CV(i^*)$; then, as $s_i(c_i, c_{-i}) = 1$, $i \in S_G$. By reporting a cost $c_i' \leq c_i$, i may be selected at an earlier iteration of the greedy algorithm. Denote by S_i (resp. S_i') the set to which i is added when reporting cost c_i (resp. c_i'). We have $S_i' \subseteq S_i$; in addition, $S_i' \subseteq S_G'$, the set selected by the greedy algorithm under (c_i', c_{-i}); if not, then greedy selection would terminate prior to selecting i also when she reports c_i, a contradiction. Moreover, we have

$$c_i' \leq c_i \leq \frac{B}{2}\frac{V(S_i \cup \{i\}) - V(S_i)}{V(S_i \cup \{i\})} \leq \frac{B}{2}\frac{V(S_i' \cup \{i\}) - V(S_i')}{V(S_i' \cup \{i\})}$$

by the monotonicity and submodularity of V. Hence $i \in S_G'$. As $L(\xi)$, is the optimal value of (8) under relaxation L when i^* is excluded from $\mathcal{N}$, reducing the costs can only increase this value, so under $c_i' \leq c_i$ the greedy set is still allocated and $s_i(c_i', c_{-i}) = 1$. Suppose now that when i reports c_i, $L(\xi) < CV(i^*)$. Then $s_i(c_i, c_{-i}) = 1$ iff $i = i^*$. Reporting $c_{i^*}' \leq c_{i^*}$ does not change $V(i^*)$ nor $L(\xi) \leq CV(i^*)$; thus $s_{i^*}(c_{i^*}', c_{-i^*}) = 1$, so the mechanism is monotone.

To show budget feasibility, suppose that $L(\xi) < CV(i^*)$. Then the mechanism selects i^*. Since the bid of i^* does not affect the above condition, the threshold payment of i^* is B and the mechanism is budget feasible. Suppose that $L(\xi) \geq CV(i^*)$. Denote by S_G the set selected by the greedy algorithm, and for $i \in S_G$, denote by S_i the subset of the solution set that was selected by the greedy algorithm just prior to the addition of i—both sets determined for the present cost vector c. Then for any submodular function V, and for all $i \in S_G$:

$$\text{if } c_i' \geq \frac{V(S_i \cup \{i\}) - V(S)}{V(S_G)} B \text{ then } s_i(c_i', c_{-i}) = 0 \qquad (24)$$

In other words, if i increases her cost to a value higher than $\frac{V(S_i \cup \{i\}) - V(S)}{V(S_G)}$, she will cease to be in the selected set S_G. As a result, (24) implies that the threshold payment of user i is bounded by the above quantity. Hence, the total payment is bounded by the telescopic sum:

$$\sum_{i \in S_G} \frac{V(S_i \cup \{i\}) - V(S_i)}{V(S_G)} B = \frac{V(S_G) - V(\emptyset)}{V(S_G)} B = B \quad \square$$