

Budget Feasible Mechanisms for Experimental Design

Thibaut Horel

École Normale Supérieure
thibaut.horel@normalesup.org

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 \mathbf{E} has access to a population of n potential experiment subjects $i \in \{1, \dots, n\}$, each associated with a vector of features $x_i \in \mathbb{R}^d$. Conducting an experiment with subject i reveals an unknown value $y_i \in \mathbb{R}$ to \mathbf{E} . \mathbf{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.g.*, through linear regression. As a proxy for various practical constraints, \mathbf{E} may select only a subset of subjects on which to conduct the experiment.

We initiate the study of budgeted mechanisms for experimental design. In this setting, \mathbf{E} has a budget B . Each subject i declares an associated cost $c_i > 0$ to be part of the experiment, and must be paid at least her cost. In particular, the *Experimental Design Problem* (EDP) is to find a set S of subjects for the experiment that maximizes $V(S) = \log \det(I_d + \sum_{i \in S} x_i x_i^T)$ under the constraint $\sum_{i \in S} c_i \leq B$; our objective function corresponds to the information gain in parameter β that is learned through linear regression methods, and is related to the so-called D -optimality criterion. Further, the subjects are *strategic* and may lie about their costs. Thus, we need to design a mechanism for EDP with suitable properties.

We present a deterministic, polynomial time, budget feasible mechanism scheme, that is approximately truthful and yields a constant (≈ 12.98) factor approximation to EDP. By applying previous work on budget feasible mechanisms with a submodular objective, one could *only* have derived either an exponential time deterministic mechanism or a randomized polynomial time mechanism. We also establish that no truthful, budget-feasible mechanism is possible within a factor 2 approximation, and show how to generalize our approach to a wide class of learning problems, beyond linear regression.

1 Introduction

In the classic setting of experimental design [26, 4], 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 measure a certain inherent property of the subjects by performing an experiment: the outcome y_i of the experiment on a subject i is unknown to E before the experiment is performed.

Typically, E has a hypothesis on the relationship between x_i 's and y_i 's. Due to its simplicity, as well as its ubiquity in statistical analysis, a large body of work has focused on linear hypotheses: *i.e.*, it is assumed that there exists a $\beta \in \mathbb{R}^d$ such that

$$y_i = \beta^T x_i + \varepsilon_i,$$

for all $i \in \{1, \dots, n\}$, where ε_i are zero-mean, i.i.d. random variables. Conducting the experiments and obtaining the measurements y_i lets E estimate β , *e.g.*, through linear regression.

The above experimental design scenario has many applications. Regression over personal data collected through surveys or experimentation is the cornerstone of marketing research, as well as research in a variety of experimental sciences such as medicine and sociology. Crucially, statistical analysis of user data is also a widely spread practice among Internet companies, which routinely use machine learning techniques over vast records of user data to perform inference and classification tasks integral to their daily operations. Beyond linear regression, there is a rich literature about estimation procedures, as well as about means of quantifying the quality of the produced estimate [26]. 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 a β that approximates the true parameter of the underlying population [16, 21, 10, 7].

We depart from this classical setup by viewing experimental design in a strategic setting, and by studying budgeted mechanism design issues. In our setting, experiments cannot be manipulated and hence measurements are reliable. E has a total budget of B to conduct all the experiments. There is a cost c_i associated with experimenting on subject i which varies from subject to subject. This cost c_i is determined by the subject i and reported to E ; subjects are strategic and may misreport these costs. Intuitively, c_i may be viewed as the cost 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 intrinsic worth of the data to the subject. The economic aspect of paying subjects has always been inherent in experimental design: experimenters often work within strict budgets and design creative incentives. Subjects often negotiate better incentives or higher payments. However, we are not aware of a principled study of this setting from a strategic point of view, when subjects declare their costs and therefore determine their payment.

Such a setting is increasingly realistic, given the growth of these experiments over the Internet.

Our contributions are as follows.

- We initiate the study of experimental design in the presence of a budget and strategic subjects. In particular, we formulate the *Experimental Design Problem* (EDP) as follows: the experimenter \mathbf{E} wishes to find a set S of subjects to maximize

$$V(S) = \log \det \left(I_d + \sum_{i \in S} x_i x_i^T \right) \quad (1)$$

subject to a budget constraint $\sum_{i \in S} c_i \leq B$, where B is \mathbf{E} 's budget. When subjects are strategic, the above problem can be naturally approached as a *budget feasible mechanism design* problem, as introduced by Singer [29].

The objective function, which is the key, is formally obtained by optimizing the information gain in β when the latter is learned through ridge regression, and is related to the so-called D -optimality criterion [26, 4].

- We present a polynomial time mechanism scheme for EDP that is approximately truthful and yields a constant factor (≈ 12.98) approximation to the optimal value of (1). In contrast to this, we show that no truthful, budget-feasible mechanisms are possible for EDP within a factor 2 approximation.

We note that the objective (1) is submodular. Using this fact, applying previous results on budget feasible mechanism design under general submodular objectives [29, 11] would yield either a deterministic, truthful, constant-approximation mechanism that requires exponential time, or a non-deterministic, (universally) truthful, poly-time mechanism that yields a constant approximation ratio only *in expectation* (i.e., its approximation guarantee for a given instance may in fact be unbounded).

From a technical perspective, we propose a convex optimization problem and establish that its optimal value is within a constant factor from the optimal value of EDP. In particular, we show our relaxed objective is within a constant factor from the so-called multi-linear extension of (1), which in turn can be related to (1) through pipage rounding. We establish the constant factor to the multi-linear extension by bounding the partial derivatives of these two functions; we achieve the latter by exploiting convexity properties of matrix functions over the convex cone of positive semidefinite matrices.

Our convex relaxation of EDP involves maximizing a self-concordant function subject to linear constraints. Its optimal value can be computed with arbitrary accuracy in polynomial time using the so-called barrier method. However, the outcome of this computation may not necessarily be monotone, a property needed in

designing a truthful mechanism. Nevertheless, we construct an algorithm that solves the above convex relaxation and is “almost” monotone; we achieve this by applying the barrier method on a set perturbed constraints, over which our objective is “sufficiently” concave. In turn, we show how to employ this algorithm to design a poly-time, δ -truthful, constant-approximation mechanism for EDP.

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. We present our convex relaxation to EDP in Section 4 and use it to construct our mechanism in Section 5. We conclude in Section 6. All proofs of our technical results are provided in the appendix.

2 Related work

Budget Feasible Mechanisms for General Submodular Functions Budget feasible mechanism design was originally proposed by Singer [29]. 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. [11] improve this result by providing a 7.91-approximate mechanism, and show a corresponding lower bound of 2 among universally truthful randomized mechanisms for submodular maximization.

The above approximation guarantees hold for the expected value of the randomized mechanism: for a given instance, the approximation ratio provided by the mechanism may in fact be unbounded. 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.*, propose 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 deterministic mechanisms, improving upon an earlier bound of 2 by Singer [29].

Budget Feasible Mechanisms on Specific Problems 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 [29]. Singer also provides a 7.32-approximate truthful mechanism for the budget feasible version of MATCHING, and a corresponding lower bound of 2 [29]. Improving an earlier result by Singer, Chen et al. [11] 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 [30].

The deterministic mechanisms for KNAPSACK [11] and COVERAGE [30] follow the same general framework, which we also employ in our mechanism for EDP. We describe this framework in detail in Section 5. Both of these mechanisms rely on approximating the optimal solution to the underlying combinatorial problem by a well-known linear program (LP) relaxation [1], which can be solved exactly in polynomial time. No such relaxation exists for EDP, which is unlikely to be approximable through an LP due to its logarithmic objective. We develop instead a convex relaxation to EDP; though, contrary to the above LP relaxations, this cannot be solved exactly, we establish that it can be incorporated in the framework of [11, 30] to yield a δ -truthful mechanism for EDP.

Beyond Submodular Objectives Beyond submodular objectives, it is known that no truthful mechanism with approximation ratio smaller than $n^{1/2-\varepsilon}$ exists for maximizing fractionally subadditive functions (a class that includes submodular functions) assuming access to a value query oracle [29]. Assuming access to a stronger oracle (the *demand* oracle), there exists a truthful, $O(\log^3 n)$ -approximate mechanism [12] as well as a universally truthful, $O(\frac{\log n}{\log \log n})$ -approximate mechanism for subadditive maximization [6]. 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 [6]. Posted price, rather than direct revelation mechanisms, are also studied in [5].

Monotone Approximations in Combinatorial Auctions Relaxations of combinatorial problems are prevalent in *combinatorial auctions*, in which an auctioneer aims at maximizing a set function which is the sum of utilities of strategic bidders (*i.e.*, the social welfare). As noted by Archer et al. [3], approximations to this maximization must preserve incentive compatibility and truthfulness. Most approximation algorithms do not preserve these properties, hence specific relaxations, and corresponding roundings to an integral solution, must be constructed. Archer et al. [3] propose a randomized rounding of the LP relaxation of the SETPACKING problem, yielding a mechanism which is *truthful-in-expectation*. Lavi and Swamy [20] construct randomized truthful-in-expectation mechanisms for several combinatorial auctions, improving the approximation ratio in [3], by treating the fractional solution of an LP as a probability distribution from which they sample integral solutions.

Beyond LP relaxations, Dughmi et al. [14] propose truthful-in-expectation mechanisms for combinatorial auctions in which the bidders’ utilities are matroid rank sum functions (applied earlier to the COMBINATORIALPUBLICPROJECTS problem [13]). Their framework relies on solving a convex optimization problem which can only be solved approximately. As in [20], they also treat the fractional solution as

a distribution over which they sample integral solutions. The authors ensure that a solver is applied to a “well-conditioned” problem, which resembles the technical challenge we face in Section 4.2. However, we seek a deterministic mechanism and δ -truthfulness, not truthfulness-in-expectation. In addition, our objective is not a matroid rank sum function. As such, both the methodology for dealing with problems that are not “well-conditioned” as well as the approximation guarantees of the convex relaxation in [14] do not readily extend to EDP.

Briest et al. [8] construct monotone FPTAS for problems that can be approximated through rounding techniques, which in turn can be used to construct truthful, deterministic, constant-approximation mechanisms for corresponding combinatorial auctions. EDP is not readily approximable through such rounding techniques; as such, we rely on a relaxation to approximate it.

δ -Truthfulness and Differential Privacy The notion of δ -truthfulness has attracted considerable attention recently in the context of differential privacy (see, *e.g.*, the survey by Pai and Roth [25]). McSherry and Talwar [22] were the first to observe that any ε -differentially private mechanism must also be δ -truthful in expectation, for $\delta = 2\varepsilon$. This property was used to construct δ -truthful (in expectation) mechanisms for a digital goods auction [22] and for α -approximate equilibrium selection [17]. Nissim et al. [24] propose a framework for converting a differentially private mechanism to a truthful-in-expectation mechanism by randomly selecting between a differentially private mechanism with good approximation guarantees, and a truthful mechanism. They apply their framework to the FACILITYLOCATION problem. We depart from the above works in seeking a deterministic mechanism for EDP, and using a stronger notion of δ -truthfulness.

3 Preliminaries

3.1 Linear Regression and Experimental Design

The theory of experimental design [26, 4, 10] considers the following formal setting. Suppose that an experimenter E wishes to conduct k among n possible experiments. Each experiment $i \in \mathcal{N} \equiv \{1, \dots, n\}$ is associated with a set of parameters (or features) $x_i \in \mathbb{R}^d$, normalized so that

$$b \leq \|x_i\|_2^2 \leq 1,$$

for some $b > 0$. 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 .

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 magnitude of the coefficient β_i captures the effect that feature i has on the measured variable, and its sign captures whether the correlation is positive or negative.

The purpose of these experiments is to allow \mathbf{E} to estimate the model β . In particular, assume that the experimenter \mathbf{E} has a *prior* distribution on β , *i.e.*, 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). Then, \mathbf{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 [15]:

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

where the last equality is obtained by setting $\nabla_{\beta} \Pr(\beta \mid y_S)$ to zero and solving the resulting linear system; in (3), $X_S \equiv [x_i]_{i \in S} \in \mathbb{R}^{|S| \times d}$ is the matrix of experiment features and $y_S \equiv [y_i]_{i \in S} \in \mathbb{R}^{|S|}$ are the observed measurements. This optimization, commonly known as *ridge regression*, includes an additional quadratic penalty term $\beta^T R \beta$ compared to the standard least squares estimation.

Let $V : 2^{\mathcal{N}} \rightarrow \mathbb{R}$ be a *value function*, quantifying how informative a set of experiments S is in estimating β . The classical 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 literature [26, 7]; one that has natural advantages is the *information gain*:

$$V(S) = I(\beta; y_S) = H(\beta) - H(\beta \mid y_S).\tag{4}$$

which is the entropy reduction on β after the revelation of y_S (also known as the mutual information between y_S and β). 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 reduction of its estimator. Under the linear model (2), and the Gaussian prior, the information gain takes the following form (see, *e.g.*, [10]):

$$I(\beta; y_S) = \frac{1}{2} \log \det(R + X_S^T X_S) - \frac{1}{2} \log \det R\tag{5}$$

Maximizing $I(\beta; y_S)$ is therefore equivalent to maximizing $\log \det(R + X_S^T X_S)$, which is known in the experimental design literature as the Bayes D -optimality criterion [26, 4, 10].

Our analysis will focus on the case of a *homotropic* prior, in which the prior covariance is the identity matrix, *i.e.*, $R = I_d \in \mathbb{R}^{d \times d}$. Intuitively, this corresponds to the simplest prior, in which no direction of \mathbb{R}^d is a priori favored; equivalently, it also corresponds to the case where ridge regression estimation (3) performed by \mathbf{E} has a penalty term $\|\beta\|_2^2$. A generalization of our results to arbitrary covariance matrices R can be found in Appendix H.

3.2 Budget-Feasible Experimental Design: Full Information Case

Instead of the cardinality constraint in classical experimental design discussed above, we consider a budget-constrained version. Each experiment is associated with a cost $c_i \in \mathbb{R}_+$. The cost c_i can capture, *e.g.*, the amount the subject i deems sufficient to incentivize her participation in the experiment. The experimenter \mathbf{E} is limited by a budget $B \in \mathbb{R}_+$. In the full-information case, experiment costs are common knowledge; as such, the experimenter wishes to solve:

EXPERIMENTALDESIGNPROBLEM (EDP)

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

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

W.l.o.g., we assume that $c_i \in [0, B]$ for all $i \in \mathcal{N}$, as no i with $c_i > B$ can be in an S satisfying (6b). Denote by

$$OPT = \max_{S \subseteq \mathcal{N}} \left\{ V(S) \mid \sum_{i \in S} c_i \leq B \right\} \quad (7)$$

the optimal value achievable in the full-information case. EDP, as defined above, is NP-hard; to see this, note that KNAPSACK reduces to EDP with dimension $d = 1$ by mapping the weight of each item, say, w_i , to an experiment with $x_i^2 = w_i$.

The value function (6a) has the following properties, which are proved in Appendix A. First, it is non-negative, *i.e.*, $V(S) \geq 0$ for all $S \subseteq \mathcal{N}$. Second, it is also monotone, *i.e.*, $V(S) \leq V(T)$ for all $S \subseteq T$, with $V(\emptyset) = 0$. Finally, it is submodular, *i.e.*, $V(S \cup \{i\}) - V(S) \geq V(T \cup \{i\}) - V(T)$ for all $S \subseteq T \subseteq \mathcal{N}$ and $i \in \mathcal{N}$. The above imply that a greedy algorithm yields a constant approximation ratio to EDP. In particular, consider the greedy algorithm in which, for $S \subseteq \mathcal{N}$ the set constructed thus far, the next element i included is the one which maximizes the *marginal-value-per-cost*, *i.e.*, $i = \arg \max_{j \in \mathcal{N} \setminus S} (V(S \cup \{j\}) - V(S))/c_j$. This is repeated until adding an element in S exceeds the budget B . Denote by S_G the

set constructed by this heuristic and let $i^* = \arg \max_{i \in \mathcal{N}} V(\{i\})$ be the element of maximum singleton value. Then, the following algorithm:

$$\text{if } V(\{i^*\}) \geq V(S_G) \text{ return } \{i^*\} \text{ else return } S_G \quad (8)$$

yields an approximation ratio of $\frac{5e}{e-1}$ [29]; this can be further improved to $\frac{e}{e-1}$ using more complicated greedy set constructions [18, 31].

3.3 Budget-Feasible Experimental Design: Strategic Case

We study the following *strategic* setting, in which the costs c_i are *not* common knowledge and their reporting can be manipulated by the experiment subjects. The latter are strategic and wish to maximize their utility, which is the difference of the payment they receive and their true cost. We note that, though subjects may misreport c_i , they cannot lie about x_i (*i.e.*, all public features are verifiable prior to the experiment) nor y_i (*i.e.*, the subject cannot falsify her measurement). In this setting, experimental design reduces to a *budget feasible reverse auction*, as introduced by Singer [29]; we review the formal definition in Appendix D. In short, given a budget B and a value function $V : 2^{\mathcal{N}} \rightarrow \mathbb{R}_+$, a *reverse auction mechanism* $\mathcal{M} = (S, p)$ comprises (a) an *allocation function*¹ $S : \mathbb{R}_+^n \rightarrow 2^{\mathcal{N}}$, determining the set of experiments to be purchased, and (b) a *payment function* $p : \mathbb{R}_+^n \rightarrow \mathbb{R}_+^n$, determining the payments $[p_i(c)]_{i \in \mathcal{N}}$ received by experiment subjects.

We seek mechanisms that are *normalized* (unallocated experiments receive zero payments), *individually rational* (payments for allocated experiments exceed costs), have *no positive transfers* (payments are non-negative), and are *budget feasible* (the sum of payments does not exceed the budget B). We relax the notion of truthfulness to δ -*truthfulness*, requiring that reporting one's true cost is an *almost-dominant* strategy: no subject increases their utility by reporting a cost that differs more than $\delta > 0$ from their true cost. Under this definition, a mechanism is truthful if $\delta = 0$. In addition, we would like the allocation $S(c)$ to be of maximal value; however, δ -truthfulness, as well as the hardness of EDP, preclude achieving this goal. Hence, we seek mechanisms with that are (α, β) -*approximate*, *i.e.*, there exist $\alpha \geq 1$ and $\beta > 0$ s.t. $OPT \leq \alpha V(S(c)) + \beta$, and are *computationally efficient*, in that S and p can be computed in polynomial time.

We note that the constant approximation algorithm (8) breaks truthfulness. Though this is not true for all submodular functions (see, *e.g.*, [29]), it is true for the objective of EDP: we show this in Appendix I. This motivates our study of more complex mechanisms.

¹Note that S would be more aptly termed as a *selection* function, as this is a reverse auction, but we retain the term “allocation” to align with the familiar term from standard auctions.

4 Approximation Results

Previous approaches towards designing truthful, budget feasible mechanisms for KNAPSACK [11] and COVERAGE [30] build upon polynomial-time algorithms that compute an approximation of OPT , the optimal value in the full information case. Crucially, to be used in designing a truthful mechanism, such algorithms need also to be *monotone*, in the sense that decreasing any cost c_i leads to an increase in the estimation of OPT ; the monotonicity property precludes using traditional approximation algorithms.

In the first part of this section, we address this issue by designing a convex relaxation of EDP, and showing that its solution can be used to approximate OPT . The objective of this relaxation is concave and self-concordant [7] and, as such, there exists an algorithm that solves this relaxed problem with arbitrary accuracy in polynomial time. Unfortunately, the output of this algorithm may not necessarily be monotone. Nevertheless, in the second part of this section, we show that a solver of the relaxed problem can be used to construct a solver that is “almost” monotone. In Section 5, we show that this algorithm can be used to design a δ -truthful mechanism for EDP.

4.1 A Convex Relaxation of EDP

A classical way of relaxing combinatorial optimization problems is *relaxing by expectation*, using the so-called *multi-linear* extension of the objective function V (see, e.g., [9, 34, 14]). This is because this extension can yield approximation guarantees for a wide class of combinatorial problems through *pipage rounding*, a technique proposed by Ageev and Sviridenko [1]. Crucially for our purposes, such relaxations in general preserve monotonicity which, as discussed, is required in mechanism design.

Formally, let $P_{\mathcal{N}}^{\lambda}$ be a probability distribution over \mathcal{N} parametrized by $\lambda \in [0, 1]^n$, where a set $S \subseteq \mathcal{N}$ sampled from $P_{\mathcal{N}}^{\lambda}$ is constructed as follows: each $i \in \mathcal{N}$ is selected to be in S independently with probability λ_i , i.e., $P_{\mathcal{N}}^{\lambda}(S) \equiv \prod_{i \in S} \lambda_i \prod_{i \in \mathcal{N} \setminus S} (1 - \lambda_i)$. Then, the *multi-linear* extension $F : [0, 1]^n \rightarrow \mathbb{R}$ of V is defined as the expectation of V under the distribution $P_{\mathcal{N}}^{\lambda}$:

$$F(\lambda) \equiv \mathbb{E}_{S \sim P_{\mathcal{N}}^{\lambda}}[V(S)] = \mathbb{E}_{S \sim P_{\mathcal{N}}^{\lambda}} \left[\log \det \left(I_d + \sum_{i \in S} x_i x_i^T \right) \right], \quad \lambda \in [0, 1]^n. \quad (9)$$

Function F is an extension of V to the domain $[0, 1]^n$, as it equals V on integer inputs: $F(\mathbb{1}_S) = V(S)$ for all $S \subseteq \mathcal{N}$, where $\mathbb{1}_S$ denotes the indicator vector of S . Contrary to problems such as KNAPSACK, the multi-linear extension (9) cannot be optimized in polynomial time for the value function V we study here, given by (6a).

Hence, we introduce an extension $L : [0, 1]^n \rightarrow \mathbb{R}$ s.t.

$$\begin{aligned} L(\lambda) &\equiv \log \det \left(I_d + \sum_{i \in \mathcal{N}} \lambda_i x_i x_i^T \right) \\ &= \log \det \left(\mathbb{E}_{S \sim P_{\mathcal{N}}^\lambda} \left[I_d + \sum_{i \in S} x_i x_i^T \right] \right), \quad \lambda \in [0, 1]^n \end{aligned} \tag{10}$$

Note that L also extends V , and follows naturally from the multi-linear extension by swapping the expectation and log det in (9). Crucially, it is *strictly concave* on $[0, 1]^n$, a fact that we exploit in the next section to maximize L subject to the budget constraint in polynomial time.

Our first technical lemma relates the concave extension L to the multi-linear extension F :

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

The proof of this lemma can be found in Appendix B.1. In short, exploiting the concavity of the log det function over the set of positive semi-definite matrices, we first bound the ratio of all partial derivatives of F and L . We then show that the bound on the ratio of the derivatives also implies a bound on the ratio F/L .

Armed with this result, we subsequently use pipage rounding to show that any λ that maximizes the multi-linear extension F can be rounded to an “almost” integral solution. More specifically, given a set of costs $c \in \mathbb{R}_+^n$, we say that a $\lambda \in [0, 1]^n$ is feasible if it belongs to the set

$$\mathcal{D}_c = \{ \lambda \in [0, 1]^n : \sum_{i \in \mathcal{N}} c_i \lambda_i \leq B \}. \tag{11}$$

Then, the following lemma holds:

Lemma 2 (Rounding). *For any feasible $\lambda \in \mathcal{D}_c$, there exists a feasible $\bar{\lambda} \in \mathcal{D}_c$ such that (a) $F(\lambda) \leq F(\bar{\lambda})$, and (b) at most one of the coordinates of $\bar{\lambda}$ is fractional.*

The proof of this lemma is in Appendix B.2, and follows the main steps of the pipage rounding method of Ageev and Sviridenko [1]. Together, Lemma 1 and Lemma 2 imply that OPT , the optimal value of EDP, can be approximated by solving the following convex optimization problem:

$$\begin{aligned} \text{Maximize:} \quad & L(\lambda) \\ \text{subject to:} \quad & \lambda \in \mathcal{D}_c \end{aligned} \tag{P_c}$$

In particular, for $L_c^* \equiv \max_{\lambda \in \mathcal{D}_c} L(\lambda)$ the optimal value of (P_c) , the following holds:

Proposition 1. $OPT \leq L_c^* \leq 2OPT + 2 \max_{i \in \mathcal{N}} V(i)$.

The proof of this proposition can be found in Appendix B.3. As we discuss in the next section, L_c^* can be computed by a poly-time algorithm at arbitrary accuracy. However, the outcome of this computation may not necessarily be monotone; we address this by converting this poly-time estimator of L_c^* to one that is “almost” monotone.

4.2 Polynomial-Time, Almost-Monotone Approximation

The log det objective function of (P_c) is strictly concave and *self-concordant* [7]. The maximization of a concave, self-concordant function subject to a set of linear constraints can be performed through the *barrier method* (see, e.g., [7] Section 11.5.5 for general self-concordant optimization as well as [33] for a detailed treatment of the log det objective). The performance of the barrier method is summarized in our case by the following lemma:

Lemma 3 (Boyd and Vandenberghe [7]). *For any $\varepsilon > 0$, the barrier method computes an approximation \hat{L}_c^* that is ε -accurate, i.e., it satisfies $|\hat{L}_c^* - L_c^*| \leq \varepsilon$, in time $O(\text{poly}(n, d, \log \log \varepsilon^{-1}))$. The same guarantees apply when maximizing L subject to an arbitrary set of $O(n)$ linear constraints.*

Clearly, the optimal value L_c^* of (P_c) is monotone in c : formally, for any two $c, c' \in \mathbb{R}_+^n$ s.t. $c \leq c'$ coordinate-wise, $\mathcal{D}_{c'} \subseteq \mathcal{D}_c$ and thus $L_c^* \geq L_{c'}^*$. Hence, the map $c \mapsto L_c^*$ is non-increasing. Unfortunately, the same is not true for the output \hat{L}_c^* of the barrier method: there is no guarantee that the ε -accurate approximation \hat{L}_c^* exhibits any kind of monotonicity.

Nevertheless, we prove that it is possible to use the barrier method to construct an approximation of L_c^* that is “almost” monotone. More specifically, given $\delta > 0$, we say that $f : \mathbb{R}^n \rightarrow \mathbb{R}$ is δ -decreasing if $f(x) \geq f(x + \mu e_i)$, for all $i \in \mathcal{N}, x \in \mathbb{R}^n, \mu \geq \delta$, where e_i is the i -th canonical basis vector of \mathbb{R}^n . In other words, f is δ -decreasing if increasing any coordinate by δ or more at input x ensures that the output will be at most $f(x)$.

Our next technical result establishes that, using the barrier method, it is possible to construct an algorithm that computes L_c^* at arbitrary accuracy in polynomial time and is δ -decreasing. We achieve this by restricting the optimization over a subset of \mathcal{D}_c at which the concave relaxation L is “sufficiently” concave. Formally, for $\alpha \geq 0$ let

$$\mathcal{D}_{c,\alpha} \equiv \{\lambda \in [\alpha, 1]^n : \sum_{i \in \mathcal{N}} c_i \lambda_i \leq B\} \subseteq \mathcal{D}_c.$$

Note that $\mathcal{D}_c = \mathcal{D}_{c,0}$. Consider the following perturbation of the concave relaxation (P_c) :

$$\begin{aligned} &\text{Maximize:} && L(\lambda) \\ &\text{subject to:} && \lambda \in \mathcal{D}_{c,\alpha} \end{aligned} \tag{P_{c,\alpha}}$$

Algorithm 1

Input: $B \in \mathbb{R}_+$, $c \in [0, B]^n$, $\delta \in (0, 1]$, $\varepsilon \in (0, 1]$

1: $\alpha \leftarrow \varepsilon(\delta/B + n^2)^{-1}$

2: Use the barrier method to solve $(P_{c,\alpha})$ with accuracy $\varepsilon' = \frac{1}{2^{n+1}B}\alpha\delta b$; denote the output by $\hat{L}_{c,\alpha}^*$

3: **return** $\hat{L}_{c,\alpha}^*$

Our construction of a δ -decreasing, ε -accurate approximator of L_c^* proceeds as follows: first, it computes an appropriately selected lower bound α ; using this bound, it solves the perturbed problem $(P_{c,\alpha})$ using the barrier method, also at an appropriately selected accuracy ε' , obtaining thus a ε' -accurate approximation of $L_{c,\alpha}^* \equiv \max_{\lambda \in \mathcal{D}_{c,\alpha}} L(\lambda)$. The corresponding output is returned as an approximation of L_c^* . A summary of the algorithm and the specific choices of α and ε' are given in Algorithm 1. The following proposition, which is proved in Appendix C, establishes that this algorithm has both properties we desire:

Proposition 2. *For any $\delta \in (0, 1]$ and any $\varepsilon \in (0, 1]$, Algorithm 1 computes a δ -decreasing, ε -accurate approximation of L_c^* . The running time of the algorithm is $O(\text{poly}(n, d, \log \log \frac{B}{b\varepsilon\delta}))$.*

We note that the execution of the barrier method on the restricted set $\mathcal{D}_{c,\alpha}$ is necessary. The algorithm's output when executed over the entire domain may not necessarily be δ -decreasing, even when the approximation accuracy is small. This is because costs become saturated when the optimal $\lambda \in \mathcal{D}_c$ lies at the boundary: increasing them has no effect on the objective. Forcing the optimization to happen “off” the boundary ensures that this does not occur, while taking α to be small ensures that this perturbation does not cost much in terms of approximation accuracy.

5 Mechanism for EDP

In this section we use the δ -decreasing, ε -accurate algorithm solving the convex optimization problem (P_c) to design a mechanism for EDP. The construction follows a methodology proposed in [29] and employed by Chen et al. [11] and Singer [30] to construct deterministic, truthful mechanisms for KNAPSACK and COVERAGE respectively. We briefly outline this below (see also Algorithm 2 in Appendix F for a detailed description).

Recall from Section 3.2 that $i^* \equiv \arg \max_{i \in \mathcal{N}} V(\{i\})$ is the element of maximum value, and S_G is a set constructed greedily, by selecting elements of maximum marginal value per cost. The general framework used by Chen et al. [11] and by Singer [30] for the KNAPSACK and COVERAGE value functions constructs an allocation as follows. First, a polynomial-time, monotone approximation of OPT is

computed over all elements excluding i^* . The outcome of this approximation is compared to $V(\{i^*\})$: if it exceeds $V(\{i^*\})$, then the mechanism constructs an allocation S_G greedily; otherwise, the only item allocated is the singleton $\{i^*\}$. Provided that the approximation used is within a constant from OPT , the above allocation can be shown to also yield a constant approximation to OPT . Furthermore, using Myerson's Theorem [23], it can be shown that this allocation combined with *threshold payments* (see Lemma 4 below) constitute a truthful mechanism.

The approximation algorithms used in [11, 30] are LP relaxations, and thus their outputs are monotone and can be computed exactly in polynomial time. We show that the convex relaxation (P_c), which can be solved by an ε -accurate, δ -decreasing algorithm, can be used to construct a δ -truthful, constant approximation mechanism, by being incorporated in the same framework.

To obtain this result, we use the following modified version of Myerson's theorem [23], whose proof we provide in Appendix E.

Lemma 4. *A normalized mechanism $\mathcal{M} = (S, p)$ for a single parameter auction is δ -truthful if: (a) S is δ -monotone, i.e., for any agent i and $c'_i \leq c_i - \delta$, for any fixed costs c_{-i} of agents in $\mathcal{N} \setminus \{i\}$, $i \in S(c_i, c_{-i})$ implies $i \in S(c'_i, c_{-i})$, and (b) agents are paid threshold payments, i.e., for all $i \in S(c)$, $p_i(c) = \inf\{c'_i : i \in S(c'_i, c_{-i})\}$.*

Lemma 4 allows us to incorporate our relaxation in the above framework, yielding the following theorem:

Theorem 1. *For any $\delta \in (0, 1]$, and any $\varepsilon \in (0, 1]$, there exists a δ -truthful, individually rational and budget feasible mechanism for EDP that runs in time $O(\text{poly}(n, d, \log \log \frac{B}{b\varepsilon\delta}))$ and allocates a set S^* such that $OPT \leq \frac{10e-3+\sqrt{64e^2-24e+9}}{2(e-1)}V(S^*) + \varepsilon \simeq 12.98V(S^*) + \varepsilon$.*

The proof of the theorem, as well as our proposed mechanism, can be found in Appendix F. In addition, we prove the following simple lower bound, proved in Appendix G.

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

6 Conclusions

We have proposed a convex relaxation for EDP, and showed that it can be used to design a δ -truthful, constant approximation mechanism that runs in polynomial time. Our objective function, commonly known as the Bayes D -optimality criterion, is motivated by linear regression, and in particular captures the information gain when experiments are used to learn a linear model.

A natural question to ask is to what extent the results we present here generalize to other machine learning tasks beyond linear regression. We outline a path in pursuing such generalizations in Appendix H. In particular, although the information

gain is not generally a submodular function, we show that for a wide class of models, in which experiments outcomes are perturbed by independent noise, the information gain indeed exhibits submodularity. Several important learning tasks fall under this category, including generalized linear regression, logistic regression, *etc.* In light of this, it would be interesting to investigate whether our convex relaxation approach generalizes to other learning tasks in this broader class.

The literature on experimental design includes several other optimality criteria [26, 4]. Our convex relaxation (10) involved swapping the log det scalarization with the expectation appearing in the multi-linear extension (9). The same swap is known to yield concave objectives for several other optimality criteria, even when the latter are not submodular (see, *e.g.*, Boyd and Vandenberghe [7]). Exploiting the convexity of such relaxations to design budget feasible mechanisms is an additional open problem of interest.

Acknowledgments

We thank Francis Bach for our helpful discussions on approximate solutions of convex optimization problems, and Yaron Singer for his comments and suggestions, and for insights into budget feasible mechanisms.

References

- [1] A. A. Ageev and M. Sviridenko. Pipage rounding: A new method of constructing algorithms with proven performance guarantee. *J. Comb. Optim.*, 8(3): 307–328, 2004. 4, 9, 10
- [2] A. G. Akritas, E. K. Akritas, and G. I. Malaschonok. Various proofs of Sylvester’s (determinant) identity. *Mathematics and Computers in Simulation*, 42(4–6):585 – 593, 1996. 17
- [3] A. Archer, C. Papadimitriou, K. Talwar, and E. Tardos. An approximate truthful mechanism for combinatorial auctions with single parameter agents. *Internet Mathematics*, 1(2):129–150, 2004. 4
- [4] A. Atkinson, A. Donev, and R. Tobias. *Optimum experimental designs, with SAS*. Oxford University Press New York, 2007. 1, 2, 5, 7, 14
- [5] A. Badanidiyuru, R. Kleinberg, and Y. Singer. Learning on a budget: posted price mechanisms for online procurement. In *EC*, 2012. 4
- [6] X. Bei, N. Chen, N. Gravin, and P. Lu. Budget feasible mechanism design: from prior-free to bayesian. In *STOC*, 2012. 4, 33

- [7] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004. 1, 6, 9, 11, 14, 17, 22
- [8] P. Briest, P. Krysta, and B. Vöcking. Approximation techniques for utilitarian mechanism design. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, page 39–48, 2005. 5
- [9] G. Calinescu, C. Chekuri, M. Pál, and J. Vondrák. Maximizing a submodular set function subject to a matroid constraint. In *Integer programming and combinatorial optimization*, pages 182–196. Springer, 2007. 9
- [10] K. Chaloner and I. Verdinelli. Bayesian experimental design: A review. *Statistical Science*, pages 273–304, 1995. 1, 5, 6, 7
- [11] N. Chen, N. Gravin, and P. Lu. On the approximability of budget feasible mechanisms. In *SODA*, 2011. 2, 3, 4, 9, 12, 13, 26, 28, 32
- [12] S. Dobzinski, C. H. Papadimitriou, and Y. Singer. Mechanisms for complement-free procurement. In *ACM EC*, 2011. 4, 33
- [13] S. Dughmi. A truthful randomized mechanism for combinatorial public projects via convex optimization. In *EC*, 2011. 4
- [14] S. Dughmi, T. Roughgarden, and Q. Yan. From convex optimization to randomized mechanisms: toward optimal combinatorial auctions. In *STOC*, 2011. 4, 5, 9
- [15] J. Friedman, T. Hastie, and R. Tibshirani. *The elements of statistical learning*, volume 1. Springer Series in Statistics, 2001. 6
- [16] J. Ginebra. On the measure of the information in a statistical experiment. *Bayesian Analysis*, 2(1):167–211, 2007. 1
- [17] M. Kearns, M. M. Pai, A. Roth, and J. Ullman. Private equilibrium release, large games, and no-regret learning, 2012. <http://arxiv.org/abs/1207.4084v1>. 5
- [18] A. Krause and C. Guestrin. A note on the budgeted maximization of submodular functions. Technical Report CMU-CALD-05-103, CMU, 2005. 8
- [19] A. Krause and C. Guestrin. Near-optimal nonmyopic value of information in graphical models. In *UAI*, 2005. 32
- [20] R. Lavi and C. Swamy. Truthful and near-optimal mechanism design via linear programming. *Journal of the ACM*, 58(6):25, 2011. 4
- [21] L. Le Cam. Comparison of experiments: a short review. *Lecture Notes-Monograph Series*, pages 127–138, 1996. 1

- [22] F. McSherry and K. Talwar. Mechanism design via differential privacy. In *FOCS*, 2007. 5
- [23] R. Myerson. Optimal auction design. *Mathematics of operations research*, 6(1): 58–73, 1981. 13, 29, 31, 34
- [24] K. Nissim, R. Smorodinsky, and M. Tennenholtz. Approximately optimal mechanism design via differential privacy. In *Innovations in Theoretical Computer Science (ITCS)*, 2012. 5
- [25] M. Pai and A. Roth. Privacy and mechanism design. *SIGecom Exchanges*, 2013. 5
- [26] F. Pukelsheim. *Optimal design of experiments*, volume 50. Society for Industrial Mathematics, 2006. 1, 2, 5, 6, 7, 14
- [27] J. Schummer. Almost-dominant strategy implementation: exchange economies. *Games and Economic Behavior*, 48(1):154–170, 2004. 25
- [28] J. Sherman and W. Morrison. Adjustment of an inverse matrix corresponding to a change in one element of a given matrix. *The Annals of Mathematical Statistics*, 21(1):124–127, 1950. 22
- [29] Y. Singer. Budget feasible mechanisms. In *FOCS*, 2010. 2, 3, 4, 8, 12, 25, 26, 32
- [30] Y. Singer. How to win friends and influence people, truthfully: influence maximization mechanisms for social networks. In *WSDM*, 2012. 4, 9, 12, 13
- [31] M. Sviridenko. A note on maximizing a submodular set function subject to a knapsack constraint. *Oper. Res. Lett.*, 32(1):41–43, 2004. 8
- [32] L. G. Valiant. A theory of the learnable. In R. A. DeMillo, editor, *STOC*, pages 436–445. ACM, 1984. 31
- [33] L. Vandenberghe, S. Boyd, and S. Wu. Determinant maximization with linear matrix inequality constraints. *SIAM journal on matrix analysis and applications*, 19(2):499–533, 1998. 11
- [34] J. Vondrak. Optimal approximation for the submodular welfare problem in the value oracle model. In *Proceedings of the 40th annual ACM symposium on Theory of computing*, pages 67–74. ACM, 2008. 9

A Properties of the Value Function V

For the sake of concreteness, we prove below the positivity, monotonicity, and submodularity of $V(S) = \log \det(I_d + X_S^T X_S)$ from basic principles. We note however that these properties hold more generally for the information gain under a wider class of models than the linear model with Gaussian noise and prior that we study here: we discuss this in more detail in Appendix H.

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. With this definition, matrix inversion is decreasing and convex over symmetric positive definite matrices (see Example 3.48 p. 110 in [7]).

Recall that the determinant of a matrix equals the product of its eigenvalues. The positivity of $V(S)$ follows from the fact that $X_S^T X_S$ is positive semi-definite and, as such $I_d + X_S^T X_S \succeq I_d$, so all its eigenvalues are larger than or equal to one, and they are all one if $S = \emptyset$. The marginal contribution of item $i \in \mathcal{N}$ to set $S \subseteq \mathcal{N}$ can be written as

$$\begin{aligned} V(S \cup \{i\}) - V(S) &= \frac{1}{2} \log \det(I_d + X_S^T X_S + x_i x_i^T) - \frac{1}{2} \log \det(I_d + X_S^T X_S) \\ &= \frac{1}{2} \log \det(I_d + x_i x_i^T (I_d + X_S^T X_S)^{-1}) = \frac{1}{2} \log(1 + x_i^T A(S)^{-1} x_i) \end{aligned} \quad (12)$$

where $A(S) \equiv I_d + X_S^T X_S$, and the last equality follows from the Sylvester's determinant identity [2]. Monotonicity therefore follows from the fact that $A(S)^{-1}$ is positive semidefinite. Finally, since the inverse is decreasing over positive definite matrices, we have

$$\forall S \subseteq \mathcal{N}, \quad A(S)^{-1} \succeq A(S \cup \{i\})^{-1}. \quad (13)$$

and submodularity also follows, as a function is submodular if and only if the marginal contributions are non-increasing in S . \square

B Proofs of Statements in Section 4.1

B.1 Proof of Lemma 1

The bound $F(\lambda) \leq L(\lambda)$ follows by the concavity of the $\log \det$ function and Jensen's inequality. 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 we use $\partial_i \cdot$ as a shorthand for $\frac{\partial}{\partial \lambda_i}$, the partial derivative with respect to the i -th variable.

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

$$\partial_i P_{\mathcal{N}}^\lambda(S) = \begin{cases} P_{\mathcal{N} \setminus \{i\}}^\lambda(S \setminus \{i\}) & \text{if } i \in S, \\ -P_{\mathcal{N} \setminus \{i\}}^\lambda(S) & \text{if } i \in \mathcal{N} \setminus S. \end{cases}$$

Hence,

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

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

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

Recall from (12) that the marginal contribution of i to S is given by

$$V(S \cup \{i\}) - V(S) = \frac{1}{2} \log(1 + x_i^T A(S)^{-1} x_i),$$

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

$$\partial_i F(\lambda) = \frac{1}{2} \sum_{\substack{S \subseteq \mathcal{N} \\ i \in \mathcal{N} \setminus S}} P_{\mathcal{N} \setminus \{i\}}^\lambda(S) \log(1 + x_i^T A(S)^{-1} x_i)$$

The computation of the derivative of L uses standard matrix calculus: writing $\tilde{A}(\lambda) \equiv I_d + \sum_{i \in \mathcal{N}} \lambda_i x_i x_i^T$,

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

Hence,

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

which implies

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

Recall from (13) that the monotonicity of the matrix inverse over positive definite matrices implies

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

as $A(S) \preceq A(S \cup \{i\})$. Observe that since $1 \leq \lambda_i \leq 1$, $P_{\mathcal{N} \setminus \{i\}}^\lambda(S) \geq P_{\mathcal{N}}^\lambda(S)$ and $P_{\mathcal{N} \setminus \{i\}}^\lambda(S) \geq P_{\mathcal{N}}^\lambda(S \cup \{i\})$ for all $S \subseteq \mathcal{N} \setminus \{i\}$. Hence,

$$\begin{aligned} \partial_i F(\lambda) &\geq \frac{1}{4} \sum_{\substack{S \subseteq \mathcal{N} \\ i \in \mathcal{N} \setminus S}} P_{\mathcal{N}}^\lambda(S) \log \left(1 + x_i^T A(S)^{-1} x_i \right) \\ &\quad + \frac{1}{4} \sum_{\substack{S \subseteq \mathcal{N} \\ i \in \mathcal{N} \setminus S}} P_{\mathcal{N}}^\lambda(S \cup \{i\}) \log \left(1 + x_i^T A(S \cup \{i\})^{-1} x_i \right) \\ &\geq \frac{1}{4} \sum_{S \subseteq \mathcal{N}} P_{\mathcal{N}}^\lambda(S) \log \left(1 + x_i^T A(S)^{-1} x_i \right). \end{aligned}$$

Using that $A(S) \succeq I_d$ we get that $x_i^T A(S)^{-1} x_i \leq \|x_i\|_2^2 \leq 1$. Moreover, $\log(1+x) \geq x$ for all $x \leq 1$. Hence,

$$\partial_i F(\lambda) \geq \frac{1}{4} x_i^T \left(\sum_{S \subseteq \mathcal{N}} P_{\mathcal{N}}^\lambda(S) A(S)^{-1} \right) x_i.$$

Finally, using that the inverse is a matrix convex function over symmetric positive definite matrices (see Appendix A):

$$\partial_i F(\lambda) \geq \frac{1}{4} x_i^T \left(\sum_{S \subseteq \mathcal{N}} P_{\mathcal{N}}^\lambda(S) A(S) \right)^{-1} x_i = \frac{1}{4} x_i^T \tilde{A}(\lambda)^{-1} x_i = \frac{1}{2} \partial_i L(\lambda).$$

Having bound the ratio between the partial derivatives, we now bound the ratio $F(\lambda)/L(\lambda)$ from below. Consider the following three cases.

First, 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 .

Second, if the minimum of the ratio $F(\lambda)/L(\lambda)$ is attained at a vertex of the hypercube $[0, 1]^n$ different from 0. F and L being relaxations of the value function V , they are equal to V on the vertices which are exactly the binary points. Hence, the minimum is equal to 1 in this case; in particular, it is greater than $1/2$.

Finally, if the minimum is attained at a point λ^* with at least one coordinate belonging to $(0, 1)$, let i be one such coordinate and consider the function G_i :

$$G_i : x \mapsto \frac{F}{L}(\lambda_1^*, \dots, \lambda_{i-1}^*, x, \lambda_{i+1}^*, \dots, \lambda_n^*).$$

Then this function attains a minimum at $\lambda_i^* \in (0, 1)$ and its derivative is zero at this point. Hence:

$$0 = G'_i(\lambda_i^*) = \partial_i \left(\frac{F}{L} \right) (\lambda^*).$$

But $\partial_i(F/L)(\lambda^*) = 0$ implies that

$$\frac{F(\lambda^*)}{L(\lambda^*)} = \frac{\partial_i F(\lambda^*)}{\partial_i L(\lambda^*)} \geq \frac{1}{2}$$

using the lower bound on the ratio of the partial derivatives. This concludes the proof of the lemma. \square

B.2 Proof of Lemma 2

We give a rounding procedure which, given a feasible λ with at least two fractional components, returns some feasible λ' with one fewer 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:

$$\forall \varepsilon \in \left[\max \left(-\lambda_i, (\lambda_j - 1) \frac{c_j}{c_i} \right), \min \left(1 - \lambda_i, \lambda_j \frac{c_j}{c_i} \right) \right], \lambda_\varepsilon \text{ is feasible} \quad (14)$$

Furthermore, the function F_λ is convex; indeed:

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

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

$$\frac{c_i}{c_j} \mathbb{E}_{S' \sim P_{\mathcal{N} \setminus \{i,j\}}^\lambda(S')} \left[V(S' \cup \{i\}) + V(S' \cup \{j\}) - V(S' \cup \{i, j\}) - V(S') \right]$$

which is positive by submodularity of V . Hence, the maximum of F_λ over the interval given in (14) is attained at one of its limits, at which either the i -th or j -th component of λ_ε becomes integral. \square

B.3 Proof of Proposition 1

The lower bound on L_c^* follows immediately from the fact that L extends V to $[0, 1]^n$. For the upper bound, let us consider a feasible point $\lambda^* \in \mathcal{D}_c$ such that $L(\lambda^*) = L_c^*$. By applying Lemma 1 and Lemma 2 we get a feasible point $\bar{\lambda}$ with at most one fractional component such that

$$L(\lambda^*) \leq 2F(\bar{\lambda}). \quad (15)$$

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 :

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

By submodularity of V , $V(S \cup \{i\}) \leq V(S) + V(\{i\})$. Hence,

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

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

$$F(\bar{\lambda}) \leq OPT + \max_{i \in \mathcal{N}} V(i). \quad (16)$$

Together, (15) and (16) imply the proposition. \square

C Proof of Proposition 2

We proceed by showing that the optimal value of $(P_{c,\alpha})$ is close to the optimal value of (P_c) (Lemma 6) while being well-behaved with respect to changes of the cost (Lemma 7). These lemmas together imply Proposition 2.

Note that the choice of α given in Algorithm 1 implies that $\alpha < \frac{1}{n}$. This in turn implies that the feasible set $\mathcal{D}_{c,\alpha}$ of $(P_{c,\alpha})$ is non-empty: it contains the strictly feasible point $\lambda = (\frac{1}{n}, \dots, \frac{1}{n})$.

Lemma 5. *Let $\partial_i L(\lambda)$ denote the i -th derivative of L , for $i \in \{1, \dots, n\}$, then:*

$$\forall \lambda \in [0, 1]^n, \quad \frac{b}{2^n} \leq \partial_i L(\lambda) \leq 1$$

Proof. Recall that we had defined:

$$\tilde{A}(\lambda) \equiv I_d + \sum_{i=1}^n \lambda_i x_i x_i^T \quad \text{and} \quad A(S) \equiv I_d + \sum_{i \in S} x_i x_i^T$$

Let us also define $A_k \equiv A(\{x_1, \dots, x_k\})$. We have $\partial_i L(\lambda) = x_i^T \tilde{A}(\lambda)^{-1} x_i$. Since $\tilde{A}(\lambda) \succeq I_d$, $\partial_i L(\lambda) \leq x_i^T x_i \leq 1$, which is the right-hand side of the lemma. For

the left-hand side, note that $\tilde{A}(\lambda) \preceq A_n$. Hence $\partial_i L(\lambda) \geq x_i^T A_n^{-1} x_i$. Using the Sherman-Morrison formula [28], for all $k \geq 1$:

$$x_i^T A_k^{-1} x_i = x_i^T A_{k-1}^{-1} x_i - \frac{(x_i^T A_{k-1}^{-1} x_k)^2}{1 + x_k^T A_{k-1}^{-1} x_k}$$

By the Cauchy-Schwarz inequality:

$$(x_i^T A_{k-1}^{-1} x_k)^2 \leq x_i^T A_{k-1}^{-1} x_i x_k^T A_{k-1}^{-1} x_k$$

Hence:

$$x_i^T A_k^{-1} x_i \geq x_i^T A_{k-1}^{-1} x_i - x_i^T A_{k-1}^{-1} x_i \frac{x_k^T A_{k-1}^{-1} x_k}{1 + x_k^T A_{k-1}^{-1} x_k}$$

But $x_k^T A_{k-1}^{-1} x_k \leq 1$ and $\frac{a}{1+a} \leq \frac{1}{2}$ if $0 \leq a \leq 1$, so:

$$x_i^T A_k^{-1} x_i \geq x_i^T A_{k-1}^{-1} x_i - \frac{1}{2} x_i^T A_{k-1}^{-1} x_i \geq \frac{x_i^T A_{k-1}^{-1} x_i}{2}$$

By induction:

$$x_i^T A_n^{-1} x_i \geq \frac{x_i^T x_i}{2^n}$$

Using that $x_i^T x_i \geq b$ concludes the proof of the left-hand side of the lemma's inequality. \square

Let us introduce the Lagrangian of problem $(P_{c,\alpha})$:

$$\mathcal{L}_{c,\alpha}(\lambda, \mu, \nu, \xi) \equiv L(\lambda) + \mu^T(\lambda - \alpha \mathbf{1}) + \nu^T(\mathbf{1} - \lambda) + \xi(B - c^T \lambda)$$

so that:

$$L_{c,\alpha}^* = \min_{\mu, \nu, \xi \geq 0} \max_{\lambda} \mathcal{L}_{c,\alpha}(\lambda, \mu, \nu, \xi)$$

Similarly, we define $\mathcal{L}_c \equiv \mathcal{L}_{c,0}$ the lagrangian of (P_c) .

Let λ^* be primal optimal for $(P_{c,\alpha})$, and (μ^*, ν^*, ξ^*) be dual optimal for the same problem. In addition to primal and dual feasibility, the Karush-Kuhn-Tucker (KKT) conditions [7] give $\forall i \in \{1, \dots, n\}$:

$$\begin{aligned} \partial_i L(\lambda^*) + \mu_i^* - \nu_i^* - \xi^* c_i &= 0 \\ \mu_i^*(\lambda_i^* - \alpha) &= 0 \\ \nu_i^*(1 - \lambda_i^*) &= 0 \end{aligned}$$

Lemma 6. *We have:*

$$L_c^* - \alpha n^2 \leq L_{c,\alpha}^* \leq L_c^*$$

In particular, $|L_c^ - L_{c,\alpha}^*| \leq \alpha n^2$.*

Proof. $\alpha \mapsto L_{c,\alpha}^*$ is a decreasing function as it is the maximum value of the L function over a set-decreasing domain, which gives the rightmost inequality.

Let μ^*, ν^*, ξ^* be dual optimal for $(P_{c,\alpha})$, that is:

$$L_{c,\alpha}^* = \max_{\lambda} \mathcal{L}_{c,\alpha}(\lambda, \mu^*, \nu^*, \xi^*)$$

Note that $\mathcal{L}_{c,\alpha}(\lambda, \mu^*, \nu^*, \xi^*) = \mathcal{L}_c(\lambda, \mu^*, \nu^*, \xi^*) - \alpha \mathbf{1}^T \mu^*$, and that for any λ feasible for problem (P_c) , $\mathcal{L}_c(\lambda, \mu^*, \nu^*, \xi^*) \geq L(\lambda)$. Hence,

$$L_{c,\alpha}^* \geq L(\lambda) - \alpha \mathbf{1}^T \mu^*$$

for any λ feasible for (P_c) . In particular, for λ primal optimal for (P_c) :

$$L_{c,\alpha}^* \geq L_c^* - \alpha \mathbf{1}^T \mu^* \quad (17)$$

Let us denote by the M the support of μ^* , that is $M \equiv \{i | \mu_i^* > 0\}$, and by λ^* a primal optimal point for $(P_{c,\alpha})$. From the KKT conditions we see that:

$$M \subseteq \{i | \lambda_i^* = \alpha\}$$

Let us first assume that $|M| = 0$, then $\mathbf{1}^T \mu^* = 0$ and the lemma follows.

We will now assume that $|M| \geq 1$. In this case $c^T \lambda^* = B$, otherwise we could increase the coordinates of λ^* in M , which would increase the value of the objective function and contradict the optimality of λ^* . Note also, that $|M| \leq n-1$, otherwise, since $\alpha < \frac{1}{n}$, we would have $c^T \lambda^* < B$, which again contradicts the optimality of λ^* . Let us write:

$$B = c^T \lambda^* = \alpha \sum_{i \in M} c_i + \sum_{i \in \bar{M}} \lambda_i^* c_i \leq \alpha |M| B + (n - |M|) \max_{i \in \bar{M}} c_i$$

That is:

$$\max_{i \in \bar{M}} c_i \geq \frac{B - B|M|\alpha}{n - |M|} > \frac{B}{n} \quad (18)$$

where the last inequality uses again that $\alpha < \frac{1}{n}$. From the KKT conditions, we see that for $i \in M$, $\nu_i^* = 0$ and:

$$\mu_i^* = \xi^* c_i - \partial_i L(\lambda^*) \leq \xi^* c_i \leq \xi^* B \quad (19)$$

since $\partial_i L(\lambda^*) \geq 0$ and $c_i \leq 1$.

Furthermore, using the KKT conditions again, we have that:

$$\xi^* \leq \inf_{i \in \bar{M}} \frac{\partial_i L(\lambda^*)}{c_i} \leq \inf_{i \in \bar{M}} \frac{1}{c_i} = \frac{1}{\max_{i \in \bar{M}} c_i} \quad (20)$$

where the last inequality uses Lemma 5. Combining (18), (19) and (20), we get that:

$$\sum_{i \in M} \mu_i^* \leq |M| \xi^* B \leq n \xi^* B \leq \frac{nB}{\max_{i \in \bar{M}} c_i} \leq n^2$$

This implies that:

$$\mathbf{1}^T \mu^* = \sum_{i=1}^n \mu_i^* = \sum_{i \in M} \mu_i^* \leq n^2$$

which along with (17) proves the lemma. \square

Lemma 7. *If $c' = (c'_i, c_{-i})$, with $c'_i \leq c_i - \delta$, we have:*

$$L_{c',\alpha}^* \geq L_{c,\alpha}^* + \frac{\alpha \delta b}{2^n B}$$

Proof. Let μ^*, ν^*, ξ^* be dual optimal for $(P_{c',\alpha})$. Noting that:

$$\mathcal{L}_{c',\alpha}(\lambda, \mu^*, \nu^*, \xi^*) \geq \mathcal{L}_{c,\alpha}(\lambda, \mu^*, \nu^*, \xi^*) + \lambda_i \xi^* \delta,$$

we get similarly to Lemma 6:

$$L_{c',\alpha}^* \geq L(\lambda) + \lambda_i \xi^* \delta$$

for any λ feasible for $(P_{c,\alpha})$. In particular, for λ^* primal optimal for $(P_{c,\alpha})$:

$$L_{c',\alpha}^* \geq L_{c,\alpha}^* + \alpha \xi^* \delta$$

since $\lambda_i^* \geq \alpha$.

Using the KKT conditions for $(P_{c',\alpha})$, we can write:

$$\xi^* = \inf_{i: \lambda_i^* > \alpha} \frac{x_i^T S(\lambda'^*)^{-1} x_i}{c'_i}$$

with λ'^* optimal for $(P_{c',\alpha})$. Since $c'_i \leq B$, using Lemma 5, we get that $\xi^* \geq \frac{b}{2^n B}$, which concludes the proof. \square

We are now ready to conclude the proof of Proposition 2. Let $\hat{L}_{c,\alpha}^*$ be the approximation computed by Algorithm 1.

1. using Lemma 6:

$$|\hat{L}_{c,\alpha}^* - L_c^*| \leq |\hat{L}_{c,\alpha}^* - L_{c,\alpha}^*| + |L_{c,\alpha}^* - L_c^*| \leq \frac{\alpha \delta}{B} + \alpha n^2 = \varepsilon$$

which proves the ε -accuracy.

2. for the δ -decreasingness, let $c' = (c'_i, c_{-i})$ with $c'_i \leq c_i - \delta$, then:

$$\hat{L}_{c',\alpha}^* \geq L_{c',\alpha}^* - \frac{\alpha \delta b}{2^{n+1} B} \geq L_{c,\alpha}^* + \frac{\alpha \delta b}{2^{n+1} B} \geq \hat{L}_{c,\alpha}^*$$

where the first and last inequalities follow from the accuracy of the approximation, and the inner inequality follows from Lemma 7.

3. the accuracy of the approximation $\hat{L}_{c,\alpha}^*$ is:

$$\varepsilon' = \frac{\varepsilon \delta b}{2^{n+1}(\delta + n^2 B)}$$

Note that:

$$\log \log(\varepsilon')^{-1} = O\left(\log \log \frac{B}{\varepsilon \delta b} + \log n\right)$$

Using Lemma 3 concludes the proof of the running time. \square

D Budget Feasible Reverse Auction Mechanisms

We review in this appendix the formal definition of a budget feasible reverse auction mechanisms, as introduced by Singer [29]. We depart from the definitions in [29] only in considering δ -truthful, rather than truthful, mechanisms.

Given a budget B and a value function $V : 2^{\mathcal{N}} \rightarrow \mathbb{R}_+$, a *mechanism* $\mathcal{M} = (S, p)$ comprises (a) an *allocation function* $S : \mathbb{R}_+^n \rightarrow 2^{\mathcal{N}}$ and (b) a *payment function* $p : \mathbb{R}_+^n \rightarrow \mathbb{R}_+^n$. Let $s_i(c) = \mathbb{1}_{i \in S(c)}$ be the binary indicator of $i \in S(c)$. We seek mechanisms that have the following properties [29]:

- *Normalization.* Unallocated experiments receive zero payments: $s_i(c) = 0$ implies $p_i(c) = 0$.
- *Individual Rationality.* Payments for allocated experiments exceed costs: $p_i(c) \geq c_i \cdot s_i(c)$.
- *No Positive Transfers.* Payments are non-negative: $p_i(c) \geq 0$.
- *δ -Truthfulness.* Reporting one's true cost is an *almost-dominant* [27] strategy. Formally, let c_{-i} be a vector of costs of all agents except i . Then, $p_i(c_i, c_{-i}) - s_i(c_i, c_{-i}) \cdot c_i \geq p_i(c'_i, c_{-i}) - s_i(c'_i, c_{-i}) \cdot c_i$, for every $i \in \mathcal{N}$ and every two cost vectors (c_i, c_{-i}) and (c'_i, c_{-i}) such that $|c_i - c'_i| > \delta$. The mechanism is *truthful* if $\delta = 0$.
- *Budget Feasibility.* The sum of the payments should not exceed the budget constraint, *i.e.* $\sum_{i \in \mathcal{N}} p_i(c) \leq B$.
- *(α, β) -approximation.* The value of the allocated set should not be too far from the optimum value of the full information case, as given by (7). Formally, there must exist some $\alpha \geq 1$ and $\beta > 0$ such that $OPT \leq \alpha V(S(p)) + \beta$, where $OPT = \max_{S \subseteq \mathcal{N}} \{V(S) \mid \sum_{i \in S} c_i \leq B\}$.
- *Computational Efficiency.* The allocation and payment function should be computable in time polynomial in various parameters.

E Proof of Lemma 4

Using the notations of Lemma 4, we want to prove that if c_i and c'_i are two different costs reported by user i with $|c_i - c'_i| \geq \delta$, and if c_{-i} is any vector of costs reported by the other users:

$$p_i(c_i, c_{-i}) - s_i(c_i, c_{-i}) \cdot c_i \geq p_i(c'_i, c_{-i}) - s_i(c'_i, c_{-i}) \cdot c_i \quad (21)$$

We distinguish four cases depending on the value of $s_i(c_i, c_{-i})$ and $s'_i(c'_i, c_{-i})$.

1. $s_i(c_i, c_{-i}) = s_i(c'_i, c_{-i}) = 0$. Since the mechanism is normalized we have $p_i(c_i, c_{-i}) = p_i(c'_i, c_{-i}) = 0$ and (21) is true.
2. $s_i(c'_i, c_{-i}) = s_i(c_i, c_{-i}) = 1$. Note that i is paid her threshold payment when allocated, and since this payment does not depend on i 's reported cost, (21) is true (and is in fact an equality).
3. $s_i(c'_i, c_{-i}) = 0$ and $s_i(c_i, c_{-i}) = 1$. We then have $p_i(c'_i, c_{-i}) = 0$ by normalization and (21) follows from individual rationality.
4. $s_i(c'_i, c_{-i}) = 1$ and $s_i(c_i, c_{-i}) = 0$. By δ -decreasingness of s_i , $c_i \geq c'_i + \delta$, and $s_i(c_i, c_{-i}) = 0$ implies that i 's threshold payment is less than c_i , i.e. $p_i(c'_i, c_{-i}) \leq c_i$. This last inequality is equivalent to (21) in this final case. \square

F Proof of Theorem 1

We now present the proof of Theorem 1. Our mechanism for EDP is composed of (a) the allocation function presented in Algorithm 2, and (b) 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 . A closed-form formula for threshold payments when S_G is the allocated set can be found in [29].

We use the notation OPT_{-i^*} to denote the optimal value of EDP when the maximum value element i^* is excluded. We also use OPT'_{-i^*} to denote the approximation computed by the δ -decreasing, ε -accurate approximation of $L_{c_{-i^*}}^*$, as defined in Algorithm 2.

The properties of δ -truthfulness and individual rationality follow from δ -monotonicity and threshold payments. δ -monotonicity and budget feasibility follow similar steps as the analysis of Chen et al. [11], adapted to account for δ -monotonicity:

Lemma 8. *Our mechanism for EDP is δ -monotone and budget feasible.*

Algorithm 2 Mechanism for EDP

Input: $B \in \mathbb{R}_+, c \in [0, B]^n, \delta \in (0, 1], \varepsilon \in (0, 1]$

- 1: $i^* \leftarrow \arg \max_{j \in \mathcal{N}} V(j)$
 - 2: $OPT'_{-i^*} \leftarrow$ using Proposition 2, compute a ε -accurate, δ -decreasing approximation of
$$L_{c_{-i^*}}^* \equiv \max_{\lambda \in [0, 1]^n} \{L(\lambda) \mid \lambda_{i^*} = 0, \sum_{i \in \mathcal{N} \setminus \{i^*\}} c_i \lambda_i \leq B\}$$
 - 3: $C \leftarrow \frac{8e-1+\sqrt{64e^2-24e+9}}{2(e-1)}$
 - 4: **if** $OPT'_{-i^*} < C \cdot V(i^*)$ **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**
-

Proof. Consider an agent i with cost c_i that is selected by the mechanism, and suppose that she reports a cost $c'_i \leq c_i - \delta$ while all other costs stay the same. Suppose that when i reports c_i , $OPT'_{-i^*} \geq CV(i^*)$; then, as $s_i(c_i, c_{-i}) = 1$, $i \in S_G$. By reporting cost 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 $i \in S'_G$. By δ -decreasingness of OPT'_{-i^*} , under $c'_i \leq c_i - \delta$ the greedy set is still allocated and $s_i(c'_i, c_{-i}) = 1$. Suppose now that when i reports c_i , $OPT'_{-i^*} < CV(i^*)$. Then $s_i(c_i, c_{-i}) = 1$ iff $i = i^*$. Reporting $c'_i \leq c_i$ does not change $V(i^*)$ nor $OPT'_{-i^*} \leq CV(i^*)$; thus $s_{i^*}(c'_{i^*}, c_{-i^*}) = 1$, so the mechanism is monotone.

To show budget feasibility, suppose that $OPT'_{-i^*} < CV(i^*)$. 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 $OPT'_{-i^*} \geq CV(i^*)$. 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 $i \in S_G$:

$$\text{if } c'_i \geq \frac{V(S_i \cup \{i\}) - V(S)}{V(S_G)} B \text{ then } s_i(c'_i, c_{-i}) = 0 \quad (22)$$

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, (22) implies that the threshold payment of user i is bounded by the above quantity. Hence, the total payment is bounded by the telescopic sum:

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

The complexity of the mechanism is given by the following lemma.

Lemma 9 (Complexity). *For any $\varepsilon > 0$ and any $\delta > 0$, the complexity of the mechanism is $O(\text{poly}(n, d, \log \log \frac{B}{b\varepsilon\delta}))$*

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 . By Proposition 2, line 2 of Algorithm 2 can be computed in time $O(\text{poly}(n, d, \log \log \frac{B}{b\varepsilon\delta}))$. Hence the allocation function's complexity is as stated. \square

Finally, we prove the approximation ratio of the mechanism. We use the following lemma from [11] which bounds OPT in terms of the value of S_G , as computed in Algorithm 2, and i^* , the element of maximum value.

Lemma 10 ([11]). *Let S_G be the set computed in Algorithm 2 and let $i^* = \arg \max_{i \in \mathcal{N}} V(\{i\})$. We have:*

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

Using Proposition 1 and Lemma 10 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:

$$OPT \leq \frac{10e-3 + \sqrt{64e^2 - 24e + 9}}{2(e-1)} V(S^*) + \varepsilon. \quad (23)$$

To see this, let $L_{c_{-i^*}}^*$ be the maximum value of L subject to $\lambda_{i^*} = 0$, $\sum_{i \in \mathcal{N} \setminus i^*} c_i \leq B$. From line 2 of Algorithm 2, we have $L_{c_{-i^*}}^* - \varepsilon \leq OPT'_{-i^*} \leq L_{c_{-i^*}}^* + \varepsilon$.

If the condition on line 4 of the algorithm holds then, from the lower bound in Proposition 1,

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

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

$$OPT \leq (1 + C)V(i^*) + \varepsilon. \quad (24)$$

If the condition on line 4 does not hold, by observing that $L_{c_{-i^*}}^* \leq L_c^*$ and the upper bound of Proposition 1, we get

$$V(i^*) \leq \frac{1}{C}L_{c_{-i^*}}^* + \frac{\varepsilon}{C} \leq \frac{1}{C}(2OPT + 2V(i^*)) + \frac{\varepsilon}{C}.$$

Applying Lemma 10,

$$V(i^*) \leq \frac{1}{C} \left(\frac{2e}{e-1} (3V(S_G) + 2V(i^*)) + 2V(i^*) \right) + \frac{\varepsilon}{C}.$$

Note that C satisfies $C(e-1) - 6e + 2 > 0$, hence

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

Finally, using Lemma 10 again, we get

$$OPT \leq \frac{3e}{e-1} \left(1 + \frac{4e}{C(e-1) - 6e + 2} \right) V(S_G) + \frac{2e\varepsilon}{C(e-1) - 6e + 2}. \quad (25)$$

Our choice of C , namely,

$$C = \frac{8e-1 + \sqrt{64e^2 - 24e + 9}}{2(e-1)}, \quad (26)$$

is precisely to minimize the maximum among the coefficients of V_{i^*} and $V(S_G)$ in (24) and (25), respectively. Indeed, consider:

$$\max \left(1 + C, \frac{3e}{e-1} \left(1 + \frac{4e}{C(e-1) - 6e + 2} \right) \right).$$

This function has two minima, only one of those is such that $C(e-1) - 6e + 2 \geq 0$. This minimum is precisely (26). For this minimum, $\frac{2e\varepsilon}{C(e-1) - 6e + 2} \leq \varepsilon$. Placing the expression of C in (24) and (25) gives the approximation ratio in (23), and concludes the proof of Theorem 1. \square

G Proof of Theorem 2

Suppose, for contradiction, that such a mechanism exists. From Myerson's Theorem [23], a single parameter auction is truthful if and only if the allocation function is monotone and agents are paid threshold payments. Consider two experiments with dimension $d = 2$, such that $x_1 = e_1 = [1, 0]$, $x_2 = e_2 = [0, 1]$ and $c_1 = c_2 = B/2 + \varepsilon$.

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 $B/2 - \varepsilon$, by monotonicity it remains in the solution; by threshold payment, it is paid at least $B/2 + \varepsilon$. 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 $2 \log 2$. \square

H Extensions

H.1 Strategic Experimental Design with Non-Homotropic Prior

In the general case where the prior distribution of the experimenter on the model β in (2) is not homotropic and has a generic covariance matrix R , the value function takes the general form given by (5).

Let us denote by λ (resp. Λ) the smallest (resp. largest) eigenvalue of R , applying the mechanism described in Algorithm 2 and adapting the analysis of the approximation ratio *mutatis mutandis*, we get the following result which extends Theorem 1.

Theorem 3. *For any $\delta \in (0, 1]$, and any $\varepsilon \in (0, 1]$, there exists a δ -truthful, individually rational and budget feasible mechanism for the objective function V given by (5) that runs in time $O(\text{poly}(n, d, \log \log \frac{B\Lambda}{\varepsilon\delta b\lambda}))$ and allocates a set S^* such that:*

$$OPT \leq \left(\frac{6e - 2}{e - 1} \frac{1/\lambda}{\log(1 + 1/\lambda)} + 4.66 \right) V(S^*) + \varepsilon.$$

H.2 Non-Bayesian Setting

In the non-bayesian setting, *i.e.* when the experimenter has no prior distribution on the model, the covariance matrix R is the zero matrix. In this case, the ridge regression estimation procedure (3) reduces to simple least squares (*i.e.*, linear regression), and the D -optimality criterion reduces to the entropy of $\hat{\beta}$, given by:

$$V(S) = \log \det(X_S^T X_S) \quad (27)$$

A natural question which arises is whether it is possible to design a deterministic mechanism in this setting. Since (27) 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 11. *For any $M > 1$, there is no M -approximate, truthful, budget feasible, individually rational mechanism for a budget feasible reverse auction with value function $V(S) = \det X_S^T X_S$.*

Proof. From Myerson’s Theorem [23], a single parameter auction is truthful if and only if the allocation function is monotone and agents are paid threshold payments. 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 $c_1 = c_2 = 0.5 + \varepsilon$, while $c_3 = c_4 = \varepsilon$, for some small $\varepsilon > 0$. 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 $B/2 - 3\varepsilon$, x_1 will still be in the solution. By threshold payments, experiment x_1 receives in this case a payment that is at least $B/2 + \varepsilon$. 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

H.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, consider the following variant of the standard PAC learning setup [32]: assume that the features $x_i, i \in \mathcal{N}$ take values in some generic set Ω , called the *query space*. Measurements $y_i \in \mathbb{R}$ are given by

$$y_i = h(x_i) + \varepsilon_i \quad (28)$$

where $h \in \mathcal{H}$ for some subset \mathcal{H} of all possible mappings $h : \Omega \rightarrow \mathbb{R}$, called the *hypothesis space*. As before, we assume that the experimenter has a prior distribution on the hypothesis $h \in \mathcal{H}$; we also assume that ε_i are random variables in \mathbb{R} , not necessarily identically distributed, that are independent *conditioned on* h . As before, the features x_i are public, and the goal of the experimenter is to (a) retrieve measurements y_i and (b) estimate h as accurately as possible.

This model is quite broad, and captures many classic machine learning tasks; we give a few concrete examples below:

1. **Generalized Linear Regression.** In this case, $\Omega = \mathbb{R}^d$, \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 variables (not necessarily identically distributed).
2. **Learning Binary Functions with Bernoulli Noise.** When learning a binary function under noise, the experimenter wishes to determine a binary function h by testing its output on different inputs; however, the output may be corrupted with probability p . Formally, $\Omega = \{0, 1\}^d$, \mathcal{H} is some subset of

binary functions $h : \Omega \rightarrow \{0, 1\}$, and

$$\varepsilon_i = \begin{cases} 0, & \text{w. prob. } 1 - p \\ \bar{h}(x_i) - h(x_i), & \text{w. prob. } p \end{cases}$$

3. **Logistic Regression.** Logistic regression aims to learn a hyperplane separating $+1$ -labeled values from -1 -labeled values; again, values can be corrupted, and the probability that a label is flipped drops with the distance from the hyperplane. Formally, $\Omega = \mathbb{R}^d$, \mathcal{H} is the set of maps $\{h(x) = \text{sign}(\beta^T x) \text{ for some } \beta \in \mathbb{R}^d\}$, and ε_i are independent conditioned on β such that

$$\varepsilon_i = \begin{cases} -2 \cdot \mathbb{1}_{\beta^T x > 0}, & \text{w. prob. } \frac{1}{1 + e^{\beta^T x}} \\ +2 \cdot \mathbb{1}_{\beta^T x < 0}, & \text{w. prob. } \frac{e^{\beta^T x}}{1 + e^{\beta^T x}} \end{cases}$$

We can again define the information gain as an objective to maximize:

$$V(S) = H(h) - H(h \mid y_S), \quad S \subseteq \mathcal{N} \quad (29)$$

This is a monotone set function, and it clearly satisfies $V(\emptyset) = 0$. In general, the information gain is not a submodular function. However, when the errors ε_i are independent conditioned on h , the following lemma holds:

Lemma 12. *The value function given by the information gain (29) is submodular.*

Proof. A more general statement for graphical models is shown in [19]; in short, using the chain rule for the conditional entropy we get:

$$V(S) = H(y_S) - H(y_S \mid h) = H(y_S) - \sum_{i \in S} H(y_i \mid h) \quad (30)$$

where the second equality comes from the independence of the y_i 's conditioned on h . Recall that the joint entropy of a set of random variables is a submodular function. Thus, the value function is written in (30) 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 by Singer [29] and Chen et al. [11] to get the following corollary:

Corollary 1. *For Bayesian experimental design with an objective given by the information gain (29), there exists a randomized, polynomial-time, budget feasible, individually rational, and universally truthful mechanism with a 7.91 approximation ratio, in expectation.*

In cases where maximizing (29) can be done in polynomial time in the full-information setup, there exists a deterministic, polynomial-time, budget feasible, individually rational, and truthful mechanism for Bayesian experimental design with an 8.34 approximation ratio.

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 [12, 6] remains an interesting open question.

I Non-Truthfulness of the Maximum Operator

We give a counterexample of the truthfulness of the maximum mechanism whose allocation rule is defined in (8) when the value function V is as defined in (1). We denote by (e_1, e_2, e_3) the canonical basis of \mathbb{R}^3 and define the following feature vectors: $x_1 = e_1$, $x_2 = \frac{1}{\sqrt{2}} \cos \frac{\pi}{5} e_2 + \frac{1}{\sqrt{2}} \sin \frac{\pi}{5} e_3$, $x_3 = \frac{1}{\sqrt{2}} e_2$ and $x_4 = \frac{1}{2} e_3$, with associated costs $c_1 = \frac{5}{2}$, $c_2 = c_3 = 1$ and $c_4 = \frac{2}{3}$. We also assume that the budget of the auctioneer is $B = \frac{5}{2}$.

Note that $V(x_i) = \log(1 + \|x_i\|^2)$, so x_1 is the point of maximum value. Let us now compute the output of the greedy heuristic. We have:

$$\frac{V(x_1)}{c_1} \simeq 0.277, \quad \frac{V(x_2)}{c_2} = \frac{V(x_3)}{c_3} \simeq 0.405, \quad \frac{V(x_4)}{c_4} \simeq 0.335 \quad (31)$$

so the greedy heuristic will start by selecting x_2 or x_3 . Without loss of generality, we can assume that it selected x_2 . From the Sherman-Morrison formula we get:

$$V(\{x_i, x_j\}) - V(x_i) = \log \left(1 + \|x_j\|^2 - \frac{\langle x_i, x_j \rangle^2}{1 + \|x_i\|^2} \right)$$

In particular, when x_i and x_j are orthogonal $V(\{x_i, x_j\}) = V(x_j)$. This allows us to compute:

$$\begin{aligned} \frac{V(\{x_2, x_3\}) - V(x_2)}{c_3} &= \log \left(1 + \frac{1}{2} - \frac{1}{6} \cos^2 \frac{\pi}{5} \right) \simeq 0.329 \\ \frac{V(\{x_2, x_4\}) - V(x_2)}{c_4} &= \frac{3}{2} \log \left(1 + \frac{1}{4} - \frac{1}{12} \sin^2 \frac{\pi}{5} \right) \simeq 0.299 \end{aligned}$$

Note that at this point x_1 cannot be selected without exceeding the budget. Hence, the greedy heuristic will add x_3 to the greedy solution and returns the set $\{x_2, x_3\}$ with value:

$$V(\{x_2, x_3\}) = V(x_2) + V(\{x_2, x_3\}) - V(x_2) \simeq 0.734$$

In contrast, $V(x_1) \simeq 0.693$ so the mechanism will allocate to $\{x_2, x_3\}$.

Let us now assume that user 3 reduces her cost. It follows from (31) that the greedy heuristic will start by adding her to the greedy solution. Furthermore:

$$\frac{V(\{x_3, x_2\}) - V(x_3)}{c_2} = \log \left(1 + \frac{1}{2} - \frac{1}{6} \cos^2 \frac{\pi}{5} \right) \simeq 0.329$$

$$\frac{V(\{x_3, x_4\}) - V(x_3)}{c_4} = \frac{3}{2} \log \left(1 + \frac{1}{4} \right) \simeq 0.334$$

Hence, the greedy solution will be $\{x_3, x_4\}$ with value:

$$V(\{x_3, x_4\}) = V(x_3) + V(\{x_3, x_4\}) - V(x_3) \simeq 0.628$$

As a consequence the mechanism will allocate to user 1 in this case. By reducing her cost, user 3, who was previously allocated, is now rejected by the mechanism. This contradicts the monotonicity of the allocation rule, hence its truthfulness by Myerson's theorem [23], which states that a single parameter auction is truthful if and only if the allocation function is monotone and agents are paid threshold payments.