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

David Ellis Hershkowitz
Carnegie Mellon University, Pittsburgh, PA, USA
dhershko@cs.cmu.edu

R. Ravi
Carnegie Mellon University, Pittsburgh, PA, USA
ravi@andrew.cmu.edu

Sahil Singla
Princeton University, Princeton, NJ, USA
Institute for Advanced Study, Princeton, NJ, USA
singla@cs.princeton.edu

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. Although computing the objective in our model might seem to require enumerating exponentially-many possibilities, we show it is well estimated by a surrogate objective which is representable by a polynomial-size 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 variants of minimum spanning and Steiner trees, minimum cuts, and facility location. Up to constants, our approximation guarantees match those of previously-studied algorithms for 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.

2012 ACM Subject Classification Theory of computation → Approximation algorithms analysis; Theory of computation → Packing and covering problems; Theory of computation → Routing and network design problems; Theory of computation → Facility location and clustering; Theory of computation → Rounding techniques

Keywords and phrases Approximation Algorithms, Optimization Under Uncertainty, Two-Stage Optimization, Expected Max

Digital Object Identifier 10.4230/LIPIcs.APPROX-RANDOM.2019.4

Category APPROX

Related Version A full version of the paper is available at `https://arxiv.org/abs/1811.11635`.

Funding *David Ellis Hershkowitz*: Supported in part by NSF grants CCF-1527110, CCF-1618280, CCF-1814603, CCF-1910588, NSF CAREER award CCF-1750808 and a Sloan Research Fellowship.
R. Ravi: Supported in part by the U. S. Office of Naval Research award N00014-18-1-2099, and the U. S. NSF award CCF-1527032.
Sahil Singla: Supported in part by Schmidt Foundation and NSF awards CCF-1319811, CCF-1536002, and CCF-1617790.

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2019).
Editors: Dimitris Achlioptas and László A. Végh; Article No. 4; pp. 4:1–4:19
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

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 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 how to efficiently 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.

1. We introduce the MinEMax model for optimization of reconfigurable resources under uncertainty.
2. We show that, although evaluating the MinEMax objective function may seem difficult, a MinEMax problem can be approximately reduced to a "TruncatedTwoStage" problem whose objective is representable by an LP.
3. Armed with 2, 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.
4. Lastly, we show that the MinEMax model captures the commonly studied two-stage models for optimization under uncertainty: the *stochastic* and *demand-robust* models. We even show that it generalizes a "Hybrid" problem that interpolates between these models.

1.1 Related Work

Significant work has been done in two-stage optimization under uncertainty. The two most commonly studied models are the stochastic model [22, 15, 24] and the demand-robust model [9, 4, 14, 13]. In the **stochastic two-stage model** a probability distribution is given over scenarios and our objective is the *expected* total cost. In the **demand-robust** two-stage model we are given scenarios and our objective is the cost of the *worst-case* scenario given our first stage solution.

Another related model is *Distributionally robust optimization* (DRO) [23, 12, 8, 5]. In DRO we are given a distribution along with a ball of "nearby" distributions, and we must pay the *worst-case expectation* over all these distributions. Similarly to our model, DRO generalizes both the stochastic and demand-robust two-stage models. Our model can be seen as a "flip" of the DRO model: while the DRO model takes the worst-case over distributions our model takes a distribution over worst cases. Like DRO, our model is also sufficiently general to capture stochastic and demand-robust optimization. A recent result [20] – which shows that approximation algorithms are possible in DRO – complements our approximation algorithms in MinEMax.

A well studied measure for risk-aversion from stochastic programming is *conditional value at risk (CVaR)* [1]. Roughly, CVaR gives the average cost in the worst-case case α tail of a distribution. A notable recent work in CVaR presents a data-driven approach to two-stage risk aversion [18]. Theorem 1 in their work is reminiscent of our reduction of MinEMax to Hybrid; this theorem shows that their objective can be reformulated as a combination of the CVaR cost and the worst-case distribution. We emphasize that while CVaR might appear similar to the TruncatedTwoStage metric studied in this work, these two metrics are distinct and not readily comparable. Two salient differences are: (1) the threshold in the TruncatedTwoStage objective is the minimizing threshold while in CVaR the threshold is fixed, and (2) the TruncatedTwoStage objective sums up the truncated cost over a set of Bernoulli random variables whereas CVaR takes a truncated average cost with respect to a single distribution. Moreover, to the best of our knowledge, CVaR has not been studied in the context of approximation algorithms.

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 [26, 25, 27] examined various models of two-stage optimization that capture risk-aversion. Notably, the model of [27] interpolates between stochastic and demand-robust while also accommodating *black-box* distributions. Other papers (e.g., [15]) studied algorithms for stochastic optimization given access to black-box distributions. There has also been work on two-stage stochastic models in which – as in our model – independent stochastic outcomes factor prominently. For example, Immorlica et al. [17] study a two-stage stochastic model in which each "client" activates independently and the realized scenario consists of all activated clients. The primary difference between their model and ours is that for us entire scenarios – rather than clients – activate independently. Moreover, reconfigurability of resources is not factored in their model.

Lastly, in our reduction from MinEMax to TruncatedTwoStage, we make use of a bound which has appeared before in other settings [19, 21, 6, 11]. For example, [6] use this bound to tightly estimate the optimum value in an optimization problem where the cost function is random and only marginal distributions for the coefficients of the cost function are known. Unlike our work, these works do not design approximation algorithms for two-stage problems.

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.

1.2.1 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 $\mathsf{cost} : 2^X \times 2^X \to \mathbb{R}$. 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 discussed below.

This paper makes the common assumption that cost is *linear*, i.e., $\mathsf{cost}(X_1, X_2^{(s)})$ equals $\mathsf{cost}(\emptyset, X_2^{(s)}) + \mathsf{cost}(X_1, \emptyset)$ 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.

1.2.2 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:

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

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:

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

1.2.3 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:

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

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. We shall assume without loss of generality that $\sum_s p_s \geq 1$ throughout this paper since one can always ensure this without affecting solutions to the problem by adding dummy scenarios of cost 0 and probability 1.

As a concrete example of 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 $\sigma \cdot c_i$ for $\sigma > 1$. Our goal is to connect r to an *unknown* vertex v_s with minimum total two-stage cost.

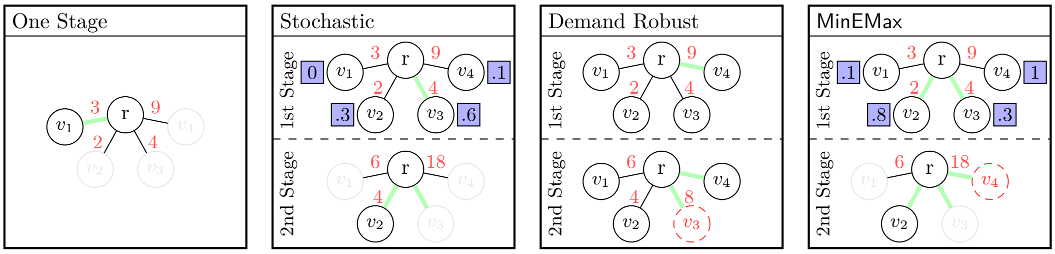

(a) If scenario to be covered is known to be v_1, problem is trivial.

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

(c) Given first stage soln., adversary chooses costliest scenario.

(d) Given first stage soln., adversary chooses 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.

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 resources 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 seems challenging: scenarios may realize in exponentially-many ways and so 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:

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

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

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 $\mathsf{cost}(X_1, X_2^{(s)}) \geq \mathsf{cost}(X_1, X_2^{(s+1)})$ 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 $B := \mathsf{cost}(X_1, X_2^{(k)})$ 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 $\mathsf{cost}_{\mathsf{Trunc}}$ 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 2.** *For two-stage covering problems there exist polynomial-time approximation algorithms with the following guarantees.*[1]

MinEMax *Problem*	*Steiner tree*	*UFL*	*MST*	*Min-cut*	*k-center*
Approximation	$\frac{30}{1-1/e}$	$\frac{8}{1-1/e}$	$O(\log n + \log m)$	$\frac{4}{1-1/e}$	$O(1)$

Our earlier Theorem 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 defer details on min-cut to the full version of our paper.

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: 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 [7].[2]

It is also worth noting that Anthony et al. [4] 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. We defer details on our k-center results to the full version of our paper.

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 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.*

We defer a formal definition and discussion of the Hybrid model as well as the intuition and proof for Theorem 3 to the full version of our paper. As a corollary of Theorems 2 and 3, we immediately recover polynomial-time approximations for Hybrid MST, UFL, Steiner tree and min-cut.[3]

[1] The $O(1)$ in the k-center approximation is roughly 57.
[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.
[3] Though not k-center since its cost function is not linear.

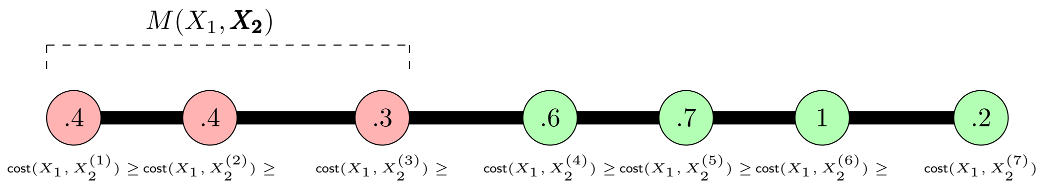

Figure 2 $M(X_1, \boldsymbol{X_2})$. $B(X_1, \boldsymbol{X_2}) = \mathsf{cost}(X_1, X_2^{(3)})$. Red circles: scenarios in $M(X_1, \boldsymbol{X_2})$. Green circles: all other scenarios. Numbers in circles: probabilities. Scenarios arranged left to right in descending order of $\mathsf{cost}(X_1, X_2^{(s)})$.

2 Reducing MinEMax to TruncatedTwoStage

In this section, we demonstrate a technique to simplify both computing and reasoning about $\mathsf{cost}_{\mathsf{EMax}}$ 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.** *Let P_{EMax} be a* MinEMax *problem and let P_{Trunc} be its* TruncatedTwoStage *version. An α-approximation algorithm for P_{Trunc} is a $\left(\frac{\alpha}{1-1/e}\right)$-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 remark that this result can be seen to follow from results regarding the "correlation gap" [2, 3] which show a similar bound where instead of max we have any sub-modular function. We give a different proof in §A for completeness that we find simpler in our setting where we consider the max and not any sub-modular function.

▶ **Lemma 4.** *Let $\boldsymbol{Y} = \{Y_1, \dots, Y_m\}$ 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 $b = \min\{a : \sum_{s=1}^{a} p_s \geq 1\}$. Then*

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

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

For a *given solution* $(X_1, \boldsymbol{X}_2)$ to MinEMax, Lemma 4 yields a computationally tractable form of $\mathsf{cost}_{\mathsf{EMax}}$. Specifically, let our scenarios be indexed such that $\mathsf{cost}(X_1, X_2^{(s)}) \geq \mathsf{cost}(X_1, X_2^{(s+1)})$ and let b be the smallest 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):

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

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 4, we can approximate $\mathsf{cost}_{\mathsf{EMax}}(X_1, \boldsymbol{X_2})$. However, we would like to estimate $\mathsf{cost}_{\mathsf{EMax}}(X_1, \boldsymbol{X}_2)$ 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 4 to estimate $\mathsf{cost}_{\mathsf{EMax}}(X_1, \boldsymbol{X}_2)$ 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 approximately the $\mathsf{cost}_{\mathsf{Trunc}}$ objective. We remind the reader that, as per Eq.(4), $\mathsf{cost}_{\mathsf{Trunc}}(X_1, \boldsymbol{X_2}) := \min_B[B + \sum_{s\in[m]} p_s \cdot (\mathsf{cost}(X_1, X_2^{(s)}) - B)^+]$. The following lemma shows that the B achieving the minimum in $\mathsf{cost}_{\mathsf{Trunc}}(X_1, \boldsymbol{X_2})$ is $B(X_1, \boldsymbol{X_2})$ and therefore shows that $\mathsf{cost}_{\mathsf{Trunc}}$ is a good approximation of $\mathsf{cost}_{\mathsf{EMax}}$.

▶ **Lemma 5.** *Let $(X_1, \boldsymbol{X_2})$ be a solution to a* TruncatedTwoStage *or* MinEMax *problem. We have*

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

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

Proof Sketch. The rough idea of the proof is to show that $B + \sum_s p_s(\mathsf{cost}(X_1, X_2^{(s)} - B)^+$ 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 $B + \sum_s p_s(\mathsf{cost}(X_1, X_2^{(s)} - B)^+$ 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 $\sum_s p_s(\mathsf{cost}(X_1, X_2^{(s)} - B)^+$ increases and so $B + \sum_s p_s(\mathsf{cost}(X_1, X_2^{(s)} - B)^+$ becomes larger overall. The case when B is made larger is symmetric. The full proof is available in §A. ◀

Using Lemma 4 and Lemma 5, 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 6.** *For feasible solution $(X_1, \boldsymbol{X}_2)$ of any P_{EMax} we have, $\mathsf{cost}_{\mathsf{EMax}}(X_1, \boldsymbol{X}_2) \leq \mathsf{cost}_{\mathsf{Trunc}}(X_1, \boldsymbol{X}_2)$.*

▶ **Lemma 7.** *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 optimal solutions to P_{EMax} and P_{Trunc} respectively. We have that $\mathsf{cost}_{\mathsf{Trunc}}(T_1, \boldsymbol{T_2}) \leq \left(\frac{1}{1-1/e}\right) \mathsf{cost}_{\mathsf{EMax}}(E_1, \boldsymbol{E_2})$.*

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. 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 6 we have $\mathsf{cost}_{\mathsf{EMax}}(\hat{T}_1, \hat{\boldsymbol{T}}_2) \leq \mathsf{cost}_{\mathsf{Trunc}}(\hat{T}_1, \hat{\boldsymbol{T}}_2)$. Since $(\hat{T}_1, \hat{\boldsymbol{T}}_2)$ is an α-approximation we have this is at most $\alpha \cdot \mathsf{cost}_{\mathsf{Trunc}}(T_1, \boldsymbol{T_2})$. Applying Lemma 7 this is at most $\left(\frac{\alpha}{1-1/e}\right) \mathsf{cost}_{\mathsf{EMax}}(E_1, \boldsymbol{E_2})$. 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 $\left(\frac{\alpha}{1-1/e}\right) \mathsf{cost}_{\mathsf{EMax}}(E_1, \boldsymbol{E_2})$, 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 [22, 15, 24] 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 $\mathsf{cost}_{\mathsf{Trunc}}$ as an LP objective. For this technique we need to extend the definition of $\mathsf{cost}_{\mathsf{Trunc}}$ 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 $\mathsf{cost}(x_1, x_2^{(s)})$ for fractional solution $(x_1, x_2^{(s)})$ to scenario s and let $\mathsf{cost}_{\mathsf{Trunc}}(x_1, \boldsymbol{x_2})$ be defined similarly to the integral case, i.e. for fractional $(x_1, \boldsymbol{x_2})$,

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

Given a minimization LP, it is easy to see that by introducing an additional variable to represent B and additional variables to represent $(\mathsf{cost}(x_1, x_2(s)) - B)^+$ for every s, we can represent $\mathsf{cost}_{\mathsf{Trunc}}(x_1, \boldsymbol{x_2})$ in an LP. For cleanliness of exposition, when we write our LPs we omit these additional variables and simply write our objective as "$\mathsf{cost}_{\mathsf{Trunc}}(x_1, \boldsymbol{x_2})$." 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 $\mathsf{cost}_{\mathsf{Trunc}}(x_1, \boldsymbol{x_2})$.

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 $\mathsf{cost}(x_1, x_2^{(s)}) \geq \mathsf{cost}(x_1, x_2^{(s+1)})$. Let b be the smallest positive integer such that $\sum_{s=1}^{b} p_s \geq 1$. For fractional $(x_1, \boldsymbol{x_2})$, we define

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

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

▶ Remark 8. It is easy to verify that the proof of Lemma 5 also holds for $\mathsf{cost}_{\mathsf{Trunc}}(x_1, \boldsymbol{x_2})$ 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 9.** *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 $\mathsf{cost}(X_1, X_2^{(s)}) \leq c \cdot \mathsf{cost}(Y_1, Y_2^{(s)})$ then*

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

Proof. We have

$$\mathsf{cost}_{\mathsf{Trunc}}(X_1, \boldsymbol{X_2}) = \min_B \left[B + \sum_s p_s \cdot (\mathsf{cost}(X_1, X_2^{(s)}) - B)^+) \right] \tag{9}$$

$$\leq c \cdot B(Y_1, \boldsymbol{Y_2}) + \sum_s p_s \cdot (\mathsf{cost}(X_1, X_2^{(s)}) - c \cdot B(Y_1, \boldsymbol{Y_2}))^+) \tag{10}$$

$$\leq c \cdot B(Y_1, \boldsymbol{Y_2}) + \sum_s p_s \cdot (c \cdot \mathsf{cost}(Y_1, Y_2^{(s)}) - c \cdot B(Y_1, \boldsymbol{Y_2}))^+) \tag{11}$$

$$= c \cdot \left(B(Y_1, \boldsymbol{Y_2}) + \sum_s p_s \cdot (\mathsf{cost}(Y_1, Y_2^{(s)}) - B(Y_1, \boldsymbol{Y_2}))^+) \right) \tag{12}$$

$$= c \cdot \mathsf{cost}_{\mathsf{Trunc}}(Y_1, \boldsymbol{Y_2}) \tag{13}$$

where Eq.(10) is by letting $B = c \cdot B(Y_1, \boldsymbol{Y_2})$, Eq.(11) is by $\mathsf{cost}(X_1, X_2^{(s)}) \leq c \cdot \mathsf{cost}(Y_1, Y_2^{(s)})$ and Eq.(13) is by Lemma 5. ◀

3.2 Steiner Tree

In this section we give a $\left(\frac{30}{1-1/e}\right)$-approximation for MinEMax rooted Steiner tree.

▶ **Definition 10** (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*

$$\mathsf{cost}(X_1, X_2^{(s)}) := \sum_{e \in X_1} c_e + \sigma_s \cdot \sum_{e \in X_2^{(s)}} c_e. \tag{14}$$

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

Our algorithm is based on an LP rounding algorithm of Gupta et al. [16] for two-stage stochastic Steiner tree. Roughly, we use Lemma 9 to argue that the first stage solution for every optimal TruncatedTwoStage solution is, up to small constants, a tree rooted at r. This structural property allows us to write an LP that approximately captures TruncatedTwoStage Steiner tree. Gupta et al. [16] showed that this LP can be rounded s.t. every scenario has a good cost. We then combine this rounding with Lemma 9 to derive an approximation algorithm for TruncatedTwoStage Steiner tree, which is sufficient for approximating MinEMax Steiner tree by Theorem 1.

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

▶ **Lemma 11.** *There exists an integral solution $(\hat{X}_1, \hat{\boldsymbol{X}}_{\boldsymbol{2}})$ to* TruncatedTwoStage *Steiner tree s.t. $G[\hat{X}_1]$ is a tree rooted at r and $\mathsf{cost}_{\mathit{Trunc}}(\hat{X}_1, \hat{\boldsymbol{X}}_{\boldsymbol{2}}) \leq 2 \cdot \mathsf{cost}_{\mathit{Trunc}}(O_1, \boldsymbol{O_2})$, where $(O_1, \boldsymbol{O_2})$ is the optimal solution to* TruncatedTwoStage *Steiner tree.*

Proof. Lemma 4.1 of Dhamdhere et al. [9] shows that given $(O_1, \boldsymbol{O_2})$ it is possible to modify it to a feasible solution $(\hat{X}_1, \hat{\boldsymbol{X}}_{\boldsymbol{2}})$ such that $G[\hat{X}_1]$ is a tree rooted at r and $\mathsf{cost}(\hat{X}_1, \hat{X}_2^{(s)}) \leq 2 \cdot \mathsf{cost}(O_1, O_2^{(s)})$ for every s. It follows by Lemma 9 that $\mathsf{cost}_{\mathsf{Trunc}}(\hat{X}_1, \hat{\boldsymbol{X}}_{\boldsymbol{2}}) \leq 2 \cdot \mathsf{cost}_{\mathsf{Trunc}}(O_1, \boldsymbol{O_2})$. ◀

We now describe how to formulate an LP that leverages the structural property in Lemma 11. 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

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

which as described by Eq.(6) also defines $\mathsf{cost}_{\mathsf{Trunc}}(x_1, \boldsymbol{x_2})$. Letting $\delta^-(v)$ and $\delta^+(v)$ stand for all directed edges going into and out of v, respectively. The following is our LP.

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

Notably, the third family of constraints enforces that terminal t is serviced by the first stage more and more as one moves closer to the root. The characteristic vector of $(\hat{X}_1, \hat{\boldsymbol{X}}_{\boldsymbol{2}})$ as described in Lemma 11 gives a feasible solution to ST LP. As a result, Lemma 11 demonstrates that ST LP has nearly optimal objective as stated in the following corollary.

▶ **Corollary 12.** *Let $(x_1, \boldsymbol{x_2})$ be the optimal solution of ST LP. We have $\mathsf{cost}_{\mathit{Trunc}}(x_1, \boldsymbol{x_2}) \leq 2 \cdot \mathsf{cost}_{\mathit{Trunc}}(O_1, \boldsymbol{O_2})$, where $(O_1, \boldsymbol{O_2})$ is the optimal solution to* TruncatedTwoStage *Steiner tree.*

Proof. Let $(\hat{x}_1, \hat{\boldsymbol{x}}_{\boldsymbol{2}})$ be the characteristic vector of $(\hat{X}_1, \hat{\boldsymbol{X}}_{\boldsymbol{2}})$ from Lemma 11. Fix an arbitrary terminal t. Let P_2 for terminal t be the shortest path from t to $\hat{X}_1$ in $G[\hat{X}_2]$. Let u_t be the sink of P_2 and let P_1 be the shortest path from u_t to r in $G[\hat{X}_1]$. Notice that $(\hat{x}_1, \hat{\boldsymbol{x}}_{\boldsymbol{2}})$ 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 for every t is a feasible solution to ST LP. Moreover, notice that cost of this solution is $\mathsf{cost}_{\mathsf{Trunc}}(\hat{x}_1, \boldsymbol{x_2}) = \mathsf{cost}_{\mathsf{Trunc}}(X_1, \boldsymbol{X_2}) \leq 2 \cdot \mathsf{cost}_{\mathsf{Trunc}}(O_1, \boldsymbol{O_2})$ by Lemma 11. ◀

Previous work of Gupta et al. [16] shows that it is possible to round a fractional solution of ST LP such that *every* scenario has a good cost.

▶ **Lemma 13** ([16]). *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. $\mathsf{cost}(X_1, X_2^{(s)}) \leq 15 \cdot \mathsf{cost}(x_1, x_2^{(s)})$ for every s.*

Since Corollary 12 gives ST LP has a good optimal solution, we can round ST LP such that every scenario has a low cost. Now Lemma 9 tells us that such a rounding preserves the cost of a solution for TruncatedTwoStage optimization. This gives the following theorem.

▶ **Theorem 14.** MinEMax *Steiner tree can be* $\left(\frac{30}{1-1/e}\right)$*-approximated in polynomial time.*

Proof. Our algorithm first solves ST LP to get fractional solution $(x_1, \boldsymbol{x_2})$. Next, we apply Lemma 13 to round $(x_1, \boldsymbol{x_2})$ in polynomial time to give $(X_1, \boldsymbol{X_2})$ as our solution. Thus, we have

$$\begin{aligned} \mathsf{cost}_{\mathsf{Trunc}}(X_1, \boldsymbol{X_2}) &\le 15 \cdot \mathsf{cost}_{\mathsf{Trunc}}(x_1, \boldsymbol{x_2}) && \text{(by Lemma 9, Lemma 13)} \\ &\le 30 \cdot \mathsf{cost}_{\mathsf{Trunc}}(O_1, \boldsymbol{O_2}), && \text{(by Corollary 12)} \end{aligned}$$

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, we have a $\left(\frac{30}{1-1/e}\right)$-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. ◀

3.3 Uncapacitated Facility Location

In this section we give a polynomial-time $\left(\frac{8}{1-1/e}\right)$-approximation algorithm for MinEMax uncapacitated facility location (UFL).

▶ **Definition 15** (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 $S_1, \dots, S_m \subseteq \mathcal{D}$, 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*

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

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

Our algorithm is based on the work of Ravi and Sinha [22] 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 9 to this rounding gives an algorithm that approximates TruncatedTwoStage UFL, which by Theorem 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, respectively. For a fractional solution $(x_1, \boldsymbol{x_2})$, we define

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

[4] This easily generalizes to more demand.

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 $\mathsf{cost}(x_1, x_2^{(s)})$ defines $\mathsf{cost}_{\mathsf{Trunc}}(x_1, \boldsymbol{x_2})$ for fractional $(x_1, \boldsymbol{x_2})$, which allows us to define our LP.

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

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 16** (Theorem 2, Lemma 1 in [22]). *Given a fractional solution $(x_1, \boldsymbol{x_2})$ to UFL LP, it is possible to round it to integral $(X_1, \boldsymbol{X_2})$ in polynomial-time s.t. for every scenario s we have $\mathsf{cost}(X_1, X_2^{(s)}) \leq 8 \cdot \mathsf{cost}(x_1, x_2^{(s)})$.*

We now give our approximation algorithm for MinEMax UFL.

▶ **Theorem 17.** MinEMax *UFL can be $\left(\frac{8}{1-1/e}\right)$-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 16 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, $\mathsf{cost}_{\mathsf{Trunc}}(o_1, \boldsymbol{o_2}) = \mathsf{cost}_{\mathsf{Trunc}}(O_1, \boldsymbol{O_2})$. Now using Lemma 9 and Lemma 16 it follows that

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

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

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

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

Lastly, notice that our algorithm is trivially polynomial-time. ◀

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 18** (MinEMax MST). *We are given a graph $G = (V, E)$ where $|V| = n$, a set of m scenarios $S_1, \dots S_m$ where each scenario S_s has an associated second-stage cost function $\mathsf{cost}_2^{(s)} : E \to \mathbb{Z}^+$ and a probability p_s. We are also given a first-stage cost function, $\mathsf{cost}_1 : E \to \mathbb{Z}^+$. 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*

$$\mathsf{cost}(X_1, X_2^{(s)}) := \sum_{e \in X_1} \mathsf{cost}_1(e) + \sum_{e \in X_2^{(s)}} \mathsf{cost}_2^{(s)}(e). \tag{15}$$

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

Our algorithm is based on the work of Dhamdhere et al. [10] 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 9, we get an approximation algorithm for TruncatedTwoStage MST, which by Theorem 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

$$\mathsf{cost}(x_1, x_2^{(s)}) := \sum_e x_1(e) \cdot \mathsf{cost}_1(e) + x_2^{(s)}(e) \cdot \mathsf{cost}_2(e), \tag{16}$$

which as described in Eq.(6), defines $\mathsf{cost}_{\mathsf{Trunc}}(x_1, \boldsymbol{x_2})$ 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.

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

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. [10].

We need the following result of Dhamdhere et al. [10] to round MST LP such that every scenario has a low cost.

▶ **Lemma 19** ([10]). *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 $\mathbb{E}[\mathsf{cost}(X_1, X_2^{(s)})] \leq \mathsf{cost}(x_1, x_2^{(s)}) \cdot (40 \log n + 16 \log m)$. 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 20.** *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 19 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

[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.

verify that $\mathsf{cost}_{\mathsf{Trunc}}(o_1, \boldsymbol{o_2}) = \mathsf{cost}_{\mathsf{Trunc}}(O_1, \boldsymbol{O_2})$. Taking expectations over the randomness of our algorithm and applying Lemma 9 and Lemma 19, we have with probability at least $1 - \frac{1}{mn^2}$ that

$$\begin{aligned}\mathbb{E}[\mathsf{cost}_{\mathsf{Trunc}}(X_1, \boldsymbol{X_2})] &\leq (40 \log n + 16 \log m) \cdot \mathsf{cost}_{\mathsf{Trunc}}(o_1, \boldsymbol{o_2}) \\ &= (40 \log n + 16 \log m) \cdot \mathsf{cost}_{\mathsf{Trunc}}(O_1, \boldsymbol{O_2}).\end{aligned}$$

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 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 19. ◀

References

1 Carlo Acerbi and Dirk Tasche. Expected shortfall: a natural coherent alternative to value at risk. *Economic notes*, 31(2):379–388, 2002.

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

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

4 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.

5 D. Bertsimas, D. Brown, and C. Caramanis. Theory and Applications of Robust Optimization. *SIAM Review*, 53(3):464–501, 2011.

6 Dimitris Bertsimas, Karthik Natarajan, and Chung-Piaw Teo. Probabilistic combinatorial optimization: Moments, semidefinite programming, and asymptotic bounds. *SIAM Journal on Optimization*, 15(1):185–209, 2004.

7 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.

8 Erick Delage and Yinyu Ye. Distributionally Robust Optimization Under Moment Uncertainty with Application to Data-Driven Problems. *Operations Research*, 58(3):595–612, 2010.

9 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.

10 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.

11 Anulekha Dhara and Karthik Natarajan. On the polynomial solvability of distributionally robust k-sum optimization. *Optimization Methods and Software*, 32(4):738–753, 2017.

12 Joel Goh and Melvyn Sim. Distributionally Robust Optimization and Its Tractable Approximations. *Operations Research*, 58(4-part-1):902–917, 2010.

13 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, February 2015.

[6] Although Theorem 1 and Lemma 9 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.

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

15 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.

16 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.

17 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.

18 Ruiwei Jiang and Yongpei Guan. Risk-averse two-stage stochastic program with distributional ambiguity. *Operations Research*, 66(5):1390–1405, 2018.

19 TL Lai and Herbert Robbins. Maximally dependent random variables. *Proceedings of the National Academy of Sciences*, 73(2):286–288, 1976.

20 Andre Linhares and Chaitanya Swamy. Approximation Algorithms for Distributionally-Robust Stochastic Optimization with Black-Box Distributions. In *Proceedings of the Symposium on the Theory of Computing (STOC)*, 2019.

21 Isaac Meilijson and Arthur Nádas. Convex majorization with an application to the length of critical paths. *Journal of Applied Probability*, 16(3):671–677, 1979.

22 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.

23 Herbert Scarf. A Min-Max Solution of an Inventory Problem. *Studies in the Mathematical Theory of Inventory and Production*, pages 201–209, 1958.

24 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.

25 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.

26 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.

27 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.

A Deferred Proofs of §2

▶ **Lemma 4.** *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 $b = \min\{a : \sum_{s=1}^{a} p_s \geq 1\}$. Then*

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

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

Proof. We begin by showing the lower bound on $\mathbb{E}_{A\sim \boldsymbol{Y}}[\max_{s\in A} v_s]$. Let $M := [b]$. Consider the new set of probabilities

$$p'_s = \begin{cases} 1 - \sum_{s<b} p_s & \text{if } s = b \\ p_s & \text{otherwise} \end{cases} \tag{17}$$

and let $\boldsymbol{Y'}$ be the corresponding Bernoulli r.v.s. Notice that $\sum_{s\in M} p'_s = 1$.

Since $p'_s \le p_s$, clearly we have that $\mathbb{E}_{A\sim \boldsymbol{Y}}[\max_{s\in A} v_s] \ge \mathbb{E}_{A\sim \boldsymbol{Y'}}[\max_{s\in A} v_s]$. Thus, we will focus on lower bounding $\mathbb{E}_{A\sim \boldsymbol{Y'}}[\max_{s\in A} v_s]$. The probability that no element of M is in A when drawn from $\boldsymbol{Y'}$ is

$$\prod_{s\in M}(1-p'_s) \quad \le \quad e^{-\sum_{s\in M} p'_s} \quad = \quad \frac{1}{e}$$

because $1 - x \le e^{-x}$ and $\sum_{s\in M} p'_s = 1$. It follows that

$$\begin{aligned}
\mathbb{E}_{A\sim \boldsymbol{Y}}\left[\max_{s\in A} v_s\right] &\ge \mathbb{E}_{A\sim \boldsymbol{Y'}}\left[\max_{s\in A} v_s\right] \\
&\ge \left(1-\frac{1}{e}\right)\mathbb{E}_{A\sim \boldsymbol{Y'}}\left[\max_{s\in A} v_s \mid \text{at least 1 element from } M \text{ in } A\right] \\
&\ge \left(1-\frac{1}{e}\right)\mathbb{E}_{A\sim \boldsymbol{Y'}}\left[\max_{s\in A} v_s \mid \text{exactly 1 element from } M \text{ in } A\right] \\
&= \left(1-\frac{1}{e}\right)\sum_{s\in M} v_s \frac{p'_s}{\sum_{i\in M} p'_s} \quad = \left(\frac{1-1/e}{1}\right)\sum_{s\in M} p'_s v_s
\end{aligned}$$

where the last line follows since $\sum_{s\in M} p'_s = 1$.

Thus, we have that

$$\begin{aligned}
\mathbb{E}_{A\sim \boldsymbol{Y}}\left[\max_{s\in A} v_s\right] &\ge \left(1-\frac{1}{e}\right)\sum_{s\in M} p'_s v_s \\
&= \left(1-\frac{1}{e}\right)\sum_{s\in M} p'_s\Big((v_s - v_b)^+ + v_b\Big) && (\text{by } v_s \ge v_b \text{ for } s\in M) \\
&\ge \left(1-\frac{1}{e}\right)\Big(v_b + \sum_{s\in M} p'_s\big((v_s-v_b)^+\big)\Big) && \Big(\text{by } 1 = \sum_{s\in M} p'_s\Big) \\
&\ge \left(1-\frac{1}{e}\right)\Big(v_b + \sum_{s\in M} p_s\big((v_s-v_b)^+\big)\Big) && (\text{by } (v_b - v_b)^+ = 0) \\
&= \left(1-\frac{1}{e}\right)\Big(v_b + \sum_{s} p_s\big((v_s-v_b)^+\big)\Big) && (\text{by } v_s > v_b \text{ iff } s\in M)
\end{aligned}$$

which gives our lower bound.

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

$$\max(x, y) \le t + (x-t)^+ + (y-t)^+. \tag{18}$$

In particular, Eq. (18) follows because the RHS in each of the following cases is always $\ge \max\{x, y\}$.

- if $t \ge \max\{x, y\}$ we get t for the RHS.
- if $t \ge x$ 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 $t + x - t + y - t = x + y - t \ge \max\{x, y\}$ 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 $\max(S) \le t + \sum_{s\in S}(s-t)^+$. Thus, we have

$$\mathbb{E}_{A\sim \boldsymbol{Y}}\Big[\max_{s\in A} v_s\Big] \le \mathbb{E}_{A\sim \boldsymbol{Y}}\Big[v_b + \sum_{s\in A}(v_s - v_b)^+\Big] \quad = \quad v_b + \mathbb{E}_{A\sim \boldsymbol{Y}}\Big[\sum_{s\in A}(v_s - v_b)^+\Big] \tag{19}$$

$$= v_b + \mathbb{E}_{A\sim \boldsymbol{Y}}\Big[\sum_{s\in A\cap M}(v_s - v_b)^+ + \sum_{s\in A\setminus M}(v_s - v_b)^+\Big] \tag{20}$$

$$= v_b + \mathbb{E}_{A\sim \boldsymbol{Y}}\Big[\sum_{s\in A\cap M}(v_s - v_b)^+\Big] \tag{21}$$

$$= v_b + \mathbb{E}_{A\sim \boldsymbol{Y}}\Big[\sum_{s\in A\cap M}(v_s - v_b)\Big] \tag{22}$$

$$= v_b + \sum_{s\in M} p_s \cdot (v_s - v_b) \tag{23}$$

$$= v_b + \sum_{s} p_s \cdot (v_s - v_b)^+, \tag{24}$$

where Eq.(19) is by Eq.(18), Eq.(21) is by $v_s > v_b$ iff $s \le b$, Eq.(22) is by $v_s \ge v_b$ for $s \in M$ and Eq.(24) is by $v_s > v_b$ iff $s \in M$. This is exactly the desired upper bound. ◀

▶ **Lemma 5.** *Let* $(X_1, \boldsymbol{X_2})$ *be a solution to a* TruncatedTwoStage *or* MinEMax *problem. We have*

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

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

Proof. To clear our notation we let $\bar{B} := B(X_1, \boldsymbol{X_2})$, $c_s := \mathsf{cost}(X_1, X_2^{(s)})$ 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 $f(\bar{B}) \le f(\bar{B} - \epsilon)$.

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

$$a^+ - b^+ \ge a - b \tag{25}$$

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

$$f(\bar{B} + \epsilon) - f(\bar{B}) = \epsilon + \sum_{s\in[m]} p_s \cdot \big((c_s - \bar{B} - \epsilon)^+ - (c_s - \bar{B})^+\big)$$

$$= \epsilon + \sum_{s\in\hat{M}} p_s \cdot \big((c_s - \bar{B} - \epsilon)^+ - (c_s - \bar{B})^+\big)\,, \tag{26}$$

where Eq.(26) follows since for $s \notin \hat{M}$ we have $c_s \le \bar{B}$ and so $\big((c_s - \bar{B} - \epsilon)^+ - (c_s - \bar{B})^+\big) = 0$ for $s \notin \hat{M}$. Now noticing that for every s we have $(c_s - \bar{B} - \epsilon) \le (c_s - \bar{B})$, applying (25) to (26) gives

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

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

$$\begin{aligned} f(\bar{B}-\epsilon) - f(\bar{B}) &= -\epsilon + \sum_s p_s \cdot \left((c_s - \bar{B} + \epsilon)^+ - (c_s - \bar{B})^+\right) \\ &\geq -\epsilon + \sum_{s \in \bar{M}} p_s \cdot \left((c_s - \bar{B} + \epsilon)^+ - (c_s - \bar{B})^+\right) && (27) \\ &\geq -\epsilon + \sum_{s \in \bar{M}} p_s \cdot \left((c_s - \bar{B} + \epsilon) - (c_s - \bar{B})\right) && (28) \\ &\geq \epsilon\Big(1 - \sum_{s \in \bar{M}} p_s\Big) \quad \geq \quad 0 && (29) \end{aligned}$$

where Eq.(27) is by $(c_s - \bar{B} + \epsilon)^+ \geq (c_s - \bar{B})^+$, Eq.(28) is by $c_s \geq \bar{B}$ for $s \in \bar{M}$ and Eq.(29) is by $\sum_{s \in \bar{M}} p_s \geq 1$. 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 6.** *For feasible solution $(X_1, \boldsymbol{X}_2)$ of any P_{EMax} we have,* $\mathsf{cost}_{EMax}(X_1, \boldsymbol{X}_2) \leq \mathsf{cost}_{Trunc}(X_1, \boldsymbol{X}_2)$.

Proof. We have

$$\begin{aligned} \mathsf{cost}_{\mathsf{EMax}}(X_1, \boldsymbol{X_2}) &= \mathbb{E}_A[\max_{s \in A}\{\mathsf{cost}(X_1, X_2^{(s)})\}] && (30) \\ &\leq B(X_1, \boldsymbol{X_2}) + \sum_s p_s \cdot \left(\mathsf{cost}(X_1, X_2^{(s)}) - B(X_1, \boldsymbol{X_2})\right)^+ && (31) \\ &= \mathsf{cost}_{\mathsf{Trunc}}(X_1, \boldsymbol{X}_2) && (32) \end{aligned}$$

where Equation (31) is by Lemma 4 and Equation (32) is by Lemma 5. ◀

▶ **Lemma 7.** *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 optimal solutions to P_{EMax} and P_{Trunc} respectively. We have that* $\mathsf{cost}_{Trunc}(T_1, \boldsymbol{T_2}) \leq \left(\frac{1}{1-1/e}\right) \mathsf{cost}_{EMax}(E_1, \boldsymbol{E_2})$.

Proof. We have that

$$\begin{aligned} & \mathsf{cost}_{\mathsf{Trunc}}(T_1, \boldsymbol{T_2}) \\ & \leq \mathsf{cost}_{\mathsf{Trunc}}(E_1, \boldsymbol{E_2}) && \text{(by } (T_1, \boldsymbol{T_2}) \text{ minimizes } \mathsf{cost}_{\mathsf{Trunc}}) \\ & = \min_B \Big[B + \sum_s p_s \cdot (\mathsf{cost}(E_1, E_2^{(s)}) - B)^+\Big] \\ & \leq B(E_1, \boldsymbol{E_2}) + \sum_s p_s \cdot (\mathsf{cost}(E_1, E_2^{(s)}) - B(E_1, \boldsymbol{E_2}))^+ \\ & \leq \left(\frac{1}{1-1/e}\right) \mathbb{E}_A[\max_{s \in A}\{\mathsf{cost}(E_1, E_2^{(s)})\}] && \text{(by Lemma 4)} \\ & = \left(\frac{1}{1-1/e}\right) \mathsf{cost}_{\mathsf{EMax}}(E_1, \boldsymbol{E_2}). \end{aligned}$$

◀