Prepare for the Expected Worst: Algorithms for Reconfigurable Resources Under Uncertainty

D Ellis Hershkowitz*[1], R. Ravi†[1], and Sahil Singla‡[2]

[1]Carnegie Mellon University, Pittsburgh PA 15213, USA
[2]Princeton University, Princeton NJ 08544, USA

November 28, 2018

Abstract

In this paper we study how to optimally balance cheap inflexible resources with more expensive, reconfigurable resources despite uncertainty in the input problem. Specifically, we introduce the MinEMax model to study "build versus rent" problems. In our model different scenarios appear independently. Before knowing which scenarios appear, we may build rigid resources that cannot be changed for different scenarios. Once we know which scenarios appear, we are allowed to rent reconfigurable but expensive resources to use across scenarios. While the objective in our model is difficult to compute, we show it is well-estimated by a surrogate objective which is representable by an LP. In this surrogate objective we pay for each scenario only to the extent that it exceeds a certain threshold. Using this objective we design algorithms that approximately-optimally balance inflexible and reconfigurable resources for several NP-hard covering problems. For example, we study minimum spanning and Steiner trees, minimum cuts and facility location variants. Up to constants our approximation guarantees match those of previous algorithms for the previously-studied demand-robust and stochastic two-stage models. Lastly, we demonstrate that our problem is sufficiently general to smoothly interpolate between previous demand-robust and stochastic two-stage problems.

*Supported in part by NSF grants CCF-1527110, CCF-1618280 and NSF CAREER award CCF-1750808.

†Supported in part by the U. S. Office of Naval Research award N00014-18-1-2099, and the U. S. NSF award CCF-1527032.

‡Much of this work was done while at Carnegie Mellon University. Supported in part by Schmidt Foundation and NSF awards CCF-1319811, CCF-1536002, and CCF-1617790.

1 Introduction

Optimizing for reconfigurable resources under uncertainty formalizes the challenges of balancing expensive, flexible resources with cheap, inflexible ones. For example, such optimization problems formalize the challenges in "build versus rent" problems. Concretely, consider the algorithmic challenges faced by an Internet service provider (ISP). An ISP must provide content to its customers while balancing between rigid and reconfigurable resources. In particular, it can build out its own network—a rigid resource—or choose to support traffic on a competitor's network—a flexible resource—at a marked up premium. This latter resource is reconfigurable since an ISP can change which edges in a competitor's network it uses at any given time. To minimize the additional load on its network, the competitor charges the ISP for the maximum extra bandwidth it must support at any given moment. Furthermore, an ISP only has probabilistic knowledge of where customer demands will occur: Based on where previous demands have occurred an ISP estimates future demands, but it does not exactly know the future demands. If a demand occurs which the ISP's network cannot service, it must use the competitor's network to support it. Thus, an ISP balances between rigid and flexible resources in the face of uncertainty, and pays for the cost of its own network plus the cost of supporting the *expected maximum* traffic routed on its competitor's network.

In this paper, we introduce the MinEMax model to study the algorithmic challenges associated with optimizing reconfigurable resources under uncertainty. In our model we are given a set of *scenarios* that might occur. In the preceding example these scenarios were the sets of possible demands. We think of problems in our model as being divided between a first stage where we "build" rigid resources and a second stage where we "rent" flexible resources. In particular, in the first stage we can build non-reconfigurable resources without knowing which scenarios occur. In the second stage, each scenario independently realizes according to its specified Bernoulli probability, and we can rent reconfigurable resources at an increased cost to use among any of our scenarios. For instance, in the preceding example the ISP first built its own network and then, once it learned where demands occurred, it could rent bandwidth to support different demands over time. In fact, this example is exactly our MinEMax Steiner tree problem. Thus, the objective we minimize is the first stage cost plus the *expected maximum* cost of additional reconfigurable resources required for any realized scenario; hence the name of our model.

Since every scenario is an independent Bernoulli, there are exponentially-many ways in which scenarios realize. It is not even clear, then, how to compute the expected second-stage cost. Nonetheless, we provide techniques to simplify and reason about the MinEMax cost and, therefore, solve various MinEMax problems.

The primary contributions of our work are as follows.

(i) We introduce the MinEMax model for optimization of reconfigurable resources under uncertainty.

(ii) We show that, although evaluating the MinEMax objective function is difficult, a MinEMax problem can be approximately reduced to a "TruncatedTwoStage" problem whose objective is representable by an LP.

(iii) Armed with (ii), we adapt various rounding techniques to give approximation algorithms for a variety of two-stage MinEMax problems including spanning and Steiner trees, cuts and facility location problems.

(iv) Lastly, we show that the MinEMax model captures the commonly studied two-stage models for optimization under uncertainty: the stochastic and demand-robust models. Indeed, we show that it generalizes a "Hybrid" problem that smoothly interpolates between the stochastic and demand-robust models.

1.1 Related Work

Significant prior work has been done in two-stage optimization under uncertainty. The two most commonly studied models are the stochastic model [RS04, GPRS04, SS06] and the demand-robust model [DGRS05, AGGN08, GNR10, GGP+15a]. In the **stochastic two-stage model** a probability distribution is given over scenarios, and our objective is to minimize the *expected* total cost. In the **demand-robust** two-stage model we always pay for the *worst-case* scenario given our first stage solution.

Distributionally robust optimization (DRO) [Sca58, GS10, DY10, BBC11] captures a problem similar to our own. In the DRO model we are given a distribution along with a ball of "nearby" distributions and must pay the *worst-case expectation* over all distributions. Similarly to our own model, DRO generalizes both the stochastic and demand-robust two-stage models. Our model can be seen as dual to the DRO model: while the DRO model takes the worst-case over distributions our model takes a distribution over worst cases. While our model is also sufficiently general to capture stochastic and demand-robust optimization, unlike DRO it is not so difficult as to preclude approximation algorithms. That is, we give approximation algorithms for discrete MinEMax problems and to the best of our knowledge no such results exist for DRO.

Several additional models for optimization under uncertainty—some of which even interpolate between stochastic and demand-robust—have also been studied. A series of papers [Sri07, SZY09, Swa11] examined various models of two-stage optimization to capture risk-aversion. Notably, the model of [Swa11] interpolates between stochastic and demand-robust while also accommodating *black-box* distributions. Other papers [GPRS04] have also studied algorithms for stochastic optimization given access to black-box distributions. There has also been work on two-stage stochastic models where each scenario can arise independently of the others [IKMM04], as in our own model.

Lastly, our problems bear some similarities to *centrum* problems [Tam01, CS18] whose objectives interpolate between a sum and a max. Similarly, previous work in scheduling [BP03, KMPS05] in which the L_p norm of the vector of completion times for $1 < p < \infty$ has been proposed to trade-off between the sum of completion times objective $p = 1$ and the makespan objective $p = \infty$ is reminiscent of our own.

1.2 Models

We now formally define our new MinEMax model and the prior models that we generalize. We study two-stage covering problems, defined as follows.

Two-Stage Covering. Let U be the universe of *clients* (or demand requirements), and let X be the set of *elements* that we can purchase. Every scenario $S_1, S_2, \ldots, S_m$ is a subset of clients. Let $sol(S_s)$ for $s \in [m]$ denote the sets in 2^X which are feasible to *cover* scenario S_s. In covering problems if $A \subseteq B$ and $A \in sol(S_s)$, then $B \in sol(S_s)$. We are also given a cost function $\text{cost} : 2^X \times 2^X \to \mathbb{R}$. For a given a specification of cost, scenarios, clients, and feasibility constraints, we must find a set of elements $X_1 \subseteq X$ to be bought in the first stage, and a set of elements $X_2^{(s)} \subseteq X$ to be bought in the second stage s.t. $X_1 \cup X_2^{(s)} \in sol(S_s)$ for every s. Our goal is to find a solution of minimal cost where the cost of a solution is different for different two-stage covering models.

This paper makes the common assumption that cost is *linear*, i.e., $\text{cost}(X_1, X_2^{(s)})$ equals $\text{cost}(\emptyset, X_2^{(s)}) + \text{cost}(X_1, \emptyset)$ for any $X_1, X_2^{(s)} \subseteq 2^X$. Let $\boldsymbol{X_2} := (X_2^{(1)}, \ldots, X_2^{(m)})$; throughout the paper a bold variable denotes a vector.

We now describe and discuss how different cost functions yield different two-stage covering models.

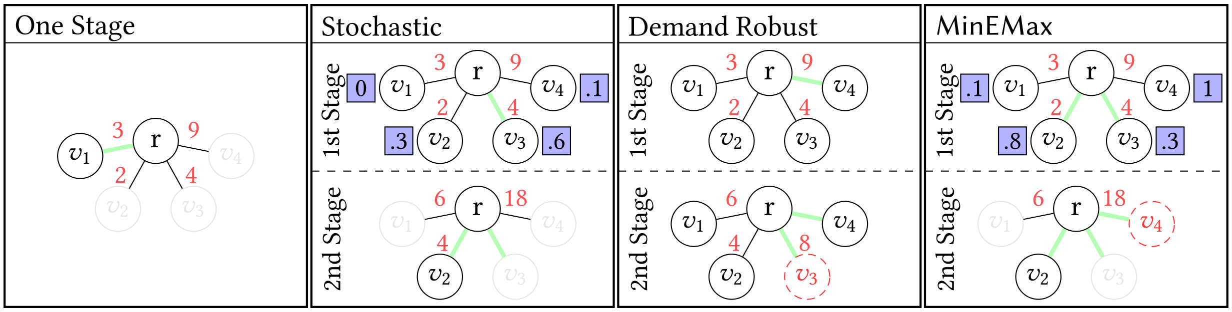

(a) If the scenario to be covered is known to be v_1, the problem is trivial.

(b) Exactly one scenario realizes according to a probability distribution.

(c) Given a first stage solution, adversary chooses the costliest scenario.

(d) Given a first stage solution, adversary chooses the costliest realized scenario.

Figure 1: Star graph MinEMax for $m = 4$. Green edges: edges bought by solution. e_i labeled by its cost in each stage for $\sigma = 2$. Non-opaque second-stage node: realized scenario. Blue square: probability of scenario. Dashed red nodes: nodes chosen by an adversary.

Prior Models. In the *demand-robust* two-stage covering model the cost of solution $(X_1, \boldsymbol{X_2})$ is the maximum cost over all the scenarios:

$$\text{cost}_{\text{Rob}}(X_1, \boldsymbol{X_2}) := \max_{s \in [m]} \left\{ \text{cost}(X_1, X_2^{(s)}) \right\}. \tag{1}$$

In the *stochastic* two-stage covering model we are given a probability distribution $\mathcal{D}$ over m scenarios with which exactly one of them realizes; i.e. $\sum_{s \in [m]} \mathcal{D}(s) = 1$. The cost of solution $(X_1, \boldsymbol{X_2})$ is the expected cost:

$$\text{cost}_{\text{Stoch}}(X_1, \boldsymbol{X_2}) := \mathbb{E}_{s \sim \mathcal{D}}[\text{cost}(X_1, X_2^{(s)})]. \tag{2}$$

Our New MinEMax Model. In the MinEMax two-stage covering model we are given probabilities $\mathbf{p} = \{p_1, \ldots, p_m\}$ with which each scenario *independently* realizes. The cost of solution $(X_1, \boldsymbol{X_2})$ is the expected maximum cost among the realized scenarios:

$$\text{cost}_{\text{EMax}}(X_1, \boldsymbol{X_2}) := \mathbb{E}_{A \sim \mathbf{p}} \left[\max_{s \in A} \left\{ \text{cost}(X_1, X_2^{(s)}) \right\} \right] \tag{3}$$

where A contains each s independently w.p. p_s. To avoid confusion, we reiterate that unlike the stochastic model, in MinEMax multiple scenarios may simultaneously appear in A because each of them independently realizes.

As a concrete example these models, consider the following star covering problem. We are given a star graph with root r and leaves $v_1, \ldots, v_m$. Each edge $e_i = (r, v_i)$ can be purchased in the first stage at cost c_i and in the second stage at an inflated cost $\sigma \cdot c_s$ for $\sigma > 1$. Our goal is to connect r to an *unknown* vertex v_s with minimum total two-stage cost. In particular, v_s is only revealed after we purchase our first-stage edges, X_1, at which point we must purchase e_s in a second stage at cost $\sigma \cdot c_s$ if e_s was not already purchased in the first stage. In all three models we initially buy some set of edges. In the stochastic version of this problem a single v_s then appears according to a distribution and we must pay to connect v_s if we have not already. In the demand-robust version of this problem, v_s is always chosen so as to maximize our second stage cost. However, in our MinEMax version of this problem several v_s appear and we must pay for a budget of reconfigurable edge resource to be reused for every v_s. See Figure 1 for an illustration.

1.3 Technical Results and Intuition

We now discuss our technical results. As earlier noted, capturing the MinEMax objective is challenging: scenarios may realize in exponentially-many ways and even computing the objective seems computationally infeasible. We solve this issue by showing that to solve a MinEMax problem, P_{EMax}, it suffices to solve its TruncatedTwoStage version, P_{Trunc}. A TruncatedTwoStage problem is identical to a MinEMax problem but the cost of a solution $(X_1, \boldsymbol{X_2})$ is its truncated sum:

$$\mathrm{cost}_{\mathsf{Trunc}}(X_1, \boldsymbol{X_2}) := \min_B \left[B + \sum_{s \in [m]} p_s \cdot (\mathrm{cost}(X_1, X_2^{(s)}) - B)^+ \right]. \tag{4}$$

We will later see that P_{Trunc} can be represented by an LP and, therefore, can be efficiently approximated by various rounding techniques. The following theorem shows that to approximate a MinEMax problem, it suffices to consider its TruncatedTwoStage version.

Theorem 1.1. *Let P_{EMax} be a* MinEMax *problem and let P_{Trunc} be its corresponding* TruncatedTwoStage *problem. An α-approximation algorithm for P_{Trunc} is a $\left(\frac{2\alpha}{1-1/e}\right)$-approximation algorithm for P_{EMax}.*

Intuition. The main observation we use to show this theorem is that a set of expensive scenarios with large total probability mass dominates the cost of a given MinEMax solution. We illustrate this observation with an example. Let $(X_1, \boldsymbol{X_2})$ be a solution for a MinEMax problem. Now WLOG let $\mathrm{cost}(X_1, X_2^{(s)}) \geq \mathrm{cost}(X_1, X_2^{(s+1)})$ for all s, i.e., the sth scenario is more expensive than the $(s+1)$th scenario for our solution. Let $M := [k]$ be the indices of the first k scenarios such that $\sum_{s \leq k} p_s$ is large; say, at least 1. Let the border $B := \mathrm{cost}(X_1, X_2^{(k)})$ be the cost of the least expensive scenario with an index in M. Because there is a great deal of probability mass among scenarios in M we know that with large probability some scenario in M will always appear. Whenever a scenario of cost less than B appears we know that with good probability something in M has also appeared of greater cost. Thus, as far as the expected max is concerned, a scenario that costs less than B can be ignored. Lastly, while it is not immediately clear how to represent $\mathrm{cost}_{\mathsf{Trunc}}$ function in an LP, we show using a simple convexity argument how this can accomplished.

Next, we design approximation algorithms for two-stage covering problems in the MinEMax model.

Theorem 1.2. *For two-stage covering problems there exist polynomial time approximation algorithms with the following guarantees where bold results indicate results given in this paper.* [1]

MinEMax *Problem*	*UFL*	*Steiner tree*	*MST*	*Min-cut*	*k-center*
Approximation	$\boldsymbol{\frac{16}{1-1/e}}$	$\boldsymbol{\frac{60}{1-1/e}}$	$\boldsymbol{O(\log n + \log m)}$	$\boldsymbol{\frac{8}{1-1/e}}$	$\boldsymbol{O(1)}$

Intuition. Our earlier Theorem 1.1 demonstrated that to solve a MinEMax problem, P_{EMax}, we need to only solve its TruncatedTwoStage version, P_{Trunc}. While it is not clear how to represent P_{EMax} with an LP, P_{Trunc} can be represented with an LP. Furthermore, by adapting previous two-stage optimization rounding techniques to the TruncatedTwoStage setting, we are able to approximately solve the TruncatedTwoStage versions of uncapacitated facility location (UFL), Steiner tree, minimum spanning tree (MST), and min-cut.

We use different techniques to give an approximation algorithm for k-center. The intuition for our k-center proof is similar to that of Theorem 1.1: Truncated costs approximate MinEMax cost. However, for k-center we truncate more aggressively. Rather than truncating costs of scenarios, we truncate distances in the input metric. To do this, we draw on methods of Chakrabarty and Swamy [CS18].[2]

[1] The $O(1)$ in the k-center approximation is roughly 114.

[2] We also note here that, unlike the previous problems we study, the cost function in k-center is not linear as described in §1.2.

It is also worth noting that Anthony et al. [AGGN08] proved hardness of approximation for a two-stage k-center problem. In particular, they show stochastic k-center where scenarios consist of *multiple* clients is as hard to approximate as dense k-subgraph. Thus, since our MinEMax model generalizes the stochastic model, we restrict our attention in k-center to scenarios consisting of *single* clients; otherwise our problem would be prohibitively hard to approximate. Since our scenarios consist of single clients the stochastic and demand-robust versions of the k-center problem we solve correspond to k-median and k-center respectively.

Our last theorem shows that MinEMax generalizes the stochastic and demand-robust models as well as a Hybrid model which smoothly interpolates between stochastic and demand-robust optimization.

Theorem 1.3. *An α-approximation for a two-stage covering algorithm in the* MinEMax *model implies an α-approximation for the corresponding two-stage covering problem in the stochastic, demand-robust, and* Hybrid *models.*

For cleanliness of exposition, we defer a formal definition and discussion of the Hybrid model as well as the intuition and proof for Theorem 1.3 to Section 5. As a corollary of Theorems 1.2 and 1.3, we immediately recover polynomial time approximations for Hybrid MST, UFL, Steiner tree and min-cut.[3]

2 Reducing MinEMax to TruncatedTwoStage

In this section, we demonstrate a technique to simplify both computing and reasoning about $\text{cost}_{\mathsf{EMax}}$ by reducing a MinEMax problem to a TruncatedTwoStage problem with only a small loss in the approximation factor. Specifically, we show the following theorem.

Theorem 1.1. *Let P_{EMax} be a* MinEMax *problem and let P_{Trunc} be its corresponding* TruncatedTwoStage *problem. An α-approximation algorithm for P_{Trunc} is a $\left(\frac{2\alpha}{1-1/e}\right)$-approximation algorithm for P_{EMax}.*

As earlier noted, we show this by observing that a set of expensive scenarios with “large” total probability mass dominates the cost of a given MinEMax solution.

We begin by observing that the expected max of a set of independent random variables is approximately bounded by the most expensive of these random variables whose probabilities sum to 1. We defer the proof of this lemma to §A. Variants of this lemma have appeared before [ADSY12, Ala14].

Lemma 2.1. *Let $\boldsymbol{Y} = \{Y_1, \ldots, Y_m\}$ be a set of independent Bernoulli r.v.s, where Y_s is 1 with probability p_s, and 0 otherwise. Let $v_s \in \mathbb{R}_{\geq 0}$ be a value associated with Y_s. WLOG assume $v_s \geq v_{s+1}$ for $s \in [m-1]$. Let $b = \min_{a:\sum_{s=1}^{a} p_s \geq 1} a$. Then*

$$\left(\frac{1-1/e}{2}\right)\left(v_b + \sum_s p_s \cdot (v_s - v_b)^+\right) \leq \mathbb{E}_{\boldsymbol{Y}}\left[\max_s \{Y_s \cdot v_s\}\right] \leq v_b + \sum_s p_s \cdot (v_s - v_b)^+,$$

where $x^+ := \max\{x, 0\}$.

For a *given solution* $(X_1, \boldsymbol{X}_2)$ to MinEMax, Lemma 2.1 yields a computationally tractable form of $\text{cost}_{\mathsf{EMax}}$. Specifically, let our scenarios be indexed such that $\text{cost}(X_1, X_2^{(s)}) \geq \text{cost}(X_1, X_2^{(s+1)})$ and let b be the smallest

[3] Though not k-center since its cost function is not linear.

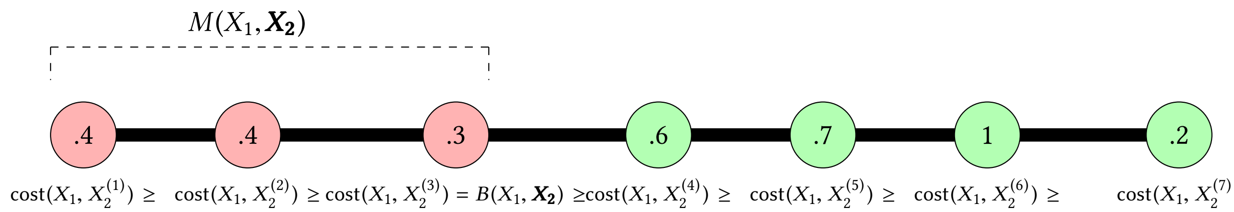

Figure 2: $M(X_1, \boldsymbol{X_2})$ and $B(X_1, \boldsymbol{X_2})$. Red circles: scenarios in $M(X_1, \boldsymbol{X_2})$. Green circles: all other scenarios. Numbers in circles: probabilities. Scenarios arranged left to right in descending order of $\text{cost}(X_1, X_2^{(s)})$.

positive integer such that $\sum_{s=1}^{b} p_s \geq 1$. We define the following terms analogous to those in the lemma (see Figure 2 for an illustration):

$$M(X_1, \boldsymbol{X}_2) := [b] \qquad \text{and} \qquad B(X_1, \boldsymbol{X}_2) := \min_{s \in M(X_1, \boldsymbol{X}_2)} \text{cost}(X_1, X_2^{(s)}). \tag{5}$$

Notice that $\sum_{s \in M(X_1, \boldsymbol{X_2})} p_s < 2$. Now, by letting $B(X_1, \boldsymbol{X_2})$ be v_b in Lemma 2.1, we can approximate $\text{cost}_{\text{EMax}}(X_1, \boldsymbol{X_2})$. However, we would like to estimate $\text{cost}_{\text{EMax}}(X_1, \boldsymbol{X}_2)$ within an LP where $(X_1, \boldsymbol{X}_2)$ are variables since our algorithms are LP based. Unfortunately, it is not clear how to capture v_b in an LP and so it is not clear how to directly use Lemma 2.1 to estimate $\text{cost}_{\text{EMax}}(X_1, \boldsymbol{X}_2)$ within an LP.

For this reason, we derive an even simpler form of the above approximation of the expected max which can be computed using an LP. In particular, we show that the expected max is the $\text{cost}_{\text{Trunc}}$ objective. We remind the reader that, as per Eq.(4), $\text{cost}_{\text{Trunc}}(X_1, \boldsymbol{X_2}) := \min_B [B + \sum_{s \in [m]} p_s \cdot (\text{cost}(X_1, X_2^{(s)}) - B)^+]$. The following lemma shows that the B achieving the minimum in $\text{cost}_{\text{Trunc}}(X_1, \boldsymbol{X_2})$ is $B(X_1, \boldsymbol{X_2})$ and therefore shows that $\text{cost}_{\text{Trunc}}$ is a good approximation of $\text{cost}_{\text{EMax}}$.

Lemma 2.2. *Let $(X_1, \boldsymbol{X_2})$ be a solution to a* TruncatedTwoStage *or* MinEMax *problem. We have*

$$B(X_1, \boldsymbol{X_2}) = \arg\min_B \left[B + \sum_{s \in [m]} p_s \cdot (\text{cost}(X_1, X_2^{(s)}) - B)^+ \right],$$

where the arg min *takes the largest B minimizing the relevant quantity.*

Proof Sketch. The rough idea of the proof is to show that $B + \sum_s p_s(\text{cost}(X_1, X_2^{(s)} - B)^+$ is convex in B and that $B(X_1, \boldsymbol{X_2})$ is a local minimum. In particular, imagine that B is currently set at $B(X_1, \boldsymbol{X_2})$ and consider what happens to $B + \sum_s p_s(\text{cost}(X_1, X_2^{(s)} - B)^+$ if we shift B to be smaller. Recall that we have at least one probability mass across elements which are larger than B by definition of $B(X_1, \boldsymbol{X_2})$. Thus, when we shift B to be smaller, B decreases slower than $\sum_s p_s(\text{cost}(X_1, X_2^{(s)} - B)^+$ increases and so $B + \sum_s p_s(\text{cost}(X_1, X_2^{(s)} - B)^+$ becomes larger overall. The case when B is made larger is symmetric. The full proof is available in §A. ◻

Using Lemma 2.1 and Lemma 2.2, it is easy to show the following two lemmas. These lemmas—proved in §A—upper and lower bound the MinEMax cost of a solution with respect to its TruncatedTwoStage solution respectively.

Lemma 2.3. *For feasible solution $(X_1, \boldsymbol{X_2})$ of any P_{EMax} we have,* $\text{cost}_{EMax}(X_1, \boldsymbol{X_2}) \leq \text{cost}_{Trunc}(X_1, \boldsymbol{X}_2)$.

Lemma 2.4. *Let P_{EMax} be a* MinEMax *problem and P_{Trunc} be its truncated version. Let $(E_1, \boldsymbol{E_2})$ and $(T_1, \boldsymbol{T_2})$ be the optimal solutions to P_{EMax} and P_{Trunc} respectively. We have* $\text{cost}_{Trunc}(T_1, \boldsymbol{T_2}) \leq \left(\frac{2}{1-1/e}\right) \text{cost}_{EMax}(E_1, \boldsymbol{E_2})$.

The preceding lemmas allow us to conclude that an α-approximation algorithm for a TruncatedTwoStage problem is an $O(\alpha)$-approximation algorithm for the corresponding MinEMax problem.

Proof of Theorem 1.1. Let $(\hat{T}_1, \hat{\boldsymbol{T}}_2)$ be the solution returned by an α-approximation algorithm for P_{Trunc}. Let $(E_1, \boldsymbol{E_2})$ and $(T_1, \boldsymbol{T_2})$ be the optimal solutions to P_{EMax} and P_{Trunc} respectively. By Lemma 2.3 we have $\text{cost}_{\mathsf{EMax}}(\hat{T}_1, \hat{\boldsymbol{T}}_2) \le \text{cost}_{\mathsf{Trunc}}(\hat{T}_1, \hat{\boldsymbol{T}}_2)$. Since $(\hat{T}_1, \hat{\boldsymbol{T}}_2)$ is an α approx we have this is at most $\alpha \cdot \text{cost}_{\mathsf{Trunc}}(T_1, \boldsymbol{T_2})$. Applying Lemma 2.4 this is at most $\left(\frac{2\alpha}{1-1/e}\right) \text{cost}_{\mathsf{EMax}}(E_1, \boldsymbol{E_2})$. Since any solution that is feasible for P_{Trunc} is also feasible for P_{EMax}, we conclude that $(\hat{T}_1, \hat{\boldsymbol{T}}_2)$ is a feasible solution for P_{EMax} with cost in P_{EMax} at most $\left(\frac{2\alpha}{1-1/e}\right) \text{cost}_{\mathsf{EMax}}(E_1, \boldsymbol{E_2})$, giving our theorem. □

3 Applications to Linear Two-Stage Covering Problems

In this section we give an $O(\log n + \log m)$-approximation algorithm for MinEMax MST and $O(1)$ approximation algorithms for MinEMax Steiner tree, MinEMax facility location, and MinEMax min-cut. Our algorithms are LP based. To derive our algorithms we use our reduction from §2 to transform a MinEMax problem into a TruncatedTwoStage problem with only a small constant loss in the approximation factor. This transformation allows us to adapt existing LP rounding techniques in which every scenario has a rounding cost close to its fractional cost [RS04, GPRS04, SS06] to solve our TruncatedTwoStage problems and, therefore, our MinEMax problems.

We first give two general techniques to solve a TruncatedTwoStage problem.

3.1 General Techniques

Our *first technique* is to represent $\text{cost}_{\mathsf{Trunc}}$ as an LP objective. For this technique we need to extend the definition of $\text{cost}_{\mathsf{Trunc}}$ from an integral solution $(X_1, \boldsymbol{X_2})$ to a fractional solution $(x_1, \boldsymbol{x_2})$. To do so, in each of our problems we locally define $\text{cost}(x_1, x_2^{(s)})$ for fractional solution $(x_1, x_2^{(s)})$ to scenario s and let $\text{cost}_{\mathsf{Trunc}}(x_1, \boldsymbol{x_2})$ be defined similarly to the integral case, i.e. for fractional $(x_1, \boldsymbol{x_2})$,

$$\text{cost}_{\mathsf{Trunc}}(x_1, \boldsymbol{x_2}) := \min_B \left[B + \sum_s p_s (\text{cost}(x_1, x_2(s)) - B)^+ \right]. \tag{6}$$

Given a minimization LP, it is easy to see that by introducing an additional variable to represent B and additional variables to represent $(\text{cost}(x_1, x_2(s)) - B)^+$ for every s, we can represent $\text{cost}_{\mathsf{Trunc}}(x_1, \boldsymbol{x_2})$ in an LP. For cleanliness of exposition, when we write our LPs we omit these additional variables and simply write our objective as "$\text{cost}_{\mathsf{Trunc}}(x_1, \boldsymbol{x_2})$." Moreover, even though some of our LPs have an exponential number of constraints, we rely on the existence of efficient separation oracles for these LPs. It is easy to verify that this holds even after one introduces the additional variables needed to represent $\text{cost}_{\mathsf{Trunc}}(x_1, \boldsymbol{x_2})$.

We also extend M and B from the integral case as defined in §2 to the fractional case in the following natural way. Given a fractional solution $(x_1, \boldsymbol{x_2})$ and a cost function on fractional solutions, cost, WLOG let our scenarios be indexed such that $\text{cost}(x_1, x_2^{(s)}) \ge \text{cost}(x_1, x_2^{(s+1)})$. Let b be the smallest positive integer such that $\sum_{s=1}^{b} p_s \ge 1$. For fractional $(x_1, \boldsymbol{x_2})$, we define

$$M(x_1, \boldsymbol{x}_2) := [b] \tag{7}$$

$$B(x_1, \boldsymbol{x}_2) := \min_{s \in M(x_1, \boldsymbol{x}_2)} \text{cost}(x_1, x_2^{(s)}). \tag{8}$$

Remark 3.1. It is easy to verify that the proof of Lemma 2.2 also holds for $\text{cost}_{\mathsf{Trunc}}(x_1, \boldsymbol{x_2})$ for fractional $(x_1, \boldsymbol{x_2})$. We will therefore invoke it on fractional $(x_1, \boldsymbol{x_2})$, even though it is stated only for integral $(X_1, \boldsymbol{X_2})$.

Our *second technique* is a generic rounding technique for TruncatedTwoStage problems. Several past works in two-stage optimization show that it is possible to round an LP solution such that the resulting integral solution has cost roughly the same as the fractional solution for *every scenario*. We prove the following lemma to make use of such rounding algorithms.

Lemma 3.2. *Let P_{Trunc} be a* TruncatedTwoStage *problem. Let $(X_1, \boldsymbol{X_2})$ and $(Y_1, \boldsymbol{Y_2})$ be integral or fractional solutions to P_{Trunc}. If for every scenario s we have $\text{cost}(X_1, X_2^{(s)}) \leq c \cdot \text{cost}(Y_1, Y_2^{(s)})$ then*

$$\text{cost}_{Trunc}(X_1, \boldsymbol{X_2}) \leq c \cdot \text{cost}_{Trunc}(Y_1, \boldsymbol{Y_2}).$$

Proof. We have

$$\begin{aligned}
&\text{cost}_{\text{Trunc}}(X_1, \boldsymbol{X_2}) = \min_B \left[B + \sum_s p_s \cdot (\text{cost}(X_1, X_2^{(s)}) - B)^+) \right] \\
&\leq c \cdot B(Y_1, \mathbf{Y_2}) + \sum_s p_s \cdot (\text{cost}(X_1, X_2^{(s)}) - c \cdot B(Y_1, \mathbf{Y_2}))^+) && \text{(by letting } B = c \cdot B(Y_1, \mathbf{Y_2})) \\
&\leq c \cdot B(Y_1, \mathbf{Y_2}) + \sum_s p_s \cdot (c \cdot \text{cost}(Y_1, Y_2^{(s)}) - c \cdot B(Y_1, \mathbf{Y_2}))^+) && \left(\text{by cost}(X_1, X_2^{(s)}) \leq c \cdot \text{cost}(Y_1, Y_2^{(s)})\right) \\
&= c \cdot \left(B(Y_1, \mathbf{Y_2}) + \sum_s p_s \cdot (\text{cost}(Y_1, Y_2^{(s)}) - B(Y_1, \mathbf{Y_2}))^+) \right) \\
&= c \cdot \text{cost}_{\text{Trunc}}(Y_1, \mathbf{Y_2}) && \text{(by Lemma 2.2).}
\end{aligned}$$

□

3.2 Uncapacitated Facility Location

In this section we give a polynomial time $\left(\frac{16}{1-1/e}\right)$-approximation algorithm for MinEMax uncapacitated facility location (UFL).

Definition 3.3 (MinEMax UFL)**.** We are given a set of facilities F and a set of clients $\mathcal{D}$ with a metric c_{ij} specifying the distances between every client j and facility i. We are also given scenarios $S_1, \ldots, S_m \subseteq \mathcal{D}$, where in scenario S_s client j has demand $d_j^s \in \{0, 1\}$[4], and a probability p_s for each scenario. Facility i's opening cost is $f_{1,i}$ in the first stage and $f_{2,i}^{(s)}$ in scenario S_s. These opening costs can be ∞, which indicates the facility cannot be opened. A feasible solution consists of a set of first and second stage facilities $(X_1, \boldsymbol{X_2})$ s.t. $X_1 \cup \bigcup_s X_2^{(s)} \neq \emptyset$. The cost for scenario s in solution $(X_1, \boldsymbol{X_2})$ is

$$\text{cost}(X_1, X_2^{(s)}) := \sum_{i \in X_1} f_{1,i} + \sum_{i \in X_2^{(s)}} f_{2,i}^{(s)} + \sum_{j \in D} \min_{i \in X_1 \cup X_2^{(s)}} c_{ij}.$$

The total cost of our solution $(X_1, \boldsymbol{X_2})$ is $\text{cost}_{\text{EMax}}(X_1, \boldsymbol{X_2}) := \mathbb{E}_{A \sim \boldsymbol{p}} \left[\max_{s \in A} \{\text{cost}(X_1, X_2^{(s)})\} \right]$.

Our algorithm is based on the work of Ravi and Sinha [RS04] on two-stage stochastic UFL. This work shows how to round an LP such that every scenario has a “good” cost after rounding. Applying Lemma 3.2 to this rounding gives an algorithm that approximates TruncatedTwoStage UFL, which by Theorem 1.1 is sufficient to approximate MinEMax UFL.

We use the following LP. Variable $z_{ij}^{(s)}$ corresponds to whether client j is served by facility i in scenario s. Variables $x_1(i)$ and $x_2^{(s)}(i)$ corresponds to whether facility i is opened in the first stage or scenario s,

[4] This easily generalizes to more demand.

respectively. For a fractional solution $(x_1, \boldsymbol{x_2})$, we define

$$\mathrm{cost}(x_1, x_2^{(s)}) := \sum_{i \in F} \left[x_1(i) \cdot f_{1,i} + x_2^{(s)}(i) \cdot f_{2,i}^{(s)} + \sum_{j \in \mathcal{D}} \hat{z}_{ij}^{(s)} \cdot c_{ij} \right],$$

where $\hat{z}_{ij}^{(s)}$ is the natural fractional assignment given fractional facilities $(x_1, x_2^{(s)})$; namely, one that sends clients to their nearest fractionally opened facilities. As described by Eq.(6), this definition of $\mathrm{cost}(x_1, x_2^{(s)})$ defines $\mathrm{cost}_{\mathsf{Trunc}}(x_1, \boldsymbol{x}_2)$ for fractional $(x_1, \boldsymbol{x_2})$, which allows us to define our LP.

$$\begin{aligned} \min \quad & \mathrm{cost}_{\mathsf{Trunc}}(x_1, \boldsymbol{x_2}) && && \text{(UFL LP)} \\ \text{s.t.} \quad & \sum_{i \in F} z_{ij}^{(s)} \geq d_j^{(s)} && \forall j \in \mathcal{D}, \forall s \\ & z_{ij}^{(s)} \leq x_1(i) + x_2^{(s)}(i) && \forall i \in F, \forall j \in \mathcal{D}, \forall s \\ & 0 \leq x_1, \boldsymbol{x_2}, \boldsymbol{z} \end{aligned}$$

Note that an integral solution to the above LP is a feasible solution for MinEMax UFL. Ravi and Sinha showed how to round this LP.

Lemma 3.4 (Theorem 2, Lemma 1 in [RS04]). *Given a fractional solution $(x_1, \boldsymbol{x_2})$ to UFL LP, it is possible to round it to integral $(X_1, \boldsymbol{X_2})$ s.t. for every scenario s we have* $\mathrm{cost}(X_1, X_2^{(s)}) \leq 8 \cdot \mathrm{cost}(x_1, x_2^{(s)})$.

We now give our approximation algorithm for MinEMax UFL.

Theorem 3.5. MinEMax *UFL can be* $\left(\frac{16}{1-1/e}\right)$*-approximated in polynomial time.*

Proof. Our algorithm starts by solving UFL LP to get a fractional $(x_1, \boldsymbol{x_2})$. Next, round $(x_1, \boldsymbol{x_2})$ using Lemma 3.4 to integral $(X_1, \boldsymbol{X_2})$. Return $(X_1, \boldsymbol{X_2})$.

Let $(O_1, \boldsymbol{O_2})$ be the optimal integral solution to the TruncatedTwoStage instance of our problem and let $(o_1, \boldsymbol{o_2})$ be its corresponding characteristic function. By definition, $\mathrm{cost}_{\mathsf{Trunc}}(o_1, \boldsymbol{o_2}) = \mathrm{cost}_{\mathsf{Trunc}}(O_1, \boldsymbol{O_2})$. Now using Lemma 3.2 and Lemma 3.4 it follows that

$$\mathrm{cost}_{\mathsf{Trunc}}(X_1, X_2) \leq 8 \cdot \mathrm{cost}_{\mathsf{Trunc}}(x_1, \boldsymbol{x_2}).$$

Since $(o_1, \boldsymbol{o_2})$ feasible for UFL LP, we get

$$\mathrm{cost}_{\mathsf{Trunc}}(X_1, X_2) \quad \leq \quad 8 \cdot \mathrm{cost}_{\mathsf{Trunc}}(o_1, \boldsymbol{o_2}) \quad = \quad 8 \cdot \mathrm{cost}_{\mathsf{Trunc}}(O_1, \boldsymbol{O_2}).$$

Thus, our algorithm is an 8-approximation for TruncatedTwoStage UFL. Applying Theorem 1.1 gives a $\left(\frac{16}{1-1/e}\right)$-approximation for MinEMax UFL.

Lastly, notice that our algorithm is trivially polynomial time. □

3.3 Steiner Tree

In this section we give a $\left(\frac{60}{1-1/e}\right)$-approximation for MinEMax rooted Steiner tree.

Definition 3.6 (MinEMax Rooted Steiner tree). We are given a graph $G = (V, E)$, a root $r \in V$, a cost c_e for each edge e. We are also given scenarios $S_1, \ldots, S_m \subseteq V$, each with an associated probability p_s and an *inflation factor* $\sigma_s > 0$. We must find a first stage solution $X_1 \subseteq E$ and a second-stage solution for every scenario, $X_2^{(j)} \subseteq E$. A solution is feasible if for every s we have $X_1 \cup X_2^{(s)}$ connects $\{r\} \cup S_s$. The cost for scenario s in solution $(X_1, \boldsymbol{X_2})$ is

$$\mathrm{cost}(X_1, X_2^{(s)}) := \sum_{e \in X_1} c_e + \sigma_s \cdot \sum_{e \in X_2^{(s)}} c_e. \tag{9}$$

The total cost we pay for solution $(X_1, \boldsymbol{X_2})$ is $\text{cost}_{\text{EMax}}(X_1, \boldsymbol{X_2}) := E_{A\sim\boldsymbol{p}}\left[\max_{s\in A}\{\text{cost}(X_1, X_2^{(s)})\}\right]$.

Our algorithm is based on an LP rounding algorithm of Gupta et al. [GRS04] for two-stage stochastic Steiner tree. Roughly, we use Lemma 3.2 to argue that the first stage solution for every optimal TruncatedTwoStage solution is a tree rooted at r. This structural property allows us to write an LP that approximately captures TruncatedTwoStage Steiner tree. Gupta et al. [GRS04] showed that this LP can be rounded s.t. every scenario has a good cost. As in the previous section, we combine this rounding with Lemma 3.2 to derive an approximation algorithm for TruncatedTwoStage Steiner tree, which is sufficient for approximating MinEMax Steiner tree by Theorem 1.1.

We begin by arguing that up to small constants, the optimal first stage solution is a tree rooted at r.

Lemma 3.7. *There exists an integral solution $(\hat{X}_1, \hat{\boldsymbol{X}}_2)$ to* TruncatedTwoStage *Steiner tree s.t. $G[X_1]$ is a tree rooted at r and $\text{cost}_{Trunc}(\hat{X}_1, \hat{\boldsymbol{X}}_2) \leq 2 \cdot \text{cost}_{Trunc}(O_1, \boldsymbol{O_2})$, where $(O_1, \boldsymbol{O_2})$ is the optimal solution to* TruncatedTwoStage *Steiner tree.*

Proof. Lemma 4.1 of Dhamdhere et al. [DGRS05] shows that given $(O_1, \boldsymbol{O_2})$ it is possible to modify it to a feasible solution $(\hat{X}_1, \hat{\boldsymbol{X}}_2)$ such that $G[\hat{X}_1]$ is a tree rooted at r and $\text{cost}(\hat{X}_1, \hat{X}_2^{(s)}) \leq 2 \cdot \text{cost}(O_1, O_2^{(s)})$ for every s. It follows by Lemma 3.2 that $\text{cost}_{\text{Trunc}}(\hat{X}_1, \hat{\boldsymbol{X}}_2) \leq 2 \cdot \text{cost}_{\text{Trunc}}(O_1, \boldsymbol{O_2})$. □

We now describe how to formulate an LP that leverages the structural property in Lemma 3.7. In particular, this indicates that as one gets closer to r, one must fractionally buy edges to a greater and greater extent. This constraint can be captured in an LP. Specifically, every node in a scenario (a.k.a. terminal) is the source of one unit of flow that is ultimately routed to r; this flow follows a path whose fractional "first stage-ness" is monotonically increasing.

More formally, we copy each edge $e = \{u, v\}$ into two directed edges (u, v) and (v, u). Let $\vec{e}$ be either one of these directed edges. Next, for each such directed edge $\vec{e}$ and every terminal in $t \in \bigcup_s S_s$, we define variables $r_1(t, \vec{e})$ and $r_2^{(s)}(t, \vec{e})$ for every s to represent how much t is connected to r by e in the first stage and in scenario s, respectively. Also, for undirected edge e, define variables $x_1(e)$ and $x_2^{(s)}(e)$ to stand for how much we buy e in the first stage and scenario s, respectively. For fractional $(x_1, \boldsymbol{x_2})$, we define

$$\text{cost}_{\text{Trunc}}(x_1, x_2^{(s)}) := \sum_e c_e \cdot x_1(e) + \sigma_s \cdot c_e \cdot x_2(e),$$

which as described by Eq.(6) also defines $\text{cost}_{\text{Trunc}}(x_1, \boldsymbol{x}_2)$. Letting $\delta^-(v)$ and $\delta^+(v)$ stand for all directed edges going into and out of v, respectively. The following is our LP.

$$\begin{aligned}
\min \quad & \text{cost}_{\text{Trunc}}(x_1, \boldsymbol{x_2}) && \text{(ST LP)} \\
\text{s.t.} \quad & \sum_{\vec{e}\in\delta^+(v)} r_1(t,\vec{e}) + r_2^{(s)}(t,\vec{e}) = \sum_{\vec{e}\in\delta^-(v)} r_1(t,\vec{e}) + r_2^{(s)}(t,\vec{e}) && \forall s, t \in S_s, v \notin \{t, r\} \\
& \sum_{\vec{e}\in\delta^+(t)} r_1(t,\vec{e}) + r_2^{(s)}(t,\vec{e}) \geq 1 && \forall s, t \in S_s \\
& \sum_{\vec{e}\in\delta^-(v)} r_1(t,\vec{e}) \leq \sum_{\vec{e}\in\delta^+(v)} r_1(t,\vec{e}) && \forall s, t \in S_s, v \notin \{t, r\} \quad (10) \\
& r_1(t,\vec{e}) \leq x_1(e); r_2^{(s)}(t,\vec{e}) \leq x_2^{(s)}(e) && \forall s, t \in S_s, \vec{e} \\
& r, x_1, \boldsymbol{x_2} \geq 0
\end{aligned}$$

Notably, Eq. (10) enforces that terminal t is serviced by the first stage more and more as one moves closer to the root. The characteristic vector of $(\hat{X}_1, \hat{\boldsymbol{X}}_2)$ as described in Lemma 3.7 gives a feasible solution to ST LP. As a result, Lemma 3.7 demonstrates that ST LP has nearly optimal objective as stated in the following corollary.

Corollary 3.8. *Let* $(x_1, \boldsymbol{x_2})$ *be the optimal solution of ST LP. We have* $\text{cost}_{Trunc}(x_1, \boldsymbol{x_2}) \leq 2 \cdot \text{cost}_{Trunc}(O_1, \boldsymbol{O_2})$, *where* $(O_1, \boldsymbol{O_2})$ *is the optimal solution to* TruncatedTwoStage *Steiner tree.*

Proof. Let $(\hat{x}_1, \hat{\boldsymbol{x}}_2)$ be the characteristic vector of $(\hat{X}_1, \hat{\boldsymbol{X}}_2)$ from Lemma 3.7. Let P_2 for terminal t be the shortest path from t to X_1 in $G[X_2]$. Let u_t be the sink of P_2 and let P_1 be the shortest path from u_t to r. Notice that $(\hat{x}_1, \hat{\boldsymbol{x}}_2)$ along with r_2 which sends one unit of flow from t to u_t along P_2 and r_1 which sends one unit of flow from u_t to r along P_1 is a feasible solution to ST LP. Moreover, notice that cost of this solution is $\text{cost}_{\text{Trunc}}(\hat{x}_1, \boldsymbol{x_2}) = \text{cost}_{\text{Trunc}}(X_1, \boldsymbol{X_2}) \leq 2 \cdot \text{cost}_{\text{Trunc}}(O_1, \boldsymbol{O_2})$ by Lemma 3.7. □

Previous work of Gupta et al. [GRS04] show that it is possible to round a fractional solution of ST LP such that *every* scenario has a good cost.

Lemma 3.9 ([GRS04])**.** *A fractional solution* $(x_1, \boldsymbol{x_2})$ *to ST LP can be rounded in polynomial time to a feasible integral solution* $(X_1, \boldsymbol{X_2})$ *s.t.* $\text{cost}(X_1, X_2^{(s)}) \leq 15 \cdot \text{cost}(x_1, x_2^{(s)})$ *for every* s.

Since Corollary 3.8 gives ST LP has a good optimal solution, we can round ST LP s.t. every scenario has a low cost. Now Lemma 3.2 tells us that such a rounding preserves the cost of a solution for TruncatedTwoStage optimization. This gives the following theorem.

Theorem 3.10. MinEMax *Steiner tree can be* $\left(\frac{60}{1-1/e}\right)$*-approximated in polynomial time.*

Proof. Our algorithm first solves ST LP to get fractional solution $(x_1, \boldsymbol{x_2})$. Next, we apply Lemma 3.9 to round $(x_1, \boldsymbol{x_2})$ in polynomial time to give $(X_1, \boldsymbol{X_2})$ as our solution. Thus, we have

$$\begin{aligned}
\text{cost}_{\text{Trunc}}(X_1, \boldsymbol{X_2}) &\leq 15 \cdot \text{cost}_{\text{Trunc}}(x_1, \boldsymbol{x_2}) && \text{(by Lemma 3.2, Lemma 3.9)} \\
&\leq 30 \cdot \text{cost}_{\text{Trunc}}(O_1, \boldsymbol{O_2}), && \text{(by Corollary 3.8)}
\end{aligned}$$

where $(O_1, \boldsymbol{O_2})$ is the optimal TruncatedTwoStage Steiner tree solution. This implies we have a 30-approximation algorithm for TruncatedTwoStage Steiner tree. Now by Theorem 1.1, we have a $\left(\frac{60}{1-1/e}\right)$-approximation for MinEMax Steiner tree.

Lastly, each of our subroutines has a polynomial runtime by previous lemmas, and so we conclude that our algorithm has a polynomial runtime. □

3.4 MST

In this section we give a randomized polynomial time algorithm which with high probability has expected cost $O(\log n + \log m)$ times the optimal MinEMax minimum spanning tree (MST) on an n-node graph with m different scenarios.

Definition 3.11 (MinEMax MST)**.** We are given a graph $G = (V, E)$ where $|V| = n$, a set of m scenarios $S_1, \ldots S_m$ where each scenario S_s has an associated second-stage cost function $\text{cost}_2^{(s)} : E \to \mathbb{Z}^+$ and a probability p_s. We are also given a first-stage cost function, $\text{cost}_1 : E \to \mathbb{Z}^+$. We must provide a first stage solution $X_1 \subseteq E$ and a solution $X_2^{(s)} \subseteq E$ for every scenario s, which is feasible if $G[X_1 \cup X_2^{(s)}]$ spans V for every s. The cost for scenario s in solution $(X_1, \boldsymbol{X_2})$ is

$$\text{cost}(X_1, X_2^{(s)}) := \sum_{e \in X_1} \text{cost}_1(e) + \sum_{e \in X_2^{(s)}} \text{cost}_2^{(s)}(e). \tag{11}$$

The total cost for solution $(X_1, \boldsymbol{X_2})$ is $\text{cost}_{\text{EMax}}(X_1, \boldsymbol{X_2}) := \mathbb{E}_{A\sim\boldsymbol{p}}\left[\max_{s\in A}\{\text{cost}(X_1, X_2^{(s)})\}\right]$.

Our algorithm is based on the work of Dhamdhere et al. [DRS05] on two-stage stochastic MST. They give a rounding technique that produces integral solutions where every scenario has a cost close to the fractional cost. Using this rounding, and applying Lemma 3.2, we get an approximation algorithm for TruncatedTwoStage MST, which by Theorem 1.1 is also sufficient to approximate MinEMax MST.

Notice that since MinEMax generalizes two-stage robust optimization, our MinEMax result gives a $O(\log n + \log m)$ approximation for two-stage robust MST as a corollary. To the best of our knowledge, this is the first non-trivial algorithm for two-stage robust MST.

Our algorithm is based on an LP. We have $m + 1$ variables for each edge e, namely $x_1(e)$ and $x_2^{(s)}(e)$ for $s \in [m]$ indicating if we take e in the first stage and in the second stage for scenario s, respectively. For a fractional solution $(x_1, \boldsymbol{x_2})$, we define

$$\text{cost}(x_1, x_2^{(s)}) := \sum_e x_1(e) \cdot \text{cost}_1(e) + x_2^{(s)}(e) \cdot \text{cost}_2(e), \tag{12}$$

which as described in Eq.(6), defines $\text{cost}_{\text{Trunc}}(x_1, \boldsymbol{x_2})$ for fractional $(x_1, \boldsymbol{x_2})$. Letting $\delta(S)$ be all edges with exactly one endpoint in $S \subseteq V$. The following is our LP.

$$\begin{aligned} \min \quad & \text{cost}_{\text{Trunc}}(x_1, \boldsymbol{x_2}) && \text{(MST LP)} \\ \text{s.t.} \quad & \sum_{e\in\delta(S)} \left(x_1(e) + x_2^{(s)}(e)\right) \geq 1 && \forall S \subset V, s \in [m] \\ & x_1, \boldsymbol{x_2} \geq 0 \end{aligned}$$

Note that an integral solution to MST LP is a feasible solution for the TruncatedTwoStage MST problem as a set of edges with at least one edge leaving every cut is a spanning tree.[5] Also, although this LP has super-polynomial constraints, it is easy to obtain an efficient separation by solving min-cut; see Dhamdhere et al. [DRS05].

We need the following result of Dhamdhere et al. [DRS05] to round MST LP such that every scenario has a low cost.

Lemma 3.12 ([DRS05])**.** *It is possible to randomly round a feasible fractional solution $(x_1, \boldsymbol{x_2})$ to MST LP to an integral solution $(X_1, \boldsymbol{X_2})$ in polynomial time s.t. with probability at least $1 - \frac{1}{mn^2}$ for every scenario s we have $\mathbb{E}[\text{cost}(X_1, X_2^{(s)})] \leq \text{cost}(x_1, x_2^{(s)}) \cdot (40 \log n + 16 \log m)$. Here the expectation is taken over the randomness of our rounding and m is the number of scenarios.*

We can now design our approximation algorithm for MinEMax MST.

Theorem 3.13. *There exists a randomized polynomial-time algorithm that with probability at least $1 - \frac{1}{mn^2}$ in expectation $O(\log n + \log m)$-approximates* MinEMax *MST where $n = |V|$ and m is the number of scenarios.*

Proof. Our algorithm starts by following MST LP to get a fractional solution $(x_1, \boldsymbol{x_2})$. Next, apply Lemma 3.12 to round $(x_1, \boldsymbol{x_2})$ to an integral solution $(X_1, \boldsymbol{X_2})$. Return $(X_1, \boldsymbol{X_2})$.

Next consider the cost of $(X_1, \boldsymbol{X_2})$. Let $(O_1, \boldsymbol{O_2})$ be the optimal integral solution to our TruncatedTwoStage MST problem and let $(o_1, \boldsymbol{o_2})$ be the corresponding characteristic vector. Notice that $(o_1, \boldsymbol{o_2})$ is a feasible solution to MST LP. Moreover, it is easy to verify that $\text{cost}_{\text{Trunc}}(o_1, \boldsymbol{o_2}) = \text{cost}_{\text{Trunc}}(O_1, \boldsymbol{O_2})$. Taking expectations over the randomness of our algorithm and applying Lemma 3.2 and Lemma 3.12, we have with

[5] If such a solution has any cycles it is not necessarily an MST, though one can always delete an edge from such a cycle and improve the cost of the solution.

probability at least $1-\frac{1}{mn^2}$ that

$$\mathbb{E}[\text{cost}_{\text{Trunc}}(X_1, \boldsymbol{X_2})] \le (40\log n + 16\log m)\cdot \text{cost}_{\text{Trunc}}(o_1, \boldsymbol{o_2})$$
$$= (40\log n + 16\log m)\cdot \text{cost}_{\text{Trunc}}(O_1, \boldsymbol{O_2}).$$

Thus, with probability at least $1-\frac{1}{mn^2}$ our algorithm's expected TruncatedTwoStage cost is within $(40\log n + 16\log m)$ of the cost of the optimal TruncatedTwoStage MST solution. We conclude by Theorem 1.1 that with high probability in expectation our algorithm $O(\log n + \log m)$-approximates MinEMax MST.[6]

Our algorithm is trivially polynomial time by the separability of our LP and Lemma 3.12. □

3.5 Min-Cut

In this section we give a polynomial time $\left(\frac{8}{1-1/e}\right)$-approximation for MinEMax min-cut.

Definition 3.14 (MinEMax min-cut)**.** We are given a graph $G=(V,E)$, a root $r \in V$, a cost c_e for edge e, and m scenarios specified by terminals $t_1,\dots,t_m \in V$. Each scenario t_s has an associated probability p_s and inflation factor $\sigma_s > 0$. We must provide a first stage solution $X_1 \subseteq E$ and a second-stage solution $X_2^{(s)} \subseteq E$ for each s. A feasible solution is one where $X_1 \cup X_2^{(s)}$ cuts r from t_s for every s. The cost for scenario s in solution $(X_1, \boldsymbol{X_2})$ is

$$\text{cost}(X_1, X_2^{(s)}) := \sum_{e\in X_1} c_e + \sigma_s \cdot \sum_{e\in X_2^{(s)}} c_e. \tag{13}$$

The total cost of solution $(X_1, \boldsymbol{X_2})$ is $\text{cost}_{\text{EMax}}(X_1, \boldsymbol{X_2}) := \mathbb{E}_{A\sim \boldsymbol{p}}\left[\max_{s\in A}\{\text{cost}(X_1, X_2^{(s)})\}\right]$.

We draw on past work on two-stage stochastic min-cut. In particular, we use the insight of Golovin et al. [GGP+15b] that, when approximating min-cut in a two-stage setting, it suffices to consider a relaxed version of the two-stage problem. In the second stage of this relaxed problem one does not pay the cost of completing their first stage solution. Rather, if the vertex corresponding to a scenario is not fully cut away in the first stage, in the second stage one must pay the full cost of cutting away that vertex in the original graph. The utility of this observation is that the relaxed problem can be captured by an LP (which is not clear in general for two-stage min-cut problems).

Thus, we first write an LP for the relaxed version of TruncatedTwoStage min-cut and then round it using ideas from Golovin et al. [GGP+15b]. We make use of Lemma 3.2 in several places in our analysis to show that the optimal solution has certain structure, ultimately showing that such an algorithm 4-approximates the TruncatedTwoStage version of min-cut. As shown in Theorem 1.1, 4-approximating TruncatedTwoStage min-cut is sufficient to $(\frac{8}{1-1/e})$-approximate MinEMax min-cut.

We let $\widehat{P_{\text{Trunc}}}$ be the previously mentioned relaxed problem; we use "$\widehat{\ }$" to indicate objects and functions in this relaxed problem. Problem $\widehat{P_{\text{Trunc}}}$ is similar to P_{Trunc} but the form of the second stage solution and cost of each scenario is different. In particular, we must give a first stage solution $X_1 \subseteq E$ and a second stage solution $\hat{\boldsymbol{X}}_{\boldsymbol{2}} \in \{0,1\}^m$ indicating if each t_s is cut away by our first stage solution. A solution $(X_1, \hat{\boldsymbol{X}}_{\boldsymbol{2}})$ is feasible if for every s we have $\hat{X}_2{}^{(s)} = 1$ iff X_1 cuts t_s from r. For scenario s in this solution we pay

$$\widehat{\text{cost}}(X_1, \hat{X}_2{}^{(s)}) := \sum_{e\in X_1} [c_e] + \sigma_s(1-\hat{X}_2^{(s)})\cdot \text{cut-cost}(r, t_s), \tag{14}$$

where cut-cost(r, t_s) is the minimum cost of a cut separating r from t_s in G.

[6] Although Theorem 1.1 and Lemma 3.2 do not explicitly account for an expectation taken over the randomness of an algorithm, it is easy to verify that the such an expectation does not affect these results.

We capture $\widehat{P_{\text{Trunc}}}$ with an LP. We have a variable $x_1(e)$ for each edge e standing for whether we cut e in the first stage and a variable $x_2^{(s)}$ for each s standing for whether or not t_s is cut from r in the first stage. For fractional $(x_1, \boldsymbol{x_2})$, we give its cost in a given scenario s as

$$\widehat{\text{cost}}(x_1, x_2^{(s)}) := \sum_e c_e \cdot x_1(e) + \sigma_s(1 - x_2^{(s)}) \cdot \text{cut-cost}(r, t_s). \tag{15}$$

As described in Eq.(6), this definition of $\widehat{\text{cost}}(x_1, x_2^{(s)})$ also defines $\widehat{\text{cost}_{\text{Trunc}}}(x_1, \boldsymbol{x_2})$ for fractional $(x_1, \boldsymbol{x_2})$. Thus, we can now give our LP with $\widehat{\text{cost}_{\text{Trunc}}}(x_1, \boldsymbol{x_2})$ as the objective value.

$$\begin{aligned}
\min \quad & \widehat{\text{cost}_{\text{Trunc}}}(x_1, \boldsymbol{x_2}) & & \text{(MC LP)} \\
\text{s.t.} \quad & \sum_{e \in P} x_1(e) \geq x_2^{(s)} & \forall P \in \mathcal{P}_G(r, t_s), \forall s \quad & (16) \\
& 0 \leq x_1(e), x_2^{(s)} \leq 1 & \forall s, e \in E, \quad & (17)
\end{aligned}$$

where $\mathcal{P}_G(r, t_s)$ gives all paths from r to t_s in G. Although this LP has super-polynomial constraints, it is easy to see that a polynomial time algorithm for $s - t$ shortest path gives an efficient separation oracle. Hence, this LP is solvable in polynomial time.

As Golovin at al. [GGP+15b] demonstrated, one can construct a feasible solution to MC LP which has cost in $\widehat{P_{\text{Trunc}}}$ roughly analogous to the optimal costs in P_{Trunc} for every scenario.

Lemma 3.15 (Lemma 3.1 in [GGP+15b])**.** *Let $(O_1, \boldsymbol{O_2})$ be the optimal integral solution to P_{Trunc}. There exists a feasible integral solution to MC LP, $(x_1, \hat{\boldsymbol{x}}_2)$, such that $\widehat{\text{cost}}(x_1, \hat{x}_2^{(s)}) \leq 2 \cdot \text{cost}(O_1, \boldsymbol{O_2})$ for every s.*

Applying Lemma 3.2 and the fact that the optimal solution to MC LP is certainly no more than $\widehat{\text{cost}_{\text{Trunc}}}(x_1, \hat{\boldsymbol{x}}_2)$ as given in Lemma 3.15, we have the following corollary. This corollary shows that the above LP has optimal value roughly the same the cost of the optimal integral TruncatedTwoStage min-cut solution.

Corollary 3.16. *Let $(o_1, \boldsymbol{o_2})$ be the optimal fractional solution to MC LP and let $(O_1, \boldsymbol{O_2})$ be the optimal solution to P_{Trunc}. We have $\widehat{\text{cost}}_{Trunc}(o_1, \boldsymbol{o_2}) \leq 2\text{cost}_{Trunc}(O_1, \boldsymbol{O_2})$.*

Having shown how MC LP has optimal cost analogous to the optimal solution to P_{Trunc}, we need only make use of the fractional solution to MC LP to construct an integral solution to P_{Trunc}. We do so with algorithm MinCutMinEMax. Roughly, this algorithm first cuts away all scenarios that were fractionally cut away by MC LP to an extent of at least $\frac{1}{2}$; its second stage solution is the minimum remaining cut for each scenario. See Algorithm 1.

Algorithm 1 MinCutMinEMax

Input: An instance of min-cut MinEMax
Output: A solution to the input instance
$(o_1, \boldsymbol{o_2}) \leftarrow$ optimal fractional solution to *MCLP*
$U = \left\{t_s \mid o_2^{(s)} \geq \frac{1}{2}, s \in [m]\right\}$
$X_1 \leftarrow$ minimum $r - U$ cut in G
$X_2^{(s)} \leftarrow$ minimum $r - t_s$ cut in $G \setminus X_1$
return $(X_1, \boldsymbol{X_2})$

Given a set of first-stage edges, X_1, we let $\hat{X}_2^{(s)}(X_1)$ be the natural way to derive a second-stage solution for $\widehat{P_{\text{Trunc}}}$ from a first-stage solution. In particular,

$$\hat{X}_2^{(s)}(X_1) := \begin{cases} 1 & \text{if } X_1 \text{ cuts } t_s \text{ from } r \\ 0 & \text{o/w} \end{cases}$$

We now argue that MinCutMinEMax solves P_{Trunc} at cost proportional to its completion to a solution to $\widehat{P_{\text{Trunc}}}$.

Lemma 3.17. *Let $(X_1, \boldsymbol{X_2})$ be the returned values of* MinCutMinEMax. *We have*

$$\text{cost}_{Trunc}(X_1, \boldsymbol{X_2}) \le \widehat{\text{cost}_{Trunc}}(X_1, \boldsymbol{\hat{X}_2(X_1)}).$$

Proof. We first argue that $\sum_{e \in X_2^{(s)}} c_e \le (1 - \hat{X}_2^{(s)}(X_1)) \cdot \text{cut-cost}_G(r, t_s)$ for any s. We case on whether X_1 cuts t_s from r.

- If X_1 cuts t_s from r then we have $\hat{X}_2^{(s)}(X_1) = 1$ and so trivially $\sum_{e \in X_2^{(s)}} c_e = 0$ meaning $\sum_{e \in X_2^{(s)}} c_e \le (1 - \hat{X}_2^{(s)}(X_1)) \cdot \text{cut-cost}_G(r, t_s)$.
- If X_1 does not cut t_s from r then we have $\hat{X}_2^{(s)}(X_1) = 0$. But $X_2^{(s)}$ is a minimum $r - t_s$ cut in $G \setminus X_1$ and $\text{cut-cost}_G(r, t_s)$ is the minimum cut cost in G and so $\sum_{e \in X_2^{(s)}} c_e \le \text{cut-cost}_G(r, t_j) = (1 - \hat{X}_2^{(s)}(X_1)) \cdot \text{cut-cost}_G(r, t_s)$.

Thus, for any s we have

$$\sum_{e \in X_2^{(s)}} c_e \le (1 - X_2^{(s)}(X_1)) \cdot \text{cut-cost}_G(r, t_s). \tag{18}$$

Next, notice that it follows that for any s we have $\text{cost}(X_1, X_2^{(s)}) \le \widehat{\text{cost}}(X_1, \hat{X}_2^{(s)}(X_1))$ since

$$\begin{aligned}
\text{cost}(X_1, X_2^{(s)}) &= \sum_{e \in X_1} c_e + \sigma_s \cdot \sum_{e \in X_2^{(s)}} c_e \\
&\le \sum_{e \in X_1} c_e + \sigma_s (1 - X_2^{(s)}(X_1)) \cdot \text{cut-cost}_G(r, t_s) && \text{(by Eq. (18))} \\
&= \widehat{\text{cost}}(X_1, \hat{X}_2^{(s)}(X_1))
\end{aligned}$$

Thus for any s we have $\text{cost}(X_1, X_2^{(s)}) \le \widehat{\text{cost}}(X_1, \hat{X}_2^{(s)}(X_1))$; applying Lemma 3.2 gives Lemma 3.17. □

Finally, combining previous lemmas we can prove that MinCutMinEMax efficiently approximates MinEMax min-cut.

Theorem 3.18. MinCutMinEMax *is a polynomial time* $\left(\frac{8}{1-1/e}\right)$*-approximation for* MinEMax *min-cut.*

Proof. First, notice that by Lemma 3.17 we have

$$\text{cost}_{\text{Trunc}}(X_1, \boldsymbol{X_2}) \le \widehat{\text{cost}_{\text{Trunc}}}(X_1, \boldsymbol{\hat{X}_2(X_1)}) \tag{19}$$

Thus, for the remainder of this proof it will suffice to upper bound $\widehat{\text{cost}_{\text{Trunc}}}(X_1, \boldsymbol{\hat{X}_2(X_1)})$. We begin by showing that $(X_1, \boldsymbol{\hat{X}_2(X_1)})$ has cost in $\widehat{P_{\text{Trunc}}}$ roughly the same as the optimal solution to $\widehat{P_{\text{Trunc}}}$. In particular, we will show that $\widehat{\text{cost}_{\text{Trunc}}}(X_1, \boldsymbol{\hat{X}_2(X_1)}) \le 2\widehat{\text{cost}_{\text{Trunc}}}(o_1, \boldsymbol{o_2})$ where $(o_1, \boldsymbol{o_2})$ is the optimal fractional solution to MC LP. Let $\bar{o}_1 = 2o_1$.

First, we upper bound the cost of X_1 in $\widehat{P_{\text{Trunc}}}$ relative to o_1. We do so by first arguing that $\bar{o}_1$ is a fractional $r - U$ cut. For $t_s \in U$ we have $o_2^{(s)} \ge \frac{1}{2}$. It follows by Equation (16) that for every path P from r to t_s we have $\sum_{e \in P} o_1(e) \ge o_2^{(s)} \ge \frac{1}{2}$ and so we have for every path P from r to t_s that $\sum_{e \in P} \bar{o}_1(e) = \sum_{e \in P} 2o_1(e) \ge 2o_2^{(s)} \ge 1$. Thus, $\bar{o}_1$ is a fractional $r - U$ cut. Next notice that since the minimum cut is a lower bound on any fractional cut and $\bar{o}_1$ is a fractional $r - U$ cut and X_1 is a minimum $r - U$ cut, we have

$$\sum_{e \in X_1} c_e \le \sum_{e} \bar{o}_1(e) \cdot c_e = \sum_{e} 2o_1(e) \cdot c_e \tag{20}$$

Next, we upper bound the cost of $\hat{\boldsymbol{X}}_{\boldsymbol{2}}(\boldsymbol{X_1})$ in $\widehat{P_{\text{Trunc}}}$ relative to $\boldsymbol{o_2}$. First recall that $\hat{X}_2^{(s)}(X_1)$ is 1 if X_1 cuts t_s from r and 0 otherwise. If $\hat{X}_2^{(s)}(X_1) = 1$ we trivially have $\text{cut-cost}_G(r, t_s) \cdot (1 - \hat{X}_2^{(s)}(X_1)) \leq \text{cut-cost}_G(r, t_s) \cdot (1 - o_2^{(s)})$. If $\hat{X}_2^{(s)}(X_1) = 0$ then we have X_1 does not cut t_s from r which means that $t_s \notin U$ and so $o_2^{(s)} < \frac{1}{2}$. It follows that in this case $\text{cut-cost}_G(r, t_s) \cdot (1 - \hat{X}_2^{(s)}(X_1)) = \text{cut-cost}_G(r, t_s) \leq 2 \cdot \text{cut-cost}_G(r, t_s) \cdot (1 - o_2^{(s)})$. Thus, for any s we have

$$\text{cut-cost}_G(r, t_s) \cdot (1 - \hat{X}_2^{(s)}(X_1)) \leq 2 \cdot \text{cut-cost}_G(r, t_s)(1 - o_2^{(s)}) \tag{21}$$

Since we have upper bound the first and second stage costs of $(X_1, \hat{\boldsymbol{X}}_{\boldsymbol{2}}(\boldsymbol{X_1}))$ in $\widehat{P_{\text{Trunc}}}$ we can upper bound its total cost in $\widehat{P_{\text{Trunc}}}$. In particular, we have for every s it holds that

$$\begin{aligned}
\widehat{\text{cost}}(X_1, \hat{X}_2^{(s)}(X_1)) &= \sum_{e \in X_1} [c_e] + (1 - \hat{X}_2^{(s)}) \cdot \text{cut-cost}(r, t_s) && \left(\text{by dfn. of } \widehat{\text{cost}}(X_1, \hat{X}_2^{(s)}(X_1))\right) \\
&\leq \sum_{e} 2o_1(e) \cdot c_e + 2 \cdot \text{cut-cost}_G(r, t_s)(1 - o_2^{(s)}) && \text{(by Equations 20, 21)} \\
&= 2\widehat{\text{cost}}(o_1, o_2^{(s)}) && \left(\text{by dfn. of } \widehat{\text{cost}}(o_1, o_2^{(s)})\right)
\end{aligned}$$

Thus, for any scenario s we have $\widehat{\text{cost}}(X_1, \hat{X}_2^{(s)}(X_1)) \leq 2\widehat{\text{cost}}(o_1, o_2^{(s)})$ and so applying Lemma 3.2 we get

$$\widehat{\text{cost}_{\text{Trunc}}}(X_1, \hat{\boldsymbol{X}}_{\boldsymbol{2}}(\boldsymbol{X_1})) \leq 2\widehat{\text{cost}_{\text{Trunc}}}(o_1, \boldsymbol{o_2}). \tag{22}$$

To complete our proof we notice that by Corollary 3.16 we have

$$2\widehat{\text{cost}_{\text{Trunc}}}(o_1, \boldsymbol{o_2}) \leq 4\text{cost}_{\text{Trunc}}(O_1, \boldsymbol{O_2}). \tag{23}$$

Combining Equation (19), Equation (22) and Equation (23) we conclude that

$$\text{cost}_{\text{Trunc}}(X_1, \boldsymbol{X_2}) \leq 4\text{cost}_{\text{Trunc}}(O_1, \boldsymbol{O_2}).$$

Thus, our algorithm, MinCutMinEMax, is a 4-approximation for TruncatedTwoStage min-cut. Applying Theorem 1.1, we conclude MinCutMinEMax is a $\left(\frac{8}{1-1/e}\right)$-approximation for MinEMax min-cut.

A polynomial runtime follows from the fact that our algorithm needs to only solve LP MC LP, compute U and then compute polynomially many min-cuts. □

4 MinEMax k-Center

In this section we give a constant approximation for MinEMax k-center.

Definition 4.1 (MinEMax k-center)**.** We are given a metric space $(\mathcal{D}, \{c_{ij}\})$ over points $\mathcal{D}$, a set of scenarios $\{S_s\}_{s=1}^m$ where S_s corresponds to client $s \in \mathcal{D}$, and a probability p_s for each scenario. We must output $X \subseteq \mathcal{D}$ which is feasible if $|X| \leq k$. The cost we pay for solution X is

$$\mathbb{E}_{A \sim \boldsymbol{p}}\left[\max_{s \in A} d(X, s)\right],$$

where $d(X, s) := \min_{i \in X} c_{is}$.

Notice that unlike the preceding MinEMax problems, here we only provide a first stage solution. MinEMax k-center can be phrased as a two-stage covering problem with a non-linear cost if we set the cost of any solution that opens a facility in the second stage to ∞. For this reason, we let $B(X)$ and $M(X)$ stand for $B(X, \emptyset)$ and $M(X, \emptyset)$ respectively as defined in Eq.(5).

Roughly, our algorithm works as follows. We draw on the intuition behind Theorem 1.1 that the expected max is well-approximated by truncating and summing values, and therefore truncate distances in the metric. We then solve a k-center-like LP on this truncated metric and use our LP solution to cluster together nearby clients. Finally, in this clustered version of our problem we run a k-median algorithm and return its solution as our solution for the MinEMax k-center problem. Our techniques follow Chakrabarty and Swamy [CS18] combined with careful applications of Lemma 2.1 among other techniques to handle challenges unique to MinEMax k-center.

We begin by describing the LP. Define f_B as the function that truncates at B, i.e.,

$$f_B(d) := \begin{cases} d & \text{if } d \geq B \\ 0 & \text{otherwise.} \end{cases}$$

Let X^* be an optimal solution to MinEMax k-center, let X_T^* be the optimal solution to the corresponding TruncatedTwoStage problem. Given B as a guess of $B(X_T^*)$, our LP has a variable x_i indicating if i is a center and a variable z_{is} indicating the extent to which we assign s to i.

$$\begin{aligned} \min \quad & \sum_s p_s \cdot \Big(\sum_i f_B(c_{is}) \cdot z_{is} \Big) && (P_B) \\ \text{s.t.} \quad & \sum_i z_{is} \geq 1, \quad \forall s \\ & 0 \leq z_{is} \leq x(i), \quad \forall i, \forall s \\ & \sum_i x_i \leq k \end{aligned}$$

Let val(B) be the optimal value of P_B given parameter B. Let OPT be the cost of the optimal solution for the input MinEMax k-center problem.

Although we would like to use $B(X_T^*)$ as our value for B in P_B, we do not know $B(X_T^*)$. For this reason, in the following lemma we argue how to efficiently compute a value that, up to constants, works as well. The proof of this lemmas is deferred to §B; roughly, the idea is to take B as the best power of $(1+\epsilon)$.

Lemma 4.2. *There exists a polynomial-time algorithm which for a given $\epsilon > 0$ and an input instance of MinEMax k-center, returns a $\hat{B}$ such that $\hat{B} \leq (1+\epsilon)\left(\frac{6 \cdot OPT}{1-1/e}\right)$ and val$(\hat{B}) \leq \left(\frac{6 \cdot OPT}{1-1/e}\right)$. The algorithm's runtime is polynomial in n, m, and $\log_{1+\epsilon}$ OPT.*

Given a good value of B, we can describe our algorithm in full. Our algorithm first computes $\hat{B}$ as in Lemma 4.2. It next uses $P_{\hat{B}}$ to cluster clients. Let

$$P_{\hat{B}}(s) := \sum_i f_{\hat{B}}(c_{is}) z_{is}$$

be the cost of the scenario with client s in $P_{\hat{B}}$. We sort scenarios in increasing order of $P_{\hat{B}}(s)$. For each client s' initialize $\mathcal{P}_{s'}$ to 0. Next, iterate through the clients. For client s', if there exists a client s s.t. $c_{ss'} \leq 2\hat{B}$ with $\mathcal{P}_s > 0$, then increment $\mathcal{P}_s$ by $p_{s'}$. Otherwise, set $\mathcal{P}_{s'}$ to $p_{s'}$. Let $\mathcal{D}' := \{s \in \mathcal{D} : \mathcal{P}_s > 0\}$ and let $\sigma : \mathcal{D} \to \mathcal{D}'$ be a function where $\sigma(s')$ denotes the client to whom we move client s''s probability mass.

Now consider the weighted k-median instance consisting of clients $\mathcal{D}'$ with distances $\{c_{ij}\}$, where $s' \in \mathcal{D}'$ has weight $\mathcal{P}_{s'}$ and where one can choose centers only at points in $\mathcal{D}'$. Call this instance $\mathcal{MED}$. Notice that weighted k-median can be reduced to unweighted k-median by just duplicating points and scaling costs by the appropriate amount. Run any α-approximation for k-median on $\mathcal{MED}$ and return the output as our solution to the input MinEMax k-center problem.

We now prove that our algorithm achieves a constant approximation in polynomial-time. Henceforth, let (x, z) denote an optimal solution to $P_{\hat{B}}$ and as before let val$(\hat{B})$ denote the value of this solution. Moreover, let val$'(\hat{B}) := \sum_{s' \in \mathcal{D}'} \mathcal{P}_{s'} \cdot P_{\hat{B}}(s')$ be the cost of (x, z) applied to $\mathcal{D}'$.

We first show that clients in $\mathcal{MED}$ are far apart which will allow us to argue that truncating distances at $\hat{B}$ does not affect distances.

Lemma 4.3. *If $s'_1, s'_2 \in \mathcal{D}'$ then $c_{s'_1 s'_2} > 2\hat{B}$.*

Proof. WLOG suppose that s'_1 is considered before s'_2 in the clustering above. Moreover, suppose for the sake of contradiction that $c_{s'_1 s'_2} \leq 2\hat{B}$. Notice that when we examine s'_2 we will assign s'_2 to s'_1; i.e., $\sigma(s'_2) = s'_1$. However, it follows that $\mathcal{P}_{s'_1} = 0$ and as such $s'_1 \notin \mathcal{D}'$ by definition of $\mathcal{D}'$, a contradiction. □

We next show that the value of our LP solution to $P_{\hat{B}}$ only decreases in cost when applied to $\mathcal{D}'$.

Lemma 4.4. *$val'(\hat{B}) \leq val(\hat{B})$.*

Proof. When we cluster points in increasing order of $P_{\hat{B}}(s)$, we have that the cluster to which any given client is reassigned is always of lesser cost in $P_{\hat{B}}$; i.e., if $s \in \sigma^{-1}(s')$ then

$$\sum_i f_{\hat{B}}(c_{is'}) \cdot z_{is'} \leq \sum_i f_{\hat{B}}(c_{is}) \cdot z_{is}. \tag{24}$$

Thus, we can get val($\hat{B}$) equals

$$\begin{aligned}
\sum_{s \in \mathcal{D}} p_s \cdot \sum_i f_{\hat{B}}(c_{is}) \cdot z_{is} &= \sum_{s' \in \mathcal{D}'} \Big[\sum_{s \in \sigma^{-1}(s')} p_s \cdot \sum_i f_{\hat{B}}(c_{is}) \cdot z_{is} \Big] \\
&\geq \sum_{s' \in \mathcal{D}'} \Big[\sum_{s \in \sigma^{-1}(s')} p_s \sum_i f_{\hat{B}}(c_{is'}) \cdot z_{is'} \Big] && \text{(by Eq. (24))} \\
&= \sum_{s' \in \mathcal{D}'} \Big[\Big(\sum_i f_{\hat{B}}(c_{is'}) \cdot z_{is'} \Big) \cdot \sum_{s \in \sigma^{-1}(s')} p_s \Big] \\
&= \sum_{s' \in \mathcal{D}'} \Big[\Big(\sum_i f_{\hat{B}}(c_{is'}) \cdot z_{is'} \Big) \cdot \mathcal{P}_{s'} \Big] \\
&= \text{val}'(\hat{B}).
\end{aligned}$$

□

Next we show that there exists a solution to $\mathcal{MED}$ of cost about val($\hat{B}$).

Lemma 4.5. *The optimal solution to $\mathcal{MED}$ has cost at most $O(val(\hat{B}))$.*

Proof. Again, let (x, z) be the optimal solution to $P_{\hat{B}}$. We prove this lemma by constructing a fractional solution (x', z') to the LP of $\mathcal{MED}$ of value at most $2 \cdot \text{val}(\hat{B})$. The LP of $\mathcal{MED}$ is as follows and has variables analogous to P_B (we overload (x', z') here to stand for the variables in our LP along with the feasible solution for our LP that we will construct).

$$\min \quad \sum_{s'_1 \in \mathcal{D}'} \mathcal{P}_s \Big(\sum_{s'_2 \in \mathcal{D}'} c_{s'_2 s'_1} \cdot z'_{s'_2 s'_1} \Big) \tag{k-M LP}$$

$$\text{s.t.} \quad \sum_{s'_2 \in \mathcal{D}'} z'_{s'_2 s'_1} \geq 1, \qquad \forall s'_1 \tag{25}$$

$$0 \leq z'_{s'_2 s'_1} \leq x'(s'_2), \qquad \forall s'_2, \forall s'_1 \tag{26}$$

$$\sum_{s' \in \mathcal{D}'} x'(s') \leq k \tag{27}$$

To construct (x', z'), we do the following. We first define a new clustering such that every point in $\mathcal{D}$ goes to the closest point in $\mathcal{D}'$. Formally, for $s' \in \mathcal{D}'$, we define $F_{s'} := \{i \in \mathcal{D} : s' = \arg\min_{s' \in \mathcal{D}'} c_{is'}\}$. Intuitively, our solution (x', z') reroutes services that were provided by facilities in $F_{s'}$ to s's.

Let $x'(s') := \sum_{i \in F_{s'}} x(i)$. That is, we open a facility at s' by summing up the facilities in $P_{\hat{B}}$ that were clustered to s'. We let $z'_{s'_1 s'_2} = \sum_{i \in F_{s'_1}} z_{is'_2}$. That is, we assign s'_2 to s'_1 to the extent that $P_{\hat{B}}$ assigned s'_2 to points clustered with s'_1.

We now prove the feasibility of (x', z') for k-M LP. Since for every client s in $P_{\hat{B}}$ we know $\sum_i z_{is} \geq 1$, we have that every client in $\mathcal{D}'$ is serviced: For every client $s'_1 \in \mathcal{D}'$ it holds that $\sum_{s'_2 \in \mathcal{D}'} z'_{s'_2 s'_1} = \sum_{s'_2 \in \mathcal{D}'} \sum_{i \in F_{s'_2}} z_{is'_1} = \sum_i z_{is'_1} \geq 1$. Moreover, we open no more than k centers fractionally since in $P_{\hat{B}}$ we have that $\sum_i x(i) \leq k$ and

$$\sum_{s' \in \mathcal{D}'} x'(s') \quad = \quad \sum_{s' \in \mathcal{D}'} \sum_{i \in F_{s'}} x(i) \quad = \quad \sum_i x(i) \quad \leq \quad k.$$

Also, no client is serviced by an unopened facility: in $P_{\hat{B}}$ we have that $0 \leq z_{is} \leq x(i)$. Hence, for any $s'_1, s'_2 \in \mathcal{D}'$ we have $z'_{s'_1 s'_2} = \sum_{i \in F_{s'_1}} z_{is'_2} \leq \sum_{i \in F_{s'_1}} x(i) = x'(s'_1)$.

Lastly, we bound the objective value of (x', y'). By Lemma 4.3 we know that for $s'_1 \neq s'_2$, where $s'_1, s'_2 \in \mathcal{D}'$, it holds that $c_{s'_1 s'_2} > 2\hat{B}$. We first show that if $i \in F_{s'_2}$ then $c_{is'_1} > \hat{B}$. This is because if $c_{is'_1} \leq \hat{B}$ then $c_{is'_2} \leq \hat{B}$ since s'_2 is the closest point in $\mathcal{D}'$ to i by $i \in F_{s'_2}$. Now by triangle inequality, we get $c_{s'_1 s'_2} \leq c_{s'_1 i} + c_{is'_2} \leq 2\hat{B}$, which is a contradiction to Lemma 4.3. Hence, $c_{is'_1} > \hat{B}$ implies $f_{\hat{B}}(c_{is'_1}) = c_{is'_1}$, and

$$c_{s'_1 s'_2} \leq 2c_{is'_1} = 2 \cdot f_{\hat{B}}(c_{is'_1}), \tag{28}$$

where $i \in F_{s'_2}$ for $s'_2 \neq s'_1$ and $s'_1, s'_2 \in \mathcal{D}'$.

Thus, the value of (x', z') in k-M LP is

$$\begin{aligned}
\sum_{s'_1 \in \mathcal{D}'} \mathcal{P}_{s'_1} \left(\sum_{s'_2 \in \mathcal{D}'} c_{s'_2 s'_1} \cdot z'_{s'_2 s'_1} \right) &= \sum_{s'_1, s'_2 \in \mathcal{D}'} \mathcal{P}_{s'_1} \cdot c_{s'_1 s'_2} \cdot z'_{s'_2 s'_1} \\
&\leq \sum_{s'_1, s'_2 \in \mathcal{D}' : s'_1 \neq s'_2} \mathcal{P}_{s'_1} \cdot c_{s'_1 s'_2} \cdot z'_{s'_2 s'_1} && (\text{by } \ c_{ss} = 0) \\
&= \sum_{s'_1, s'_2 \in \mathcal{D}' : s'_1 \neq s'_2} \sum_{i \in F_{s'_2}} \mathcal{P}_{s'_1} \cdot c_{s'_1 s'_2} \cdot z_{is'_1} && \left(\text{by definition of } z'_{s'_1 s'_2}\right) \\
&\leq 2 \cdot \sum_{s'_1, s'_2 \in \mathcal{D}' : s'_1 \neq s'_2} \sum_{i \in F_{s'_2}} \mathcal{P}_{s'_1} \cdot f_{\hat{B}}(c_{is'_1}) \cdot z_{is'_1} && (\text{by Eq. (28)}) \\
&\leq 2 \cdot \sum_{s'_1 \in \mathcal{D}'} \sum_i \mathcal{P}_{s'_1} \cdot f_{\hat{B}}(c_{is'_1}) \cdot z_{is'_1} \\
&= 2 \cdot \mathrm{val}'(\hat{B}) \\
&\leq 2 \cdot \mathrm{val}(\hat{B}) && (\text{by Lemma 4.4}).
\end{aligned}$$

Thus, since (x', z') is feasible and has cost at most $2\mathrm{val}(\hat{B})$ we know that the optimal solution to k-M LP has cost at most $2\mathrm{val}(\hat{B})$. Moreover, since past work has demonstrated that k-M LP has a constant integrality gap—e.g.[CGTS02] shows it is at most 20/3—we conclude that the optimal solution to $\mathcal{MED}$ has cost at most $O(\mathrm{val}(\hat{B}))$. □

Next, we show that any solution to $\mathcal{MED}$ is a good solution to our MinEMax k-center problem.

Lemma 4.6. *An integer solution to $\mathcal{MED}$ of cost C solves the input* MinEMax *k-center problem with cost at most $C + 4 \cdot \hat{B}$.*

Proof. Observe that the primary difference between MinEMax k-center problem and $\mathcal{MED}$ is that the former is on $\mathcal{D}$ while the latter is on $\mathcal{D}'$. Roughly, this lemma is true because any point in $\mathcal{D}$ is at most $2\hat{B}$ from a point in $\mathcal{D}'$.

Let X be our $\mathcal{MED}$ solution of cost C. We have that the cost of X as an MinEMax solution is

$$\begin{aligned}
\mathbb{E}_A[\max_{s \in A}\{d(s,F)\}] &\leq \sum_{s \in M(F)} p_s \cdot d(s,F) && \text{(by Lemma 2.1)} \\
&= \sum_{s' \in \mathcal{D}'} \sum_{s \in \sigma^{-1}(s') \cap M(F)} p_s \cdot d(s,F) \\
&\leq \sum_{s' \in \mathcal{D}'} \sum_{s \in \sigma^{-1}(s') \cap M(F)} p_s \cdot (c_{ss'} + d(s',F)) && \text{(by triangle inequality)} \\
&\leq \sum_{s' \in \mathcal{D}'} \sum_{s \in \sigma^{-1}(s') \cap M(F)} p_s \cdot \left(2\hat{B} + d(s',F)\right) && \left(\text{by } c_{ss'} \leq 2\hat{B} \text{ if } \sigma(s) = s'\right) \\
&\leq 4\hat{B} + \sum_{s' \in \mathcal{D}'} (d(s',F)) \sum_{s \in \sigma^{-1}(s') \cap M(F)} p_s && \left(\text{by } \sum_{s \in M(F)} p_s < 2\right) \\
&\leq 4\hat{B} + \sum_{s' \in \mathcal{D}'} (d(s',F))\, \mathcal{P}_{s'} && \text{(by definition of } \mathcal{P}_{s'}) \\
&\leq 4\hat{B} + C && \text{(by definition of } C). \quad \square
\end{aligned}$$

Lastly, we conclude the approximation factor of our algorithm.

Theorem 4.7. MinEMax *k-center can be $O(1)$-approximated in polynomial time.*

Proof. By Lemma 4.5 the optimal solution to $\mathcal{MED}$ has cost at most $O(\text{val}(\hat{B}))$. Now, applying our α-approximation algorithm for k-median to $\mathcal{MED}$ results in an integer solution of cost at most $\alpha \cdot O(\text{val}(\hat{B}))$. By Lemma 4.6 such an integer solution solves MinEMax k-center with cost at most $\alpha \cdot O(\text{val}(\hat{B})) + 4\hat{B}$. Applying Lemma 4.2 we then have that our solution costs at most

$$\alpha \cdot O\left(\left(\frac{6}{1-1/e}\right) \cdot \text{OPT}\right) + 4(1+\epsilon)\left(\frac{6 \cdot \text{OPT}}{1-1/e}\right).$$

Letting ϵ be any constant > 0 and using an $\alpha = O(1)$ approximation algorithm for k-median—e.g., [CGTS02]—we conclude that our solution has cost at most $O(OPT)$.

Lastly, we argue that our algorithm runs in polynomial time. Solving our LPs and performing clustering are trivially poly-time. Running our α-approximation for k-median is poly-time by assumption, so we conclude the polynomial runtime of our algorithm. □

5 Reducing Stochastic, Demand-Robust, and Hybrid to MinEMax

In this section we show how to use an α-approximation algorithm in the MinEMax model to design an α-approximation in the stochastic, the demand-robust, and the Hybrid models (Theorem 1.3). We believe our Hybrid model gives a clean way of modeling inaccuracy of the input distribution in a stochastic two-stage setting.

Our New Hybrid **Model.** In our Hybrid two-stage covering model—as in our stochastic model—we are given a distribution $\mathcal{D}$ over scenarios from which exactly one scenario realizes. However, we are also given a *caution parameter* ρ which specifies the inaccuracy of $\mathcal{D}$. A low value of ρ signals that the second-stage realization is expected to be close to the stochastic model while a large ρ signals that any scenario may occur. Specifically, in the second-stage w.p. $(1-\rho)$ the scenario we must pay for realizes from the input distribution, and w.p. ρ it is chosen adversarially. Thus, the cost for solution $(X_1, \boldsymbol{X_2})$ is a convex combination of the stochastic and robust objectives using ρ:

$$\mathrm{cost}_{\mathsf{Hyb}}(X_1, \boldsymbol{X_2}) := \rho \cdot \mathrm{cost}_{\mathsf{Rob}}(X_1, \boldsymbol{X_2}) + (1-\rho) \cdot \mathrm{cost}_{\mathsf{Stoch}}(X_1, \boldsymbol{X_2}). \tag{29}$$

Note that for $\rho = 0$ we recover the stochastic model and for $\rho = 1$ we recover the demand-robust model. Thus, for $\rho \in (0,1)$, the Hybrid model smoothly interpolates between these two models and so it suffices to prove Theorem 1.3 for only the Hybrid model. Recall that the theorem that we will prove is as follows.

Theorem 1.3. *An α-approximation for a two-stage covering algorithm in the* MinEMax *model implies an α-approximation for the corresponding two-stage covering problem in the stochastic, demand-robust, and* Hybrid *models.*

Theorem 1.3 Intuition. The main idea of our reduction from a Hybrid problem, P_{Hyb}, to a MinEMax problem, P_{EMax}, is as follows. For each scenario in our original Hybrid problem we create two scenarios in our MinEMax problem; one to represent the demand-robust cost of this scenario and one to represent its stochastic cost. The scenario to represent the robust cost has probability one but its cost is dampened by ρ. The scenario to represent the stochastic cost has a sufficiently low probability so that the independent Bernoulli trials of these scenarios are effectively disjoint. Moreover, the costs of the scenarios to represent the stochastic costs are inflated to make up for the dampened probability, and to ensure that they are more expensive than the demand-robust scenarios.

Our Reduction

Let P_{Hyb} be a Hybrid problem with distribution $\mathcal{D}$ over scenarios $\{\overline{S}_1, \dots, \overline{S}_m\}$, a caution parameter ρ, and a linear cost function

$$\mathrm{cost}_{\mathsf{Hyb}}(X_1, X_2) = \mathrm{cost}^1_{\mathsf{Hyb}}(X_1) + \mathrm{cost}^2_{\mathsf{Hyb}}(X_2),$$

where $\mathrm{cost}^1_{\mathsf{Hyb}}(X_1) := \mathrm{cost}_{\mathsf{Hyb}}(X_1, \emptyset)$ and $\mathrm{cost}^2_{\mathsf{Hyb}}(X_2) := \mathrm{cost}_{\mathsf{Hyb}}(\emptyset, X_2)$. We produce a MinEMax instance P_{EMax} with $2m$ scenarios $S_1, S_2, \dots, S_{2m}$ where the first m scenarios encode the demand-robust cost and the last m scenarios encode the stochastic cost. For $s \in [m]$, the covering constraints for S_s and S_{m+s} are the same as $\overline{S}_s$. We set the cost function

$$\mathrm{cost}_{\mathsf{EMax}}(X_1, X_2^{(s)}) := \mathrm{cost}^1_{\mathsf{EMax}}(X_1) + \mathrm{cost}^2_{\mathsf{EMax}}(X_2^{(s)})$$

such that the first stage cost is the same as the Hybrid first stage, i.e., $\mathrm{cost}^1_{\mathsf{EMax}}(X_1) = \mathrm{cost}^1_{\mathsf{Hyb}}(X_1)$. In the second stage for $s \in [m]$, we set

$$\mathrm{cost}^2_{\mathsf{EMax}}(X_2^{(s)}) := \rho \cdot \mathrm{cost}^2_{\mathsf{Hyb}}(X_2^{(s)})$$

$$\text{and} \quad \mathrm{cost}^2_{\mathsf{EMax}}(X_2^{(m+s)}) := \gamma(1-\rho) \cdot \mathrm{cost}^2_{\mathsf{Hyb}}(X_2^{(s)}),$$

where $\gamma \geq 1$ is a sufficiently large scaling factor for our stochastic scenarios. In particular, γ satisfies that for all $s, s' \in [m]$,

$$\gamma(1-\rho) \cdot \mathrm{cost}^2_{\mathsf{Hyb}}(X_2^{(s)}) > \rho \cdot \mathrm{cost}^2_{\mathsf{Hyb}}(X_2^{(s')}). \tag{30}$$

For $s \in [m]$, we set the probabilities in P_{EMax} to be $p_s := 1$ and $p_{m+s} = \frac{\mathcal{D}(s)}{\gamma}$.

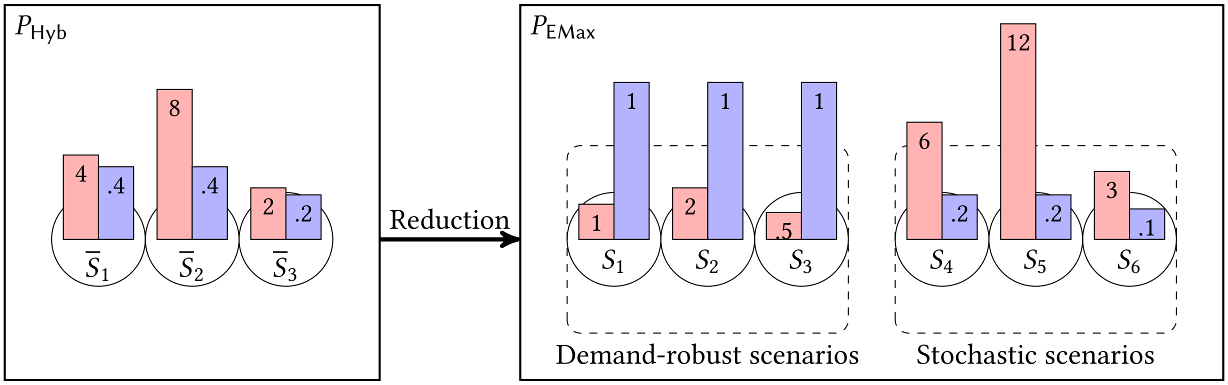

Figure 3: Reduction from P_{Hyb} for $\rho = .25$ to P_{EMax} for $m = 3$. Circles: scenarios. Costs are drawn for fixed solution $(X_1, \overline{\boldsymbol{X}}_{\boldsymbol{2}})$ for P_{Hyb} and its corresponding solution for P_{EMax}, $(X_1, \boldsymbol{X}_{\boldsymbol{2}})$ where for $s \equiv s' \mod m$ we define $X_2^{(s)} := \overline{X}_2^{(s')}$. Red rectangles in P_{Hyb}: $\mathrm{cost}_{\mathsf{Hyb}}(X_1, \overline{X}_2^{(s)})$. Red rectangles in P_{EMax}: $\mathrm{cost}_{\mathsf{EMax}}(X_1, X_2^{(s)})$. Blue rectangles in P_{Hyb}: $\mathcal{D}(s)$ of $\overline{S}_s$. Blue rectangles in P_{EMax}: p_s. $\gamma = 2$ and $\mathrm{cost}^1_{\mathsf{Hyb}}(X_1) = \mathrm{cost}^1_{\mathsf{EMax}}(X_1) = 0$. Demand-robust scenarios have high probability and stochastic scenarios have high cost.

Notice that the stochastic copy of a scenario always costs more than the demand-robust one, i.e., $\mathrm{cost}_{\mathsf{EMax}}(X_1, X_2^{(s')}) > \mathrm{cost}_{\mathsf{EMax}}(X_1, X_2^{(s)})$ for $s \le m$ and $s' > m$. See Figure 3 for an illustration.

Our Proof of Theorem 1.3

Observe that by linearity of costs in MinEMax, our first stage and second stage costs can be separated. This implies for any solution $(X_1, \boldsymbol{X}_{\boldsymbol{2}})$ to P_{EMax}, we have

$$\mathrm{cost}_{\mathsf{EMax}}(X_1, \boldsymbol{X}_{\boldsymbol{2}}) = \mathrm{cost}^1_{\mathsf{EMax}}(X_1) + \mathbb{E}_A\Big[\max_{s \in A}\Big\{\mathrm{cost}^2_{\mathsf{EMax}}(X_2^{(s)})\Big\}\Big]. \tag{31}$$

To prove Theorem 1.3, in Lemma 5.1 we show that for our reduction the optimal P_{EMax} solution costs less than the optimal Hybrid solution. We defer the proof of this lemma to §C. In Lemma 5.2 we show the converse direction (up to a small factor).

Lemma 5.1. *Let $(H_1, \overline{\boldsymbol{H}}_{\boldsymbol{2}})$ be the optimal solution to P_{Hyb} and let $(H_1, \boldsymbol{H}_{\boldsymbol{2}})$ be its natural interpretation in P_{EMax}, i.e., for $s \equiv s' \mod m$ we define $H_2^{(s)} := \overline{H}_2^{(s')}$. We have $\mathrm{cost}_{EMax}(H_1, \boldsymbol{H}_{\boldsymbol{2}}) \le \mathrm{cost}_{Hyb}(H_1, \overline{\boldsymbol{H}}_{\boldsymbol{2}})$.*

The main idea of the proof is to use a union bound over the stochastic scenarios.

Next, we show that any solution costs more in P_{EMax} than in P_{Hyb} (up to small multiplicative factors).

Lemma 5.2. *Let $(X_1, \boldsymbol{X}_{\boldsymbol{2}})$ be any solution to P_{EMax} and let $(X_1, \overline{\boldsymbol{X}}_{\boldsymbol{2}})$ be its natural interpretation in P_{Hyb}, i.e., $\overline{X}_2^{(s)} := \arg\min\Big\{\mathrm{cost}^2_{Hyb}(X_2^{(s)}), \mathrm{cost}^2_{Hyb}(X_2^{(s+m)})\Big\}$. We have $\mathrm{cost}_{Hyb}(X_1, \overline{\boldsymbol{X}}_{\boldsymbol{2}}) \le \Big(1 - \frac{m}{\gamma}\Big)^{-1} \cdot \mathrm{cost}_{EMax}(X_1, \boldsymbol{X}_{\boldsymbol{2}})$.*

Proof. Notice that by the definition of $\overline{X}_2$, we know

$$\mathrm{cost}_{\mathsf{Hyb}}(X_1, \overline{X}_2^{(s)}) = \min\Big\{\mathrm{cost}_{\mathsf{Hyb}}(X_1, X_2^{(s)}), \mathrm{cost}_{\mathsf{Hyb}}(X_1, X_2^{(s+m)})\Big\}. \tag{32}$$

We first lower bound the second stage cost of MinEMax using the inclusion-exclusion principle on the number of stochastic scenarios that realize into A. Recollect, A always contains the first m demand-robust scenarios because their probability is 1. Since a stochastic scenario (if it realizes) always costs more than all the demand-robust scenarios by Eq. (30), we get

$$\mathbb{E}_A\Big[\max_{s\in A}\Big\{\mathrm{cost}^2_{\mathsf{EMax}}(X_2^{(s)})\Big\}\Big] \geq \Pr_A[A=[m]]\cdot\max_{s\in[m]}\Big\{\mathrm{cost}^2_{\mathsf{EMax}}(X_2^{(s)})\Big\} + \sum_{s>m}\Pr_A[s\in A]\cdot\mathrm{cost}^2_{\mathsf{EMax}}(X_2^{(s)})$$
$$-\sum_{s,s'>m}\Pr_A[s\in A]\Pr_A[s'\in A]\cdot\max\Big\{\mathrm{cost}^2_{\mathsf{EMax}}(X_2^{(s)}),\mathrm{cost}^2_{\mathsf{EMax}}(X_2^{(s')})\Big\}.$$

Now using $\Pr_A[A=[m]]=\prod_s\left(1-\frac{\mathcal{D}(s)}{\gamma}\right)\geq(1-\frac{1}{\gamma})^m$ and the definition of $\mathrm{cost}^2_{\mathsf{EMax}}$, we get

$$\mathbb{E}_A\Big[\max_{s\in A}\Big\{\mathrm{cost}^2_{\mathsf{EMax}}(X_2^{(s)})\Big\}\Big] \geq \left(1-\frac{1}{\gamma}\right)^m\cdot\rho\max_{s\in[m]}\Big\{\mathrm{cost}^2_{\mathsf{Hyb}}(X_2^{(s)})\Big\}+\gamma(1-\rho)\sum_{s>m}p_s\cdot\mathrm{cost}^2_{\mathsf{Hyb}}(X_2^{(s)})$$
$$-\gamma(1-\rho)\cdot\sum_{s,s'>m}p_sp_{s'}\cdot\max\Big\{\mathrm{cost}^2_{\mathsf{Hyb}}(X_2^{(s)}),\mathrm{cost}^2_{\mathsf{Hyb}}(X_2^{(s')})\Big\}. \tag{33}$$

Since for any $a,b\geq 0$, we have $\max\{a,b\}\leq a+b$, we can bound

$$\sum_{s,s'>m}p_sp_{s'}\max\Big\{\mathrm{cost}^2_{\mathsf{Hyb}}(X_2^{(s)}),\mathrm{cost}^2_{\mathsf{Hyb}}(X_2^{(s')})\Big\}\leq 2\sum_{s,s'>m}p_sp_{s'}\mathrm{cost}^2_{\mathsf{Hyb}}(X_2^{(s)})=\frac{2}{\gamma}\cdot\sum_{s>m}p_s\mathrm{cost}^2_{\mathsf{Hyb}}(X_2^{(s)}),$$

where the last equality uses $\sum_{s'>m}p_{s'}=\sum_{s'>m}\frac{\mathcal{D}(s')}{\gamma}=\frac{1}{\gamma}$. Combining this with Eq. (32) and Eq. (33),

$$\begin{aligned}
&\mathbb{E}_A\Big[\max_{s\in A}\Big\{\mathrm{cost}^2_{\mathsf{EMax}}(X_2^{(s)})\Big\}\Big]\\
&\geq\left(1-\frac{1}{\gamma}\right)^m\cdot\rho\max_{s\in[m]}\Big\{\mathrm{cost}^2_{\mathsf{Hyb}}(X_2^{(s)})\Big\}+\gamma(1-\rho)\Big(\sum_{s>m}p_s\cdot\mathrm{cost}^2_{\mathsf{Hyb}}(X_2^{(s)})-\frac{2}{\gamma}\sum_{s>m}p_s\cdot\mathrm{cost}^2_{\mathsf{Hyb}}(X_2^{(s)})\Big)\\
&\geq\left(1-\frac{m}{\gamma}\right)\cdot\rho\max_{s\in[m]}\Big\{\mathrm{cost}^2_{\mathsf{Hyb}}(\overline{X}_2^{(s)})\Big\}+(1-\rho)\Big(1-\frac{m}{\gamma}\Big)\cdot\mathbb{E}_{s\sim\mathcal{D}}\Big[\mathrm{cost}^2_{\mathsf{Hyb}}(\overline{X}_2^{(s)})\Big],
\end{aligned}$$

where the last inequality uses $(1-\frac{1}{\gamma})^m\geq(1-\frac{m}{\gamma})$ along with $(1-\frac{2}{\gamma})\geq(1-\frac{m}{\gamma})$ and $\gamma p_s=\mathcal{D}(s)$. Our upper bound follows when we combine the last inequality with Eq. (31) to get

$$\begin{aligned}
\mathrm{cost}_{\mathsf{EMax}}(X_1,\boldsymbol{X_2})&\geq\mathrm{cost}^1_{\mathsf{EMax}}(X_1)+\left(1-\frac{m}{\gamma}\right)\Big(\rho\max_{s\in[m]}\Big\{\mathrm{cost}^2_{\mathsf{Hyb}}(\overline{X}_2^{(s)})\Big\}+(1-\rho)\mathbb{E}_{s\sim\mathcal{D}}\Big[\mathrm{cost}^2_{\mathsf{Hyb}}(\overline{X}_2^{(s)})\Big]\Big)\\
&\geq\left(1-\frac{m}{\gamma}\right)\Big(\rho\max_{i\in[m]}\Big\{\mathrm{cost}_{\mathsf{Hyb}}(X_1,\overline{X}_2^{(s)})\Big\}+(1-\rho)\cdot\mathbb{E}_{s\sim\mathcal{D}}\Big[\mathrm{cost}_{\mathsf{Hyb}}(X_1,\overline{X}_2^{(s)})\Big]\Big)\\
&=\left(1-\frac{m}{\gamma}\right)\cdot\mathrm{cost}_{\mathsf{Hyb}}(X_1,\overline{\boldsymbol{X_2}}).
\end{aligned}$$

□

We now conclude the proof of the main theorem by combining the α-approximation algorithm with Lemma 5.1 and Lemma 5.2 which allow us to move between MinEMax and Hybrid via our reduction.

Proof of Theorem 1.3. Since Hybrid captures the stochastic and demand-robust models it suffices to prove the theorem only for Hybrid.

Consider an input Hybrid problem P_{Hyb} with optimal solution $(H_1,\overline{\boldsymbol{H_2}})$. Suppose we have an α-approx algorithm for MinEMax. To design an α-approx for P_{Hyb}, we simply run our reduction to get an instance of MinEMax problem P_{EMax}, run our MinEMax approximation algorithm on P_{EMax} to get back $(A_1,\boldsymbol{A_2})$, and then return $(A_1,\overline{\boldsymbol{A_2}})$ where $\overline{A_2}$ is the natural interpretation of A_2 as a solution for P_{Hyb}. In particular,

$$\overline{A}_2^{(s)}:=\arg\min\Big\{\mathrm{cost}^2_{\mathsf{Hyb}}(A_2^{(s)}),\mathrm{cost}^2_{\mathsf{Hyb}}(A_2^{(s+m)})\Big\}.$$

Now consider the cost of our returned solution for P_{Hyb}. Let $(H_1,\boldsymbol{H_2})$ be its natural interpretation in P_{EMax}, i.e., for $s \equiv s' \mod m$ we define $H_2^{(s)} := \overline{H}_2^{(s')}$. By Lemma 5.2 we know

$$\mathsf{cost}_{\mathsf{Hyb}}(A_1, \overline{\boldsymbol{A_2}}) \leq \left(1 - \frac{m}{\gamma}\right)^{-1} \mathsf{cost}_{\mathsf{EMax}}(A_1, \boldsymbol{A_2}).$$

Let $(E_1, \boldsymbol{E_2})$ be the optimal solution of MinEMax. Since $(A_1, \boldsymbol{A_2})$ is an α-approximation,

$$\mathsf{cost}_{\mathsf{EMax}}(A_1, \boldsymbol{A_2}) \leq \alpha \cdot \mathsf{cost}_{\mathsf{EMax}}(E_1, \boldsymbol{E_2}).$$

Combining the above two equations, we get

$$\mathsf{cost}_{\mathsf{Hyb}}(A_1, \overline{\boldsymbol{A_2}}) \leq \alpha \cdot \left(1 - \frac{m}{\gamma}\right)^{-1} \mathsf{cost}_{\mathsf{EMax}}(E_1, \boldsymbol{E_2}),$$

Since $(E_1, \boldsymbol{E_2})$ is the optimal solution of MinEMax, $\mathsf{cost}_{\mathsf{EMax}}(E_1, \boldsymbol{E_2}) \leq \mathsf{cost}_{\mathsf{EMax}}(H_1, \boldsymbol{H_2})$. By Lemma 5.1,

$$\mathsf{cost}_{\mathsf{EMax}}(H_1, \boldsymbol{H_2}) \leq \mathsf{cost}_{\mathsf{Hyb}}(H_1, \boldsymbol{H_2}).$$

Thus, $(A_1, \overline{\boldsymbol{A_2}})$ is an $\alpha \cdot \left(1 - \frac{m}{\gamma}\right)^{-1}$-approx for P_{Hyb}. Choosing $\gamma \to \infty$ and noting that the reduction is polynomial-time, the theorem follows. □

Appendix

A Deferred Proofs of §2

Lemma 2.1. *Let $\boldsymbol{Y} = \{Y_1, \ldots, Y_m\}$ be a set of independent Bernoulli r.v.s, where Y_s is 1 with probability p_s, and 0 otherwise. Let $v_s \in \mathbb{R}_{\geq 0}$ be a value associated with Y_s. WLOG assume $v_s \geq v_{s+1}$ for $s \in [m-1]$. Let $b = \min_{a: \sum_{s=1}^{a} p_s \geq 1} a$. Then*

$$\left(\frac{1 - 1/e}{2}\right)\left(v_b + \sum_s p_s \cdot (v_s - v_b)^+\right) \leq \mathbb{E}_{\boldsymbol{Y}}\left[\max_s \{Y_s \cdot v_s\}\right] \leq v_b + \sum_s p_s \cdot (v_s - v_b)^+,$$

where $x^+ := \max\{x, 0\}$.

Proof. We begin by showing the lower bound on $\mathbb{E}_{A \sim \boldsymbol{Y}}\left[\max_{s \in A} v_s\right]$. Notice that $1 \leq \sum_{s \leq b} p_s < 2$.

Let $M := [b]$. The probability that no element of M is in A is

$$\prod_{s \in M} (1 - p_s) \quad \leq \quad e^{-\sum_{s \in M} p_s} \quad = \quad \frac{1}{e}$$

because $1 - x \leq e^{-x}$ and $\sum_{s \in M} p_s \geq 1$. It follows that

$$\begin{aligned}
\mathbb{E}_{A \sim \boldsymbol{Y}}\left[\max_{s \in A} v_s\right] &\geq \left(1 - \frac{1}{e}\right)\mathbb{E}_{A \sim \boldsymbol{Y}}\left[\max_{s \in A} v_s \mid \text{at least 1 element from } M \text{ in } A\right] \\
&\geq \left(1 - \frac{1}{e}\right)E_{A \sim \boldsymbol{Y}}\left[\max_{s \in A} v_s \mid \text{exactly 1 element from } M \text{ in } A\right] \\
&= \left(1 - \frac{1}{e}\right)\sum_{s \in M} v_s \frac{p_s}{\sum_{i \in M} p_s} \quad > \left(\frac{1 - 1/e}{2}\right)\sum_{s \in M} p_s v_s \qquad \left(\text{by } \sum_{s \in M} p_s < 2\right).
\end{aligned}$$

Thus, we have that

$$\begin{aligned}
\mathbb{E}_{A \sim \boldsymbol{Y}}\left[\max_{s \in A} v_s\right] &> \left(\frac{1 - 1/e}{2}\right)\sum_{s \in M} p_s v_s \\
&= \left(\frac{1 - 1/e}{2}\right)\sum_{s \in M} p_s\left((v_s - v_b)^+ + v_b\right) \qquad (\text{by } v_s \geq v_b \text{ for } s \in M)
\end{aligned}$$

$$\geq \left(\frac{1-1/e}{2}\right)\left(v_b + \sum_{s\in M} p_s\left((v_s - v_b)^+\right)\right) \qquad \left(\text{by } 1 \leq \sum_{s\in M} p_s\right)$$

$$= \left(\frac{1-1/e}{2}\right)\left(v_b + \sum_{s} p_s\left((v_s - v_b)^+\right)\right) \qquad (\text{by } v_s > v_b \text{ iff } s \in M)$$

as our lower bound.

We now show the upper bound. Recall $x^+ := \max(x, 0)$. Notice that we have for any t,

$$\max(x, y) \leq t + (x-t)^+ + (y-t)^+. \tag{34}$$

In particular, Eq. (34) follows because the RHS in each of the following cases is always $\geq \max\{x, y\}$.

- if $t \geq \max\{x, y\}$ we get t for the RHS.
- if $t \geq x$ and $t < y$ we get $t + y - t = y = \max\{x, y\}$ for the RHS; the symmetric case also holds.
- if $t < x$ and $t < y$ we get $t + x - t + y - t = x + y - t \geq \max\{x, y\}$ for the RHS.

It is easy to verify that this holds for a max of more than two inputs; i.e. for a set S of reals we have $\max(S) \leq t + \sum_{s\in S}(s-t)^+$. Thus, we have

$$\mathbb{E}_{A\sim \boldsymbol{Y}}\left[\max_{s\in A} v_s\right] \leq \mathbb{E}_{A\sim \boldsymbol{Y}}\left[v_b + \sum_{s\in A}(v_s - v_b)^+\right] = v_b + \mathbb{E}_{A\sim \boldsymbol{Y}}\left[\sum_{s\in A}(v_s - v_b)^+\right] \qquad (\text{by Eq. (34)})$$

$$= v_b + \mathbb{E}_{A\sim \boldsymbol{Y}}\left[\sum_{s\in A\cap M}(v_s - v_b)^+ + \sum_{s\in A\cap (X\setminus M)}(v_s - v_b)^+\right]$$

$$= v_b + \mathbb{E}_{A\sim \boldsymbol{Y}}\left[\sum_{s\in A\cap M}(v_s - v_b)^+\right] \qquad (\text{by } v_s > v_b \text{ iff } s \leq b)$$

$$= v_b + \mathbb{E}_{A\sim \boldsymbol{Y}}\left[\sum_{s\in A\cap M}(v_s - v_b)\right] \qquad (\text{by } v_s \geq v_b \text{ for } s \in M)$$

$$= v_b + \sum_{s\in M} p_s \cdot (v_s - v_b)$$

$$= v_b + \sum_{s} p_s \cdot (v_s - v_b)^+ \qquad (\text{by } v_s > v_b \text{ iff } s \in M),$$

which is exactly the desired upper bound. □

Lemma 2.2. *Let* $(X_1, \boldsymbol{X_2})$ *be a solution to a* TruncatedTwoStage *or* MinEMax *problem. We have*

$$B(X_1, \boldsymbol{X_2}) = \arg\min_B \left[B + \sum_{s\in[m]} p_s \cdot (\text{cost}(X_1, X_2^{(s)}) - B)^+\right],$$

where the arg min *takes the largest* B *minimizing the relevant quantity.*

Proof. To clear our notation we let $\bar{B} := B(X_1, \boldsymbol{X_2})$, $c_s := \text{cost}(X_1, X_2^{(s)})$ and $\bar{M} := M(X_1, \boldsymbol{X_2})$. Let $f(B) := B + \sum_{s\in[m]} p_s \cdot (c_s - B)^+$. We argue that $\bar{B}$ is the largest global minimum of f by showing that for any $\epsilon > 0$ we know that $f(\bar{B}) < f(\bar{B} + \epsilon)$ and $f(\bar{B}) \leq f(\bar{B} - \epsilon)$.

We begin by noting that for any reals $a \leq b$ we have

$$a^+ - b^+ \geq a - b \tag{35}$$

by casing on which of a and b are larger than 0.

Let $\hat{M} := \{s \in \bar{M} : c_s > \bar{B}\}$. Notice that $\sum_{s\in\hat{M}} p_s < 1$. For fixed and arbitrary $\epsilon > 0$ consider the relative values of $f(\bar{B})$ and $f(\bar{B} + \epsilon)$. We have

$$f(\bar{B} + \epsilon) - f(\bar{B}) = \epsilon + \sum_{s} p_s \cdot \left((c_s - \bar{B} - \epsilon)^+ - (c_s - \bar{B})^+\right)$$

$$= \epsilon + \sum_{s \in \hat{M}} p_s \cdot \left((c_s - \bar{B} - \epsilon)^+ - (c_s - \bar{B})^+ \right), \tag{36}$$

where (36) follows since for $s \notin \hat{M}$ we have $c_s \leq \bar{B}$ and so $\left((c_s - \bar{B} - \epsilon)^+ - (c_s - \bar{B})^+ \right) = 0$ for $s \notin \hat{M}$. Now noticing that for every s we have $(c_s - \bar{B} - \epsilon) \leq (c_s - \bar{B})$, applying (35) to (36) gives

$$f(\bar{B} + \epsilon) - f(\bar{B}) \geq \epsilon + \sum_{s \in \hat{M}} p_s \cdot (-\epsilon) = \epsilon \left(1 - \sum_{s \in \hat{M}} p_s \right) > 0,$$

where the last inequality uses $\sum_{s \in \hat{M}} p_s < 1$. Thus, we have $f(\bar{B} + \epsilon) > f(\bar{B})$.

Now consider the relative values of $f(\bar{B})$ and $f(\bar{B} - \epsilon)$. We have

$$\begin{aligned}
f(\bar{B} - \epsilon) - f(\bar{B}) &= -\epsilon + \sum_{s} p_s \cdot \left((c_s - \bar{B} + \epsilon)^+ - (c_s - \bar{B})^+ \right) \\
&\geq -\epsilon + \sum_{s \in \bar{M}} p_s \cdot \left((c_s - \bar{B} + \epsilon)^+ - (c_s - \bar{B})^+ \right) && \left(\text{by } (c_s - \bar{B} + \epsilon)^+ \geq (c_s - \bar{B})^+ \right) \\
&\geq -\epsilon + \sum_{s \in \bar{M}} p_s \cdot \left((c_s - \bar{B} + \epsilon) - (c_s - \bar{B}) \right) && \left(\text{by } c_s \geq \bar{B} \text{ for } s \in \bar{M} \right) \\
&\geq \epsilon \left(1 - \sum_{s \in \bar{M}} p_s \right) \quad \geq \quad 0 && \left(\text{by } \sum_{s \in \bar{M}} p_s \geq 1 \right).
\end{aligned}$$

Thus, for any $\epsilon > 0$ we know that $f(\bar{B}) < f(\bar{B} + \epsilon)$ and $f(\bar{B}) \leq f(\bar{B} - \epsilon)$. It follows that, not only is $\bar{B}$ a global minimum of f but it is the largest global minimum. The lemma follows immediately. □

Lemma 2.3. *For feasible solution* $(X_1, \boldsymbol{X}_2)$ *of any* P_{EMax} *we have,* $\text{cost}_{EMax}(X_1, \boldsymbol{X}_2) \leq \text{cost}_{Trunc}(X_1, \boldsymbol{X}_2)$.

Proof. We have

$$\begin{aligned}
\text{cost}_{\text{EMax}}(X_1, \boldsymbol{X_2}) &= \mathbb{E}_A[\max_{s \in A}\{\text{cost}(X_1, X_2^{(s)})\}] \\
&\leq B(X_1, \boldsymbol{X_2}) + \sum_s p_s \cdot \left(\text{cost}(X_1, X_2^{(s)}) - B(X_1, \boldsymbol{X_2}) \right)^+ && \text{(by Lemma 2.1)} \\
&= \text{cost}_{\text{Trunc}}(X_1, \boldsymbol{X}_2) && \text{(by Lemma 2.2)}.
\end{aligned}$$

□

Lemma 2.4. *Let* P_{EMax} *be a* MinEMax *problem and* P_{Trunc} *be its truncated version. Let* $(E_1, \boldsymbol{E_2})$ *and* $(T_1, \boldsymbol{T_2})$ *be the optimal solutions to* P_{EMax} *and* P_{Trunc} *respectively. We have* $\text{cost}_{Trunc}(T_1, \boldsymbol{T_2}) \leq \left(\frac{2}{1-1/e} \right) \text{cost}_{EMax}(E_1, \boldsymbol{E_2})$.

Proof. We have

$$\begin{aligned}
&\text{cost}_{\text{Trunc}}(T_1, \boldsymbol{T_2}) \\
&\leq \text{cost}_{\text{Trunc}}(E_1, \boldsymbol{E_2}) && \text{(by } (T_1, \boldsymbol{T_2}) \text{ minimizes cost}_{\text{Trunc}}) \\
&= \min_B \left[B + \sum_s p_s \cdot (\text{cost}(E_1, E_2^{(s)}) - B)^+ \right] \\
&\leq B(E_1, \boldsymbol{E_2}) + \sum_s p_s \cdot (\text{cost}(E_1, E_2^{(s)}) - B(E_1, \boldsymbol{E_2}))^+ \\
&\leq \left(\frac{2}{1 - 1/e} \right) \mathbb{E}_A[\max_{s \in A}\{\text{cost}(E_1, E_2^{(s)})\}] && \text{(by Lemma 2.1)}
\end{aligned}$$

$$= \left(\frac{2}{1-1/e}\right) \text{cost}_{\text{EMax}}(E_1, \boldsymbol{E_2}).$$

□

B Deferred Proofs of §4

We use the following lemma in the proof of Lemma 4.4.

Lemma B.1. *If $B \geq OPT$ then $val\left(\frac{2B}{1-1/e}\right) \leq \left(\frac{6}{1-1/e}\right) \cdot OPT \leq \left(\frac{6}{1-1/e}\right) \cdot B$.*

Proof. Again, let X_T^* be the optimal solution to the TruncatedTwoStage problem corresponding to our MinEMax k-center problem. First, notice that by Lemma 2.4 we have $B(X_T^*) + \sum_s p_s \cdot (d(s, X_T^*) - B(X_T^*))^+ = \text{cost}_T(X_T^*) \leq \left(\frac{2}{1-1/e}\right) \text{cost}_{\text{EMax}}(X^*) = \left(\frac{2}{1-1/e}\right) \text{OPT}$ and so we have $B(X_T^*) \leq \text{OPT} \cdot \left(\frac{2}{1-1/e}\right)$.

Now suppose $B \geq \text{OPT}$. Let z_{is}^* be 1 if X_T^* assigns s to i and 0 otherwise and let $x^*(i)$ be 1 if $i \in X_T^*$ and 0 otherwise. Since (x^*, z^*) is feasible for $P_{\left(\frac{2B}{1-1/e}\right)}$ we have

$$\begin{aligned}
\text{val}\left(\frac{2B}{1-1/e}\right) &\leq \sum_s p_s \left(\sum_i f_{\left(\frac{2B}{1-1/e}\right)}(c_{is}) \cdot z_{is}^*\right) \\
&\leq \sum_s p_s \left(\sum_i f_{B(X_T^*)}(c_{is}) \cdot z_{is}^*\right) && (37) \\
&= \sum_{s \in M(X_T^*)} p_s \cdot (d(s, X_T^*)) && (38) \\
&\leq \sum_{s \in M(X_T^*)} p_s \cdot ((d(s, X_T^*) - B(X_T^*))^+ + B(X_T^*)) && (39) \\
&< 2B(X_T^*) + \text{cost}_{\text{Trunc}}(X_T^*) && (40) \\
&\leq 2B(X_T^*) + \left(\frac{2}{1-1/e}\right) \cdot \text{OPT}, && (41)
\end{aligned}$$

where Eq.(37) follows since $\left(\frac{2}{1-1/e}\right) \cdot B \geq \left(\frac{2}{1-1/e}\right) \cdot \text{OPT} \geq B(X_T^*)$, Eq.(38) follows by definition of $M(X_T^*)$, Eq.(39) follows by $x \leq (x-t)^+ + t$, Eq.(40) follows since $\sum_{s \in M(X_T^*)} p_s < 2$ and Eq.(41) follows by Lemma 2.4. Lastly, since $B(X_T^*) \leq \text{OPT} \cdot \left(\frac{2}{1-1/e}\right)$ by assumption, we have $\text{val}\left(\frac{2B}{1-1/e}\right) \leq \left(\frac{6}{1-1/e}\right) \cdot \text{OPT}$. □

Lemma 4.2. *There exists a polynomial-time algorithm which for a given $\epsilon > 0$ and an input instance of MinEMax k-center, returns a $\hat{B}$ such that $\hat{B} \leq (1+\epsilon)\left(\frac{6 \cdot OPT}{1-1/e}\right)$ and $val(\hat{B}) \leq \left(\frac{6 \cdot OPT}{1-1/e}\right)$. The algorithm's runtime is polynomial in n, m, and $\log_{1+\epsilon}$ OPT.*

Proof. Our algorithm to return $\hat{B}$ is as follows: let $\bar{B}$ be $(1+\epsilon)^i$ where i is the smallest $i \in \mathbb{Z}^+$ such that $\text{val}\left(\frac{2(1+\epsilon)^i}{1-1/e}\right) \leq \left(\frac{6}{1-1/e}\right) \cdot (1+\epsilon)^i$. Thus, we have

1. $\bar{B} \leq (1+\epsilon) \cdot \text{OPT}$ by Lemma B.1 and how we choose $2\bar{B}$;
2. $\text{val}\left(\frac{6\bar{B}}{1-1/e}\right) \leq \frac{6}{1-1/e} \cdot \bar{B}$ trivially by how we choose $\bar{B}$.

Lastly, we return $\hat{B} := \frac{6\bar{B}}{1-1/e}$. By (1) and the definition of $\hat{B}$ we have that $\hat{B} \leq (1+\epsilon)\left(\frac{6 \cdot \text{OPT}}{1-1/e}\right)$. Moreover, if $\bar{B} < \text{OPT}$ then by (2) we have $\text{val}(\hat{B}) \leq \left(\frac{6}{1-1/e}\right)\text{OPT}$ and if $\bar{B} \geq \text{OPT}$ then by Lemma B.1 we have $\text{val}(\hat{B}) \leq \left(\frac{6}{1-1/e}\right)\text{OPT}$. □

C Deferred Proofs of §5

Lemma 5.1. *Let* $(H_1, \overline{\boldsymbol{H_2}})$ *be the optimal solution to* P_{Hyb} *and let* $(H_1, \boldsymbol{H_2})$ *be its natural interpretation in* P_{EMax}*, i.e., for* $s \equiv s' \mod m$ *we define* $H_2^{(s)} := \overline{H}_2^{(s')}$*. We have* $\text{cost}_{EMax}(H_1, \boldsymbol{H_2}) \leq \text{cost}_{Hyb}(H_1, \overline{\boldsymbol{H_2}})$*.*

Proof. We first bound the second stage cost of MinEMax depending on whether a stochastic scenario ($s > m$) realizes into A or not. Recollect, A always contains the first m demand-robust scenarios because their probability is 1. Since a stochastic scenario (if it realizes) always costs more than all the demand-robust scenarios by Eq. (30), we get

$$\begin{aligned}
\mathbb{E}_A\left[\max_{s\in A}\left\{\text{cost}^2_{\text{EMax}}(H_2^{(s)})\right\}\right] &\leq \Pr_A[A=[m]]\cdot\max_{s\in[m]}\left\{\text{cost}^2_{\text{EMax}}(H_2^{(s)})\right\} + \sum_{s>m}\Pr_A[s\in A]\cdot\text{cost}^2_{\text{EMax}}(H_2^{(s)}) \\
&\leq 1\cdot\max_{s\in[m]}\left\{\text{cost}^2_{\text{EMax}}(\overline{H}_2^{(s)})\right\} + \sum_{s\in[m]} p_s\cdot\text{cost}^2_{\text{EMax}}(\overline{H}_2^{(s)}) \\
&= \max_{s\in[m]}\left\{\rho\cdot\text{cost}^2_{\text{Hyb}}(\overline{H}_2^{(s)})\right\} + \frac{1}{\gamma}\mathbb{E}_{s\sim\mathcal{D}}\left[\gamma(1-\rho)\cdot\text{cost}^2_{\text{Hyb}}(\overline{H}_2^{(s)})\right],
\end{aligned}$$

where the last equality uses the definition of $\text{cost}^2_{\text{EMax}}$ and that $p_s = \frac{\mathcal{D}(s)}{\gamma}$. Now using Eq. (31),

$$\begin{aligned}
\text{cost}_{\text{EMax}}(H_1, \boldsymbol{H_2}) &\leq \text{cost}^1_{\text{Hyb}}(H_1) + \rho\cdot\max_{s\in[m]}\left\{\text{cost}^2_{\text{Hyb}}(\overline{H}_2^{(s)})\right\} + (1-\rho)\cdot\mathbb{E}_{s\sim\mathcal{D}}\left[\text{cost}^2_{\text{Hyb}}(\overline{H}_2^{(s)})\right] \\
&= \text{cost}_{\text{Hyb}}(H_1, \overline{\boldsymbol{H_2}}).
\end{aligned}$$

□

References

[ADSY12] Shipra Agrawal, Yichuan Ding, Amin Saberi, and Yinyu Ye. Price of correlations in stochastic optimization. *Operations Research*, 60(1):150–162, 2012.

[AGGN08] Barbara M. Anthony, Vineet Goyal, Anupam Gupta, and Viswanath Nagarajan. A plant location guide for the unsure. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2008.

[Ala14] Saeed Alaei. Bayesian combinatorial auctions: Expanding single buyer mechanisms to many buyers. *SIAM Journal on Computing (SICOMP)*, 43(2):930–972, 2014.

[BBC11] D. Bertsimas, D. Brown, and C. Caramanis. Theory and applications of robust optimization. *SIAM Review*, 53(3):464–501, 2011.

[BP03] Nikhil Bansal and Kirk Pruhs. Server scheduling in the lp norm: A rising tide lifts all boat. In *Proceedings of the Symposium on the Theory of Computing (STOC)*, pages 242–250, 2003.

[CGTS02] Moses Charikar, Sudipto Guha, Éva Tardos, and David B Shmoys. A constant-factor approximation algorithm for the k-median problem. *Journal of Computer and System Sciences (JCSS)*, 65(1):129–149, 2002.

[CS18] Deeparnab Chakrabarty and Chaitanya Swamy. Interpolating between k-median and k-center: Approximation algorithms for ordered k-median. In *Proceedings of the International Colloquium on Automata, Languages and Programming (ICALP)*, pages 29:1–29:14, 2018.

[DGRS05] Kedar Dhamdhere, Vineet Goyal, R. Ravi, and Mohit Singh. How to pay, come what may: Approximation algorithms for demand-robust covering problems. In *Proceedings of the Symposium on the Foundations of Computer Science (FOCS)*, pages 367–378, 2005.

[DRS05] Kedar Dhamdhere, R Ravi, and Mohit Singh. On two-stage stochastic minimum spanning trees. In *Proceedings of International Conference on Integer Programming and Combinatorial Optimization (IPCO)*, pages 321–334, 2005.

[DY10] Erick Delage and Yinyu Ye. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research*, 58(3):595–612, 2010.

[GGP+15a] Daniel Golovin, Vineet Goyal, Valentin Polishchuk, R. Ravi, and Mikko Sysikaski. Improved approximations for two-stage min-cut and shortest path problems under uncertainty. *Mathematical Programming*, 149(1):167–194, Feb 2015.

[GGP+15b] Daniel Golovin, Vineet Goyal, Valentin Polishchuk, R Ravi, and Mikko Sysikaski. Improved approximations for two-stage min-cut and shortest path problems under uncertainty. *Mathematical Programming*, 149(1-2):167–194, 2015.

[GNR10] Anupam Gupta, Viswanath Nagarajan, and R Ravi. Thresholded covering algorithms for robust and max-min optimization. *Automata, Languages and Programming*, pages 262–274, 2010.

[GPRS04] Anupam Gupta, Martin Pal, R Ravi, and Amitabh Sinha. Boosted sampling: approximation algorithms for stochastic optimization. In *Proceedings of the Symposium on the Theory of Computing (STOC)*, pages 417–426, 2004.

[GRS04] Anupam Gupta, R Ravi, and Amitabh Sinha. An edge in time saves nine: Lp rounding approximation algorithms for stochastic network design. In *Proceedings of the Symposium on the Foundations of Computer Science (FOCS)*, pages 218–227, 2004.

[GS10] Joel Goh and Melvyn Sim. Distributionally robust optimization and its tractable approximations. *Operations Research*, 58(4-part-1):902–917, 2010.

[IKMM04] Nicole Immorlica, David Karger, Maria Minkoff, and Vahab S Mirrokni. On the costs and benefits of procrastination: Approximation algorithms for stochastic combinatorial optimization problems. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 691–700, 2004.

[KMPS05] V. S. Anil Kumar, Madhav Marathe, Srinivasan Parthasarathy, and Aravind Srinivasan. Approximation algorithms for scheduling on multiple machines. In *Proceedings of the Symposium on the Foundations of Computer Science (FOCS)*, pages 254–263, 2005.

[RS04] R Ravi and Amitabh Sinha. Hedging uncertainty: Approximation algorithms for stochastic optimization problems. In *Proceedings of International Conference on Integer Programming and Combinatorial Optimization (IPCO)*, pages 101–115, 2004.

[Sca58] Herbert Scarf. A min-max solution of an inventory problem. *Studies in the Mathematical Theory of Inventory and Production*, pages 201–209, 1958.

[Sri07] Aravind Srinivasan. Approximation algorithms for stochastic and risk-averse optimization. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1305–1313, 2007.

[SS06] David B. Shmoys and Chaitanya Swamy. An approximation scheme for stochastic linear programming and its application to stochastic integer programs. *Journal of the ACM (JACM)*, 53(6):978–1012, November 2006.

[Swa11] Chaitanya Swamy. Risk-averse stochastic optimization: probabilistically-constrained models and algorithms for black-box distributions. In *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1627–1646, 2011.

[SZY09] Anthony Man-Cho So, Jiawei Zhang, and Yinyu Ye. Stochastic combinatorial optimization with controllable risk aversion level. *Mathematics of Operations Research*, 34(3):522–537, 2009.

[Tam01] Arie Tamir. The k-centrum multi-faculty location problem. *Discrete Applied Mathematics*, 109(3):293–307, May 2001.