Dual Dynamic Programming for Data-driven Distributionally Robust Multistage Convex Optimization

Shixuan Zhang[1] Xu Andy Sun[2]

[1]Institute for Computational and Experimental Research in Mathematics, Brown University

[2]Operations Research and Statistics, Sloan School of Management, Massachusetts Institute of Technology

INFORMS, Oct 18, 2022

Outline for Part 1

Multistage Distributionally Robust Optimization

Consider a multistage distributionally robust optimization (MDRO) problem

$$\begin{aligned}\min_{x_1\in\mathcal{X}_1} \quad & f_1(x_0,x_1;\xi_1) + \sup_{p_2\in\mathcal{P}_2} \mathbb{E}_{\xi_2\sim p_2} \min_{x_2\in\mathcal{X}_2} \Bigg[f_2(x_1,x_2;\xi_2) + \\ & + \sup_{p_3\in\mathcal{P}_3} \mathbb{E}_{\xi_3\sim p_3} \min_{x_3\in\mathcal{X}_3} \Bigg[f_3(x_2,x_3;\xi_3) + \cdots \\ & + \sup_{p_T\in\mathcal{P}_T} \mathbb{E}_{\xi_T\sim p_T} \min_{x_T\in\mathcal{X}_T} f_T(x_{T-1},x_T;\xi_T) \Bigg]\Bigg].\end{aligned}$$

Here, for each stage $t = 1, \ldots, T$,

- $\mathcal{X}_t \subset \mathbb{R}^{d_t}$ is a compact set of feasible decisions x_t and f_t is a nonnegative lsc cost function;
- ξ_t is the uncertainty vector, which obeys a Borel probability distribution p_t over all possible realizations Ξ_t in a given ambiguity set $\mathcal{P}_t$.

This framework encompasses multistage stochastic optimization (MSO) when $|\mathcal{P}_t| = 1$ and multistage robust optimization (MRO) when $\delta_\xi \in \mathcal{P}_t$ for all $\xi \in \Xi_t$.

Dynamic Programming Recursion

The MDRO problem can be defined recursively using the deterministic *(worst-case) expected cost-to-go functions*

$$\mathcal{Q}_{t-1}(x_{t-1}) := \sup_{p_t \in \mathcal{P}_t} \mathbb{E}_{\xi_t \sim p_t} \left(\min_{x_t \in \mathcal{X}_t} f_t(x_{t-1}, x_t; \xi_t) + \mathcal{Q}_t(x_t) \right),$$

for $t = T, \ldots, 2$, where $\mathcal{Q}_T \equiv 0$. We can also define the *value functions* to simplify the notation

$$Q_t(x_{t-1}; \xi_t) := \min_{x_t \in \mathcal{X}_t} f_t(x_{t-1}, x_t; \xi_t) + \mathcal{Q}_t(x_t),$$

for each $t = 1, 2, \ldots, T$ so

$$\mathcal{Q}_{t-1}(x_{t-1}) = \sup_{p_t \in \mathcal{P}_t} \mathbb{E}_{\xi_t \sim p_t} Q_t(x_{t-1}; \xi_t)$$

for each $t \geq 2$, and the optimal value of the MDRO is $Q_1(x_0; \xi_1)$.

Dual dynamic programming (DDP) algorithms construct under-approximations $\underline{\mathcal{Q}}_t^i$, and optionally over-approximations $\overline{\mathcal{Q}}_t^i$, for $\mathcal{Q}_t$ iteratively.

Illustration of Dual Dynamic Programming Algorithms

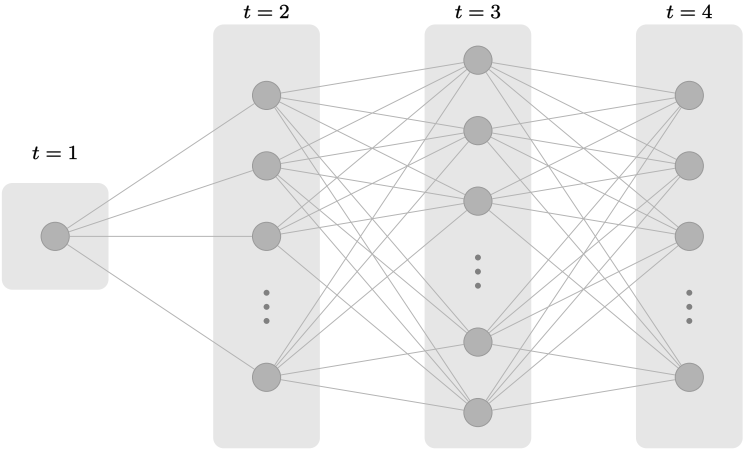

Figure: Illustration of DDP on a 4-stage problem

Illustration of Dual Dynamic Programming Algorithms

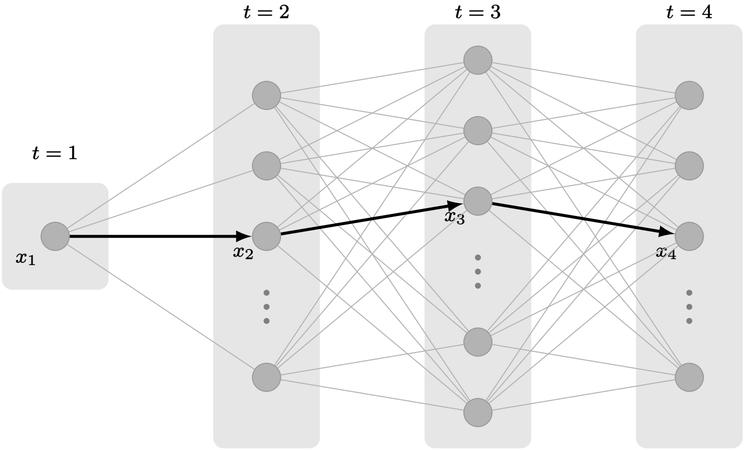

Figure: Illustration of DDP on a 4-stage problem

Illustration of Dual Dynamic Programming Algorithms

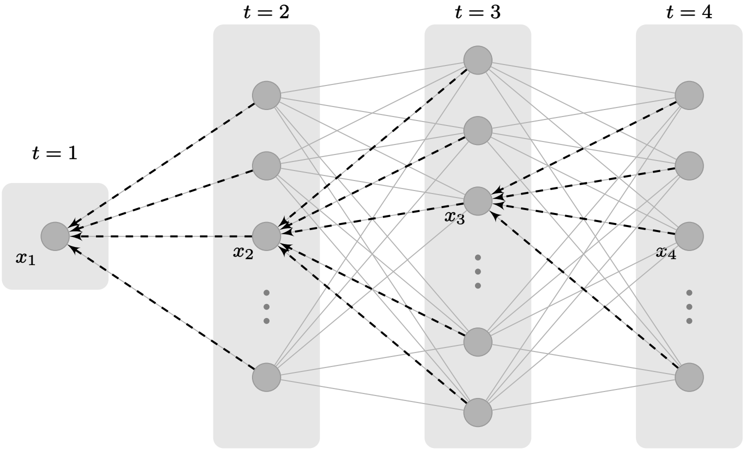

Figure: Illustration of DDP on a 4-stage problem

Literature Review and Motivational Questions

- Stochastic DDP (SDDP) for linear MSO: Pereira and Pinto (1991)
- SDDP Convergence: Shapiro (2011), Girardeau et al. (2015), Baucke et al. (2017)
- DDP for mixed-integer linear MSO: (Zou et al. 2019, Ahmed et al. 2020)
- Robust DDP (RDDP) for linear MRO: Georghiou et al. (2019)
- DDP for risk-averse MSO and MDRO: Philpott et al. (2013), Shapiro et al. (2013), Philpott et al. (2018), Duque and Morton (2020).
- DDP complexity for MSO: Lan (2020), Zhang and Sun (2022b)

Motivational Questions

1. Are the DDP complexity results valid for MDRO problems?
   - In particular, does the MSO linear scalability still hold?
2. How do we solve MDRO-DDP recursion problems, especially for data-driven models?

Outline for Part 2

New DDP Framework: Single Stage Subproblems

We define single stage subproblem oracles (SSSO) because

- there could be infinite outcomes of the uncertainty in each stage;
- the information passing through stages need to be quantified precisely.

For the initial stage, the SSSO is required to provide an optimal solution for the problem $x_1^* \in \min_{x_1 \in \mathcal{X}_1} f_1(x_0, x_1; \xi_1) + \underline{\mathcal{Q}}_1(x_1)$ and the associated gap $\gamma_1 := \overline{\mathcal{Q}}_1(x_1^*) - \underline{\mathcal{Q}}_1(x_1^*)$.

Definition (Noninitial stage subproblem oracle)

Given a feasible state $x_{t-1} \in \mathcal{X}_{t-1}$, the noninitial SSSO provides a feasible state $x_t \in \mathcal{X}_t$, an L_t-Lipschitz continuous cut $\mathcal{V}_{t-1}(\cdot)$, and an over-estimate value v_{t-1} such that

1. they are valid, i.e., $\mathcal{V}_{t-1}(x) \leq \mathcal{Q}_{t-1}(x)$ for any $x \in \mathcal{X}_{t-1}$ and $v_{t-1} \geq \mathcal{Q}_{t-1}(x_{t-1})$;
2. the gap is controlled, i.e., $v_{t-1} - \mathcal{V}_{t-1}(x_{t-1}) \leq \gamma_t := \overline{\mathcal{Q}}_t(x_t) - \underline{\mathcal{Q}}_t(x_t)$.

New DDP Framework: Complexity Analysis

We show that the total number of SSSO evaluations before termination is bounded by (Zhang and Sun 2020)

$$\texttt{\#Eval} \leq 1 + T \cdot \inf_{\delta} \left\{ \sum_{t=1}^{T-1} \left(1 + \frac{2L_t D_t}{\delta_t - \delta_{t+1}}\right)^{d_t} : \varepsilon = \delta_1 > \delta_2 > \cdots > \delta_T = 0 \right\}.$$

- In a relative optimality scale, this upper bound can be specialized to $\mathcal{O}(T^2)$;
- and further improved to $\mathcal{O}(T)$ if we allow nonconsecutive updates in the DDP.

It is also at least

$$\texttt{\#Eval} \geq \frac{d}{d-1} \sqrt{\frac{\pi}{2}(d^2 - 4)} \left(\frac{DL(T-2)}{16\varepsilon}\right)^{d/2-1}.$$

- This reconfirms the well-known "curse of dimensionality" in the state space even for convex problems.

Data-driven DDP: Wasserstein Ambiguity

Given a distance function $d_t : \Xi_t \to \mathbb{R}$, Wasserstein (1-)distance is defined by

$$W_t(\mu, \nu) := \inf_{\pi \in \mathcal{M}^{\mathrm{Prob}}(\Xi_t \times \Xi_t)} \left\{ \int_{\Xi_t \times \Xi_t} d_t(\xi^1, \xi^2) \, \mathrm{d}\pi(\xi^1, \xi^2) : P_*^1(\pi) = \mu, \; P_*^2(\pi) = \nu \right\},$$

for any two probability measures μ, ν on Ξ_t. We define the Wasserstein ambiguity set $\mathcal{P}_t$ centered at $\hat{\nu}_t$ as

$$\mathcal{P}_t := \left\{ p \in \mathcal{W}_t : \; W_t(p, \hat{\nu}_t) \le \rho_{t,0}, \; \langle g_{t,j}, p \rangle \le \rho_{t,j}, \; j = 1, \ldots, m_t \right\}$$

with any given vector $\rho_t \in \mathbb{R}^{m_t+1}$ and continuous functions $g_{t,j} : \Xi_t \to \mathbb{R}$ for $j = 1, \ldots, m_t$.

Data-driven DDP: Out-of-Sample Performance Guarantee

Suppose $\delta_t \geq 3$. We have the out-of-sample performance guarantee with probability at least $\alpha \in (0,1)$ if either of the following conditions holds for each $t \in \mathcal{T}$ (Zhang and Sun 2022a):

1. the probability measure ν_t is sub-Gaussian, and
$$n_t \cdot \rho_{t,0}^{\delta_t} \geq \frac{1}{C'_t}\left[\ln C_t - \ln\left(1 - \alpha^{1/(T-1)}\right)\right],$$
2. the probability measure ν_t has finite third moments and
$$n_t \cdot \rho_{t,0}^{2} \geq \frac{C''_t}{1 - \alpha^{1/(T-1)}},$$

where C_t, C'_t, and C''_t are the positive constants depending only on ν_t, $t = 2, \dots, T$.

Data-driven DDP: Finite-Dimensional Dual Recursion

We show that the expected cost-to-go function can be equivalently rewritten as

$$\mathcal{Q}_{t-1}(x_{t-1}) = \min_{\lambda \geq 0} \left\{ \sum_{j=0}^{m_t} \rho_{t,j}\lambda_j + \frac{1}{n_t} \sum_{k=1}^{n_t} \sup_{\xi_k \in \Xi_t} \left\{ Q_t(x_{t-1};\xi_k) - \lambda_0 \, d_{t,k}(\xi_k) - \sum_{j=1}^{m_t} \lambda_j g_{t,j}(\xi_k) \right\} \right\}.$$

We design exact SSSO implementations based on this recursion assuming either

- the cost function $f_t(x_{t-1}, x_t; \xi_t)$ is usc, concave in ξ_t for any $x_{t-1} \in \mathcal{X}_{t-1}$, $x_t \in \mathcal{X}_t$ with $g_{t,j}$'s being convex; or
- f_t is jointly convex in (x_t, ξ_t) for any $x_{t-1} \in \mathcal{X}_{t-1}$ and Ξ_t, d_t are polyhedral, with $g_{t,j}$'s being concave.

Data-driven DDP: Finite-Dimensional Dual Recursion

We show that the expected cost-to-go function can be equivalently rewritten as

$$\mathcal{Q}_{t-1}(x_{t-1}) = \min_{\lambda \geq 0} \left\{ \sum_{j=0}^{m_t} \rho_{t,j}\lambda_j + \frac{1}{n_t} \sum_{k=1}^{n_t} \sup_{\xi_k \in \Xi_t} \left\{ Q_t(x_{t-1};\xi_k) - \lambda_0 d_{t,k}(\xi_k) - \sum_{j=1}^{m_t} \lambda_j g_{t,j}(\xi_k) \right\} \right\}.$$

We design exact SSSO implementations based on this recursion assuming either

- the cost function $f_t(x_{t-1}, x_t; \xi_t)$ is usc, concave in ξ_t for any $x_{t-1} \in \mathcal{X}_{t-1}$, $x_t \in \mathcal{X}_t$ with $g_{t,j}$'s being convex; or
- f_t is jointly convex in (x_t, ξ_t) for any $x_{t-1} \in \mathcal{X}_{t-1}$ and Ξ_t, d_t are polyhedral, with $g_{t,j}$'s being concave.

Outline for Part 3

Multi-commodity Inventory Problem

Now we consider a distributionally robust multi-commodity inventory problem defined recursively as

$$\mathcal{Q}_{t-1}(x_{t-1}) := \sup_{p_t \in \mathcal{P}_t} \int_{\Xi_t} \left[f_t(x_{t-1}, x_t; \xi_t) + \mathcal{Q}_t(x_t) \right] \mathrm{d}p_t,$$

where

$$\begin{aligned} f_t(x_{t-1}, x_t; \xi_t) := \min \quad & C^F + \sum_{j \in \mathcal{J}} \left(C^a_j y^a_{t,j} + C^b_j x^b_{t,j} + C^H_j [x^l_{t,j}]_+ + C^B_j [x^l_{t,j}]_- + C^r_j y^r_{t,j} \right) + \mathcal{Q}_t(x_t) \\ \text{s.t.} \quad & \sum_{j \in \mathcal{J}} y^a_{t,j} \le B^c, \\ & x^l_{t,j} - y^a_{t,j} - x^b_{t-1,j} \ge x^l_{t-1,j} - D_{t,j} + y^r_{t,j}, && \forall j \in \mathcal{J}, \\ & y^a_{t,j} \in [0, B^a_j],\ x^b_{t,j} \in [0, B^b_j], && \forall j \in \mathcal{J}, \\ & y^r_{t,j} \in [0, D_{t,j}],\ x^l_{t,j} \in [B^{l,-}_j, B^{l,+}_j], && \forall j \in \mathcal{J}. \end{aligned}$$

Multi-commodity Inventory Problem

Experiment Procedure

1. Draw $n_t \in \{5, 10, 20, 40\}$ data points from ν_t to get $\hat{\nu}_t$ for each $t \geq 2$;
2. Solve the nominal risk-neutral and risk-averse MSO using $\hat{\nu}_t$ (and MRO if applicable);
3. Construct Wasserstein ambiguity sets using $\hat{\nu}_t$ and solve the MDRO.
4. Evaluate the policies via $N = 100,000$ independent sample paths.

For inventory problems with uncertain prices, we set

$$D_{t,j} := D_0 \left[1 + \cos\left(\frac{2\pi(t+j)}{\tau} \right) \right] + \bar{D}, \quad C^b_{t,j}(\xi_t) := \xi_{t,j}, \quad C^a_{t,j}(\xi_t) := C_1 \cdot \xi_{t,j}, \quad j \in \mathcal{J},$$

and

$$\xi_t := \max \left\{ \text{Normal}(\mu_t, \bar{C} \cdot \Sigma_t), \underline{C} \right\}, \quad \mu_t := C_0 \left[1 + \sin\left(\frac{2\pi(t+j)}{\tau} \right) \right],$$

with Σ_t randomly generated for all $t \geq 2$. We use $|\mathcal{J}| = 5$, $T = 10$ and $\tau = 5$.

Multi-commodity Inventory Problem

Experiment Procedure

1. Draw $n_t \in \{5, 10, 20, 40\}$ data points from ν_t to get $\hat{\nu}_t$ for each $t \geq 2$;
2. Solve the nominal risk-neutral and risk-averse MSO using $\hat{\nu}_t$ (and MRO if applicable);
3. Construct Wasserstein ambiguity sets using $\hat{\nu}_t$ and solve the MDRO.
4. Evaluate the policies via $N = 100,000$ independent sample paths.

For inventory problems with uncertain prices, we set

$$D_{t,j} := D_0 \left[1 + \cos\left(\frac{2\pi(t+j)}{\tau} \right) \right] + \bar{D}, \quad C^b_{t,j}(\xi_t) := \xi_{t,j}, \quad C^a_{t,j}(\xi_t) := C_1 \cdot \xi_{t,j}, \quad j \in \mathcal{J},$$

and

$$\xi_t := \max \left\{ \mathrm{Normal}(\mu_t, \bar{C} \cdot \Sigma_t), \underline{C} \right\}, \quad \mu_t := C_0 \left[1 + \sin\left(\frac{2\pi(t+j)}{\tau} \right) \right],$$

with Σ_t randomly generated for all $t \geq 2$. We use $|\mathcal{J}| = 5$, $T = 10$ and $\tau = 5$.

Multi-commodity Inventory with Uncertain Prices

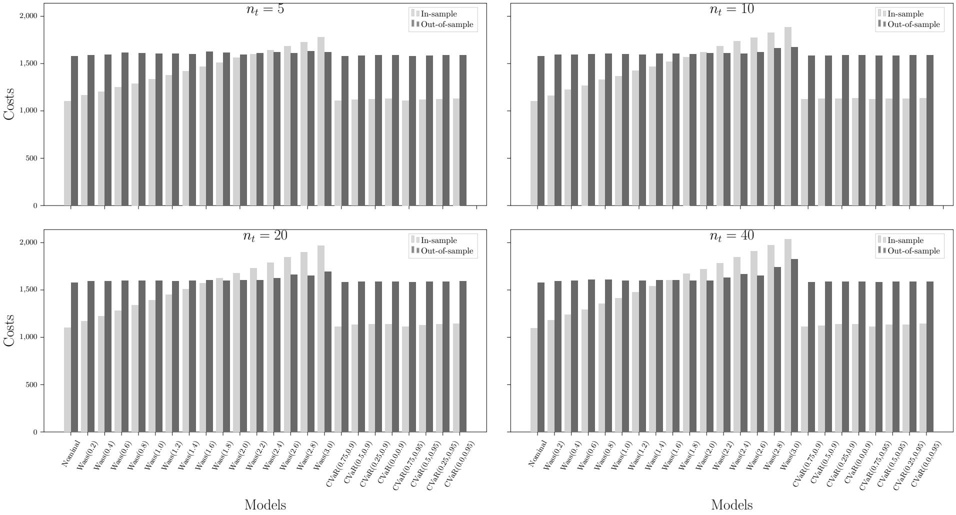

Figure: In-sample and Out-of-sample Costs Comparison for Inventory with Uncertain Prices

Multi-commodity Inventory with Uncertain Demands

Now we set the prices to be constants and consider uncertain demands modeled by

$$D_{t,j}(\xi_t) := D_0 \left[1 + \cos\left(\frac{2\pi(t+j)}{\tau} \right) \right] + \bar{D} \cdot \xi_{t,j}, \quad j \in \mathcal{J}.$$

The uncertain vector $\xi_t \in [0,1]^J$ has its components described as follows: $\xi_{t,1} \sim \mathrm{Uniform}(0,1)$, and for $j = 2, \ldots, J$, we have

$$\xi_{t,j} \mid \xi_{t,j-1} \sim \begin{cases} \mathrm{Uniform}(0, (1+\xi_{t,j-1})/2), & \text{if } \xi_{t,j-1} \leq \frac{1}{2}, \\ \mathrm{Uniform}(\xi_{t,j-1}/2, 1), & \text{otherwise.} \end{cases}$$

We use $J = 3$, $T = \tau = 5$ for the tests.

Multi-commodity Inventory with Uncertain Demands

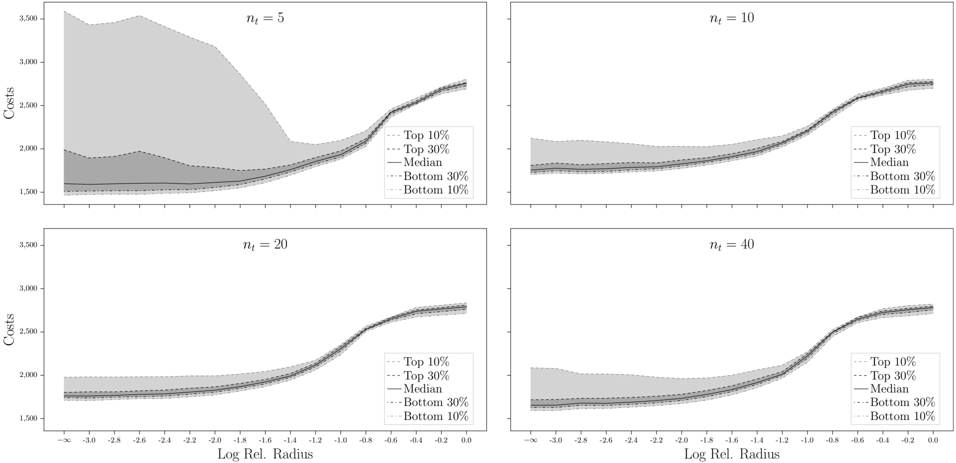

Figure: Out-of-sample Performance for Inventory with Uncertain Demands with Different Radii

Multi-commodity Inventory with Uncertain Demands

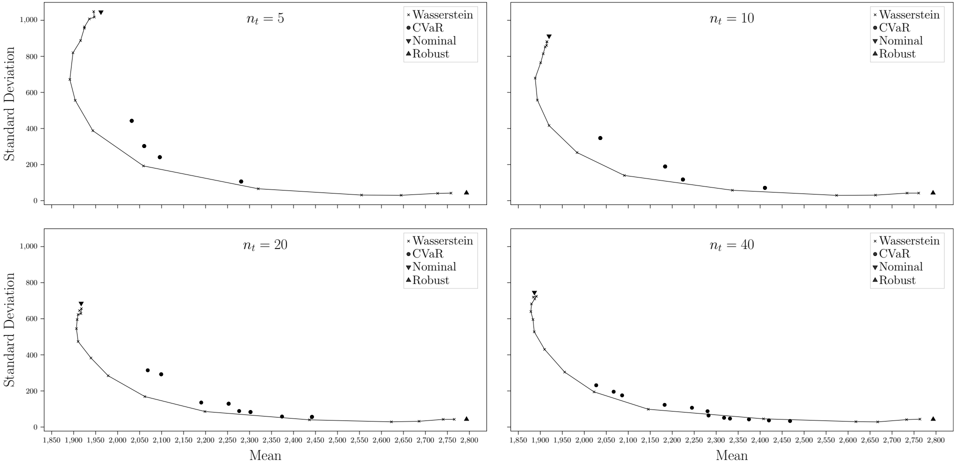

Figure: Comparison against Baseline Models for Inventory with Uncertain Demands

Hydro-thermal Power Planning Problem

We consider a distributionally robust hydro-thermal power planning problem defined recursively as

$$\mathcal{Q}_{t-1}(x_{t-1}) := \sup_{p_t \in \mathcal{P}_t} \int_{\Xi_t} \left[f_t(x_{t-1}, x_t; \xi_t) + \mathcal{Q}_t(x_t) \right] \mathrm{d}p_t,$$

where

$$\begin{aligned}
f_t(x_{t-1}, x_t; \xi_t) := \min_{y_t} \quad & \sum_{j \in \mathcal{J}} \left(C^s y^s_{t,j} + \sum_{l \in \mathcal{L}_j} C^g_l y^g_{t,l} + \sum_{j' \neq j} \left(C^e_{j,j'} y^e_{t,j,j'} + C^a_{j,j'} y^a_{t,j,j'} \right) \right) \\
\text{s.t.} \quad & x^l_{t,j} + y^h_{t,j} + y^s_{t,j} = x^l_{t-1,j} + \xi_{t,j}, && \forall\, j \in \mathcal{J}, \\
& y^h_{t,j} + \sum_{l \in \mathcal{L}_j} y^g_{t,l} + \sum_{j' \neq j} (y^a_{t,j,j'} - y^e_{t,j,j'} + y^e_{t,j',j}) = D_{t,j}, && \forall\, j \in \mathcal{J}, \\
& y^g_{t,l} \in [B^{g,-}_l, B^{g,+}_l], && \forall\, l \in \mathcal{L}, \\
& x^l_{t,j} \in [0, B^l_j], \;\; y^h_{t,j} \in [0, B^h_j], && \forall\, j \in \mathcal{J}, \\
& y^a_{t,j,j'} \in [0, B^a_{j,j'}], \;\; y^e_{t,j,j'} \in [0, B^e_{j,j'}], && \forall\, j, j' \in \mathcal{J}.
\end{aligned}$$

Hydro-thermal Power Planning Problem

The uncertain energy inflow ξ_t is assumed to follow

$$\ln \xi_t - \mu_t = \varphi_t(\ln \xi_{t-1} - \mu_{t-1}) + \varepsilon_t, \quad \varepsilon_t \sim \mathrm{Normal}(0, \Sigma_t),$$

but we may still formulate our MDRO using the marginal empirical measures $\hat{\nu}_t$'s.

Experiment Procedure

1. Draw $n_t \in \{5, 10, 20, 40\}$ data points from ν_t to get $\hat{\nu}_t$ for each $t \geq 2$;
2. Take additional samples from the fitted multivariate normal distributions to get $\tilde{\nu}_t$, and solve the approximate risk-neutral and risk-averse MSO using $\tilde{\nu}_t$ with at most 1000 iterations;
3. Construct Wasserstein ambiguity sets using $\hat{\nu}_t$ with the radius set to be relative to $W_t(\hat{\nu}_t, \tilde{\nu}_t)$, and solve the MDRO with at most 1000 iterations.
4. Evaluate the policies via $N = 100,000$ independent sample paths.

Hydro-thermal Power Planning Problem

The uncertain energy inflow ξ_t is assumed to follow

$$\ln \xi_t - \mu_t = \varphi_t(\ln \xi_{t-1} - \mu_{t-1}) + \varepsilon_t, \quad \varepsilon_t \sim \mathrm{Normal}(0, \Sigma_t),$$

but we may still formulate our MDRO using the marginal empirical measures $\hat{\nu}_t$'s.

Experiment Procedure

1. Draw $n_t \in \{5, 10, 20, 40\}$ data points from ν_t to get $\hat{\nu}_t$ for each $t \geq 2$;
2. Take additional samples from the fitted multivariate normal distributions to get $\tilde{\nu}_t$, and solve the approximate risk-neutral and risk-averse MSO using $\tilde{\nu}_t$ with at most 1000 iterations;
3. Construct Wasserstein ambiguity sets using $\hat{\nu}_t$ with the radius set to be relative to $W_t(\hat{\nu}_t, \tilde{\nu}_t)$, and solve the MDRO with at most 1000 iterations.
4. Evaluate the policies via $N = 100,000$ independent sample paths.

Hydro-thermal Power Planning Problem

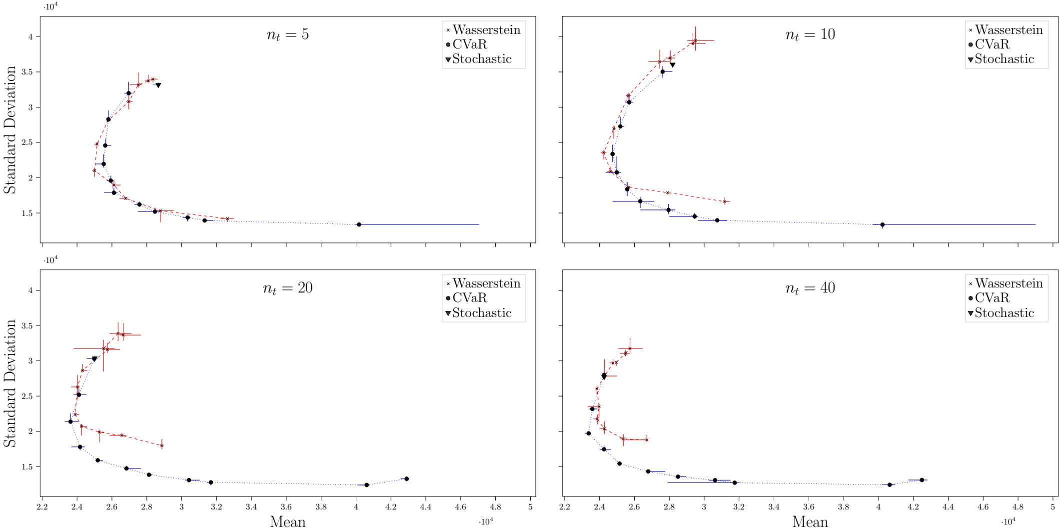

Figure: Comparison against Baseline Models on Hydro-thermal Power Planning for $T = 13$

Outline for Part 4

Introduction and Motivational Questions
- Multistage Distributionally Robust Optimization (MDRO)
- Dual Dynamic Programming (DDP) Algorithms

A New DDP Framework with Applications in Data-driven MDRO
- New DDP Framework with Complexity Analysis
- Data-driven DDP using Wasserstein Ambiguity

Numerical Experiments
- Multi-commodity Inventory Problem
- Hydro-thermal Power Planning Problem

Concluding Remarks

Concluding Remarks

- We study the oracle complexity of DDP algorithms for MDRO problems for the first time, with matching upper and lower bounds;
- We design exact SSSO implementations that allow us to tackle MDRO problems with data-driven Wasserstein ambiguity sets;
- We compare the out-of-sample performance of the MDRO model against risk-neutral and risk-averse MSO models and the MRO model using multi-commodity inventory problems and hydro-thermal power planning problem.

Our related works:

Shixuan Zhang and Xu Andy Sun. On distributionally robust multistage convex optimization: New algorithms and complexity analysis. *arXiv preprint arXiv:2010.06759*, 2020.

Shixuan Zhang and Xu Andy Sun. On distributionally robust multistage convex optimization: Data-driven models and performance. *arXiv preprint arXiv:2210.08433*, 2022a.

Shixuan Zhang and Xu Andy Sun. Stochastic dual dynamic programming for multistage stochastic mixed-integer nonlinear optimization. *Mathematical Programming*, pages 1–51, 2022b.

References I

Shabbir Ahmed, Filipe Goulart Cabral, and Bernardo Freitas Paulo da Costa. Stochastic lipschitz dynamic programming. *Mathematical Programming*, pages 1–39, 2020.

R Baucke, A Downward, and G Zakeri. A deterministic algorithm for solving multistage stochastic programming problems. *Optimization Online*, 2017.

Daniel Duque and David P Morton. Distributionally robust stochastic dual dynamic programming. *SIAM Journal on Optimization*, 30(4):2841–2865, 2020.

A Georghiou, A Tsoukalas, and W Wiesemann. Robust dual dynamic programming. *Operations Research*, 67(3):813–830, 2019.

P. Girardeau, V. Leclere, and A. B. Philpott. On the Convergence of Decomposition Methods for Multistage Stochastic Convex Programs. *Mathematics of Operations Research*, 40(1):130–145, February 2015. ISSN 0364-765X, 1526-5471. doi: 10.1287/moor.2014.0664.

Guanghui Lan. Complexity of stochastic dual dynamic programming. *Mathematical Programming*, pages 1–38, 2020.

M. V. F. Pereira and L. M. V. G. Pinto. Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming*, 52(1-3):359–375, May 1991. ISSN 0025-5610, 1436-4646. doi: 10.1007/BF01582895.

References II

A. B. Philpott, V. L. de Matos, and L. Kapelevich. Distributionally robust SDDP. *Computational Management Science*, 15(3-4):431–454, October 2018. ISSN 1619-697X, 1619-6988. doi: 10.1007/s10287-018-0314-0.

Andy Philpott, Vitor de Matos, and Erlon Finardi. On Solving Multistage Stochastic Programs with Coherent Risk Measures. *Operations Research*, 61(4):957–970, August 2013. ISSN 0030-364X, 1526-5463. doi: 10.1287/opre.2013.1175.

A Shapiro. Analysis of stochastic dual dynamic programming method. *European Journal of Operational Research*, 209(1):63–72, February 2011. ISSN 03772217. doi: 10.1016/j.ejor.2010.08.007.

A Shapiro, W Tekaya, J P da Costa, and M P Soares. Risk neutral and risk averse Stochastic Dual Dynamic Programming method. *European Journal of Operational Research*, 224(2):375–391, January 2013. ISSN 03772217. doi: 10.1016/j.ejor.2012.08.022.

J Zou, S Ahmed, and X A Sun. Stochastic dual dynamic integer programming. *Mathematical Programming*, 175(1-2):461–502, May 2019. ISSN 0025-5610, 1436-4646. doi: 10.1007/s10107-018-1249-5.

Illustration of DDP Complexity Proofs I

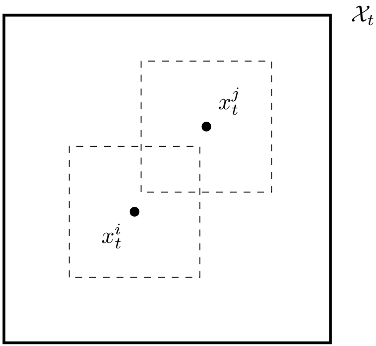

(a) Norm-ball covering

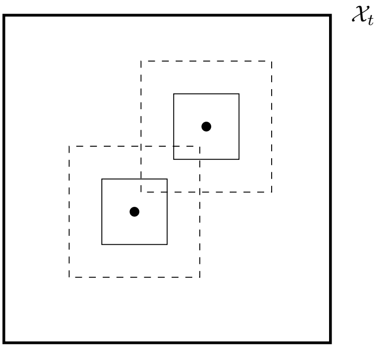

(b) Norm-ball non-overlapping

Illustration of DDP Complexity Proofs II

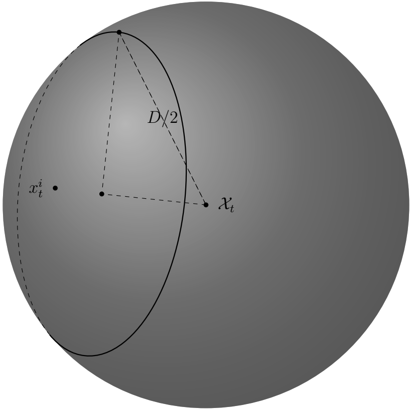

Figure: Norm-ball on a sphere