Randomized Minmax Regret for Combinatorial Optimization Under Uncertainty⋆

Andrew Mastin[1], Patrick Jaillet[2], and Sang Chin[3]

[1] Lawrence Livermore National Laboratory⋆⋆, Livermore CA 94550, USA; mastin1@llnl.gov
[2] Massachusetts Institute of Technology, Cambridge MA 02139, USA; jaillet@mit.edu
[3] Draper Laboratory, Cambridge MA 02139, USA; schin@draper.com

Abstract. The minmax regret problem for combinatorial optimization under uncertainty can be viewed as a zero-sum game played between an optimizing player and an adversary, where the optimizing player selects a solution and the adversary selects costs with the intention of maximizing the regret of the player. The conventional minmax regret model considers only deterministic solutions/strategies, and minmax regret versions of most polynomial solvable problems are NP-hard. In this paper, we consider a randomized model where the optimizing player selects a probability distribution (corresponding to a mixed strategy) over solutions and the adversary selects costs with knowledge of the player's distribution, but not its realization. We show that under this randomized model, the minmax regret version of any polynomial solvable combinatorial problem becomes polynomial solvable. This holds true for both interval and discrete scenario representations of uncertainty. Using the randomized model, we show new proofs of existing approximation algorithms for the deterministic model based on primal-dual approaches. We also determine integrality gaps of minmax regret formulations, giving tight bounds on the limits of performance gains from randomization. Finally, we prove that minmax regret problems are NP-hard under general convex uncertainty.

Keywords: Robust optimization · Approximation algorithms · Game theory

1 Introduction

Many optimization applications involve cost coefficients that are not fully known. When probability distributions are available for cost coefficients (e.g. from historical data or other estimates), stochastic programming is often an appropriate

⋆ Research supported in part by NASA ESTOs Advanced Information System Technology (AIST) program under grant number NNX12H81G. Also supported by NSF grant 1029603, ONR grant N00014-12-1-0033, and AFOSR grant FA9550-12-1-0136.

⋆⋆ Lawrence Livermore National Laboratory is operated by Lawrence Livermore National Security, LLC, for the U.S. Department of Energy, National Nuclear Security Administration under Contract DE-AC52-07NA27344.

modeling choice [13, 26]. In other cases, costs may only be known to be contained in intervals (i.e. each cost has a known lower and upper bound) or to be a member of a finite set of scenarios, and one is more interested in worst-case performance. Robust optimization formulations are desirable here as they employ a minmax-type objective [12, 17, 23].

In a general robust optimization problem with cost uncertainty, one must select a set of items from some feasible *solution set* such that item costs are contained in some *uncertainty set*. The basic problem of selecting an optimal solution from the solution set when costs are known is referred to as the *nominal problem*. When only the uncertainty set is known, the goal under the *minmax* criterion (also referred to as absolute robustness) is to select a solution that gives the best upper bound on objective cost over all possible costs from the uncertainty set [27] (assuming that the nominal problem is a minimization problem). That is, one must select the solution that, when item costs are chosen to maximize the cost of the selected solution, is minimum. Under the *minmax regret* criterion (sometimes called the robust deviation model), the goal is instead to select the solution that minimizes the maximum possible regret, defined as the difference between the cost of the selected solution and the optimal solution [25].

A problem under the minmax regret criterion can be viewed as a two-stage game played between an optimizing player and an adversary. In the first stage, the optimizing player selects a deterministic solution. In the second stage, an adversary observes the selected solution and chooses values/costs from the uncertainty set with the intention of maximizing the player's regret. The goal of the optimizing player is thus to select a solution that least allows the adversary to generate regret. For both interval and discrete scenario representations of cost uncertainty, the minmax regret versions of most polynomial solvable problems are NP-hard [1,2,5,8,23,28]. A variation on this model, first suggested by Bertsimas et al. [10] for minmax robust optimization, is to allow the optimizing player to select a probability distribution over solutions and require the adversary to select costs based on knowledge of the player's distribution, but not its realization. In this paper, we study this randomized model under the minmax regret criterion instead of the minmax criterion.

We show that under this randomized model, the minmax regret version of any polynomial solvable 0–1 integer linear programming problem is polynomial solvable. This holds true for both interval and discrete scenario representations of uncertainty. Our observation is that the randomized model corresponds to the linear programming relaxation of the mixed integer program for the deterministic model. While the relaxation may have an exponential number of constraints, an efficient separation oracle is given by the nominal problem.

The linear programming formulation leads to further insights. We show that currently known approximation algorithms for deterministic minmax regret problems [3, 19], which have been proved using combinatorial arguments, can be proved using simpler primal-dual methods. This analysis also yields integrality gaps for deterministic minmax regret problems. The integrality gaps are shown to be equal to k for discrete scenario uncertainty, where k is the number of scenarios, and equal to 2 for interval uncertainty. Both gaps match the ratios of the approximation algorithms, showing that these algorithms are optimal. The integrality gaps also establish lower bounds on performance when moving from the deterministic model to the randomized model. Letting Z_{D} and Z_{R} denote the deterministic and randomized minmax regret values for a common nominal problem, we effectively show that $Z_{\mathrm{D}}/k \leq Z_{\mathrm{R}} \leq Z_{\mathrm{D}}$ for discrete scenario uncertainty and $Z_{\mathrm{D}}/2 \leq Z_{\mathrm{R}} \leq Z_{\mathrm{D}}$ for interval uncertainty.

Given that the randomized model makes many minmax regret problems polynomial solvable for interval uncertainty and discrete scenario uncertainty, it is natural to ask if polynomial solvability remains in the presence of slightly more elaborate uncertainty sets. We show that for general convex uncertainty sets, however, deterministic and randomized minmax regret problems are NP-hard, even for polynomial solvable nominal problems.

The paper is structured as follows. In the remainder of this section we review related work. Section 2 introduces notation and definitions. Section 3 highlights our most important results for discrete scenario uncertainty, interval uncertainty, and convex uncertainty. A full development of these topics, with proofs, is contained in the full version of the paper [24]. A conclusion is given in Section 4.

Related work. One of the first studies of minmax regret from both an algorithmic and complexity perspective was that of Averbakh [7]. He looked at the minmax regret version of the simple problem of selecting k items out of n total items, where the cost of each item is uncertain and the goal is to select the set of items with minimum total cost. For interval uncertainty, he derived a polynomial-time algorithm based on interchange arguments. He demonstrated that for the discrete scenario representation of uncertainty, however, the minmax regret problem becomes NP-hard, even for the case of only two scenarios. It is interesting to contrast these results with the case of general minmax regret linear programming which, as shown by Averbakh and Lebedev [9], is NP-hard for interval uncertainty but polynomial solvable for discrete scenario uncertainty.

Apart from the item selection problem, most polynomial solvable minmax regret combinatorial problems are NP-hard, both for interval and discrete scenario uncertainty [1, 2, 8, 23, 28]. The survey paper of Aissi et al. [5] provides a comprehensive summary of results related to both minmax and minmax regret combinatorial problems. For problems that are already NP-complete, most of their minmax regret versions are Σ_2^p-complete [16]. To solve minmax regret problems in practice, the book by Kasperski reviews standard mixed integer program (MIP) formulations for both interval and discrete scenario uncertainty [17].

General approximation algorithms for deterministic minmax regret problems are known for both types of uncertainty. Kasperski and Zieliński [19] proved a 2-approximation algorithm based on midpoint costs under interval uncertainty, and Aissi et al. [3] gave a k-approximation algorithm using mean costs under discrete scenario uncertainty, where k is the number of scenarios. Under interval uncertainty, fully polynomial-time approximation schemes are known for many problems [17, 20]. For discrete scenario uncertainty, Kasperski et al. [18] looked at the minmax regret item selection problem, which models special cases of many combinatorial problems. They showed that for a non-constant number of scenarios, the problem is not approximable within any constant factor unless P=NP. If the number of scenarios is constant, fully polynomial-time approximation schemes are known for some problems [4, 6].

The application of a game-theoretic model with mixed strategies to robust optimization problems was introduced by Bertsimas et al. [10]. They focused on the minmax robust model, and their analysis was motivated by adversarial models used for online optimization algorithms. They showed that if it is possible to optimize over both the solution set and the uncertainty set in polynomial time, then an optimal mixed strategy solution can be calculated in polynomial time, and that the expected cost under the randomized model is no greater than the cost for the deterministic model. They also gave bounds on the improvement gained from randomization for various uncertainty sets. Our work is similar to theirs, but we focus on the minmax regret criterion instead of the minmax criterion.

Other related areas of research are Stackelberg security games [21, 22], network interdiction games [11], and dueling algorithms [15]. A common feature of many of these works, as well as ours, is that they involve games that at first glance have exponential size but can be solved efficiently using the appropriate reductions.

2 Definitions

We consider a general combinatorial optimization problem where we are given a set of n items $E = \{e_1, e_2, \ldots, e_n\}$ and a set $\mathcal{F}$ of feasible subsets of E. Each item $e \in E$ has a cost $c_e \in \mathbb{R}$. Given the vector $c = (c_1, \ldots, c_n)$, the goal of the optimization problem is to select a feasible subset of items that minimizes the total cost; we refer to this as the *nominal problem*:

$$F^*(c) := \min_{T \in \mathcal{F}} \sum_{e \in T} c_e. \tag{1}$$

Let $x = (x_1, \ldots, x_n)$ be a characteristic vector for some set T, so that $x_e = 1$ if $e \in T$ and $x_e = 0$ otherwise. Also let $\mathcal{X} \subseteq \{0, 1\}^n$ denote the set of all characteristic vectors corresponding to feasible sets $T \in \mathcal{F}$. We assume that $\mathcal{X}$ is described in size m (e.g. with m linear inequalities). We can equivalently write the nominal problem with a linear objective function

$$F^*(c) = \min_{x \in \mathcal{X}} \sum_{e \in E} c_e x_e. \tag{2}$$

Throughout the paper, we use both set notation and characteristic vectors for ease of presentation.

We review the conventional definitions for the deterministic minmax regret framework and then present the analogous definitions for our randomized model.

For some cost vector $c \in \mathcal{C}$, the deterministic cost of a solution $T \in \mathcal{F}$ is

$$F(T, c) := \sum_{e \in T} c_e. \tag{3}$$

The regret of a solution T under some cost vector c is the difference between the cost of the solution and the optimal cost:

$$R(T, c) := F(T, c) - F^*(c). \tag{4}$$

The *maximum regret problem* for a solution T is

$$R_{\max}(T) := \max_{c \in \mathcal{C}} R(T, c) = \max_{c \in \mathcal{C}} \left(F(T, c) - F^*(c)\right). \tag{5}$$

The *deterministic minmax regret problem* is then

$$Z_{\mathrm{D}} := \min_{T \in \mathcal{F}} R_{\max}(T) = \min_{T \in \mathcal{F}} \max_{c \in \mathcal{C}} (F(T, c) - F^*(c)). \tag{6}$$

We abuse the notation $F(\cdot, c)$, $R(\cdot, c)$ and $R_{\max}(\cdot)$ by replacing set arguments with vectors (e.g. $F(x, c)$ in place of $F(T, c)$), but we generally use capital letters for sets and lowercase letters for vectors.

We now move to the randomized framework, where the optimizing player selects a distribution over solutions and the adversary selects a distribution over costs. Starting with the optimizing player, for some set $T \in \mathcal{F}$, let y_T denote the probability that the optimizing player selects set T. Let $y = (y_T)_{T \in \mathcal{F}}$ be the vector of length $|\mathcal{F}|$ specifying the set selection distribution; we refer to y simply as a *solution*. Define the feasible region for y as

$$\mathcal{Y} := \{y \mid y \geq \mathbf{0}, \mathbf{1}^\top y = 1\}, \tag{7}$$

where the notation $\mathbf{0}$ and $\mathbf{1}$ indicates a full vector of zeros and ones, respectively. We similarly define a distribution over costs for the adversary. The set $\mathcal{C}$ may in general be infinite, but we only consider strategies with finite support; for now we assume that such strategies are sufficient. Thus consider a finite set $\mathcal{C}_f \subseteq \mathcal{C}$, and for some $c \in \mathcal{C}_f$, let w_c denote the probability that the adversary selects costs c. Then let $w = (w_c)_{c \in \mathcal{C}_f}$ and define the feasible region

$$\mathcal{W} := \{w \mid w \geq \mathbf{0}, \mathbf{1}^\top w = 1\}. \tag{8}$$

We are interested in succinct descriptions of strategies for both players. We define for the optimizing player a *mixed strategy encoding* $\mathcal{M} = (\Theta, Y)$ as a set of deterministic solutions $\Theta = \{T_i \in \mathcal{F} \mid i = 1, \ldots, \mu\}$ that should be selected with nonzero probability and the corresponding probabilities $Y = \{y_{T_i} \in [0, 1] \mid i = 1, \ldots, \mu\}$ that satisfy $\sum_{i=1}^{\mu} y_{T_i} = 1$. Here μ is the support size of the mixed strategy (i.e. the number of deterministic solutions with nonzero probability). Likewise, define an adversarial mixed strategy encoding $\mathcal{L} = (C, W)$ as a set of costs $C = \{c^j \in \mathcal{C}_f \mid j = 1, \ldots, \eta\}$ to be selected with corresponding probabilities $W = \{w_{c^j} \in [0, 1] \mid j = 1, \ldots, \eta\}$ satisfying $\sum_{j=1}^{\eta} w_{c^j} = 1$.

The expected regret under y and w is simply

$$\overline{R}(y,w) := \sum_{T\in\mathcal{F}}\sum_{c\in\mathcal{C}_f} y_T w_c R(T,c) = \sum_{T\in\mathcal{F}}\sum_{c\in\mathcal{C}_f} y_T w_c (F(T,c) - F^*(c)). \tag{9}$$

For a given y, the *maximum expected regret problem* is

$$\overline{R}_{\max}(y) := \max_{w\in\mathcal{W}} \sum_{c\in\mathcal{C}_f} w_c \sum_{T\in\mathcal{F}} y_T R(T,c) = \max_{c\in\mathcal{C}_f} \sum_{T\in\mathcal{F}} y_T R(T,c). \tag{10}$$

The above equality follows from a standard observation in game theory: the optimization of $w \in \mathcal{W}$ is maximization of the function $G(y,c) = \sum_{T\in\mathcal{F}} y_T R(T,c)$ over the convex hull of $\mathcal{C}_f$, which is equivalent to optimizing over $\mathcal{C}_f$ itself. The minmax expected regret problem, which we refer to as the *randomized minmax regret problem*, is

$$Z_{\mathrm{R}} := \min_{y\in\mathcal{Y}} \overline{R}_{\max}(y) = \min_{y\in\mathcal{Y}} \max_{c\in\mathcal{C}} \left(\sum_{T\in\mathcal{F}} y_T (F(T,c) - F^*(c)) \right), \tag{11}$$

where we have replaced $\mathcal{C}_f$ with $\mathcal{C}$ under the assumption that $\mathcal{C}_f$ contains the maximizing cost vector.

3 Results

We only consider the perspective of the optimizing player here; analogous results for the adversary are given in the full version of the paper [24].

3.1 Discrete Scenario Uncertainty

Under discrete scenario uncertainty, we are given a finite set $\mathcal{S}$ of $|\mathcal{S}| = k$ scenarios. For each $S \in \mathcal{S}$, there exists a cost vector $c^S = (c_e^S)_{e\in E}$. Our solvability result for the optimizing player is the following.

Theorem 1. *For discrete scenario uncertainty, if the nominal problem $F^*(c)$ can be solved in time polynomial in n and m, then the corresponding randomized minmax regret problem $\min_{y\in\mathcal{Y}} \max_{S\in\mathcal{S}}(\overline{F}(y,c^S) - F^*(c^S))$ can be solved in time polynomial in n, m, and k.*

Recall that the feasible region $\mathcal{X}$ is described in size m. The algorithm for determining the optimizing player's mixed strategy, shown in Algorithm 1, solves two linear programs. The first is a linear programming relaxation of the deterministic minmax regret problem,

$$\begin{aligned} \min_{p,z} \quad & z \\ \text{s.t.} \quad & \sum_{e\in E} c_e^S p_e - F^*(c^S) \le z, \qquad \forall S \in \mathcal{S}, \\ & p \in \mathrm{CH}(\mathcal{X}),\ z \text{ free}, \end{aligned} \tag{LPD}$$

where $\mathrm{CH}(\mathcal{X})$ denotes the convex hull of $\mathcal{X}$ and $p \in [0,1]^n$. We refer to the vector $p = (p_1, \dots, p_n)$ as the *marginal probability vector*; it indicates the total probability that each item should be selected in the mixed strategy. Given the optimal vector p^* from solving (LPD), the second linear program maps the marginal probabilities to an optimal mixed strategy:

$$\begin{aligned} \max_{u,w} \quad & w - \sum_{e \in E} p_e u_e \\ \text{s.t.} \quad & w - \sum_{e \in T} u_e \leq 0, \qquad \forall T \in \mathcal{F}, \\ & u, w \text{ free}, \end{aligned} \tag{LPM}$$

where $u = (u_1, \dots, u_n)$. While both (LPD) and (LPM) potentially have an exponential number of constraints, a separation oracle is given by solvability of the nominal problem. Solvability of the latter program (LPM) is a known result from [14].

Algorithm 1 Rand-Minmax-Regret (discrete scenario uncertainty)

Input: Nominal combinatorial problem, cost vectors $(c^S)_{S \in \mathcal{S}}$

Output: Optimizing player's optimal mixed strategy $\mathcal{M}^* = (\Theta^*, Y^*)$ where $\Theta^* = (T_1, \dots, T_\mu)$ and $Y^* = (y_{T_1}, \dots, y_{T_\mu})$

1: Solve linear program (LPD) to get probability vector $p^* = (p_1^*, \dots, p_n^*)$.

2: Solve linear program (LPM) with $p = p^*$ to generate constraints indexed $i = 1, \dots, \mu$. Each constraint i corresponds to a set $T_i \in \mathcal{F}$ and dual variable y_{T_i}, indicating that T_i is an element in the optimal mixed strategy and has probability y_{T_i}.

A k-approximation algorithm for the deterministic minmax regret problem was introduced by Aissi et al. [3] and is shown in Algorithm 2. Using a new primal-dual interpretation with the formulation (LPD), as well as some arguments from [3], we show a simple proof of Theorem 2.

Algorithm 2 Mean-Cost-Approximation (Aissi et al. [3])

Input: Nominal combinatorial problem, cost vectors $(c^S)_{S \in \mathcal{S}}$

Output: Feasible solution $M \in \mathcal{F}$ satisfying $R_{\max}(M) \leq k Z_{\mathrm{D}}$.

1: Determine mean costs for each item: $d_e \leftarrow \frac{1}{k} \sum_{S \in \mathcal{S}} c_e^S, \quad \forall e \in E$.

2: Solve nominal problem with mean costs: $M \leftarrow \operatorname{argmin}_{T \in \mathcal{F}} \sum_{e \in T} d_e$.

Theorem 2. *For discrete scenario uncertainty, the solution to the nominal problem with mean costs is a k-approximation for the deterministic minmax regret problem.*

Since the randomized minmax regret problem corresponds to a linear programming relaxation of the deterministic minmax regret problem (specifically, the deterministic formulation is given by replacing the constraint $p \in \mathrm{CH}(\mathcal{X})$ with $x \in \mathcal{X}$ in (LPD)), it follows that $Z_{\mathrm{R}} \leq Z_{\mathrm{D}}$. Additionally, the primal-dual interpretation allows us to prove a new lower bound on Z_{R}, stated in Theorem 3 below. We show that this bound is tight.

Theorem 3. *For discrete scenario uncertainty and all nominal problems,*

$$Z_{\mathrm{R}} \geq \frac{Z_{\mathrm{D}}}{k}, \tag{12}$$

where $k = |\mathcal{S}|$ is the number of scenarios. Equivalently, the integrality gap of the mixed integer program corresponding to (LPD) *is equal to k.*

3.2 Interval Uncertainty

For interval uncertainty, each item cost is independently contained within known lower and upper bounds:

$$c_e \in [c_e^-, c_e^+], \quad \forall e \in E. \tag{13}$$

Define the region

$$\mathcal{I} := \{c \mid c_e \in [c_e^-, c_e^+], e \in E\}. \tag{14}$$

Our solvability result for interval uncertainty is the following.

Theorem 4. *For interval uncertainty, if the nominal problem $F^*(c)$ can be solved in time polynomial in n and m, then the corresponding randomized minmax regret problem $\min_{y \in \mathcal{Y}} \max_{c \in \mathcal{I}} (\overline{F}(y, c) - F^*(c))$ can be solved in time polynomial in n and m.*

The algorithm for determining the optimizing player's mixed strategy is shown in Algorithm 3. This is the same algorithm that is used for the discrete scenario uncertainty case, except the linear program (LPI) is used instead of (LPD),

$$\begin{aligned} \min_{p,z} \quad & z \\ \text{s.t.} \quad & \sum_{e \in E \setminus T} c_e^+ p_e - \sum_{e \in T} c_e^- (1 - p_e) \leq z, \qquad \forall T \in \mathcal{F}, \\ & p \in \mathrm{CH}(\mathcal{X}), \ z \text{ free}. \end{aligned} \tag{LPI}$$

For the deterministic minmax regret problem under interval uncertainty, the known 2-approximation algorithm of Kasperski and Zieliński [19] uses midpoint costs and is shown in Algorithm 4. Using primal-dual methods, we show a new proof for this algorithm as stated by Theorem 5. We also prove Theorem 6, establishing the integrality gap for interval uncertainty, and we show that the corresponding bound is tight.

Algorithm 3 RAND-MINMAX-REGRET (interval uncertainty)

Input: Nominal combinatorial problem, item cost bounds $(c_e^-, c_e^+),\ e \in E$.
Output: Optimizing player's optimal mixed strategy $\mathcal{M}^* = (\Theta^*, Y^*)$ where $\Theta^* = (T_1, \ldots, T_\mu)$ and $Y^* = (y_{T_1}, \ldots, y_{T_\mu})$
1: Solve linear program (LPI) to get probability vector $p^* = (p_1^*, \ldots, p_n^*)$.
2: Solve linear program (LPM) with $p = p^*$ to generate constraints indexed $i = 1, \ldots, \mu$. Each constraint i corresponds to a set $T_i \in \mathcal{F}$ and dual variable y_{T_i}, indicating that T_i is an element in the optimal mixed strategy and has probability y_{T_i}.

Algorithm 4 MIDPOINT-COST-APPROXIMATION (Kasperski and Zieliński [19])

Input: Nominal combinatorial problem, item cost bounds $(c_e^-, c_e^+),\ e \in E$.
Output: Feasible solution $M \in \mathcal{F}$ satisfying $R_{\max}(M) \leq 2Z_{\mathrm{D}}$.
1: Determine midpoint costs for each item: $d_e \leftarrow \left(\frac{c_e^- + c_e^+}{2}\right), \quad \forall e \in E.$
2: Solve nominal problem with midpoint costs: $M \leftarrow \underset{T \in \mathcal{F}}{\operatorname{argmin}} \sum_{e \in T} d_e$.

Theorem 5. *For interval uncertainty, the solution to the nominal problem with midpoint costs is a* 2*-approximation for the deterministic minmax regret problem.*

Theorem 6. *For interval uncertainty and all nominal problems,*

$$Z_{\mathrm{R}} \geq \frac{Z_{\mathrm{D}}}{2}. \tag{15}$$

Equivalently, the integrality gap of the mixed integer program corresponding to (LPI) *is equal to* 2*.*

3.3 Convex Uncertainty

If the uncertainty set $\mathcal{C}$ is allowed to be a general nonnegative convex set and the nominal problem is polynomial solvable, we show that the maximum expected regret problem becomes NP-hard. This result implies that both randomized and deterministic minmax regret problems are NP-hard under convex uncertainty, since both are at least as hard as the maximum expected regret problem.

Theorem 7. *For polynomial solvable nominal problems* $F^*(c) = \min_{x \in \mathcal{X}} \sum_{e \in E} c_e x_e$ *and nonnegative convex uncertainty sets* $\mathcal{C}$*, the maximum expected regret problem* $\max_{c \in \mathcal{C}} \left(\sum_{e \in E} c_e p_e - F^*(c)\right)$ *where* $p \in \mathrm{CH}(\mathcal{X})$ *is* NP*-hard.*

4 Conclusion

Our results on lower bounds for randomized minmax regret in relation to deterministic minmax regret, specifically Theorem 3 and Theorem 6, have important implications for approximating deterministic minmax regret problems. Theorem 3 indicates that the integrality gap for the minmax regret problem under discrete scenario uncertainty is equal to k, and it is easy to create instances of nearly all nominal problems that achieve this gap. This also holds true for the integrality gap of 2 under interval uncertainty. In Kasperski [17], it is posed as an open problem whether or not there exist approximation algorithms under interval uncertainty that, for some specific nominal problems, achieve an approximation ratio better than 2. We have answered this question in the negative for approximation schemes based on our linear programming relaxations.

An important future step with randomized minmax regret research is to develop approximation algorithms for dealing with nominal problems that are already NP-hard. This problem is non-trivial: an algorithm with an approximation factor α for a nominal problem does not immediately yield an algorithm to approximate the randomized minmax regret problem with a factor α.

References

1. Aissi, H., Bazgan, C., Vanderpooten, D.: Complexity of the min–max and min–max regret assignment problems. Operations Research Letters 33(6), 634–640 (2005)
2. Aissi, H., Bazgan, C., Vanderpooten, D.: Complexity of the min-max (regret) versions of cut problems. In: Algorithms and Computation, pp. 789–798. Springer (2005)
3. Aissi, H., Bazgan, C., Vanderpooten, D.: Approximating min-max (regret) versions of some polynomial problems. In: Computing and Combinatorics, pp. 428–438. Springer (2006)
4. Aissi, H., Bazgan, C., Vanderpooten, D.: Approximation of min–max and min–max regret versions of some combinatorial optimization problems. European Journal of Operational Research 179(2), 281–290 (2007)
5. Aissi, H., Bazgan, C., Vanderpooten, D.: Min–max and min–max regret versions of combinatorial optimization problems: A survey. European Journal of Operational Research 197(2), 427–438 (2009)
6. Aissi, H., Bazgan, C., Vanderpooten, D.: General approximation schemes for min–max (regret) versions of some (pseudo-) polynomial problems. Discrete Optimization 7(3), 136–148 (2010)
7. Averbakh, I.: On the complexity of a class of combinatorial optimization problems with uncertainty. Mathematical Programming 90(2), 263–272 (2001)
8. Averbakh, I., Lebedev, V.: Interval data minmax regret network optimization problems. Discrete Applied Mathematics 138(3), 289–301 (2004)
9. Averbakh, I., Lebedev, V.: On the complexity of minmax regret linear programming. European Journal of Operational Research 160(1), 227–231 (2005)
10. Bertsimas, D., Nasrabadi, E., Orlin, J.B.: On the power of nature in robust discrete optimization. Submitted (2013)
11. Bertsimas, D., Nasrabadi, E., Orlin, J.B.: On the power of randomization in network interdiction. arXiv preprint (2013), http://arxiv.org/abs/1312.3478
12. Bertsimas, D., Sim, M.: The price of robustness. Operations Research 52(1), 35–53 (2004)
13. Birge, J.R., Louveaux, F.V.: Introduction to Stochastic Programming. Springer (1997)

14. Grötschel, M., Lovász, L., Schrijver, A.: The ellipsoid method and its consequences in combinatorial optimization. Combinatorica 1(2), 169–197 (1981)
15. Immorlica, N., Kalai, A.T., Lucier, B., Moitra, A., Postlewaite, A., Tennenholtz, M.: Dueling algorithms. In: Proceedings of the Forty-Third annual ACM Symposium on Theory of Computing. pp. 215–224. ACM (2011)
16. Johannes, B., Orlin, J.B.: Minimax regret problems are harder than minimax problems. Submitted (2012)
17. Kasperski, A.: Discrete Optimization with Interval Data: Minmax Regret and Fuzzy Approach. Springer (2008)
18. Kasperski, A., Kurpisz, A., Zieliński, P.: Approximating the min–max (regret) selecting items problem. Information Processing Letters 113(1), 23–29 (2013)
19. Kasperski, A., Zieliński, P.: An approximation algorithm for interval data minmax regret combinatorial optimization problems. Information Processing Letters 97(5), 177–180 (2006)
20. Kasperski, A., Zieliński, P.: On the existence of an fptas for minmax regret combinatorial optimization problems with interval data. Operations Research Letters 35(4), 525–532 (2007)
21. Kiekintveld, C., Jain, M., Tsai, J., Pita, J., Ordóñez, F., Tambe, M.: Computing optimal randomized resource allocations for massive security games. In: Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems-Volume 1. pp. 689–696. International Foundation for Autonomous Agents and Multiagent Systems (2009)
22. Korzhyk, D., Conitzer, V., Parr, R.: Complexity of computing optimal stackelberg strategies in security resource allocation games. In: Twenty-Fourth AAAI Conference on Artificial Intelligence (2010)
23. Kouvelis, P., Yu, G.: Robust Discrete Optimization and Its Applications. Springer (1997)
24. Mastin, A., Jaillet, P., Chin, S.: Randomized minmax regret for combinatorial optimization under uncertainty. arXiv preprint (2014), `http://arxiv.org/abs/1401.7043`
25. Savage, L.J.: The theory of statistical decision. Journal of the American Statistical Association 46(253), 55–67 (1951)
26. Shapiro, A., Dentcheva, D., Ruszczyński, A.P.: Lectures on Stochastic Programming: Modeling and Theory. SIAM (2009)
27. Wald, A.: Contributions to the theory of statistical estimation and testing hypotheses. The Annals of Mathematical Statistics 10(4), 299–326 (1939)
28. Yu, G., Yang, J.: On the robust shortest path problem. Computers & Operations Research 25(6), 457–468 (1998)