Distributionally Robust Linear and Discrete Optimization with Marginals

Louis Chen[1] Will Ma[1] Karthik Natarajan[3] James Orlin[1]
David Simchi-Levi[1,2] Zhenzhen Yan[4]

[1]Operations Research Center
Massachusetts Institute of Technology

[2]Institute for Data, Systems, and Society
Massachusetts Institute of Technology

[3]Singapore University of Technology and Design

[4]Nanyang Technological University

October 2018

Outline of Talk

1. Motivation: Distributionally Robust Max Flow as an Example

2. Main Results and Connection to Other Results

Optimization under Uncertainty

- Stochastic Optimization:

$$\inf_{r \in C} E_{\tilde{u} \sim \theta} \left[Z(r, \tilde{u}) \right]$$

- Robust Optimization:

$$\inf_{r \in C} \sup_{u \in \Omega} Z(r, u)$$

- Distributionally Robust Optimization:

$$\inf_{r \in C} \sup_{\theta \in \mathcal{P}} E_{\tilde{u} \sim \theta} \left[Z(r, \tilde{u}) \right]$$

- Specification of $\mathcal{P}$: Moments, Marginals, Distributions within some distance around a "reference" distribution - often guided by expressive power, computational and calibration issues

Problem of Interest

- $\mathcal{P} := \Gamma(\mu_1, \ldots, \mu_n)$ is the set of joint distributions consistent with given marginals $\mu_1, \ldots, \mu_n$.
- Suppose $\mathcal{X}$ is some finite set (potentially very large and might not be explicitly specified) where:

$$Z(\tilde{u}) = \max_{\chi \in \mathcal{X}} \tilde{u}^{\mathsf{T}} \chi$$

We are interested in:

$$\sup_{\theta \in \mathcal{P}} E_{\tilde{u} \sim \theta} \left[Z(\tilde{u}) \right]$$

- Consider a distributionally robust max flow problem where r is a capacity decision vector which needs to be made on arcs in addition to a random capacity vector $\tilde{u}$ and $Z(\tilde{u} + r)$ is the random maximum flow on the network:

$$\max_{r \in C} \inf_{\theta \in \mathcal{P}} E_{\tilde{u} \sim \theta} \left[Z(\tilde{u} + r) \right]$$

Distributionally Robust Max Flow

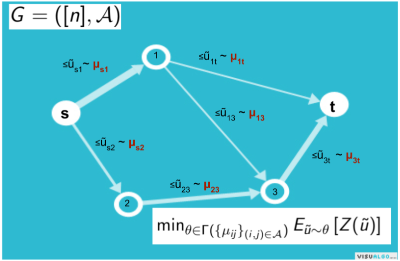

$$
\begin{aligned}
\text{(Max-flow)}\ Z(\tilde{u}) = \max \quad & v \\
\text{s.t.} \quad & \sum_{j:(i,j)\in\mathcal{A}} x_{ij} - \sum_{j:(j,i)\in\mathcal{A}} x_{ji} = \begin{cases} v, & i = s \\ 0, & i \neq s,t \\ -v, & i = t \end{cases} \\
& 0 \leq x_{ij} \leq \tilde{u}_{ij},\ (i,j) \in \mathcal{A} \\
\text{(Min-cut)}\ Z(\tilde{u}) = \min \quad & \tilde{u}^{\mathsf{T}}\chi \\
\text{s.t.} \quad & \chi \in \mathcal{X}_{cut}.
\end{aligned}
$$

Distributionally Robust Max Flow with Marginals

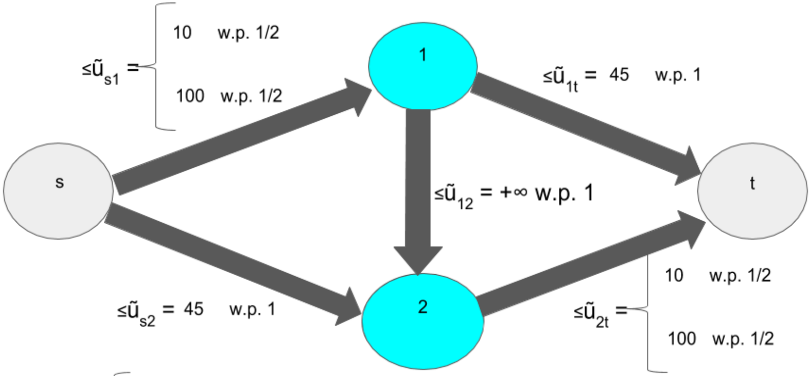

$(\tilde{u}_{s1}, \tilde{u}_{2t}) = \begin{cases} (10,10) & \text{w.p. } 1/2 \\ (100,100) & \text{w.p. } 1/2 \end{cases}$ ⇒ $EZ(\tilde{u}) = ½(20) + ½\,(145) = 82.5$

Naive Approach: assign as much probability mass to joint realization of lowest elements in the supports?

Distributionally Robust Max Flow with Marginals

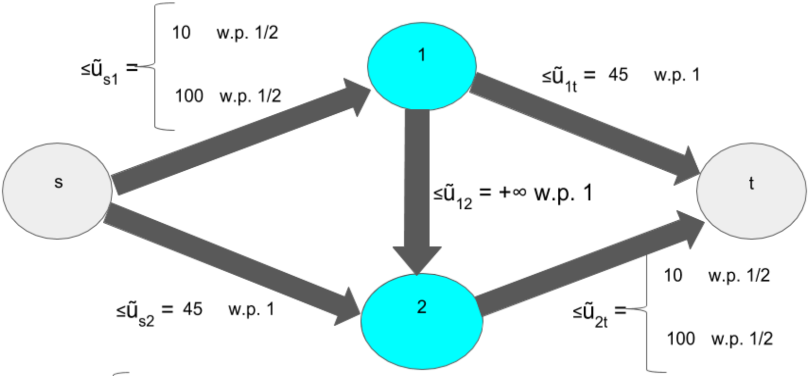

$$(\tilde{u}_{s1}, \tilde{u}_{2t}) = \begin{cases} (10,100) & \text{w.p. } 1/2 \\ (100,10) & \text{w.p. } 1/2 \end{cases}$$

$\Rightarrow$ $EZ(\tilde{u}) = \frac{1}{2}(55) + \frac{1}{2}(55) = 55$

Better than Naive!

Note that independence would give an expected value of 68.75

Dual Form: A Lower Bound

- Let w be a vector of arc capacities, θ be a probability measure in the set Γ.
- We obtain a lower bound on the expected maximum flow:

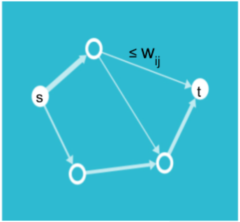

VISUALGO

$Z(w)$

$-$

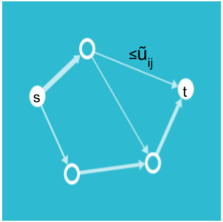

$Z(\tilde{u})$

$$\leq \sum_{(i,j)\in\mathcal{A}} max(w_{ij} - \tilde{u}_{ij}, 0)$$

$$E_{\tilde{u}\sim\theta}\left[Z(w) - Z(\tilde{u})\right] \leq \sum_{(i,j)\in\mathcal{A}} \int max(w_{ij} - \tilde{u}_{ij}, 0) d\mu_{ij}$$

$$\implies E_{\tilde{u}\sim\theta}\left[Z(\tilde{u})\right] \geq \max_{w} \left\{ Z(w) - \sum_{(i,j)\in\mathcal{A}} \int max(w_{ij} - \tilde{u}_{ij}, 0) d\mu_{ij} \right\}$$

Primal Form: An Upper Bound

- Let $(\tilde{u}, \tilde{\chi})$ be a random vector, where $\tilde{u}$ is consistent with marginals and $\tilde{\chi}$ is a random cut-set incidence vector.

$$\inf_{\theta \in \Gamma} E_{\tilde{u} \sim \theta} [Z(\tilde{u})] = \inf_{\theta \in \Gamma} E_{\tilde{u} \sim \theta} \left[\min_{\chi \in \mathcal{X}_{cut}} \tilde{u}^\intercal \chi \right] \\ \leq E_{(\tilde{u}, \tilde{\chi})} [\tilde{u}^\intercal \tilde{\chi}]$$

- Search over $(\tilde{u}, \tilde{\chi})$ for the tightest bound:
  1. Find a distribution $\nu \in \mathcal{P}(\mathcal{X}_{cut})$ for $\tilde{\chi}$. $(conv(\mathcal{X}_{cut}) \equiv \mathcal{P}(\mathcal{X}_{cut}))$
  2. Construct $\tilde{u}||_{\tilde{\chi}}$(in minimal fashion) s.t. $\tilde{u}$ is a coupling consistent with the marginals
- How does the construction work in step 2?
  Use $\tilde{u}||_{\tilde{\chi}} = \otimes\, \tilde{u}_{ij}||_{\tilde{\chi}_{ij}}$ (conditionally independent) where $\tilde{\chi}_{ij} = 1$ w.p. $\Pi_{ij}\nu(1)$ and $\tilde{\chi}_{ij} = 0$ w.p. $\Pi_{ij}\nu(0)$.
- How do we define $\tilde{u}_{ij}||_{\tilde{\chi}_{ij}}$ in a minimizing fashion?

Primal Form: An Upper Bound

Such feasible couplings yields an upper bound of the form:

$$\begin{aligned}
\inf_{\theta\in\Gamma} E_{\tilde{u}\sim\theta}\left[Z(\tilde{u})\right] &\leq E_{(\tilde{u},\tilde{\chi})}\left[\tilde{u}^\intercal\tilde{\chi}\right] \\
&= E_{\tilde{\chi}}\left[\sum_{(i,j)\in\mathcal{A}} E\left[\tilde{u}_{ij}\tilde{\chi}_{ij}||\tilde{\chi}\right]\right] \\
&= E_{\tilde{\chi}}\left[\sum_{(i,j)\in\mathcal{A}} E\left[\tilde{u}_{ij}\tilde{\chi}_{ij}||\tilde{\chi}_{ij}\right]\right] \\
&= \sum_{(i,j)\in\mathcal{A}} \Pi_{ij}\nu(1)E\left[\tilde{u}_{ij}||\tilde{\chi}_{ij}=1\right]
\end{aligned}$$

Primal Form: An Upper Bound

$\tilde{u}_{ij}||_{\tilde{\chi}_{ij}=0}$

$\tilde{u}_{ij}||_{\tilde{\chi}_{ij}=1}$

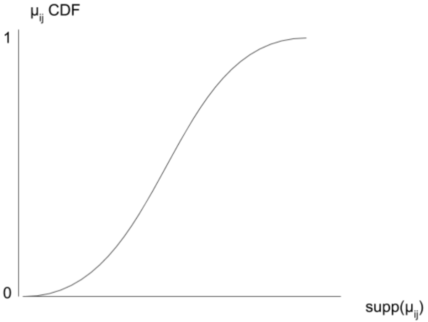

$$\inf_{\theta\in\Gamma} E_{\tilde{u}\sim\theta}\left[Z(\tilde{u})\right] \leq \min_{\nu\in\mathcal{P}(\mathcal{X}_{cut})} \sum_{(i,j)\in\mathcal{A}} \int_0^{\Pi_{ij}\nu(1)} F_{\mu_{ij}}^{-1}(q)\ dq.$$

Primal Form: An Upper Bound

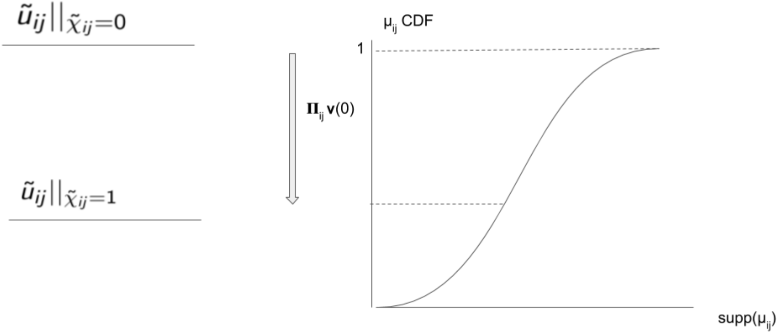

$$\inf_{\theta\in\Gamma} E_{\tilde{u}\sim\theta}\left[Z(\tilde{u})\right] \le \min_{\nu\in\mathcal{P}(\mathcal{X}_{cut})} \sum_{(i,j)\in\mathcal{A}} \int_0^{\Pi_{ij}\nu(1)} F_{\mu_{ij}}^{-1}(q) \; dq.$$

Primal Form: An Upper Bound

$$\tilde{u}_{ij}||_{\tilde{\chi}_{ij}=0}$$

$$\frac{\mu_{ij} \cdot \mathbb{1}_{Upper_{ij}}}{\Pi_{ij}\nu(0)}$$

$$\tilde{u}_{ij}||_{\tilde{\chi}_{ij}=1}$$

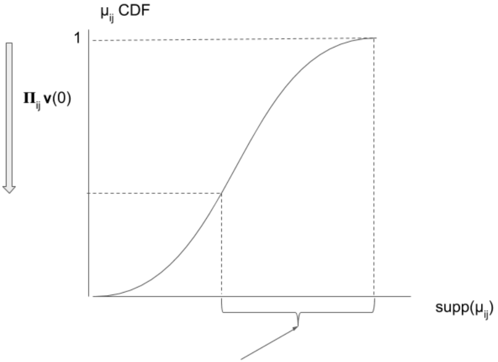

$$Upper_{ij} := \{\tilde{u}_{ij} : F_{\mu_{ij}}^{-1}(1 - \Pi_{ij}\nu(0)) \leq \tilde{u}_{ij} \leq F_{\mu_{ij}}^{-1}(1)\}$$

$$\inf_{\theta \in \Gamma} E_{\tilde{u} \sim \theta}\left[Z(\tilde{u})\right] \leq \min_{\nu \in \mathcal{P}(\mathcal{X}_{cut})} \sum_{(i,j) \in \mathcal{A}} \int_0^{\Pi_{ij}\nu(1)} F_{\mu_{ij}}^{-1}(q) \;\; dq.$$

Primal Form: An Upper Bound

$\tilde{u}_{ij}||_{\tilde{\chi}_{ij}=0}$

$$\frac{\mu_{ij}\cdot \mathbb{1}_{Upper_{ij}}}{\Pi_{ij}\nu(0)}$$

$\tilde{u}_{ij}||_{\tilde{\chi}_{ij}=1}$

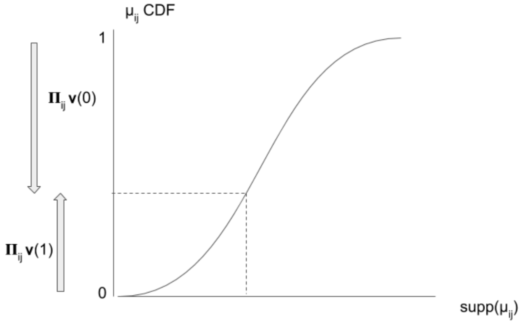

$$\inf_{\theta\in\Gamma} E_{\tilde{u}\sim\theta}\left[Z(\tilde{u})\right] \leq \min_{\nu\in\mathcal{P}(\mathcal{X}_{cut})} \sum_{(i,j)\in\mathcal{A}} \int_0^{\Pi_{ij}\nu(1)} F_{\mu_{ij}}^{-1}(q)\ dq.$$

Primal Form: An Upper Bound

$$\frac{\tilde{u}_{ij} || _{\tilde{\chi}_{ij}=0}}{\frac{\mu_{ij} \cdot \mathbb{1}_{Upper_{ij}}}{\Pi_{ij}\nu(0)}}$$

$$\frac{\tilde{u}_{ij} || _{\tilde{\chi}_{ij}=1}}{\frac{\mu_{ij} \cdot \mathbb{1}_{Lower_{ij}}}{\Pi_{ij}\nu(1)}}$$

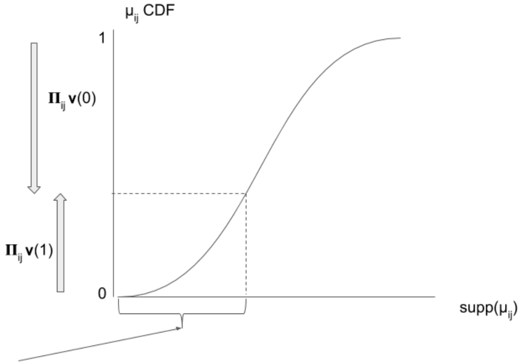

$Lower_{ij} := \{\tilde{u}_{ij} : 0 \leq \tilde{u}_{ij} \leq F^{-1}_{\mu_{ij}}(\Pi_{ij}\nu(1))\}$

$$\inf_{\theta \in \Gamma} E_{\tilde{u} \sim \theta}\left[Z(\tilde{u})\right] \leq \min_{\nu \in \mathcal{P}(\mathcal{X}_{cut})} \sum_{(i,j) \in \mathcal{A}} \int_0^{\Pi_{ij}\nu(1)} F^{-1}_{\mu_{ij}}(q) \; dq.$$

Primal-Dual Formulations

- Dual: $\max_w \left\{ Z(w) - \sum_{(i,j)\in\mathcal{A}} \int max(w_{ij} - \tilde{u}_{ij}, 0) d\mu_{ij} \right\}$
  - Concave maximization with a max-flow problem with deterministic capacities that needs to be optimized
  - Univariate expectations
- Primal: $\min_{\nu\in\mathcal{P}(\mathcal{X}_{cut})} \left\{ \sum_{(i,j)\in\mathcal{A}} \int_0^{\Pi_{ij}\nu(1)} F_{\mu_{ij}}^{-1}(q) \ \ dq \right\}$
  - Convex minimization over the s-t cut polytope
  - Suffices to find a distribution ν^* over the family of s-t cuts
  - $\nu^* \mapsto \theta^*$
  - How to find ν^*?

Finding ν^* Under Finite-Supported Marginals

Primal

$$\max_w \left\{ Z(w) - \sum_{(i,j)\in\mathcal{A}} \int \max(w_{ij} - \tilde{u}_{ij}, 0) d\mu_{ij} \right\}$$

=

Dual

$$\min_{\nu\in\mathcal{P}(\mathcal{X}_{cut})} \sum_{(i,j)\in\mathcal{A}} \int_0^{\Pi_{ij}\nu(1)} F_{\mu_{ij}}^{-1}(q) \; dq$$

Network Transformations/Red.
+
LP Duality

$$\min_{\pi,\lambda} \quad \sum_{(i,j)\in A} \sum_{k=1}^{m_{ij}} (u_{ij}^k - u_{ij}^{k-1}) \cdot \lambda_{ij}^k$$

subject to $\pi_t = 0$

$\pi_s = 1$

$0 \le \pi_i \le 1; \quad \forall i \in N \setminus \{s,t\}$

$-\sum_{\tau=1}^{k-1} p_{ij}^{\tau} + \pi_i - \pi_j \le \lambda_{ij}^k \qquad (i,j) \in A, k = 1, \ldots, m_{ij}$

$\lambda_{ij}^k \ge 0; \quad (i,j) \in A, k = 1, \ldots, m_{ij}$

Optimal π^* encodes the optimal ν^*

Decoding ν^* from π^*

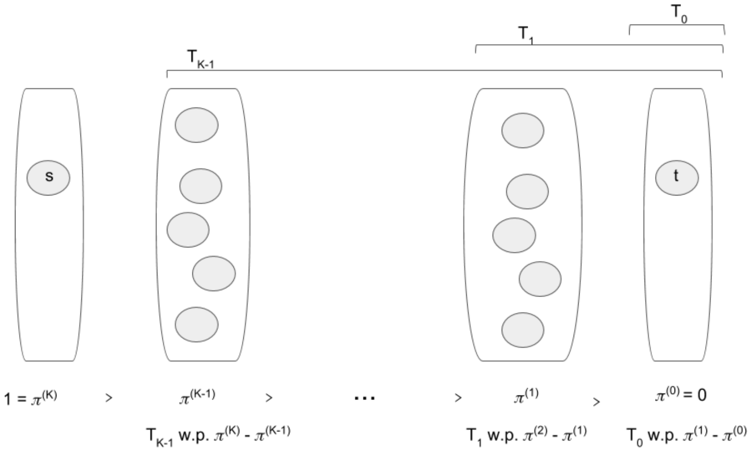

We can restrict to probability measures supported on "nested" collections of (s-t)-cuts (submodularity of cut-capacity function).

Main Results

Theorem (Maximization Form)

Let $Z^* := \max_{\theta \in \Gamma(\mu_1,\dots,\mu_n)} E_{\tilde{c}\sim\theta}\left[\max_{x\in\mathcal{X}} \tilde{c}^\intercal x\right]$ *where* $E_{\mu_i}|\tilde{c}_i| < \infty$ *for all* i*. Then,*

$$Z^* = \max_{\nu \in \mathcal{P}(\mathcal{X})} \sum_{i=1}^{n} \sum_{x_i \in \mathcal{X}_i} x_i \int_{\Pi_i \nu((-\infty, x_i - e_i])}^{\Pi_i \nu((-\infty, x_i])} F_{\mu_i}^{-1}(t)\,dt,$$

$$= \min_{\{\psi_i : \mathcal{X}_i \to \mathbb{R}\}_{i=1}^n} \left(\max_{\nu \in \mathcal{P}(\mathcal{X})} \sum_{x \in \mathcal{X}} \left[\sum_{i=1}^{n} \psi_i(x_i)\right] \nu(x) + \sum_{i=1}^{n} \int \psi_i^*(\tilde{c}_i)\,d\mu_i \right),$$

where $\psi_i^*(c_i) = \max_{x_i}(c_i x_i - \psi_i(x_i))$*. Let* ν_{rel} *denote the optimal solution. If* $\mu_1, \dots, \mu_n$ *are absolutely continuous w.r.t. the Lebesgue measure, there exists some suitably defined measurable selection* $x^* : \mathbb{R}^n \to \mathcal{X}$ *of* x^{OPT} *s.t. the "persistence values"* $P_{\tilde{c}\sim\theta^*}(x_i^*(\tilde{c}) = x_i)$ *are given by:*

$$P_{\tilde{c}\sim\theta^*}(x_i^*(\tilde{c}_i) = x_i) = \Pi_i \nu_{rel}(x_i), \ \ \forall x_i \in \mathcal{X}_i, \ \ \forall i \in [n].$$

Connection to Previous Results

- Natarajan, Song and Teo (2009) provided the primal formulation for absolutely continuous random variables - the result can be generalized to arbitrary marginals using techniques from optimal transport theory.
- Under the assumption of absolutely continuous marginals, one can show that if ν and τ both solve the primal formulation, then $\Pi_i \nu = \Pi_i \tau$ for all i.
- This provides justification to use it in choice modeling - for example the model can recreate the multinomial logit choice probabilities for an appropriate choice of Γ (see Mishra, Natarajan, Padmanabhan, Teo and Li (2014)).
- Natarajan, Song and Teo (2009) study Z^* where $\mathcal{X}$ is the feasible region to a bounded integer program (such as integer knapsack). Using a binary reformulation, they obtain tractable upper bounds, but the complexity of finding Z^* is not discussed.

Main Results: Hardness

Proposition (Related to Mangasarian and Shiau (1986))

Computing Z^* for the class of linear optimization problems given discrete marginal distributions and a H-polytope is NP-hard.

Related results:

- Computing the worst-case expected value of a function of binary random variables with fixed marginal probabilities is NP-hard, even when the function is submodular - MAX-CUT problem (see Agrawal, Ding, Saberi and Ye (2012)).
- Computing the worst-case expected value in distributionally robust linear optimization problems with a given mean and covariance matrix is NP-hard - 2-norm maximization over a polytope (see Bertsimas, Doan, Natarajan and Teo (2010)).

Main Results: Tractable Instances

Assume that the expected value of the univariate convex functions $\max_{x_i}(c_i x_i - \psi_i(x_i))$ and the subgradients are efficiently computable.

- $\mathcal{X}$ describes as V-polytope:

$$\max_{\nu \in \mathcal{P}(\mathcal{X})} \sum_{x \in \mathcal{X}} \left[\sum_{i=1}^{n} \psi_i(x_i) \right] \nu(x) \quad = \quad \max_{x \in \mathcal{X}} \sum_{i=1}^{n} \psi_i(x_i).$$

- $\mathcal{X}$ describes the extreme points to a 0-1 H-polytope ($P = conv(\mathcal{X})$):

$$\max_{\nu \in \mathcal{P}(\mathcal{X})} \sum_{x \in \mathcal{X}} \left[\sum_{i=1}^{n} \psi_i(x_i) \right] \nu(x) = \max_{x \in P} \sum_{i=1}^{n} \psi_i(1) x_i.$$

  Example: Max-flow (min-cut)

- Prior research: Meilijson and Nadas (1979) - longest path on a directed acylic graph (PERT), Birge and Maddox (1995) - PERT with marginal moments, Bertsimas, Natarajan and Teo (2004, 2006) - 0/1 optimization problem with marginal moments

Main Results: Tractable Instances

Theorem

Suppose there exists a compact extended formulation of conv$(\mathcal{X})$ as:

$$\Pi_x\left(\left\{(x,y) : y \in P,\;\; x_i = \sum_{\bar{x}_i \in \mathcal{X}_i} \bar{x}_i \sum_{j=1}^{n_{\bar{x}_i}} y_{F_j^{\bar{x}_i}} \text{ for } \{F_j^{\bar{x}_i}\} \in B,\;\; \forall j,\;\; \forall i\right\}\right),$$

where $n_{\bar{x}_i}$ is a finite integer for each i, $\bar{x}_i \in \mathcal{X}_i$ and P is a 0-1 polytope of the form:

$$P \subseteq \left\{y \in [0,1]^B : \sum_{\bar{x}_i \in \mathcal{X}_i} \sum_{j=1}^{n_{\bar{x}_i}} y_{F_j^{\bar{x}_i}} = 1,\;\; \forall i \in [n]\right\},$$

then Z^* is efficiently computable.

Main Results: Appointment Scheduling

- n patients, service time of patient i is random with $\tilde{c}_i \sim \mu_i$, s_i is the service time scheduled for patient i:

$$\min_{s \in \mathcal{S}} \sup_{\theta \in \Gamma(\mu_1, \ldots, \mu_n)} E_{\tilde{c} \sim \theta} \left[Z(s, \tilde{c}) \right],$$

where:

$$\mathcal{S} = \left\{ s \in \mathbb{R}^n : \sum_{i=1}^{n} s_i \leq T, \quad s_i \geq 0 \;\; \forall i \in [n] \right\}$$

$$\begin{aligned} Z(s, \tilde{c}) \quad = \max \quad & \sum_{i=1}^{n} (\tilde{c}_i - s_i) x_i \\ \text{s.t.} \quad & x_i - x_{i-1} \geq -1, \quad \forall i = 2, \ldots, n, \\ & x_n \leq 1, \\ & x_i \geq 0, \qquad\qquad \forall i = 1 \ldots, n \end{aligned}$$

- In this case, you can find such a compact extended formulation (this was first identified by Mak, Rong and Zhang (2016) with mean and variance information).

Main Results: Scheduling with Ranking

- n jobs, single machine, duration of job i is random with $\tilde{c}_i \sim \mu_i$, t_i is the amount by which job duration is reduced. This is decided before knowing the true realization or arrival sequence and the objective is to minimize sum of completion times:

$$\min_{t \in \mathcal{T}} \sup_{\theta \in \Gamma(\mu_1,\ldots,\mu_n)} E_{\tilde{c} \sim \theta}\left[Z(t,\tilde{c})\right],$$

where:

$$\begin{aligned} Z(t,\tilde{c}) \quad &:= \max \quad \sum_{i=1}^{n} (\tilde{c}_i - t_i) x_i \\ &\quad \text{s.t.} \quad x \in \mathcal{X}_{perm}. \end{aligned}$$

- In this case, you can find such a compact extended formulation using the Bikrhoff polytope for permutations.

Main Results: Scheduling with Random Irregular Starting Time Costs

- n jobs that need to be scheduled within a fixed time horizon $\{0, 1, \ldots, T\}$ where job $j \in N$ incurs a random cost $\tilde{c}_j(S_j) = c_j^0(S_j)\tilde{\epsilon}_j$ if it is started at time S_j.
- Precedence constraints among two jobs i and j: $S_j \geq S_i + d_{ij}$ where d_{ij} is an integer number imposing a time lag between the jobs. Precedence among jobs results in directed graph with no cycles of positive length.
- A lower bound on the total cost:

$$\inf_{\theta \in \Gamma(\mu_1,\ldots,\mu_n)} E\left[\min_{S \in \mathcal{S}} \sum_{j=1}^{n} \tilde{c}_j(S_j)\right].$$

- In this case, you can find such a compact extended formulation using the time-indexed polytope.

Main Results: (Near) Sufficiency of Tractability Conditions

- Given $m \leq n$ linearly independent vectors x^i, define:

$$\text{(Parallelotope) } Q = \sum_{i=1}^{m} [-x^i, x^i].$$

- Suppose $x^i \in \{-1, 0, 1\}^n$. Extreme point entries in $\{-m, \ldots, m\}$:

$$\text{Extr}(Q) = \sum_{i=1}^{m} \epsilon_i x^i \text{ with } |\epsilon_i| = 1 \text{ for all } i.$$

$$\Pi_x \left\{ (x, \epsilon) : x = \sum_{i=1}^{m} \epsilon_i x^i, -1 \leq \epsilon_i \leq 1, \forall i \right\}.$$

Corollary (Related to Bodlaender, Grizmann, Klee and Leeuwen (1990))

Computing Z^* for linear optimization problems over a parallelotope of the form $Q = \sum_{i=1}^{m} [-x^i, x^i]$, in which all $x^i \in \{-1, 0, 1\}^n$ is NP-hard.

Connection to Known Results: Dependence

- Price of correlations where $Z(\cdot,\cdot) \geq 0$ (see Agrawal, Ding, Saberi and Ye (2012)):

$$\frac{\sup_{\theta \in \Gamma(\mu_1,\ldots,\mu_n)} E_\theta\left[Z(\hat{r},\tilde{u})\right]}{\min_{r \in C} \sup_{\theta \in \Gamma(\mu_1,\ldots,\mu_n)} E_\theta\left[Z(r,\tilde{u})\right]} = \frac{\text{“From Independent Coupling”}}{\text{“From Worst-case Coupling”}}$$

  where

$$\hat{r} \in \arg\min_{r \in C} E_{\tilde{u} \sim \mu_1 \otimes \ldots \otimes \mu_n}\left[Z(r,\tilde{u})\right]$$

- $Z(\cdot, u)$ - monotone and submodular in u, POC $\leq e/(e-1)$ (other generalizations discussed in their paper).
- $Z(\cdot, u)$ - monotone and supermodular in u, POC can be very large.

Connection to Known Results: Independence

- $E[Z(\tilde{u})]$ might however be #P-hard to compute for discrete independent distributions for many types of functions:
  - Submodular functions ($Z(\tilde{u}) = \min\{\tilde{u}^\intercal \chi : \chi \in \mathcal{X}_{cut}\}$)
  - Supermodular functions ($Z(\tilde{u}) = (\sum_i u_i - K)^+$).
- Given a network where the edges are subject to random failure, independently and each with equal probability p, the probability that the failed edges does not contain a s-t cut is #P-hard to compute (see Provan and Ball (1983)).
- #P - Set of the counting problems associated with the decision problems in the set NP. This class was introduced by Valiant in 1979.
- Counting version of NP-hard problems are #P-hard.
- However, there are easy decision problems, for which the counting versions can be hard (number of perfect matchings in a bipartite graph).

Known Results: Independence

- A project is specified by precedence relations among tasks. Task durations are <u>independent</u> random variables with discrete, finite ranges. Then,
  - Computing a value of the cumulative distribution function of project duration is #P-hard.
  - Computing the mean of the distribution is at least as hard.
  - Neither of the problems can be computed in time polynomial in the number of points in the range of the project duration unless P = NP (see Hagstrom 1988).
- Only in special cases such as series-parallel graphs, with restricted assumptions on the randomness (such as binary random variables), these problems can be solved in polynomial time (see Ball, Colbourn and Provan (1995), Möhring (2001)).

Concluding Remarks

- Bounds under this set of distributions have been extensively studied in risk and insurance - $P(\sum_i \tilde{c}_i \geq T)$, $E(\sum \tilde{c}_i - T)^+$, $P(\max_i \tilde{c}_i \geq T)$, $E(\max \tilde{c}_i - T)^+$.
- However, in operations research and operations management, our interest is often in more complicated decision-making problems with constraints - interplay of optimization and probability is often a challenge and needs to be carefully analyzed.
- Some applications where this model has been studied includes facility location design (Lu, Ran and Shen (2014)), appointment scheduling (Mak, Rong and Zhang (2016)), traffic equilibrium (Arikan, Ahipasaolgu and Natarajan (2018) and multi-product pricing (Yan, Cheng, Natarajan and Teo (2018)).