
Robust Budget Allocation via Continuous Submodular Functions

Matthew Staib¹ Stefanie Jegelka¹

Abstract

The optimal allocation of resources for maximizing influence, spread of information or coverage, has gained attention in the past years, in particular in machine learning and data mining. But in applications, the parameters of the problem are rarely known exactly, and using wrong parameters can lead to undesirable outcomes. We hence revisit a continuous version of the Budget Allocation or Bipartite Influence Maximization problem introduced by Alon et al. (2012) from a robust optimization perspective, where an adversary may choose the least favorable parameters within a confidence set. The resulting problem is a nonconvex-concave saddle point problem (or game). We show that this nonconvex problem can be solved exactly by leveraging connections to continuous submodular functions, and by solving a constrained submodular minimization problem. Although constrained submodular minimization is hard in general, here, we establish conditions under which such a problem can be solved to arbitrary precision ϵ .

1. Introduction

The optimal allocation of resources for maximizing influence, spread of information or coverage, has gained attention in the past few years, in particular in machine learning and data mining (Domingos & Richardson, 2001; Kempe et al., 2003; Chen et al., 2009; Gomez Rodriguez & Schölkopf, 2012; Borgs et al., 2014).

In the *Budget Allocation Problem*, one is given a bipartite influence graph between channels S and people T , and the task is to assign a budget $y(s)$ to each channel s in S with the goal of maximizing the expected number of influenced people $\mathcal{I}(y)$. Each edge $(s, t) \in E$ between channel s and

person t is weighted with a probability p_{st} that, e.g., an advertisement on radio station s will influence person t to buy some product. The budget $y(s)$ controls how many independent attempts are made via the channel s to influence the people in T . The probability that a customer t is influenced when the advertising budget is y is

$$I_t(y) = 1 - \prod_{(s,t) \in E} [1 - p_{st}]^{y(s)}, \quad (1)$$

and hence the expected number of influenced people is $\mathcal{I}(y) = \sum_{t \in T} I_t(y)$. We write $\mathcal{I}(y; p) = \mathcal{I}(y)$ to make the dependence on the probabilities p_{st} explicit. The total budget y must remain within some feasible set \mathcal{Y} which may encode e.g. a total budget limit $\sum_{s \in S} y(s) \leq C$. We allow the budgets y to be continuous, as in (Bian et al., 2017).

Since its introduction by Alon et al. (2012), several works have extended the formulation of Budget Allocation and provided algorithms (Bian et al., 2017; Hatano et al., 2015; Maehara et al., 2015; Soma et al., 2014; Soma & Yoshida, 2015). Budget Allocation may also be viewed as influence maximization on a bipartite graph, where information spreads as in the Independent Cascade model. For integer y , Budget Allocation and Influence Maximization are NP-hard. Yet, constant-factor approximations are possible, and build on the fact that the influence function is submodular in the binary case, and *DR-submodular* in the integer case (Soma et al., 2014; Hatano et al., 2015). If y is continuous, the problem is a concave maximization problem.

The formulation of Budget Allocation assumes that the transmission probabilities are known exactly. But this is rarely true in practice. Typically, the probabilities p_{st} , and possibly the graph itself, must be inferred from observations (Gomez Rodriguez et al., 2010; Du et al., 2013; Narasimhan et al., 2015; Du et al., 2014; Netrapalli & Sanghavi, 2012). In Section 4 we will see that a misspecification or point estimate of parameters p_{st} can lead to much reduced outcomes. A more realistic assumption is to know *confidence intervals* for the p_{st} . Realizing this severe deficiency, recent work studied robust versions of Influence Maximization, where a budget y must be chosen that maximizes the worst-case approximation ratio over a set of possible influence functions (He & Kempe, 2016; Chen et al., 2016; Lowalekar et al., 2016). The resulting optimization problem is hard but admits bicriteria approximations.

¹Massachusetts Institute of Technology. Correspondence to: Matthew Staib <mstaib@mit.edu>, Stefanie Jegelka <stefje@mit.edu>.

In this work, we revisit Budget Allocation under uncertainty from the perspective of robust optimization (Bertsimas et al., 2011; Ben-Tal et al., 2009). We maximize the worst-case influence – not approximation ratio – for p in a confidence set centered around the “best guess” (e.g., posterior mean). This avoids pitfalls of the approximation ratio formulation (which can be misled to return poor worst-case budgets, as demonstrated in Appendix A), while also allowing us to formulate the problem as a max-min game:

$$\max_{y \in \mathcal{Y}} \min_{p \in \mathcal{P}} \mathcal{I}(y; p), \quad (2)$$

where an “adversary” can arbitrarily manipulate p within the confidence set \mathcal{P} . With p fixed, $\mathcal{I}(y; p)$ is concave in y . However, the influence function $\mathcal{I}(y; p)$ is not convex, and not even quasiconvex, in the adversary’s variables p_{st} .

The new, key insight we exploit in this work is that $\mathcal{I}(y; p)$ has the property of *continuous submodularity* in p – in contrast to previously exploited submodular maximization in y – and can hence be minimized by generalizing techniques from discrete submodular optimization (Bach, 2015). The techniques in (Bach, 2015), however, are restricted to box constraints, and do not directly apply to our confidence sets. In fact, general constrained submodular minimization is hard (Svitkina & Fleischer, 2011; Goel et al., 2009; Iwata & Nagano, 2009). We make the following contributions:

1. We present an algorithm with optimality bounds for Robust Budget Allocation in the nonconvex adversarial scenario (2).
2. We provide the first results for continuous submodular minimization with box constraints and one more “nice” constraint, and conditions under which the algorithm is guaranteed to return a global optimum.

1.1. Background and Related Work

We begin with some background material and, along the way, discuss related work.

1.1.1. SUBMODULARITY OVER THE INTEGER LATTICE AND CONTINUOUS DOMAINS

Submodularity is perhaps best known as a property of set functions. A function $F : 2^V \rightarrow \mathbb{R}$ defined on subsets $S \subseteq V$ of a ground set V is *submodular* if for all sets $S, T \subseteq V$, it holds that $F(S) + F(T) \geq F(S \cap T) + F(S \cup T)$. A similar definition extends to functions defined over a distributive lattice \mathcal{L} , e.g. the integer lattice. Such a function f is submodular if for all $x, y \in \mathcal{L}$, it holds that

$$f(x) + f(y) \geq f(x \vee y) + f(x \wedge y). \quad (3)$$

For the integer lattice and vectors x, y , $x \vee y$ denotes the coordinate-wise maximum and $x \wedge y$ the coordinate-wise

minimum. Submodularity has also been considered on continuous domains $\mathcal{X} \subset \mathbb{R}^d$, where, if f is also twice-differentiable, the property of submodularity means that all off-diagonal entries of the the Hessian are nonpositive, i.e., $\frac{\partial^2 f(x)}{\partial x_i \partial x_j} \leq 0$ for all $i \neq j$ (Topkis, 1978, Theorem 3.2). These functions may be convex, concave, or neither.

Submodular functions on lattices can be minimized by a reduction to set functions, more precisely, ring families (Birkhoff, 1937). Combinatorial algorithms for submodular optimization on lattices are discussed in (Khachaturov et al., 2012). More recently, Bach (2015) extended results based on the convex Lovász extension, by building on connections to optimal transport. The subclass of L^1 -convex functions admits strongly polynomial time minimization (Murota, 2003; Kolmogorov & Shioura, 2009; Murota & Shioura, 2014), but does not apply in our setting.

Similarly, results for submodular maximization extend to integer lattices, e.g. (Gottschalk & Peis, 2015). Stronger results are possible if the submodular function also satisfies *diminishing returns*: for all $x \leq y$ (coordinate-wise) and i such that $y + e_i \in \mathcal{X}$, it holds that $f(x + e_i) - f(x) \geq f(y + e_i) - f(y)$. For such *DR-submodular* functions, many approximation results for the set function case extend (Bian et al., 2017; Soma & Yoshida, 2015; Soma et al., 2014). In particular, Ene & Nguyen (2016) show a generic reduction to set function optimization that they apply to maximization. In fact, it also applies to minimization:

Proposition 1.1. *A DR-submodular function f defined on $\prod_{i=1}^n [k_i]$ can be minimized in strongly polynomial time $O(n^4 \log^4 k \cdot \log^2(n \log k) \cdot EO + n^4 \log^4 k \cdot \log^{O(1)}(n \log k))$, where $k = \max_i k_i$ and EO is the time complexity of evaluating f . Here, $[k_i] = \{0, 1, \dots, k_i - 1\}$.*

In particular, the time complexity is logarithmic in k . For general lattice submodular functions, this is not possible without further assumptions.

1.1.2. RELATED PROBLEMS

A sister problem of Budget Allocation is *Influence Maximization* on general graphs, where a set of seed nodes is selected to start a propagation process. The influence function is still monotone submodular and amenable to the greedy algorithm (Kempe et al., 2003), but it cannot be evaluated explicitly and requires approximation (Chen et al., 2010). *Stochastic Coverage* (Goemans & Vondrák, 2006) is a version of Set Cover where the covering sets $S_i \subseteq V$ are random. A variant of Budget Allocation can be written as stochastic coverage with multiplicity. Stochastic Coverage has mainly been studied in the online or adaptive setting, where logarithmic approximation factors can be achieved (Golovin & Krause, 2011; Deshpande et al., 2016; Adamczyk et al., 2016).

Our objective function (2) is a *signomial* in p , i.e., a linear combination of monomials of the form $\prod_i x_i^{c_i}$. General signomial optimization is NP-hard (Chiang, 2005), but certain subclasses are tractable: *posynomials* with all non-negative coefficients can be minimized via Geometric Programming (Boyd et al., 2007), and signomials with a single negative coefficient admit sum of squares-like relaxations (Chandrasekaran & Shah, 2016). Our problem, a constrained posynomial maximization, is not in general a geometric program. Some work addresses this setting via monomial approximation (Pascual & Ben-Israel, 1970; Ecker, 1980), but, to our knowledge, our algorithm is the first that solves this problem to arbitrary accuracy.

1.1.3. ROBUST OPTIMIZATION

Two prominent strategies of addressing uncertainty in parameters of optimization problems are stochastic and robust optimization. If the distribution of the parameters is known (stochastic optimization), formulations such as value-at-risk (VaR) and conditional value-at-risk (CVaR) (Rockafellar & Uryasev, 2000; 2002) apply. In contrast, robust optimization (Ben-Tal et al., 2009; Bertsimas et al., 2011) assumes that the parameters (of the cost function and constraints) can vary arbitrarily within a known confidence set U , and the aim is to optimize the worst-case setting, i.e.,

$$\min_y \sup_{u, A, b \in U} \{g(y; u) \text{ s.t. } Ay \leq b\}. \quad (4)$$

Here, we will only have uncertainty in the cost function.

In this paper we are principally concerned with robust maximization of the continuous influence function $\mathcal{I}(y)$, but mention some results for the discrete case. While there exist results for robust and CVaR optimization of modular (linear) functions (Nikolova, 2010; Bertsimas & Sim, 2003), submodular objectives do not in general admit such optimization (Maehara, 2015), but variants admit approximations (Zhang et al., 2014). The brittleness of submodular optimization under noise has been studied in (Balkanski et al., 2016; 2017; Hassidim & Singer, 2016).

Approximations for robust submodular and influence optimization have been studied in (Krause et al., 2008; He & Kempe, 2016; Chen et al., 2016; Lowalekar et al., 2016), where an adversary can pick among a *finite* set of objective functions or remove selected elements (Orlin et al., 2016).

2. Robust and Stochastic Budget Allocation

The unknown parameters in Budget Allocation are the transmission probabilities p_{st} or edge weights in a graph. If these are estimated from data, we may have posterior distributions or, a weaker assumption, confidence sets for the parameters. For ease of notation, we will work with the failure probabilities $x_{st} = 1 - p_{st}$ instead of the p_{st}

directly, and write $\mathcal{I}(y; x)$ instead of $\mathcal{I}(y; p)$.

2.1. Stochastic Optimization

If a (posterior) distribution of the parameters is known, a simple strategy is to use expectations. We place a uniform prior on x_{st} , and observe n_{st} independent observations drawn from $\text{Ber}(x_{st})$. If we observe α_{st} failures and β_{st} successes, the resulting posterior distribution on the variable X_{st} is $\text{Beta}(1 + \alpha_{st}, 1 + \beta_{st})$. Given such a posterior, we may optimize

$$\max_{y \in \mathcal{Y}} \mathcal{I}(y; \mathbb{E}[X]), \text{ or} \quad (5)$$

$$\max_{y \in \mathcal{Y}} \mathbb{E}[\mathcal{I}(y; X)]. \quad (6)$$

Proposition 2.1. *Problems (5) and (6) are concave maximization problems over the (convex) set \mathcal{Y} and can be solved exactly.*

Concavity of (6) follows since it is an expectation over concave functions, and the problem can be solved by stochastic gradient ascent or by explicitly computing gradients.

Merely maximizing expectation does not explicitly account for volatility and hence risk. One option is to include variance (Ben-Tal & Nemirovski, 2000; Bertsimas et al., 2011; Atamtürk & Narayanan, 2008):

$$\min_{y \in \mathcal{Y}} -\mathbb{E}[\mathcal{I}(y; X)] + \varepsilon \sqrt{\text{Var}(\mathcal{I}(y; X))}, \quad (7)$$

but in our case this CVaR formulation seems difficult:

Fact 2.1. *For y in the nonnegative orthant, the term $\sqrt{\text{Var}(\mathcal{I}(y; X))}$ need not be convex or concave, and need not be submodular or supermodular.*

This observation does not rule out a solution, but the apparent difficulties further motivate a robust formulation that, as we will see, is amenable to optimization.

2.2. Robust Optimization

The focus of this work is the robust version of Budget Allocation, where we allow an adversary to arbitrarily set the parameters x within an uncertainty set \mathcal{X} . This uncertainty set may result, for instance, from a known distribution, or simply assumed bounds. Formally, we solve

$$\max_{y \in \mathcal{Y}} \min_{x \in \mathcal{X}} \mathcal{I}(y; x), \quad (8)$$

where $\mathcal{Y} \subset \mathbb{R}_+^S$ is a convex set with an efficient projection oracle, and \mathcal{X} is an uncertainty set containing an estimate \hat{x} . In the sequel, we use uncertainty sets $\mathcal{X} = \{x \in \text{Box}(l, u) : R(x) \leq B\}$, where R is a distance (or divergence) from the estimate \hat{x} , and $\text{Box}(l, u)$ is the box $\prod_{(s,t) \in E} [l_{st}, u_{st}]$. The intervals $[l_{st}, u_{st}]$ can be thought of

as either confidence intervals around \hat{x} , or, if $[l_{st}, u_{st}] = [0, 1]$, enforce that each x_{st} is a valid probability.

Common examples of uncertainty sets used in Robust Optimization are *Ellipsoidal* and *D-norm uncertainty sets* (Bertsimas et al., 2011). Our algorithm in Section 3.1 applies to both.

Ellipsoidal uncertainty. The ellipsoidal or quadratic uncertainty set is defined by

$$\mathcal{X}^Q(\gamma) = \{x \in \text{Box}(0, 1) : (x - \hat{x})^T \Sigma^{-1} (x - \hat{x}) \leq \gamma\},$$

where Σ is the covariance of the random vector X of probabilities distributed according to our Beta posteriors. In our case, since the distributions on each x_{st} are independent, Σ^{-1} is actually diagonal. Writing $\Sigma = \text{diag}(\sigma^2)$, we have

$$\mathcal{X}^Q(\gamma) = \left\{ x \in \text{Box}(0, 1) : \sum_{(s,t) \in E} R_{st}(x_{st}) \leq \gamma \right\},$$

where $R_{st}(x) = (x_{st} - \hat{x}_{st})^2 \sigma_{st}^{-2}$.

D-norm uncertainty. The D-norm uncertainty set is similar to an ℓ_1 -ball around \hat{x} , and is defined as

$$\mathcal{X}^D(\gamma) = \left\{ x : \exists c \in \text{Box}(0, 1) \text{ s.t.} \right.$$

$$\left. x_{st} = \hat{x}_{st} + (u_{st} - \hat{x}_{st})c_{st}, \sum_{(s,t) \in E} c_{st} \leq \gamma \right\}.$$

Essentially, we allow an adversary to increase \hat{x}_{st} up to some upper bound u_{st} , subject to some total budget γ across all terms x_{st} . The set $\mathcal{X}^D(\gamma)$ can be rewritten as

$$\mathcal{X}^D(\gamma) = \left\{ x \in \text{Box}(\hat{x}, u) : \sum_{(s,t) \in E} R_{st}(x_{st}) \leq \gamma \right\},$$

where $R_{st}(x_{st}) = (x_{st} - \hat{x}_{st}) / (u_{st} - \hat{x}_{st})$ is the fraction of the interval $[\hat{x}_{st}, u_{st}]$ we have used up in increasing x_{st} .

The min-max formulation $\max_{y \in \mathcal{Y}} \min_{x \in \mathcal{X}} \mathcal{I}(y; x)$ has several merits: the model is not tied to a specific learning algorithm for the probabilities x as long as we can choose a suitable confidence set. Moreover, this formulation allows to fully hedge against a worst-case scenario.

3. Optimization Algorithm

As noted above, the function $\mathcal{I}(y; x)$ is concave as a function of y for fixed x . As a pointwise minimum of concave functions, $F(y) := \min_{x \in \mathcal{X}} \mathcal{I}(y; x)$ is concave. Hence, if we can compute subgradients of $F(y)$, we can solve our max-min-problem via the subgradient method, as outlined in Algorithm 1.

A subgradient $g_y \in \partial F(y)$ at y is given by the gradient of $\mathcal{I}(y; x^*)$ for the minimizing $x^* \in \arg \min_{x \in \mathcal{X}} \mathcal{I}(y; x)$, i.e.,

Algorithm 1 Subgradient Ascent

Input: suboptimality tolerance $\varepsilon > 0$, initial feasible budget $y^{(0)} \in \mathcal{Y}$

Output: ε -optimal budget y for Problem (8)

repeat

$$x^{(k)} \leftarrow \arg \min_{x \in \mathcal{X}} \mathcal{I}(y^{(k)}; x)$$

$$g^{(k)} \leftarrow \nabla_y \mathcal{I}(y^{(k)}; x^{(k)})$$

$$L^{(k)} \leftarrow \mathcal{I}(y^{(k)}; x^{(k)})$$

$$U^{(k)} \leftarrow \max_{y \in \mathcal{Y}} \mathcal{I}(y; x^{(k)})$$

$$\gamma^{(k)} \leftarrow (U^{(k)} - L^{(k)}) / \|g^{(k)}\|_2^2$$

$$y^{(k+1)} \leftarrow \text{proj}_{\mathcal{Y}}(y^{(k)} + \gamma^{(k)} g^{(k)})$$

$$k \leftarrow k + 1$$

until $U^{(k)} - L^{(k)} \leq \varepsilon$

$g_y = \nabla_y \mathcal{I}(y; x^*)$. Hence, we must be able to compute x^* for any y . We also obtain a duality gap: for any x', y' we have

$$\min_{x \in \mathcal{X}} \mathcal{I}(y'; x) \leq \max_{y \in \mathcal{Y}} \min_{x \in \mathcal{X}} \mathcal{I}(y; x) \leq \max_{y \in \mathcal{Y}} \mathcal{I}(y; x'). \quad (9)$$

This means we can estimate the optimal value \mathcal{I}^* and use it in Polyak's stepsize rule for the subgradient method (Polyak, 1987).

But $\mathcal{I}(y; x)$ is not convex in x , and not even quasiconvex. For example, standard methods (Wainwright & Chiang, 2004, Chapter 12) imply that $f(x_1, x_2, x_3) = 1 - x_1 x_2 - \sqrt{x_3}$ is not quasiconvex on \mathbb{R}_+^3 . Moreover, the above-mentioned signomial optimization techniques do not apply for an exact solution either. So, it is not immediately clear that we can solve the inner optimization problem.

The key insight we will be using is that $\mathcal{I}(y; x)$ has a different beneficial property: while not convex, $\mathcal{I}(y; x)$ as a function of x is *continuous submodular*.

Lemma 3.1. *Suppose we have $n \geq 1$ differentiable functions $f_i : \mathbb{R} \rightarrow \mathbb{R}_+$, for $i = 1, \dots, n$, either all nonincreasing or all nondecreasing. Then, $f(x) = \prod_{i=1}^n f_i(x_i)$ is a continuous supermodular function from \mathbb{R}^n to \mathbb{R}_+ .*

Proof. For $n = 1$, the resulting function is modular and therefore supermodular. In the case $n \geq 2$, we simply need to compute derivatives. The mixed derivatives are

$$\frac{\partial f}{\partial x_i \partial x_j} = f'_i(x_i) f'_j(x_j) \cdot \prod_{k \neq i, j} f_k(x_k). \quad (10)$$

By monotonicity, f'_i and f'_j have the same sign, so their product is nonnegative, and since each f_k is nonnegative, the entire expression is nonnegative. Hence, $f(x)$ is continuous supermodular by Theorem 3.2 of (Topkis, 1978). \square

Corollary 3.1. *The influence function $\mathcal{I}(y; x)$ defined in Section 2 is continuous submodular in x over the nonnegative orthant, for each $y \geq 0$.*

Proof. Since submodularity is preserved under summation, it suffices to show that each function $I_t(y)$ is continuous submodular. By Lemma 3.1, since $f_s(z) = z^{y(s)}$ is non-negative and monotone nondecreasing for $y(s) \geq 0$, the product $\prod_{(s,t) \in E} x_{st}^{y(s)}$ is continuous supermodular in x . Flipping the sign and adding a constant term yields $I_t(y)$, which is hence continuous submodular. \square

Conjecture 3.1. *Strong duality holds, i.e.*

$$\max_{y \in \mathcal{Y}} \min_{x \in \mathcal{X}} \mathcal{I}(y; x) = \min_{x \in \mathcal{X}} \max_{y \in \mathcal{Y}} \mathcal{I}(y; x). \quad (11)$$

If strong duality holds, then the duality gap $\max_{y \in \mathcal{Y}} \mathcal{I}(y; x^*) - \min_{x \in \mathcal{X}} \mathcal{I}(y^*; x)$ in Equation (9) is zero at optimality. If $\mathcal{I}(y; x)$ were quasiconvex in x , strong duality would hold by Sion's min-max theorem, but this is not the case. In practice, we observe that the duality gap always converges to zero.

Bach (2015) demonstrates how to minimize a continuous submodular function $H(x)$ subject to box constraints $x \in \text{Box}(l, u)$, up to an arbitrary suboptimality gap $\varepsilon > 0$. The constraint set \mathcal{X} in our Robust Budget Allocation problem, however, has box constraints with an additional constraint $R(x) \leq B$. This case is not addressed in any previous work. Fortunately, for a large class of functions R , there is still an efficient algorithm for continuous submodular minimization, which we present in the next section.

3.1. Constrained Continuous Submodular Function Minimization

We next address an algorithm for minimizing a monotone continuous submodular function $H(x)$ subject to box constraints $x \in \text{Box}(l, u)$ and a constraint $R(x) \leq B$:

$$\begin{aligned} & \text{minimize} && H(x) \\ & \text{s.t.} && R(x) \leq B \\ & && x \in \text{Box}(l, u). \end{aligned} \quad (12)$$

If H and R were convex, the constrained problem would be equivalent to solving, with the right Lagrange multiplier $\lambda^* \geq 0$:

$$\begin{aligned} & \text{minimize} && H(x) + \lambda^* R(x) \\ & \text{s.t.} && x \in \text{Box}(l, u). \end{aligned} \quad (13)$$

Although H and R are not necessarily convex here, it turns out that a similar approach indeed applies. The main idea of our approach bears similarity with (Nagano et al., 2011) for the set function case, but our setting with continuous functions and various uncertainty sets is more general, and requires more argumentation. We outline our theoretical results here, and defer further implementation details and proofs to the appendix.

Following (Bach, 2015), we discretize the problem; for a sufficiently fine discretization, we will achieve arbitrary accuracy. Let A be an *interpolation mapping* that maps the

discrete set $\prod_{i=1}^n [k_i]$ into $\text{Box}(l, u) = \prod_{i=1}^n [l_i, u_i]$ via the componentwise interpolation functions $A_i : [k_i] \rightarrow [l_i, u_i]$. We say A_i is δ -fine if $A_i(x_i + 1) - A_i(x_i) \leq \delta$ for all $x_i \in \{0, 1, \dots, k_i - 2\}$, and we say the full interpolation function A is δ -fine if each A_i is δ -fine.

This mapping yields functions $H^\delta : \prod_{i=1}^n [k_i] \rightarrow \mathbb{R}$ and $R^\delta : \prod_{i=1}^n [k_i] \rightarrow \mathbb{R}$ via $H^\delta(x) = H(A(x))$ and $R^\delta(x) = R(A(x))$. H^δ is lattice submodular (on the integer lattice). This construction leads to a reduction of Problem (12) to a submodular minimization problem over the integer lattice:

$$\begin{aligned} & \text{minimize} && H^\delta(x) + \lambda R^\delta(x) \\ & \text{s.t.} && x \in \prod_{i=1}^n [k_i]. \end{aligned} \quad (14)$$

Ideally, there should then exist a λ such that the associated minimizer $x(\lambda)$ yields a close to optimal solution for the constrained problem. Theorem 3.1 below states that this is indeed the case.

Moreover, a second benefit of submodularity is that we can find the entire solution path for Problem (14) by solving a single optimization problem.

Lemma 3.2. *Suppose H is continuous submodular, and suppose the regularizer R is strictly increasing and separable: $R(x) = \sum_{i=1}^n R_i(x_i)$. Then we can recover a minimizer $x(\lambda)$ for the induced discrete Problem (14) for any $\lambda \in \mathbb{R}$ by solving a single convex optimization problem.*

The problem in question arises from a relaxation h_\downarrow that extends H^δ in each coordinate i to a function on distributions over the domain $[k_i]$. These distributions are represented via their inverse cumulative distribution functions ρ_i , which take the coordinate x_i as input, and output the probability of exceeding x_i . The function h_\downarrow is an analogue of the *Lovász extension* of set functions to continuous submodular functions (Bach, 2015), it is convex and coincides with H^δ on lattice points.

Formally, this resulting single optimization problem is:

$$\begin{aligned} & \text{minimize} && h_\downarrow(\rho) + \sum_{i=1}^n \sum_{j_i=1}^{k_i-1} a_{ix_i}(\rho_i(x_i)) \\ & \text{s.t.} && \rho \in \prod_{i=1}^n \mathbb{R}_{\downarrow}^{k_i-1} \end{aligned} \quad (15)$$

where $\mathbb{R}_{\downarrow}^k$ refers to the set of ordered vectors $z \in \mathbb{R}^k$ that satisfy $z_1 \geq z_2 \geq \dots \geq z_k$, the notation $\rho_i(x_i)$ denotes the x_i -th coordinate of the vector ρ_i , and the a_{ix_i} are strictly convex functions given by

$$a_{ix_i}(t) = \frac{1}{2} t^2 \cdot [R_i^\delta(x_i) - R_i^\delta(x_i - 1)]. \quad (16)$$

Problem (15) can be solved by Frank-Wolfe methods (Frank & Wolfe, 1956; Dunn & Harshbarger, 1978; Lacoste-Julien, 2016; Jaggi, 2013). This is because the greedy algorithm for computing subgradients of the Lovász

extension can be generalized, and yields a linear optimization oracle for the dual of Problem (15). We detail the relationship between Problems (14) and (15), as well as how to implement the Frank-Wolfe methods, in Appendix C.

Let ρ^* be the optimal solution for Problem (15). For any λ , we obtain a rounded solution $x(\lambda)$ for Problem (14) by thresholding: we set $x(\lambda)_i = \max\{j \mid 1 \leq j \leq k_i - 1, \rho_i^*(j) \geq \lambda\}$, or zero if $\rho_i^*(j) < \lambda$ for all j . Each $x(\lambda')$ is the optimal solution for Problem (14) with $\lambda = \lambda'$. We use the largest parameterized solution $x(\lambda)$ that is still feasible, i.e. the solution $x(\lambda^*)$ where λ^* solves

$$\begin{aligned} \min \quad & H^\delta(x(\lambda)) \\ \text{s.t.} \quad & \lambda \geq 0 \\ & R^\delta(x(\lambda)) \leq B. \end{aligned} \quad (17)$$

This λ^* can be found efficiently via binary search or a linear scan.

Theorem 3.1. *Let H be continuous submodular and monotone decreasing, with ℓ_∞ -Lipschitz constant G , and let R be strictly increasing and separable. Assume all entries $\rho_i^*(j)$ of the optimal solution ρ^* of Problem (15) are distinct. Let $x' = A(x(\lambda^*))$ be the thresholding corresponding to the optimal solution λ^* of Problem (17), mapped back into the original continuous domain \mathcal{X} . Then x' is feasible for the continuous Problem (12), and is a $2G\delta$ -approximate solution:*

$$H(x') \leq 2G\delta + \min_{x \in \text{Box}(l, u), R(x) \leq B} H(x).$$

Theorem 3.1 implies an algorithm for solving Problem (12) to ε -optimality: (1) set $\delta = \varepsilon/G$, (2) compute ρ^* which solves Problem (15), (3) find the optimal thresholding of ρ^* by determining the smallest λ^* for which $R^\delta(x(\lambda^*)) \leq B$, and (4) map $x(\lambda^*)$ back into continuous space via the interpolation mapping A .

Optimality Bounds. Theorem 3.1 is proved by comparing x' and x^* to the optimal solution on the discretized mesh

$$x_d^* \in \underset{x \in \prod_{i=1}^n [k_i]: R^\delta(x) \leq B}{\text{argmin}} H^\delta(x).$$

Beyond the theoretical guarantee of Theorem 3.1, for any problem instance and candidate solution x' , we can compute a bound on the gap between $H(x')$ and $H^\delta(x_d^*)$. The following two bounds are proved in the appendix:

1. We can generate a discrete point $x(\lambda_+)$ satisfying

$$H(x') \leq [H(x') - H^\delta(x(\lambda_+))] + H^\delta(x_d^*).$$

2. The Lagrangian yields the bound

$$H(x') \leq \lambda^*(B - R(x')) + H^\delta(x_d^*).$$

Improvements. The requirement in Theorem 3.1 that the elements of ρ^* be distinct may seem somewhat restrictive, but as long as ρ^* has distinct elements in the neighborhood of our particular λ^* , this bound still holds. We see in Section 4.1.1 that in practice, ρ^* almost always has distinct elements in the regime we care about, and the bounds of Remark 3.1 are very good.

If H is DR-submodular and R is affine in each coordinate, then Problem (14) can be represented more compactly via the reduction of Ene & Nguyen (2016), and hence problem (12) can be solved more efficiently. In particular, the influence function $\mathcal{I}(y; x)$ is DR-submodular in x when for each s , $y(s) = 0$ or $y(s) \geq 1$.

3.2. Application to Robust Budget Allocation

The above algorithm directly applies to Robust Allocation with the uncertainty sets in Section 2.2. The ellipsoidal uncertainty set \mathcal{X}^Q corresponds to the constraint that $\sum_{(s,t) \in E} R_{st}(x_{st}) \leq \gamma$ with $R_{st}(x) = (x_{st} - \hat{x}_{st})^2 \sigma_{st}^{-2}$, and $x \in \text{Box}(0, 1)$. By the monotonicity of $\mathcal{I}(x, y)$, there is never incentive to reduce any x_{st} below \hat{x}_{st} , so we can replace $\text{Box}(0, 1)$ with $\text{Box}(\hat{x}, 1)$. On this interval, each R_{st} is strictly increasing, and Theorem 3.1 applies.

For D-norm sets, we have $R_{st}(x_{st}) = (x_{st} - \hat{x}_{st}) / (u_{st} - \hat{x}_{st})$. Since each R_{st} is monotone, Theorem 3.1 applies.

Runtime and Alternatives. Since the core algorithm is Frank-Wolfe, it is straightforward to show that Problem (15) can be solved to ε -suboptimality in time $O(\varepsilon^{-1} n^2 \delta^{-3} \alpha^{-1} |T|^2 \log n \delta^{-1})$, where α is the minimum derivative of the functions R_i . If ρ^* has distinct elements separated by η , then choosing $\varepsilon = \eta^2 \alpha \delta / 8$ results in an exact solution to (14) in time $O(\eta^{-2} n^2 \delta^{-4} \alpha^{-2} |T|^2 \log n \delta^{-1})$.

Noting that $H^\delta + \lambda R^\delta$ is submodular for all λ , one could instead perform binary search over λ , each time converting the objective into a submodular set function via Birkhoff's theorem and solving submodular minimization e.g. via one of the recent fast methods (Chakrabarty et al., 2017; Lee et al., 2015). However, we are not aware of a practical implementation of the algorithm in (Lee et al., 2015). The algorithm in (Chakrabarty et al., 2017) yields a solution in expectation. This approach also requires care in the precision of the search over λ , whereas our approach searches directly over the $O(n\delta^{-1})$ elements of ρ^* .

4. Experiments

We evaluate our Robust Budget Allocation algorithm on both synthetic test data and a real-world bidding dataset from Yahoo! Webscope (yah) to demonstrate that our method yields real improvements. For all experiments, we




Figure 1. Visualization of the sorted values of $\rho_i^*(j)$ (blue dots) with comparison to the particular Lagrange multiplier λ^* (orange line). In most regimes there are no duplicate values, so that Theorem 3.1 applies. The theorem only needs distinctness at λ^* .

used Algorithm 1 as the outer loop. For the inner submodular minimization step, we implemented the pairwise Frank-Wolfe algorithm of (Lacoste-Julien & Jaggi, 2015). In all cases, the feasible set of budgets \mathcal{Y} is $\{y \in \mathbb{R}_+^S : \sum_{s \in S} y(s) \leq C\}$ where the specific budget C depends on the experiment. Our code is available at git.io/vHXkO.

4.1. Synthetic

On the synthetic data, we probe two questions: (1) how often does the distinctness condition of Theorem 3.1 hold, so that we are guaranteed an optimal solution; and (2) what is the gain of using a robust versus non-robust solution in an adversarial setting? For both settings, we set $|S| = 6$ and $|T| = 2$ and discretize with $\delta = 0.001$. We generated true probabilities p_{st} , created Beta posteriors, and built both Ellipsoidal uncertainty sets $\mathcal{X}^Q(\gamma)$ and D-norm sets $\mathcal{X}^D(\gamma)$.

4.1.1. OPTIMALITY

Theorem 3.1 and Remark 3.1 demand that the values $\rho_i^*(j)$ be distinct at our chosen Lagrange multiplier λ^* and, under this condition, guarantee optimality. We illustrate this in four examples: for Ellipsoidal or a D-norm uncertainty set, and a total influence budget $C \in \{0.4, 4\}$. Figure 3 shows all elements of ρ^* in sorted order, as well as a horizontal line indicating our Lagrange multiplier λ^* which serves as a threshold. Despite some plateaus, the entries $\rho_i^*(j)$ are distinct in most regimes, in particular around λ^* , the regime that is needed for our results. Moreover, in practice (on the Yahoo data) we observe later in Figure 3 that both solution-dependent bounds from Remark 3.1 are very good, and all solutions are optimal within a very small gap.

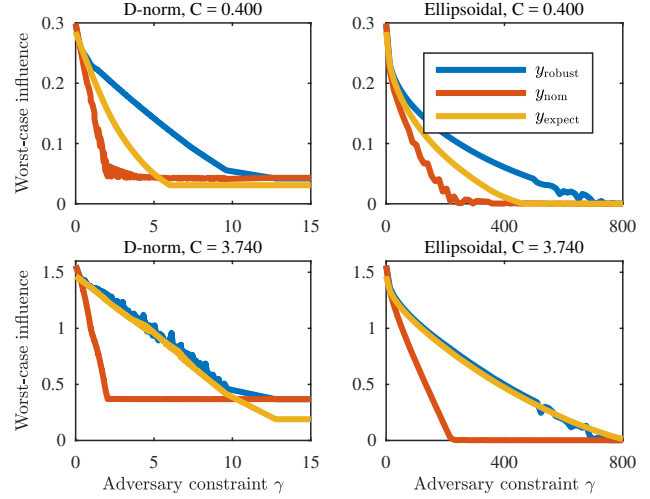


Figure 2. Comparison of worst-case expected influences for D-norm uncertainty sets $\mathcal{X}^D(\gamma)$ (left) and ellipsoidal uncertainty sets $\mathcal{X}^Q(\gamma)$ (right), for different total budget bounds C . For any particular adversary budget γ , we compare $\min_{x \in \mathcal{X}(\gamma)} \mathcal{I}(y; x)$ for each candidate allocation y .

4.1.2. ROBUSTNESS AND QUALITY

Next, we probe the effect of a robust versus non-robust solution for different uncertainty sets and budgets γ of the adversary. We compare our robust solution with using a point estimate for x , i.e., $y_{\text{nom}} \in \arg\max_{y \in \mathcal{Y}} \mathcal{I}(y; \hat{x})$, treating estimates as ground truth, and the stochastic solution $y_{\text{expect}} \in \arg\max_{y \in \mathcal{Y}} \mathbb{E}[\mathcal{I}(y; X)]$ as per Section 2.1. These two optimization problems were solved via standard first-order methods using TFOCS (Becker et al., 2011).

Figure 2 demonstrates that indeed, the alternative budgets are sensitive to the adversary and the robustly-chosen budget y_{robust} performs better, even in cases where the other budgets achieve zero influence. When the total budget C is large, y_{expect} performs nearly as well as y_{robust} , but when resources are scarce (C is small) and the actual choice seems to matter more, y_{robust} performs far better.

4.2. Yahoo! data

To evaluate our method on real-world data, we formulate a Budget Allocation instance on advertiser bidding data from Yahoo! Webscope (yah). This dataset logs bids on 1000 different phrases by advertising accounts. We map the phrases to channels S and the accounts to customers T , with an edge between s and t if a corresponding bid was made. For each pair (s, t) , we draw the associated transmission probability p_{st} uniformly from $[0, 0.4]$. We bias these towards zero because we expect people not to be easily influenced by advertising in the real world. We then generate an estimate \hat{p}_{st} and build up a posterior by gener-




Figure 3. Convergence properties of our algorithm on real data. In the first plot, ‘p’ and ‘d’ refer to primal and dual values, with dual gap shown on the second plot. The third plot demonstrates that the problem-dependent suboptimality bounds of Remark 3.1 (x for $x(\lambda_+)$ and L for Lagrangian) are very small (good) for all inner iterations of this run.

ating n_{st} samples from $\text{Ber}(p_{st})$, where n_{st} is the number of bids between s and t in the dataset.

This transformation yields a bipartite graph with $|S| = 1000$, $|T| = 10475$, and more than 50,000 edges that we use for Budget Allocation. In our experiments, the typical gap between the naive y_{nom} and robust y_{robust} was 100-500 expected influenced people. We plot convergence of the outer loop in Figure 3, where we observe fast convergence of both primal influence value and the dual bound.

4.3. Comparison to first-order methods

Given the success of first-order methods on nonconvex problems in practice, it is natural to compare these to our method for finding the worst-case vector x . On one of our Yahoo problem instances with D-norm uncertainty set, we compared our submodular minimization scheme to Frank-Wolfe with fixed stepsize as in (Lacoste-Julien, 2016), implementing the linear oracle using MOSEK (MOSEK ApS, 2015). Interestingly, from various initializations, Frank-Wolfe finds an optimal solution, as verified by comparing to the guaranteed solution of our algorithm. Note that, due to non-convexity, there are no formal guarantees for Frank-Wolfe to be optimal here, motivating the question of global convergence properties of Frank-Wolfe in the presence of submodularity.

It is important to note that there are many cases where first-order methods are inefficient or do not apply to our setup. These methods require either a projection oracle (PO) onto or linear optimization oracle (LO) over the feasible set \mathcal{X} defined by ℓ , u and $R(x)$. The D-norm set admits a LO via linear programming, but we are not aware of any efficient LO for Ellipsoidal uncertainty, nor PO for either set, that does not require quadratic programming. Even more, our algorithm applies for nonconvex functions $R(x)$ which induce nonconvex feasible sets \mathcal{X} . Such nonconvex sets may not even admit a unique projection, while our algorithm achieves provable solutions.




Figure 4. Convergence properties of Frank-Wolfe (FW), versus the optimal value attained with our scheme (SFM).

5. Conclusion

We address the issue of uncertain parameters (or, model misspecification) in Budget Allocation or Bipartite Influence Maximization (Alon et al., 2012) from a robust optimization perspective. The resulting *Robust Budget Allocation* is a nonconvex-concave saddle point problem. Although the inner optimization problem is nonconvex, we show how continuous submodularity can be leveraged to solve the problem to arbitrary accuracy ε , as can be verified with the proposed bounds on the duality gap. In particular, our approach extends continuous submodular minimization methods (Bach, 2015) to more general constraint sets, introducing a mechanism to solve a new class of constrained nonconvex optimization problems. We confirm on synthetic and real data that our method finds high-quality solutions that are robust to parameters varying arbitrarily in an uncertainty set, and scales up to graphs with over 50,000 edges.

There are many compelling directions for further study. The uncertainty sets we use are standard in the robust optimization literature, but have not been applied to e.g. Robust Influence Maximization; it would be interesting to generalize our ideas to general graphs. Finally, despite the inherent nonconvexity of our problem, first-order methods are often able to find a globally optimal solution. Explaining this phenomenon requires further study of the geometry of constrained monotone submodular minimization.

Acknowledgements

We thank the anonymous reviewers for their helpful suggestions. We also thank MIT Supercloud and the Lincoln Laboratory Supercomputing Center for providing computational resources. This research was conducted with Government support under and awarded by DoD, Air Force Office of Scientific Research, National Defense Science and Engineering Graduate (NDSEG) Fellowship, 32 CFR 168a, and also supported by NSF CAREER award 1553284.

References

- Yahoo! Webscope dataset ydata-ysm-advertiser-bids-v1.0. URL http://research.yahoo.com/Academic_Relations.
- Adamczyk, Marek, Sviridenko, Maxim, and Ward, Justin. Submodular Stochastic Probing on Matroids. *Mathematics of Operations Research*, 41(3):1022–1038, 2016.
- Alon, Noga, Gamzu, Iftah, and Tennenholtz, Moshe. Optimizing Budget Allocation Among Channels and Influencers. In *WWW*, 2012.
- Atamtürk, Alper and Narayanan, Vishnu. Polymatroids and mean-risk minimization in discrete optimization. *Operations Research Letters*, 36(5):618–622, 2008.
- Bach, Francis. Submodular Functions: From Discrete to Continuous Domains. *arXiv:1511.00394*, 2015.
- Balkanski, Eric, Rubinstein, Aviad, and Singer, Yaron. The power of optimization from samples. In *NIPS*, 2016.
- Balkanski, Eric, Rubinstein, Aviad, and Singer, Yaron. The limitations of optimization from samples. In *STOC*, 2017.
- Becker, Stephen R., Candès, Emmanuel J., and Grant, Michael C. Templates for convex cone problems with applications to sparse signal recovery. *Mathematical programming computation*, 3(3):165–218, 2011.
- Ben-Tal, Aharon and Nemirovski, Arkadi. Robust solutions of Linear Programming problems contaminated with uncertain data. *Mathematical Programming*, 88(3):411–424, 2000.
- Ben-Tal, Aharon, El Ghaoui, Laurent, and Nemirovski, Arkadi. *Robust Optimization*. Princeton University Press, 2009.
- Bertsimas, Dimitris and Sim, Melvyn. Robust discrete optimization and network flows. *Mathematical programming*, 98(1):49–71, 2003.
- Bertsimas, Dimitris, Brown, David B., and Caramanis, Constantine. Theory and Applications of Robust Optimization. *SIAM Review*, 53(3):464–501, 2011.
- Best, Michael J. and Chakravarti, Nilotpal. Active set algorithms for isotonic regression; A unifying framework. *Mathematical Programming*, 47(1-3):425–439, 1990.
- Bian, Andrew An, Mirzasoleiman, Baharan, Buhmann, Joachim M., and Krause, Andreas. Guaranteed Non-convex Optimization: Submodular Maximization over Continuous Domains. In *AISTATS*, 2017.
- Birkhoff, Garrett. Rings of sets. *Duke Mathematical Journal*, 3(3):443–454, 1937.
- Borgs, Christian, Brautbar, Michael, Chayes, Jennifer, and Lucier, Brendan. Maximizing Social Influence in Nearly Optimal Time. In *SODA*, 2014.
- Boyd, Stephen, Kim, Seung-Jean, Vandenberghe, Lieven, and Hassibi, Arash. A tutorial on geometric programming. *Optimization and engineering*, 8(1):67–127, 2007.
- Chakrabarty, Deeparnab, Lee, Yin Tat, Sidford, Aaron, and Wong, Sam Chiu-wai. Subquadratic submodular function minimization. In *STOC*, 2017.
- Chandrasekaran, Venkat and Shah, Parikshit. Relative Entropy Relaxations for Signomial Optimization. *SIAM Journal on Optimization*, 26(2):1147–1173, 2016.
- Chen, Wei, Wang, Yajun, and Yang, Siyu. Efficient influence maximization in social networks. In *KDD*, 2009.
- Chen, Wei, Wang, Chi, and Wang, Yajun. Scalable Influence Maximization for Prevalent Viral Marketing in Large-scale Social Networks. In *KDD*, 2010.
- Chen, Wei, Lin, Tian, Tan, Zihan, Zhao, Mingfei, and Zhou, Xuren. Robust influence maximization. In *KDD*, 2016.
- Chiang, Mung. Geometric Programming for Communication Systems. *Commun. Inf. Theory*, 2(1/2):1–154, 2005.
- Deshpande, Amol, Hellerstein, Lisa, and Kletenik, Deborah. Approximation Algorithms for Stochastic Submodular Set Cover with Applications to Boolean Function Evaluation and Min-Knapsack. *ACM Trans. Algorithms*, 12(3):42:1–42:28, 2016.
- Domingos, Pedro and Richardson, Matt. Mining the network value of customers. In *KDD*, 2001.
- Du, Nan, Song, Le, Gomez Rodriguez, Manuel, and Zha, Hongyuan. Scalable influence estimation in continuous-time diffusion networks. In *NIPS*, 2013.
- Du, Nan, Liang, Yingyu, Balcan, Maria-Florina, and Song, Le. Influence function learning in information diffusion networks. In *ICML*, 2014.
- Dunn, Joseph C. and Harshbarger, S. Conditional gradient algorithms with open loop step size rules. *Journal of Mathematical Analysis and Applications*, 62(2):432–444, 1978.
- Ecker, Joseph. Geometric Programming: Methods, Computations and Applications. *SIAM Review*, 22(3):338–362, 1980.
- Ene, Alina and Nguyen, Huy L. A Reduction for Optimizing Lattice Submodular Functions with Diminishing Returns. *arXiv:1606.08362*, 2016.
- Frank, Marguerite and Wolfe, Philip. An algorithm for quadratic programming. *Naval Research Logistics Quarterly*, 3(1-2):95–110, 1956.
- Goel, Gagan, Karande, Chinmay, Tripathi, Pushkar, and Wang, Lei. Approximability of combinatorial problems with multi-agent submodular cost functions. In *FOCS*, 2009.
- Goemans, Michel and Vondrák, Jan. Stochastic Covering and Adaptivity. In *LATIN 2006: Theoretical Informatics*. Springer Berlin Heidelberg, 2006.
- Golovin, Daniel and Krause, Andreas. Adaptive Submodularity: Theory and Applications in Active Learning and Stochastic Optimization. *Journal of Artificial Intelligence*, 42:427–486, 2011.

- Gomez Rodriguez, Manuel and Schölkopf, Bernhard. Influence maximization in continuous time diffusion networks. In *ICML*, 2012.
- Gomez Rodriguez, Manuel, Leskovec, Jure, and Krause, Andreas. Inferring networks of diffusion and influence. In *KDD*, 2010.
- Gottschalk, Corinna and Peis, Britta. Submodular function maximization on the bounded integer lattice. In *Approximation and Online Algorithms: 13th International Workshop (WAOA)*, 2015.
- Hassidim, Avinatan and Singer, Yaron. Submodular optimization under noise. *arXiv preprint arXiv:1601.03095*, 2016.
- Hatano, Daisuke, Fukunaga, Takuro, Maehara, Takanori, and Kawarabayashi, Ken-ichi. Lagrangian Decomposition Algorithm for Allocating Marketing Channels. In *AAAI*, 2015.
- He, Xinran and Kempe, David. Robust influence maximization. In *KDD*, 2016.
- Iwata, Satoru and Nagano, Kiyohito. Submodular function minimization under covering constraints. In *FOCS*, 2009.
- Jaggi, Martin. Revisiting Frank-Wolfe: Projection-Free Sparse Convex Optimization. In *ICML*, 2013.
- Kempe, David, Kleinberg, Jon, and Tardos, Éva. Maximizing the Spread of Influence Through a Social Network. In *KDD*, 2003.
- Khachaturov, Vladimir R., Khachaturov, Roman V., and Khachaturov, Ruben V. Supermodular programming on finite lattices. *Computational Mathematics and Mathematical Physics*, 52(6):855–878, 2012.
- Kolmogorov, Vladimir and Shioura, Akiyoshi. New algorithms for convex cost tension problem with application to computer vision. *Discrete Optimization*, 6:378–393, 2009.
- Krause, Andreas, McMahan, H Brendan, Guestrin, Carlos, and Gupta, Anupam. Robust submodular observation selection. *Journal of Machine Learning Research*, 9(Dec): 2761–2801, 2008.
- Lacoste-Julien, Simon. Convergence Rate of Frank-Wolfe for Non-Convex Objectives. *arXiv:1607.00345*, 2016.
- Lacoste-Julien, Simon and Jaggi, Martin. On the global linear convergence of Frank-Wolfe optimization variants. In *NIPS*, 2015.
- Lee, Yin Tat, Sidford, Aaron, and Wong, Sam Chiu-wai. A faster cutting plane method and its implications for combinatorial and convex optimization. In *FOCS*, 2015.
- Lowalekar, Meghna, Varakantham, Pradeep, and Kumar, Akshat. Robust Influence Maximization: (Extended Abstract). In *AAMAS*, 2016.
- Maehara, Takanori. Risk averse submodular utility maximization. *Operations Research Letters*, 43(5):526–529, 2015.
- Maehara, Takanori, Yabe, Akihiro, and Kawarabayashi, Ken-ichi. Budget Allocation Problem with Multiple Advertisers: A Game Theoretic View. In *ICML*, 2015.
- MOSEK ApS. *MOSEK MATLAB Toolbox 8.0.0.57*, 2015. URL <http://docs.mosek.com/8.0/toolbox/index.html>.
- Murota, Kazuo. *Discrete convex analysis*. SIAM, 2003.
- Murota, Kazuo and Shioura, Akiyoshi. Exact bounds for steepest descent algorithms of l -convex function minimization. *Operations Research Letters*, 42:361–366, 2014.
- Nagano, Kiyohito, Kawahara, Yoshinobu, and Aihara, Kazuyuki. Size-constrained submodular minimization through minimum norm base. In *ICML*, 2011.
- Narasimhan, Harikrishna, Parkes, David C, and Singer, Yaron. Learnability of influence in networks. In *NIPS*, 2015.
- Netrapalli, Praneeth and Sanghavi, Sujay. Learning the graph of epidemic cascades. In *SIGMETRICS*, 2012.
- Nikolova, Evdokia. Approximation algorithms for reliable stochastic combinatorial optimization. In *APPROX*, 2010.
- Orlin, James B., Schulz, Andreas, and Udvani, Rajan. Robust monotone submodular function maximization. In *IPCO*, 2016.
- Pascual, Luis D. and Ben-Israel, Adi. Constrained maximization of posynomials by geometric programming. *Journal of Optimization Theory and Applications*, 5(2): 73–80, 1970.
- Polyak, Boris T. *Introduction to Optimization*. Number 04; QA402. 5, P6. 1987.
- Rockafellar, R Tyrrell and Uryasev, Stanislav. Optimization of conditional value-at-risk. *Journal of risk*, 2:21–42, 2000.
- Rockafellar, R Tyrrell and Uryasev, Stanislav. Conditional value-at-risk for general loss distributions. *Journal of banking & finance*, 26(7):1443–1471, 2002.
- Soma, Tasuku and Yoshida, Yuichi. A Generalization of Submodular Cover via the Diminishing Return Property on the Integer Lattice. In *NIPS*, 2015.
- Soma, Tasuku, Kakimura, Naonori, Inaba, Kazuhiro, and Kawarabayashi, Ken-ichi. Optimal Budget Allocation: Theoretical Guarantee and Efficient Algorithm. In *ICML*, 2014.
- Svitkina, Zoya and Fleischer, Lisa. Submodular approximation: Sampling-based algorithms and lower bounds. *SIAM Journal on Computing*, 40(6):1715–1737, 2011.
- Topkis, Donald M. Minimizing a submodular function on a lattice. *Operations research*, 26(2):305–321, 1978.
- Wainwright, Kevin and Chiang, Alpha. *Fundamental Methods of Mathematical Economics*. McGraw-Hill Education, 2004.
- Zhang, Peng, Chen, Wei, Sun, Xiaoming, Wang, Yajun, and Zhang, Jialin. Minimizing seed set selection with

probabilistic coverage guarantee in a social network. In *KDD*, 2014.

A. Worst-Case Approximation Ratio versus True Worst-Case

Consider the function $f(x; \theta)$ defined on $\{0, 1\} \times \{0, 1\}$, with values given by:

$$f(x; 0) = \begin{cases} 1 & x = 0 \\ 0.6 & x = 1, \end{cases} \quad f(x; 1) = \begin{cases} 1 & x = 0 \\ 2 & x = 1. \end{cases} \quad (18)$$

We wish to choose x to maximize $f(x; \theta)$ robustly with respect to adversarial choices of θ . If θ were fixed, we could directly choose x_θ^* to maximize $f(x; \theta)$. In particular, $x_0^* = 0$ and $x_1^* = 1$. Of course, we want to deal with worst-case θ . One option is to maximize the worst-case approximation ratio:

$$\max_x \min_\theta \frac{f(x; \theta)}{f(x_\theta^*; \theta)}. \quad (19)$$

One can verify that the best x according to this criterion is $x = 1$, with worst-case approximation ratio 0.6 and worst-case function value 0.6. In this paper, we optimize the worst-case of the actual function value:

$$\max_x \min_\theta f(x; \theta). \quad (20)$$

This criterion will select $x = 0$, which has a worse worst-case approximation ratio of 0.5, but actually guarantees a function value of 1, significantly better than the 0.6 achieved by the other formulation of robustness.

B. DR-submodularity

B.1. Proof of Proposition 1.1

Proof. The function f can be reduced to a submodular set function $g : 2^V \rightarrow \mathbb{R}$ via (Ene & Nguyen, 2016), where $|V| = O(n \log k)$. The function g can be evaluated via mapping from 2^V to the domain of f , and then evaluating f , in time $O(n \log k \cdot EO)$. We can directly substitute these complexities into the runtime bound from (Lee et al., 2015). \square

B.2. DR-submodularity and L^{\natural} -convexity

A function is L^{\natural} -convex if it satisfies a discrete version of midpoint convexity, i.e. for all x, y it holds that

$$f(x) + f(y) \geq f\left(\left\lceil \frac{x+y}{2} \right\rceil\right) + f\left(\left\lfloor \frac{x+y}{2} \right\rfloor\right). \quad (21)$$

Remark B.1. An L^{\natural} -convex function need not be DR-submodular, and vice-versa. Hence algorithms for optimizing one type may not apply for the other.

Proof. Consider $f_1(x_1, x_2) = -x_1^2 - 2x_1x_2$ and $f_2(x_1, x_2) = x_1^2 + x_2^2$, both defined on $\{0, 1, 2\} \times \{0, 1, 2\}$. The function f_1 is DR-submodular but violates discrete midpoint convexity for the pair of points $(0, 0)$ and $(2, 2)$, while f_2 is L^{\natural} -convex but does not have diminishing returns in either dimension. \square

Intuitively-speaking, L^{\natural} -convex functions look like discretizations of convex functions. The continuous objective function $\mathcal{I}(x, y)$ we consider need not be convex, hence its discretization need not be L^{\natural} -convex, and we cannot use those tools. However, in some regimes (namely if each $y(s) \in \{0\} \cup [1, \infty)$), it happens that $\mathcal{I}(x, y)$ is DR-submodular in x .

C. Constrained Continuous Submodular Function Minimization

Define $\mathbb{R}_{\downarrow}^n$ to be the set of vectors ρ in \mathbb{R}^n which are monotone nonincreasing, i.e. $\rho(1) \geq \rho(2) \geq \dots \geq \rho(n)$. As in the main text, define $[k] = \{0, 1, \dots, k-1\}$. One of the key results from (Bach, 2015) is that an arbitrary submodular function $H(x)$ defined on $\prod_{i=1}^n [k_i]$ can be extended to a particular convex function $h_{\downarrow}(\rho)$ so that

$$\begin{aligned} \text{minimize} \quad & H(x) \\ \text{s.t.} \quad & x \in \prod_{i=1}^n [k_i] \end{aligned} \Leftrightarrow \begin{aligned} \text{minimize} \quad & h_{\downarrow}(\rho) \\ \text{s.t.} \quad & \rho \in \prod_{i=1}^n \mathbb{R}_{\downarrow}^{k_i-1}. \end{aligned} \quad (22)$$

Moreover, Theorem 4 from (Bach, 2015) states that, if a_{iy_i} are strictly convex functions for all $i = 1, \dots, n$ and each $y_i \in [k_i]$, then the two problems

$$\begin{aligned} & \text{minimize} && H(x) + \sum_{i=1}^n \sum_{y_i=1}^{x_i} a'_{iy_i}(\lambda) \\ & \text{s.t.} && x \in \prod_{i=1}^n [k_i]. \end{aligned} \quad (23)$$

and

$$\begin{aligned} & \text{minimize} && h_{\downarrow}(\rho) + \sum_{i=1}^n \sum_{x_i=1}^{k_i-1} a_{ix_i}[\rho_i(x_i)] \\ & \text{s.t.} && \rho \in \prod_{i=1}^n \mathbb{R}_{\downarrow}^{k_i-1} \end{aligned} \quad (24)$$

are equivalent. In particular, one recovers a solution to Problem (23) for any λ just as alluded to in Lemma 3.2: find ρ^* which solves Problem (24) and, for each component i , choose x_i to be the maximal value for which $\rho_i^*(x_i) \geq \lambda$.

C.1. Proof of Lemma 3.2

Proof. The discretized form of the regularizer R^δ is also separable and can be written $R^\delta(x) = \sum_{i=1}^n R_i^\delta(x)$. For each $i = 1, \dots, n$ and each $y_i \in [k_i]$ with $y_i \geq 1$, define $a_{iy_i}(t) = \frac{1}{2}t^2 \cdot [R_i^\delta(y_i) - R_i^\delta(y_i - 1)]$, so that $a'_{iy_i}(t) = t \cdot [R_i^\delta(y_i) - R_i^\delta(y_i - 1)]$. Since we assumed $R(x)$ is strictly increasing, the coefficient of t^2 in each $a_{iy_i}(t)$ is strictly positive, so that each $a_{iy_i}(t)$ is strictly convex. Then,

$$\lambda R_i^\delta(x_i) = \lambda \cdot \left[R_i^\delta(0) + \sum_{y_i=1}^{x_i} (R_i^\delta(y_i) - R_i^\delta(y_i - 1)) \right] \quad (25)$$

$$= \lambda R_i^\delta(0) + \sum_{y_i=1}^{x_i} a'_{iy_i}(\lambda), \quad (26)$$

so that the discretized version of the minimization problem can be written as

$$\begin{aligned} & \text{minimize} && H^\delta(x) + \lambda R^\delta(0) + \sum_{i=1}^n \sum_{y_i=1}^{x_i} a'_{iy_i}(\lambda) \\ & \text{s.t.} && x \in \prod_{i=1}^n [k_i]. \end{aligned} \quad (27)$$

Since the term $R^\delta(0)$ does not depend on the variable x , this minimization is equivalent to

$$\begin{aligned} & \text{minimize} && H^\delta(x) + \sum_{i=1}^n \sum_{y_i=1}^{x_i} a'_{iy_i}(\lambda) \\ & \text{s.t.} && x \in \prod_{i=1}^n [k_i]. \end{aligned} \quad (28)$$

This problem is in the precise form where we can apply the preceding equivalence result between Problems (23) and (24), so we are done. \square

C.2. Proof of Theorem 3.1

Proof. The general idea of this proof is to first show that the integer-valued point x_d^* which solves

$$x_d^* \in \underset{x \in \prod_{i=1}^n [k_i]: R^\delta(x) \leq B}{\operatorname{argmin}} H^\delta(x)$$

is also nearly a minimizer of the continuous version of the problem, due to the fineness of the discretization. Then, we show that the solutions traced out by $x(\lambda)$ get very close to x_d^* . These two results are simply combined via the triangle inequality.

C.2.1. CONTINUOUS AND DISCRETE PROBLEMS

We begin by proving that

$$H^\delta(x_d^*) \leq G\delta + \min_{x \in \mathcal{X}: R(x) \leq B} H(x). \quad (29)$$

Consider $x^* \in \operatorname{argmin}_{x \in \mathcal{X}: R(x) \leq B} H(x)$. If x^* corresponds to an integral point in the discretized domain, then $H(x^*) = H^\delta(x_d^*)$ and we are done. Else, x^* has at least one non-integral coordinate. By rounding coordinatewise, we can construct a set $X = \{x_1, \dots, x_m\} \subseteq \prod_{i=1}^n [k_i]$ so that $x^* \in \operatorname{conv}(\{A(x_1), \dots, A(x_m)\})$. By monotonicity, there must be some $x_i \in$

X with $R^\delta(x_i) \leq B$, i.e. $A(x_i)$ is feasible for the original continuous problem. By construction, since the discretization given by A is δ -fine, we must have $\|x^* - A(x_i)\|_\infty \leq \delta$. Applying the Lipschitz property of H and the optimality of x^* , we have

$$G\delta \geq H(A(x_i)) - H(x^*) = H^\delta(x_i) - H(x^*) \geq H^\delta(x_d^*) - H(x^*),$$

from which (29) follows.

C.2.2. DISCRETE AND PARAMETERIZED DISCRETE PROBLEMS

Define λ_- and λ_+ by

$$\begin{aligned} \lambda_- &\in \operatorname{argmin}_{\lambda \geq 0: R^\delta(x(\lambda)) \leq B} H^\delta(x(\lambda)) \quad \text{and} \\ \lambda_+ &\in \operatorname{argmax}_{\lambda \geq 0: R^\delta(x(\lambda)) \geq B} H^\delta(x(\lambda)). \end{aligned}$$

The next step in proving our suboptimality bound is to prove that

$$H^\delta(x(\lambda_+)) \leq H^\delta(x_d^*) \leq H^\delta(x(\lambda_-)), \quad (30)$$

from which it will follow that

$$H^\delta(x(\lambda_-)) \leq G\delta + H^\delta(x_d^*).$$

We begin by stating the min-max inequality, i.e. weak duality:

$$\min_{x \in \prod_{i=1}^n [k_i]: R^\delta(x) \leq B} H^\delta(x) = \min_{x \in \prod_{i=1}^n [k_i]} \max_{\lambda \geq 0} \{H^\delta(x) + \lambda(R^\delta(x) - B)\} \quad (31)$$

$$\geq \max_{\lambda \geq 0} \min_{x \in \prod_{i=1}^n [k_i]} \{H^\delta(x) + \lambda(R^\delta(x) - B)\} \quad (32)$$

$$= \max_{\lambda \geq 0} \{H^\delta(x(\lambda)) + \lambda(R^\delta(x(\lambda)) - B)\} \quad (33)$$

$$\geq \max_{\lambda \geq 0: R^\delta(x(\lambda)) \geq B} \{H^\delta(x(\lambda)) + \lambda(R^\delta(x(\lambda)) - B)\} \quad (34)$$

$$\geq \max_{\lambda \geq 0: R^\delta(x(\lambda)) \geq B} H^\delta(x(\lambda)) \quad (35)$$

$$= H^\delta(x(\lambda_+)). \quad (36)$$

We can also bound the optimal value of $H^\delta(x_d^*)$ from the other side:

$$H^\delta(x_d^*) = \min_{x \in \prod_{i=1}^n [k_i]: R^\delta(x) \leq B} H^\delta(x) \leq \min_{\lambda \geq 0: R^\delta(x(\lambda)) \leq B} H^\delta(x) = H^\delta(x(\lambda_-)) \quad (37)$$

because the set of $x(\lambda)$ parameterized by λ is a subset of the full set $\{x \in \prod_{i=1}^n [k_i] : R^\delta(x) \leq B\}$.

We have now bounded the optimal value of $H^\delta(x_d^*)$ on either side by optimization problems where we seek an optimal $\lambda \geq 0$ for the parameterization $x(\lambda)$:

$$H^\delta(x(\lambda_+)) \leq H^\delta(x_d^*) \leq H^\delta(x(\lambda_-)). \quad (38)$$

Recall that $x(\lambda)$ comes from thresholding the values of ρ^* by λ , and that we assume that the elements of ρ^* are unique. Hence, as we increase λ , the components of x decrease by 1 each time. Combining this with the strict monotonicity of R , we see that $\|x(\lambda_+) - x(\lambda_-)\|_\infty \leq 1$. By the Lipschitz properties of H^δ , it follows that $|H^\delta(x(\lambda_+)) - H^\delta(x(\lambda_-))| \leq G\delta$. Since $H^\delta(x_d^*)$ lies in the interval between $H^\delta(x(\lambda_+))$ and $H^\delta(x(\lambda_-))$, it follows that $|H^\delta(x_d^*) - H^\delta(x(\lambda_-))| \leq G\delta$. \square

C.3. Proof of Remark 3.1

Define $\lambda^* = \lambda_-$ as in the previous section, so that $x' = A(x(\lambda^*))$. The $x(\lambda_+)$ bound is a simple consequence from the above result that

$$H^\delta(x(\lambda_+)) \leq H^\delta(x_d^*) \leq H^\delta(x(\lambda_-)) = H(x').$$

As for the Lagrangian bound, since $x(\lambda^*)$ is a minimizer for the regularized function $H^\delta(x) + \lambda^*(R^\delta(x) - B)$, it follows that

$$H^\delta(x(\lambda^*)) + \lambda^*(R^\delta(x(\lambda^*)) - B) \leq H^\delta(x_d^*) + \lambda^*(R^\delta(x_d^*) - B). \quad (39)$$

Rearranging, and observing that $R^\delta(x_d^*) \leq B$ because x_d^* is feasible, it holds that

$$H(x') = H^\delta(x(\lambda^*)) \leq H^\delta(x_d^*) + \lambda^*(R^\delta(x_d^*) - R^\delta(x(\lambda^*))) \leq H^\delta(x_d^*) + \lambda^*(B - R(x')). \quad (40)$$

One can also combine either of these bounds with the result from the proof of Theorem 3.1 that $H^\delta(x_d^*) \leq G\delta + H(x^*)$ yielding e.g.

$$H(x') \leq G\delta + \lambda^*(B - R(x')) + H^\delta(x_d^*).$$

C.4. Solving the Optimization Problem

Now that we have proven equivalence results between the constrained problem we want to solve and the convex problem (24), we need to actually solve the convex problem. At the beginning of Section 5.2 in (Bach, 2015), it is stated that this surrogate problem can be optimized via the Frank-Wolfe method and its variants, but only the version of Problem (24) without the extra functions a_{ix_i} is elaborated upon. Here we detail how Frank-Wolfe algorithms can be used to solve the more general parametric regularized problem. Our aim is to spell out very clearly the applicability of Frank-Wolfe to this problem, for the ease of future practitioners.

Bach (2015) notes that by duality, Problem (24) is equivalent to:

$$\begin{aligned} \min_{\rho \in \prod_{i=1}^n \mathbb{R}_{\downarrow}^{k_i-1}} h_{\downarrow}(\rho) - H(0) + \sum_{i=1}^n \sum_{x_i=1}^{k_i-1} a_{ix_i}[\rho_i(x_i)] &= \min_{\rho \in \prod_{i=1}^n \mathbb{R}_{\downarrow}^{k_i-1}} \max_{w \in B(H)} \langle \rho, w \rangle + \sum_{i=1}^n \sum_{x_i=1}^{k_i-1} a_{ix_i}[\rho_i(x_i)] \\ &= \max_{w \in B(H)} \left\{ \min_{\rho \in \prod_{i=1}^n \mathbb{R}_{\downarrow}^{k_i-1}} \langle \rho, w \rangle + \sum_{i=1}^n \sum_{x_i=1}^{k_i-1} a_{ix_i}[\rho_i(x_i)] \right\} \\ &:= \max_{w \in B(H)} f(w). \end{aligned}$$

Here, the base polytope $B(H)$ happens to be the convex hull of all vectors w which could be output by the greedy algorithm in (Bach, 2015).

It is the dual problem, where we maximize over w , which is amenable to Frank-Wolfe. For Frank-Wolfe methods, we need two oracles: an oracle which, given w , returns $\nabla f(w)$; and an oracle which, given $\nabla f(w)$, produces a point s which solves the linear optimization problem $\max_{s \in B(H)} \langle s, \nabla f(w) \rangle$.

Per Bach (2015), an optimizer of the linear problem can be computed directly from the greedy algorithm. For the gradient oracle, recall that we can find a subgradient of $g(x) = \min_y h(x, y)$ at the point x_0 by finding $y(x_0)$ which is optimal for the inner problem, and then computing $\nabla_x h(x, y(x_0))$. Moreover, if such $y(x_0)$ is the unique optimizer, then the resulting vector is indeed the *gradient* of $g(x)$ at x_0 . Hence, in our case, it suffices to first find $\rho(w)$ which solves the inner problem, and then $\nabla f(w)$ is simply $\rho(w)$ because the inner function is linear in w . Since each function a_{ix_i} is strictly convex, the minimizer $\rho(w)$ is unique, confirming that we indeed get a gradient of f , and that f is differentiable.

Of course, we still need to compute the minimizer $\rho(w)$. For a given w , we want to solve

$$\min_{\rho \in \prod_{i=1}^n \mathbb{R}_{\downarrow}^{k_i-1}} \langle \rho, w \rangle + \sum_{i=1}^n \sum_{x_i=1}^{k_i-1} a_{ix_i}[\rho_i(x_i)]$$

There are no constraints coupling the vectors ρ_i , and the objective is similarly separable, so we can independently solve n problems of the form

$$\min_{\rho \in \mathbb{R}_{\downarrow}^{k-1}} \langle \rho, w \rangle + \sum_{j=1}^{k-1} a_j(\rho_j).$$

Recall that each function $a_{iy_i}(t)$ takes the form $\frac{1}{2}t^2 r_{iy_i}$ for some $r_{iy_i} > 0$. Let $D = \text{diag}(r)$, the $(k-1) \times (k-1)$ matrix with diagonal entries r_j . Our problem can then be written as

$$\begin{aligned} \min_{\rho \in \mathbb{R}_{\downarrow}^{k-1}} \langle \rho, w \rangle + \frac{1}{2} \sum_{j=1}^{k-1} r_j \rho_j^2 &= \min_{\rho \in \mathbb{R}_{\downarrow}^{k-1}} \langle \rho, w \rangle + \frac{1}{2} \langle D\rho, \rho \rangle \\ &= \min_{\rho \in \mathbb{R}_{\downarrow}^{k-1}} \langle D^{1/2}\rho, D^{-1/2}w \rangle + \frac{1}{2} \langle D^{1/2}\rho, D^{1/2}\rho \rangle. \end{aligned}$$

Completing the square, the above problem is equivalent to

$$\begin{aligned} \min_{\rho \in \mathbb{R}_{\downarrow}^{k-1}} \|D^{1/2}\rho + D^{-1/2}w\|_2^2 &= \min_{\rho \in \mathbb{R}_{\downarrow}^{k-1}} \sum_{j=1}^{k-1} (r_j^{1/2}\rho_j + r_j^{-1/2}w_j)^2 \\ &= \min_{\rho \in \mathbb{R}_{\downarrow}^{k-1}} \sum_{j=1}^{k-1} r_j(\rho_j + r_j^{-1}w_j)^2. \end{aligned}$$

This last expression is precisely the problem which is called weighted isotonic regression: we are fitting ρ to $\text{diag}(r^{-1})w$, with weights r , subject to a monotonicity constraint. Weighted isotonic regression is solved efficiently via the Pool Adjacent Violators algorithm of (Best & Chakravarti, 1990).

C.5. Runtime

Frank-Wolfe returns an ε -suboptimal solution in $O(\varepsilon^{-1}D^2L)$ iterations, where D is the diameter of the feasible region, and L is the Lipschitz constant for the gradient of the objective (Jaggi, 2013). Our optimization problem is $\max_{w \in B(H)} f(w)$ as defined in the previous section. Each $w \in B(H)$ has $O(n\delta^{-1})$ coordinates of the form $H^\delta(x + e_i) - H^\delta(x)$. Since H^δ is an expected influence in the range $[0, T]$, we can bound the magnitude of each coordinate of w by T and hence D^2 by $O(n\delta^{-1}T^2)$. If α is the minimum derivative of the functions R_i , then the smallest coefficient of the functions $a_{ix_i}(t)$ is bounded below by $\alpha\delta$. Hence the objective is the conjugate of an $\alpha\delta$ -strongly convex function, and therefore has $\alpha^{-1}\delta^{-1}$ -Lipschitz gradient. Combining these, we arrive at the $O(\varepsilon^{-1}n\delta^{-2}\alpha^{-1}T^2)$ iteration bound. The most expensive step in each iteration is computing the subgradient, which requires sorting the $O(n\delta^{-1})$ elements of ρ in time $O(n\delta^{-1} \log n\delta^{-1})$. Hence the total runtime of Frank-Wolfe is $O(\varepsilon^{-1}n^2\delta^{-3}\alpha^{-1}T^2 \log n\delta^{-1})$.

As specified in the main text, relating an approximate solution of (15) to a solution of (14) is nontrivial. Assume ρ^* has distinct elements separated by η , and chose ε to be less than $\eta^2\alpha\delta/8$. If ρ is ε -suboptimal, then by $\alpha\delta$ -strong convexity we must have $\|\rho - \rho^*\|_2 < \eta/2$, and therefore $\|\rho - \rho^*\|_\infty < \eta/2$. Since the smallest consecutive gap between elements of ρ^* is η , this implies that ρ and ρ^* have the same ordering, and therefore admit the same solution x after thresholding. Accounting for this choice in ε , we have an exact solution to (14) in total runtime of $O(\eta^{-2}n^2\delta^{-4}\alpha^{-2}T^2 \log n\delta^{-1})$.

D. Expectation and Variance of the Influence Function

We wish to study the influence $\mathcal{I}(X, y)$, its expectation and its variance as a function of y . By definition, the influence function is given by

$$\mathcal{I}(X, y) = \sum_{t \in T} \left(1 - \prod_{(s,t) \in E} X_{st}^{y(s)} \right). \quad (41)$$

Before we prove the stated results, we will simplify the functions involved.

Maximizing $\mathcal{I}(X, y)$ is equivalent to minimizing the function

$$\sum_{t \in T} \prod_{(s,t) \in E} X_{st}^{y(s)} \quad (42)$$

and vice-versa. The particular properties we are interested in, namely convexity and submodularity, are preserved under sums. Moreover, expectation is linear and variances add, so for our purposes we can focus on only one term of the above

sum. After reindexing in terms of $i = 1, \dots, n$ instead of $(s, t) \in E$, we are left studying functions of the form

$$f(y) = \prod_{i=1}^n X_i^{y_i}. \quad (43)$$

If $f(y)$ is always convex (or supermodular), then $\mathcal{I}(X, y)$ is always concave (submodular) in y , and similarly for their expectations and variances.

Expectation By independence,

$$\mathbb{E}[f(y)] = \prod_{i=1}^n \mathbb{E}[X_i^{y_i}]. \quad (44)$$

Suppose that each $X_i \sim \text{Beta}(\alpha_i, \beta_i)$, so that

$$\mathbb{E}[X_i^{y_i}] = \frac{\Gamma(\alpha_i + \beta_i)\Gamma(\alpha_i + y_i)}{\Gamma(\alpha_i + \beta_i + y_i)} \quad (45)$$

$$= \frac{\Gamma(\alpha_i + \beta_i)\Gamma(\alpha_i + y_i)\Gamma(\beta_i)}{\Gamma(\alpha_i + \beta_i + y_i)\Gamma(\beta_i)} \quad (46)$$

$$= \frac{B(\alpha_i + y_i, \beta_i)}{B(\alpha_i, \beta_i)}. \quad (47)$$

Then,

$$\mathbb{E}[f(y)] = \prod_{i=1}^n \frac{B(\alpha_i + y_i, \beta_i)}{B(\alpha_i, \beta_i)} \propto \prod_{i=1}^n B(\alpha_i + y_i, \beta_i), \quad (48)$$

where by \propto we mean that the product of the denominators is a positive constant, dependent on the problem data but independent of y .

Variance The variance of $f(y)$ can be written as

$$\text{Var} \left[\prod_{i=1}^n X_i^{y_i} \right] = \mathbb{E} \left[\prod_{i=1}^n X_i^{2y_i} \right] - \mathbb{E} \left[\prod_{i=1}^n X_i^{y_i} \right]^2 \quad (49)$$

$$= \prod_{i=1}^n \mathbb{E} \left[X_i^{2y_i} \right] - \prod_{i=1}^n \mathbb{E} \left[X_i^{y_i} \right]^2 \quad (50)$$

$$= \prod_{i=1}^n \frac{B(\alpha_i + 2y_i, \beta_i)}{B(\alpha_i, \beta_i)} - \prod_{i=1}^n \left(\frac{B(\alpha_i + y_i, \beta_i)}{B(\alpha_i, \beta_i)} \right)^2. \quad (51)$$

D.1. Gradient of Expected Influence

Recall the identity

$$\frac{\partial}{\partial a} B(a, b) = B(a, b)(\psi(a) - \psi(a + b)), \quad (52)$$

where ψ is the digamma function. We can then compute each component of the gradient of $\mathbb{E}[f(y)]$:

$$\frac{\partial}{\partial y_i} (\mathbb{E}[f(y)]) = \prod_{i=1}^n \frac{1}{B(\alpha_i, \beta_i)} \cdot \prod_{j \neq i} B(\alpha_j + y_j, \beta_j) \cdot \frac{\partial}{\partial y_i} (B(\alpha_i + y_i, \beta_i)) \quad (53)$$

$$= \prod_{i=1}^n \frac{1}{B(\alpha_i, \beta_i)} \cdot \prod_{j \neq i} B(\alpha_j + y_j, \beta_j) \cdot B(\alpha_i + y_i, \beta_i) \cdot (\psi(\alpha_i + y_i) - \psi(\alpha_i + y_i + \beta_i)) \quad (54)$$

$$= \mathbb{E}[f(y)] \cdot (\psi(\alpha_i + y_i) - \psi(\alpha_i + y_i + \beta_i)). \quad (55)$$

D.2. Counterexample for Fact 2.1

We give a specific choice of parameters n, α_i, β_i and y_i for which the resulting function $\sqrt{\text{Var}(f(y))}$ is non-convex, non-concave, non-submodular and non-supermodular for various points $y \in \mathbb{R}_+^n$. For the case $T = 1$, the function $1 - f(y)$ is a valid influence function, so we have a valid counterexample for $\sqrt{\text{Var}(\mathcal{I}(X, y))}$.

Consider the case $n = 2$, with $\alpha_1 = \alpha_2 = 1$ and $\beta_1 = \beta_2 = 1$. This corresponds to the Budget Allocation problem where we have two sources each with an edge to one customer, and we have only our prior (i.e. no data) on either of the edge probabilities. Using equation (51), we can directly compute the Hessian of $\sqrt{\text{Var}(f(y))}$ at any point y , e.g. using Mathematica. In particular, for $y_1 = y_2 = 1$, the Hessian has a positive and a negative eigenvalue, so $\sqrt{\text{Var}(f(y))}$ is neither convex nor concave at this point. Also for $y_1 = y_2 = 1$, the off-diagonal element is negative, so $\sqrt{\text{Var}(f(y))}$ is not supermodular over all of \mathbb{R}_+^2 . However, for $y_1 = y_2 = 3$, the off-diagonal element is positive, so our function is also not submodular.