Asymmetric data-driven interdiction problems with cost uncertainty: a distributionally robust optimization approach (*working paper*)

Sergey S. Ketkov

(joint work with Oleg A. Prokopyev)

Department of Business Administration, University of Zurich

30 July 2025

General interdiction problem

- A deterministic interdiction problem is a zero-sum game between a *leader* (attacker) and a *follower* (defender).
- The leader allocates limited interdiction resources to target components that are crucial to the follower's operations.
- In response, the follower observes the leader's actions and aims to maximize its own profit within the impacted environment.
- The leader is informed about the follower's objective and chooses attacks that minimize the follower's profit.

General interdiction problem

Detereministic interdiction (DI) problem

$$\textbf{[DI]:} \quad \min_{x \in X} \max_{y \in Y(x)} c^\top y, \tag{1}$$

where

$$X = \{x \in \{0,1\}^m : Hx \leq h\} \text{ and } Y(x) = \{y \in \mathbb{R}^n_+ : Fy + Lx \leq f\} \tag{2}$$

are, respectively, the leader's and the follower's feasible sets, and c is a deterministic profit vector.

Example: min-cost flow interdiction [1]

- Let $G = (N, A)$ be a directed graph, where N and A are, respectively, the sets of nodes and directed arcs.

- Assume that the leader has a single budget constraint, i.e.,

$$X = \Big\{\mathrm{x} \in \{0,1\}^{|A|} : \sum_{(i,j)\in A} r_{ij} x_{ij} \leq R\Big\}.$$

- The follower aims to maximize $-\sum_{(i,j)\in A} c_{ij} y_{ij}$ s.t.

$$\mathrm{y} \in Y(\mathrm{x}) = \Big\{\mathrm{y} \in \mathbb{R}_+^{|A|} : \sum_{j:\,(i,j)\in A} y_{ij} - \sum_{j:\,(j,i)\in A} y_{sj} = d_i \quad \forall i \in N,$$
$$0 \leq y_{ij} \leq u_{ij}(1 - x_{ij})\Big\}, \text{ where } \sum_{i\in N} d_i = 0.$$

Stochastic interdiction problem

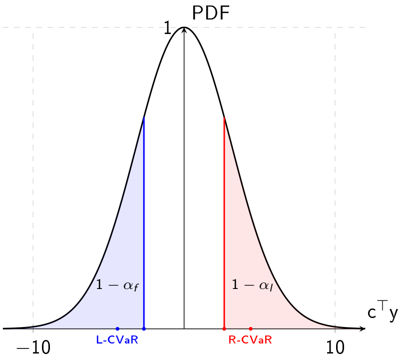

Figure: Illustration of the right-tail CVaR (red) and the left-tail CVaR (blue) for normal distribution with $\alpha_l = \alpha_f = 0.8$ [2].

- We assume that c is a random vector governed by an unknown *true* distribution $\mathbb{Q}^*$.
- If $\mathbb{Q}^*$ is given, then the decision-makers solve a stochastic interdiction (SI) problem of the form:

$$\textbf{[SI]}: \quad z^* := \min_{x,y} \ \rho_l(y, \mathbb{Q}^*)$$
$$\text{s.t. } x \in X$$
$$y \in \text{argmax}_{\tilde{y} \in Y(x)} \ \rho_f(\tilde{y}, \mathbb{Q}^*),$$

where

$$\rho_l(y, \mathbb{Q}^*) := \min_{t_l \in \mathbb{R}} \Big\{ t_l + \frac{1}{1-\alpha_l} \mathbb{E}_{\mathbb{Q}^*}\{(c^\top y - t_l)^+\} \Big\} \text{ and}$$
$$\rho_f(y, \mathbb{Q}^*) := \max_{t_f \in \mathbb{R}} \Big\{ t_f + \frac{1}{1-\alpha_f} \mathbb{E}_{\mathbb{Q}^*}\{(c^\top y - t_f)^-\} \Big\}.$$

Distributionally robust interdiction model

A1. The feasible sets X and $Y(\mathrm{x})$ for each $\mathrm{x} \in X$ are non-empty and bounded.

A2. The *support* of $\mathbb{Q}^*$ is known and given by a non-empty compact polyhedral set

$$S = \Big\{ \mathrm{c} \in \mathbb{R}^n : \mathrm{Bc} \leq \mathrm{b} \Big\}.$$

A3. The leader and the follower have access to two i.i.d. training data sets, generated according to $\mathbb{Q}^*$,

$$\hat{\mathrm{C}}_l = \Big\{ \hat{\mathrm{c}}_l^{(k)} \in S,\; k \in \{1, \ldots, k_l\} \Big\} \text{ and } \hat{\mathrm{C}}_f = \Big\{ \hat{\mathrm{c}}_f^{(k)} \in S,\; k \in \{1, \ldots, k_f\} \Big\}.$$

Furthermore, $\hat{\mathrm{C}}_f \subseteq \hat{\mathrm{C}}_l$ or, equivalently, $k_f \leq k_l$ and $\hat{\mathrm{c}}_f^{(k)} = \hat{\mathrm{c}}_l^{(k)}$ for each $k \in \{1, \ldots, k_f\}$.

Distributionally robust interdiction model

Basic DRI model

$$\textbf{[DRI]:} \quad \hat{z}_b^* := \min_{\mathrm{x},\mathrm{y}} \left\{ \max_{\mathbb{Q}_l \in \mathcal{Q}_l} \rho_l(\mathrm{y}, \mathbb{Q}_l) \right\}$$

$$\text{s.t. } \mathrm{x} \in X$$

$$\mathrm{y} \in \operatorname{argmax}_{\tilde{\mathrm{y}} \in Y(\mathrm{x})} \left\{ \min_{\mathbb{Q}_f \in \mathcal{Q}_f} \rho_f(\tilde{\mathrm{y}}, \mathbb{Q}_f) \right\},$$

where

$$\mathcal{Q}_i := \left\{ \mathbb{Q} \in \mathcal{Q}_0(S) : W^p(\hat{\mathbb{Q}}_i, \mathbb{Q}) \leq \varepsilon_i \right\} \quad i \in \{l, f\}$$

are the Wasserstein ambiguity sets, $\hat{\mathbb{Q}}_i(\hat{\mathrm{C}}_i)$ is an empirical distribution of the data and $W^p(\hat{\mathbb{Q}}_i, \mathbb{Q})$ is the type-1 Wasserstein distance w.r.t. ℓ_p-norm.

Basic DRI model

- Notably, **[DRI]** is strongly *NP*-hard, as even its deterministic version, **[DI]**, is known to be strongly *NP*-hard.
- If $\mathcal{Q}_l$ and $\mathcal{Q}_f$ are defined in terms of ℓ_1 or ℓ_∞ norm, then **[DRI]** admits a single-level MILP reformulation of polynomial size. This reformulation builds on LP reformulations of the worst-case CVaR problems for both decision-makers [3] and a strong duality-based reformulation of the follower's problem:

$$\hat{z}_b^* = \min_{\mathrm{x},\mathrm{y},\boldsymbol{\nu},\boldsymbol{s},\boldsymbol{\mu},\boldsymbol{\beta},\boldsymbol{\gamma},\lambda,t} \Big\{ t_l + \frac{1}{1-\alpha_l}\big(\lambda_l \varepsilon_l + \frac{1}{k_l}\sum_{k\in K_l} s_l^{(k)}\big)\Big\}$$

$$\begin{aligned}
\text{s.t. } & \text{primal and dual feasibility (follower)},\\
& \mathrm{x} \in X\\
& \left.\begin{aligned}
& \hat{\mathrm{c}}_l^{(k)\top}\mathrm{y} - t_l + \boldsymbol{\Delta}_l^{(k)\top}\boldsymbol{\nu}_l^{(k)} \leq s_l^{(k)}\\
& \|\mathrm{B}^\top\boldsymbol{\nu}_l^{(k)} - \mathrm{y}\|_* \leq \lambda_l\\
& \boldsymbol{\nu}_l^{(k)} \geq 0,\; s_l^{(k)} \geq 0
\end{aligned}\right\} \forall k \in K_l\\
& \color{red}{(-\mathrm{L}\mathrm{x}+\mathrm{f})^\top\boldsymbol{\beta}_f} = t_f - \frac{1}{1-\alpha_f}\big(\varepsilon_f\lambda_f + \frac{1}{k_f}\sum_{k\in K_f} s_f^{(k)}\big).
\end{aligned}$$

Basic DRI model

- As for the related one-stage DRO model [3], we demonstrate that, under mild assumptions, the basic DRI model is *asymptotically consistent*.

- That is, when the leader and the follower acquire more data, their optimal solutions and objective function values converge, in a sense, to those of the underlying stochastic programming problem [**SI**].

- Given $x \in X$, the follower's convergence, as $k_f \to \infty$, is due to Theorem 3.6 in [3].

- The leader's convergence, as $k_l \to \infty$ and $k_f \to \infty$, is based on the following assumption:
  - For every fixed $x \in X$, the full information follower's problem $\max_{y \in Y(x)} \rho_f(y, \mathbb{Q}^*)$ has a unique optimal solution.

  The proof also exploits discreteness of X, closedeness of the follower's optimal solution set and the bounded convergence theorem.

Example: packing interdiction

- $n = 4$ items, leader can block ≤ 2 items, follower selects ≤ 1 item.
- Let
$$S = \Big\{ \mathrm{c} \in \mathbb{R}^4 \, : \, c_1 \in [0,10], \; c_2 \in [1,11], \; c_3 \in [6,12], \; c_4 \in [7,13] \Big\},$$
with the true distribution $\mathbb{Q}^*$ being a *discrete uniform distribution* with expected costs $\bar{\mathrm{c}}^* = (5,6,9,10)^\top$.
- Full-information **[SI]** model:
$$z^* = \min_{\mathrm{x} \in X} \max_{\mathrm{y} \in Y(\mathrm{x})} \; \bar{\mathrm{c}}^{*\top} \mathrm{y},$$
where
$$X = \Big\{ \mathrm{x} \in \{0,1\}^4 \, : \, \sum_{i=1}^{4} x_i \leq 2 \Big\} \; \text{ and } \; Y(\mathrm{x}) = \Big\{ \mathrm{y} \in [0,1]^4 \, : \, \mathrm{y} \leq \mathbf{1} - \mathrm{x}, \; \sum_{i=1}^{4} y_i \leq 1 \Big\}$$

Figure: Optimal solution: $\mathrm{x}^* = (0,0,1,1)^\top$, $\mathrm{y}^* = (0,1,0,0)^\top$, $z^* = 6$.

Example: packing interdiction

- We assume that both decision-makers implement a myopic *sample average approximation* based on $k_l = k_f = 5$ samples given by:

$$\hat{C}_l = \begin{pmatrix} 10 & 11 & 8 & 7 \\ 6 & 7 & 7 & 9 \\ 10 & 10 & 7 & 9 \\ 9 & 11 & 6 & 7 \\ 10 & 11 & 7 & 8 \end{pmatrix} \quad \text{and} \quad \hat{C}_f = \begin{pmatrix} 3 & 3 & 8 & 13 \\ 1 & 2 & 9 & 10 \\ 3 & 11 & 11 & 13 \\ 10 & 3 & 10 & 12 \\ 4 & 1 & 12 & 7 \end{pmatrix},$$

with average profits $\bar{c}_l := (9, 10, 7, 8)^\top$ and $\bar{c}_f := (4, 5, 10, 11)^\top$.

- Notably, Assumption **A3** is violated.

Example: packing interdiction

- **True basic model:** Hypothetical benchmark assuming the leader knows the follower's data and its optimal policy.
- **Pessimistic approximation:** Assumes no knowledge of the follower's data; leader optimizes against the worst-case feasible policy.
- **Semi-pessimistic approximation:** Leader has partial information about the follower's data, e.g., it is aware of the first two columns in $\hat{C}_f$. Then, the leader constructs an uncertainty set to estimate the follower's average profits and makes a robust decision.
- **Augmented basic model:** The leader replaces the missing follower's data with its own, effectively assuming both players use the same data.

Model	**Leader's Decision**	**Leader Expected Obj.**	**Leader's True Obj.**
Full information	$(0,0,1,1)^\top$	6	6
True basic	$(0,0,0,1)^\top$	7	9
Pessimistic	$(1,1,0,0)^\top$	8	10
Semi-pessimistic	$(0,0,0,1)^\top$	7	9
Augmented basic	$(1,1,0,0)^\top$	8	10

True basic model

True basic model

$$[\mathbf{DRI}^*]: \quad z_b^* := \min_{\mathrm{x},\mathrm{y}} \left\{ \max_{\mathbb{Q}_l \in \mathcal{Q}_l} \rho_l(\mathrm{y}, \mathbb{Q}_l) \right\}$$

$$\text{s.t. } \mathrm{x} \in X$$

$$\mathrm{y} \in \operatorname{argmax}_{\tilde{\mathrm{y}} \in Y(\mathrm{x})} \left\{ \min_{\mathbb{Q}_f \in \mathcal{Q}_f(\hat{\mathrm{C}}_f^*)} \rho_f(\tilde{\mathrm{y}}, \mathbb{Q}_f) \right\},$$

where $\hat{\mathrm{C}}_f^*$ denotes the true follower's data set, and the second part of Assumption **A3** does not necessarily hold.

Pessimistic approximation

Pessimistic approximation

$$[\mathbf{DRI\text{-}P}]: \quad \hat{z}_p^* := \min_{\mathrm{x} \in X} \max_{\mathrm{y} \in Y(\mathrm{x})} \max_{\mathbb{Q}_l \in \mathcal{Q}_l} \rho_l(\mathrm{y}, \mathbb{Q}_l).$$

- The leader in **[DRI-P]** disregards any available information about the true follower's data set $\hat{\mathrm{C}}_f^*$.
- It selects the worst-case *feasible* follower's policy in terms of the leader's objective function value.

Pessimistic approximation

- We show that **[DRI-P]** is Σ_2^p-hard by a reduction from the dominating set interdiction problem [4].
- To solve **[DRI-P]**, we design a Benders decomposition algorithm tailored to two special cases of the problem, where the leader is either *risk-neutral* or *ambiguity-free*.
- The algorithm for each case is based on the standard decomposition techniques for bilevel optimization and leverages a disjoint bilinear structure of the inner optimization problem.

Semi-pessimistic approximation

Semi-pessimistic approximation

$$[\textbf{DRI-SP}]: \quad \hat{z}^*_{sp} := \min_{x \in X} \max_{\hat{\mathrm{C}}_f \in \hat{S}_l} \min_{y} \left\{ \max_{\mathbb{Q}_l \in \mathcal{Q}_l} \rho_l(y, \mathbb{Q}_l) \right\}$$

$$\text{s.t. } y \in \operatorname{argmax}_{\tilde{y} \in Y(x)} \left\{ \min_{\mathbb{Q}_f \in \mathcal{Q}_f(\hat{\mathrm{C}}_f)} \rho_f(\tilde{y}, \mathbb{Q}_f) \right\}.$$

- We relax Assumption **A3**, i.e., for each $k \in \{1, \ldots, k_f\}$, the leader either knows that

$$\hat{c}_f^{(k)} \subseteq \hat{S}_l^{(k)} := \left\{ c \in \mathbb{R}^n : \underline{b}^{(k)} \leq c \leq \overline{b}^{(k)} \right\} \subseteq S \quad \text{or} \quad \hat{c}_f^{(k)} \in \hat{\mathrm{C}}_l.$$

- The leader in **[DRI-SP]** assumes the worst-case possible realization of $\hat{\mathrm{C}}_f^*$ set w.r.t. the uncertainty set $\hat{S}_l = \hat{S}_l^{(1)} \times \ldots \times \hat{S}_l^{(k_f)}$.

Semi-pessimistic approximation

- We show that [**DRI-SP**] is Σ_2^p-hard from the robust optimistic bilevel problem with interval uncertainty [5].
- We propose to use a discretization of the uncertainty set $\hat{S}_l$, based on a finite number of scenarios for the follower's data set $\hat{C}_f$. The discretization is shown to admit a single-level MILP formulation, whose size however is proportional to the number of scenarios.
- To justify our approach, we show that the discretized semi-pessimistic approximation is almost surely *robust* with respect to the follower's data in an asymptotic sense, i.e., when the number of scenarios tends to infinity.

Computational settings

- We consider a class of general interdiction problems defined as:

$$\min_{x \in X} \max_{y \in Y(x)} c^\top y,$$

where

$$X = \big\{x \in \{0,1\}^n : Hx \leq h\ \big\} \text{ and } Y(x) = \big\{y \in \mathbb{R}^n_+ : \tilde{F}y \leq \tilde{f},\ y \leq U(1 - x)\big\}.$$

- $d_l = \dim(h) = 1$, $\tilde{d}_f = \dim(\tilde{f}) = 10$ and $U = I$.
- All elements of H and $\tilde{F}$ are generated uniformly at random from the interval $[0.01, 1]$, whereas

$$h_j = 0.4 \sum_{i=1}^{n} H_{ji} \quad \forall j \in \{1, \dots, d_l\}, \quad f_j = 0.4 \sum_{i=1}^{n} F_{ji} \quad \forall j \in \{1, \dots, \tilde{d}_f\}.$$

Computational settings

- The support set is given by:

$$S := \{\mathrm{c} \in \mathbb{R}^n : c_i \in [0.01, 1] \quad \forall i \in N := \{1, \ldots, n\}\}.$$

- The true distribution $\mathbb{Q}^*$ is defined as a product of truncated normal distributions for each component c_i, $i \in N$, with increasing mean and variance, similar to [3].
- We set $\varepsilon_l := \frac{\delta_l}{\sqrt{k_l}}$ and $\varepsilon_f := \frac{\delta_f}{\sqrt{k_f}}$ for some $\delta_l, \delta_f \in \mathbb{R}_+$, and $\alpha_l = \alpha_f = 0.95$.
- To define $\hat{S}_l$, it is assumed that the first $k_{lf} \leq k_f$ samples in $\hat{\mathrm{C}}_f^*$ are known to the leader; for each $i \in N$ and $k \in \{k_{lf} + 1, \ldots, k_f\}$, let

$$\big(\hat{\mathrm{c}}_f^{(k)}\big)_i \in [c_{ki}^* - \kappa\Delta_{ki}, c_{ki}^* + \kappa(1 - \Delta_{ki})] \cap [0.01, 1].$$

Here, κ is a fixed noise level, and $\Delta_{ki} \in [0, 1]$ is a shift parameter (Δ_{ki} is 0 with probability 0.5, and with probability 0.5, it is selected uniformly at random from the interval $[0, 1]$).

Performance metrics

- Let $(\hat{x}^*, \hat{y}^*)$ be our obtained solution and $\hat{z}^*$ the respective optimal objective function value.
- The follower's relative out-of-sample loss is defined as:

$$\mathrm{RL}_f^{(\mathrm{out})} = \frac{\rho_f(\hat{y}^*, \mathbb{Q}^*)}{\max_{y \in Y(\hat{x}^*)} \rho_f(y, \mathbb{Q}^*)} \leq 1$$

- The leader's relative out-of-sample loss is defined as:

$$\mathrm{RL}_l^{(\mathrm{out})} = \frac{\rho_l(\hat{y}^*, \mathbb{Q}^*)}{\rho_l(\tilde{y}^*, \mathbb{Q}^*)} \geq 1,$$

  where $\tilde{y}^*$ solves the full information leader's problem.
- The relative leader's in-sample loss is defined as:

$$\mathrm{RL}_l^{(\mathrm{in})} = \frac{\hat{z}^*}{z_b^*}.$$

Selected computational results

- $n = 10$, 10 random test instances, and 10 data sets for each instance.

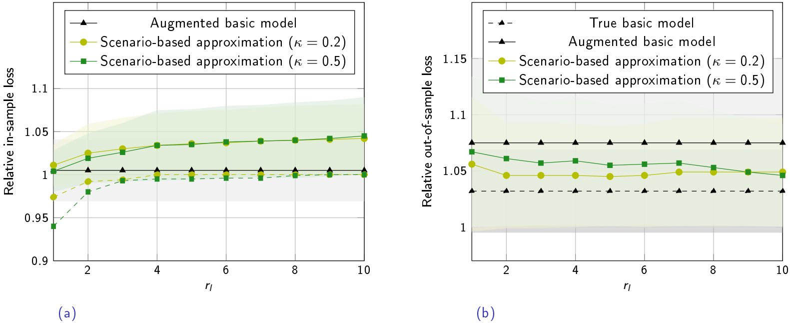

Figure: The average relative in-sample and out-of-sample loss (with MADs) as a function of the number of scenarios, r_l, for $k_l = k_f = 30$, $\delta_l = \delta_f = 0.1$ and $k_{lf} = 20$. The dashed lines in (a) correspond to the empirical 5% percentile of the relative in-sample loss.

Selected computational results

- $n = 10$, 10 random test instances, and 10 data sets for each instance.

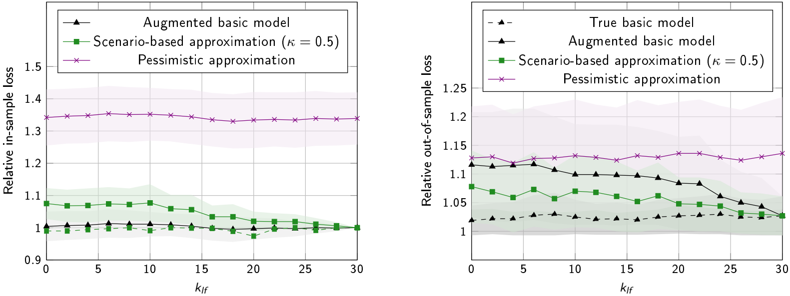

(a) (b)

Figure: The average relative in-sample and out-of-sample loss of the *ambiguity-free* leader (with MADs) as a function of k_{lf}, for $k_l = k_f = 30$, $\delta_l = \delta_f = 0.1$ and $\alpha_l = 0.9$. The follower is assumed to be *risk-averse*. The dashed line corresponds to the empirical 5% percentile of the relative in-sample loss for the scenario-based semi-pessimistic approximation.

Selected computational results

- $n = 10$, 10 random test instances, and 10 data sets for each instance.

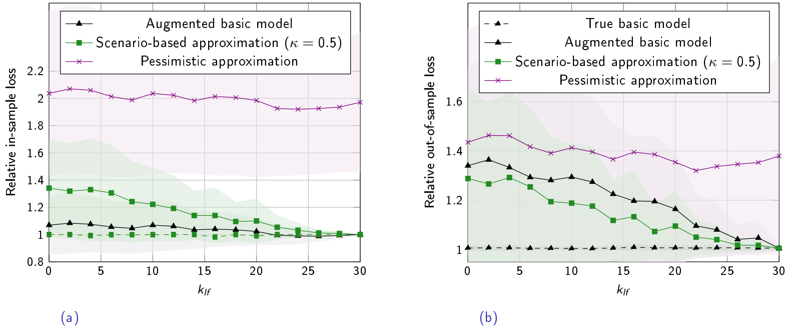

Figure: The average relative in-sample and out-of-sample loss of the *risk-neutral* leader (with MADs) as a function of k_{lf}, for $k_l = k_f = 30$ and $\delta_l = \delta_f = 0.1$. The follower is assumed to be *risk-averse*. The dashed line corresponds to the empirical 5% percentile of the relative in-sample loss for the scenario-based semi-pessimistic approximation.

Selected computational results

- $n = 10$, 10 random test instances, and 10 data sets for each instance.

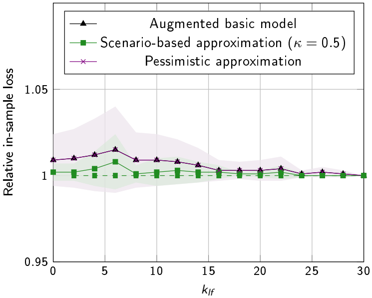

(a)

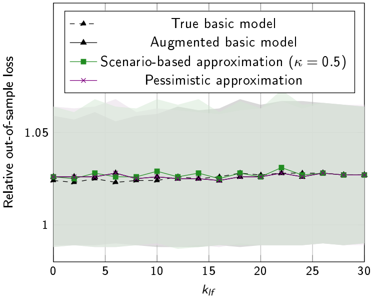

(b)

Figure: The average relative in-sample and out-of-sample loss of the *risk-neutral* leader (with MADs) as a function of k_{lf}, for $k_l = k_f = 30$ and $\delta_l = \delta_f = 0.1$. The follower is assumed to be *risk-neutral*. The dashed line corresponds to the empirical 5% percentile of the relative in-sample loss for the scenario-based semi-pessimistic approximation.

Selected computational results

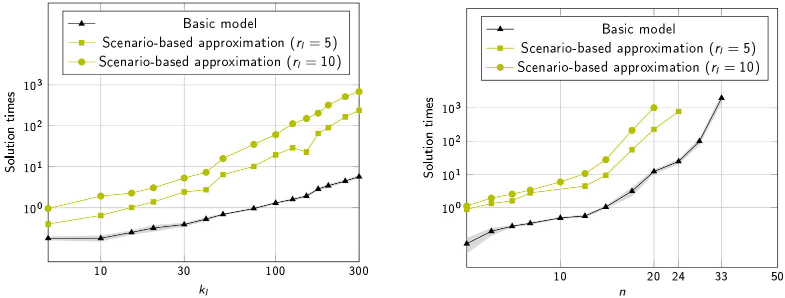

(a) $k_f = k_l, n = 10$.

(b) $k_l = k_f = 30$.

Figure: Average solution times in seconds (with MADs) of the basic and semi-pessimistic formulations as a function of $k_l = k_f$ (a) and n (b), for $\delta_l = \delta_f = 0.1$, $k_{lf} = \lfloor \frac{2}{3} k_f \rfloor$ and $\kappa = 0.2$, evaluated over 10 random test instances. The time limit is set to 60 minutes.

Selected computational results

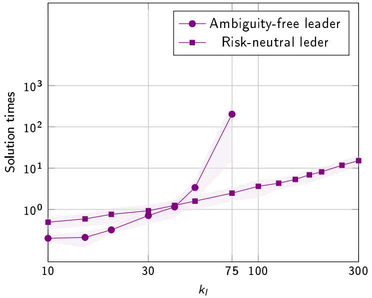

(a) $n = 10$.

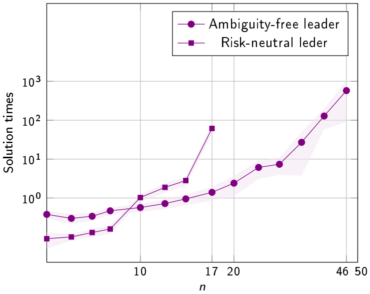

(b) $k_l = 30$.

Figure: Average solution times in seconds (with MADs) of the pessimistic formulation as a function of k_l (a) and n (b), for $\delta_l = 0.1$ and $\alpha_l = 0.9$, evaluated over 10 random test instances. The time limit is set to 60 minutes.

References

[1] J. C. Smith and C. Lim, "Algorithms for network interdiction and fortification games," in *Pareto optimality, game theory and equilibria*, pp. 609–644, Springer, 2008.

[2] R. T. Rockafellar and S. Uryasev, "Optimization of conditional value-at-risk," *Journal of Risk*, vol. 2, pp. 21–42, 2000.

[3] P. M. Esfahani and D. Kuhn, "Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations," *Mathematical Programming*, vol. 171, no. 1-2, pp. 115–166, 2018.

[4] C. Grüne and L. Wulf, "Completeness in the polynomial hierarchy for many natural problems in bilevel and robust optimization," in *International Conference on Integer Programming and Combinatorial Optimization*, pp. 256–269, Springer, 2025.

[5] C. Buchheim, D. Henke, and F. Hommelsheim, "On the complexity of robust bilevel optimization with uncertain follower's objective," *Operations Research Letters*, vol. 49, no. 5, pp. 703–707, 2021.