Multi-Stage Stochastic Programming with CVaR: Modeling, Algorithms and Robustness

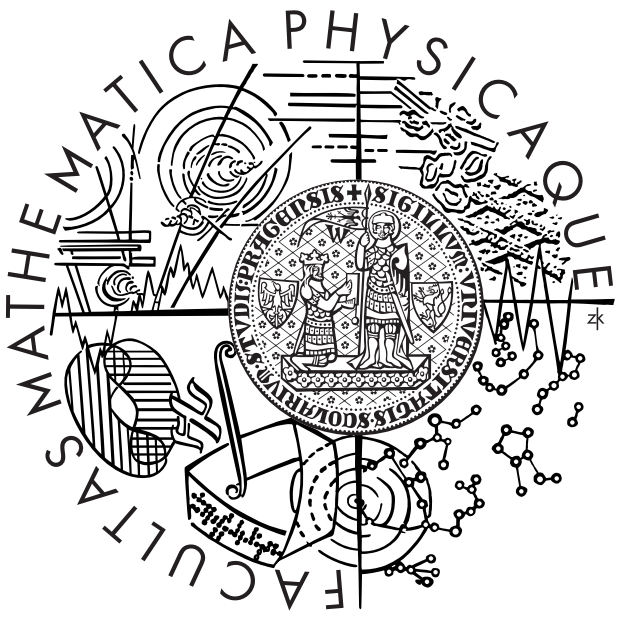

Václav Kozmík

Faculty of Mathematics and Physics
Charles University in Prague

December 11, 2014

Outline

- Introduction
  - Outline
  - Original results
- Structure of risk-averse multistage stochastic models
  - Studied less than the risk-neutral case
  - Risk measures bring multiple modeling possibilities
    - Nested models
    - Multiperiod risk measures
    - Sum of stage risks
- Stochastic Dual Dynamic Programming Algorithm
  - Originated in 1991 by Pereira & Pinto
  - Well-studied for risk-neutral case
  - Cut generation in risk-averse case is straightforward
  - Upper bound estimation is challenging in risk-averse case

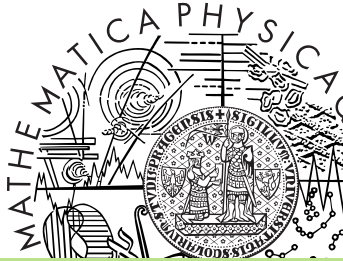

Outline

- Variance reductions schemes
  - Monte Carlo methods are widely applied, but they usually bring high variance
  - Many variance reduction schemes have been proposed: importance sampling, stratified sampling, Quasi Monte Carlo, etc.
  - Application of such schemes could be hard in practice, the proposals usually focus on the method and toy examples, not on large-scale real world applications
- Contamination technique for multistage problems
  - Captures various changes in the probability distribution
  - First results applied to two-stage setups
  - In multistage case, contamination was studied and applied for smaller problems
  - Extension to large-scale setups requires advanced algorithms and techniques

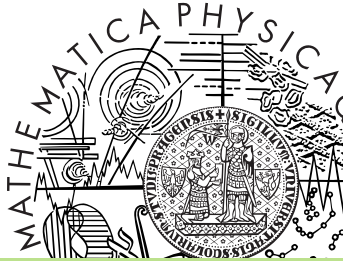

Original results

- New estimator for the policy value under the nested model setup with CVaR
  - Provides better results than the state-of-the-art estimators
  - Can be used to build valid stopping rules for the SDDP algorithm
  - General importance sampling scheme for mean-CVaR objectives
    - Closed-form solution provided for normal distribution
    - Sampling algorithm presented for any general distribution to get the suitable parameter of the variance reduction scheme
- Contamination technique for large-scale multistage programs
  - We provide an extension which accounts for the fact, that we can never compute precise optimal solution
  - Based on lower bounds from cutting-plane algorithms and upper bounds from policy value estimators
- Numerical study and comparison of two multi-stage models based on CVaR

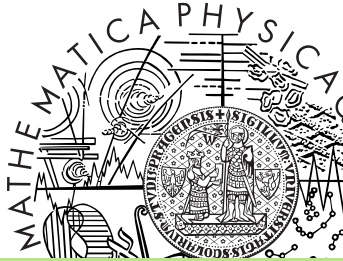

Multistage stochastic optimization

- Consider T stage stochastic program:
  - Data process $\boldsymbol{\xi} = (\boldsymbol{\xi}_1, \boldsymbol{\xi}_2, \ldots, \boldsymbol{\xi}_T)$
  - Decision process $\mathbf{x} = (\mathbf{x}_1, \ldots, \mathbf{x}_T)$
  - Stages should reflect the timing of decisions, steps can be unequal
  - Filtration $\mathcal{F}_t$ generated by the projection $\Pi_t \boldsymbol{\xi} = \boldsymbol{\xi}_{[t]} := (\boldsymbol{\xi}_1, \ldots, \boldsymbol{\xi}_t)$
  - Probability distribution of $\boldsymbol{\xi}$: $\mathbb{P}$
  - $\mathbb{P}_t$ denotes the marginal probability distribution of $\boldsymbol{\xi}_t$
  - $\mathbb{P}_t \left[\cdot | \boldsymbol{\xi}_{[t-1]} \right]$ denotes the conditional probability distribution
- The decision process is nonanticipative:
  - Decisions taken at any stage of the process do neither depend on future *realizations* of stochastic data nor on future decisions
  - $\mathbf{x}_t$ is $\mathcal{F}_t$-measurable
  - The sequence of decisions and observations is:

$$\mathbf{x}_1, \boldsymbol{\xi}_2, \mathbf{x}_2(\mathbf{x}_1, \boldsymbol{\xi}_2), \ldots, \mathbf{x}_T(\mathbf{x}_{T-1}, \boldsymbol{\xi}_2, \ldots, \boldsymbol{\xi}_T)$$

    - its random outcome $f(\mathbf{x}, \boldsymbol{\xi})$

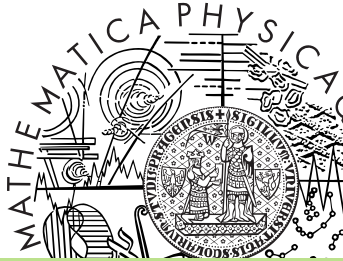

Scenario tree

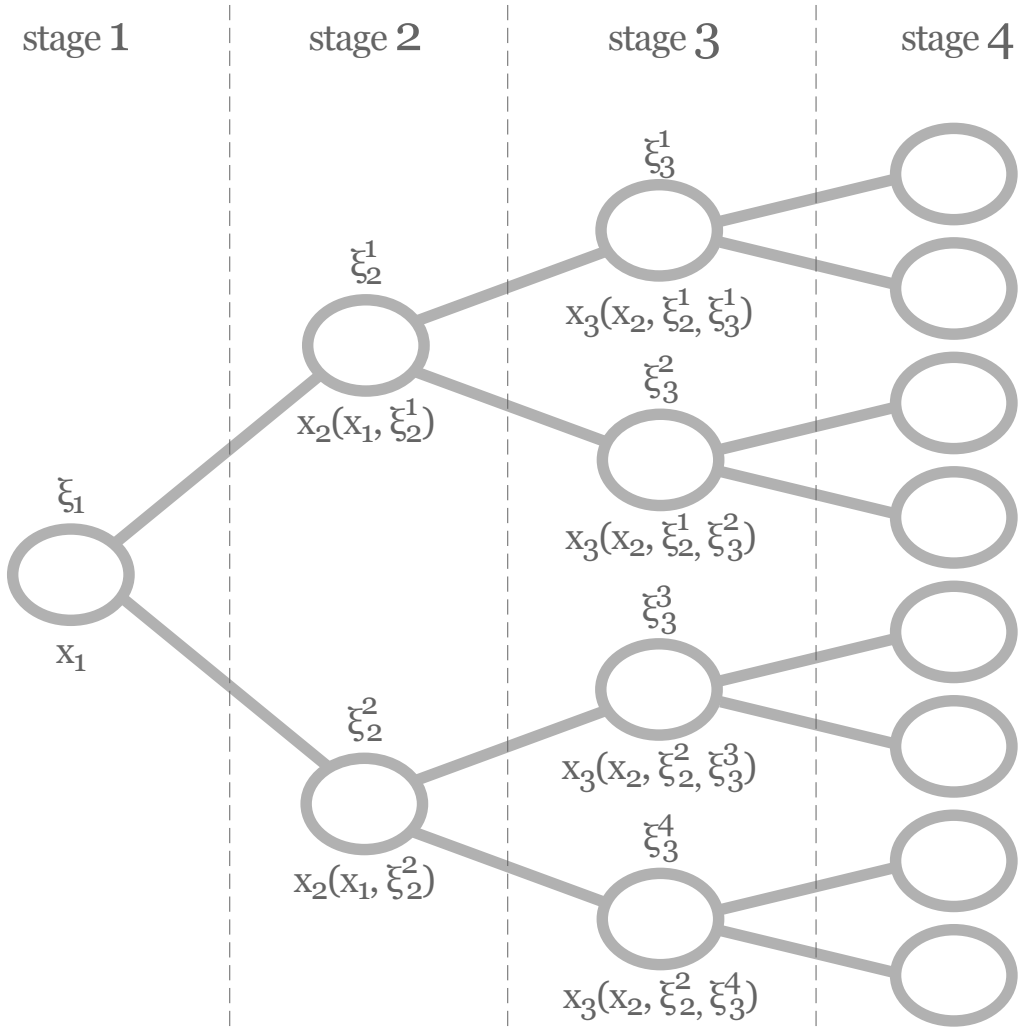

Multistage stochastic optimization

- Nested form of multistage stochastic linear program (MSLP):

$$\min_{\mathbf{x}_1 \in \mathcal{X}_1} \mathbf{c}_1^\top \mathbf{x}_1 + \mathbb{E}_{\mathbb{P}}\left[Q_2(\mathbf{x}_1, \boldsymbol{\xi}_2)\right] \text{ with } \mathcal{X}_1 := \{\mathbf{x}_1 | \, \mathbf{A}_1 \mathbf{x}_1 = \mathbf{b}_1, \, \mathbf{x}_1 \geq 0\}$$

  - With $Q_t(\mathbf{x}_{t-1}, \boldsymbol{\xi}_{[t]})$, $t = 2, \ldots, T$, defined recursively as

$$Q_t(\mathbf{x}_{t-1}, \boldsymbol{\xi}_{[t]}) = \min_{\mathbf{x}_t} \mathbf{c}_t(\boldsymbol{\xi}_{[t]})^\top \mathbf{x}_t + \mathbb{E}_{\mathbb{P}_{t+1}[\cdot | \boldsymbol{\xi}_{[t]}]}\left[Q_{t+1}(\mathbf{x}_t, \boldsymbol{\xi}_{[t+1]})\right]$$

  - $\mathcal{X}_t(\mathbf{x}_{t-1}, \boldsymbol{\xi}_{[t]})$: $\mathbf{A}_t(\boldsymbol{\xi}_{[t-1]})\mathbf{x}_t = \mathbf{b}_t(\boldsymbol{\xi}_{[t-1]}) - \mathbf{B}_t(\boldsymbol{\xi}_{[t-1]})\mathbf{x}_{t-1}$, $\mathbf{x}_t \geq 0$ a.s.,
- In the case of *stagewise independence* the conditional distributions boil down to marginal distributions $\mathbb{P}_t$ of $\boldsymbol{\xi}_t$
- We assume:
  - Constraints involving random elements hold almost surely
  - *All optimal solutions exist*, which is related with the *relatively complete recourse*
  - *All conditional expectations exist*

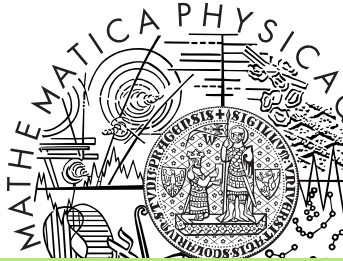

Risk-averse multistage programs

- In the risk-neutral programs possible risks are not reflected
- Risk measure is a functional which assigns a real value to the random outcome $f(\mathbf{x}, \boldsymbol{\xi})$
- Risk measures depend on decisions and probability distribution $\mathbb{P}$.
  - Filtration $\mathcal{F}_1 \subset \cdots \subset \mathcal{F}_t \cdots \subseteq \mathcal{F}$ should be taken into account
- Risk monitoring in individual stages should be incorporated

$$\min_{\mathbf{x}_1} \mathbf{c}_1^\top \mathbf{x}_1 + \rho_2 \left(\min_{\mathbf{x}_2} \mathbf{c}_2(\boldsymbol{\xi}_{[1]})^\top \mathbf{x}_2 + \cdots + \rho_T \left(\min_{\mathbf{x}_T} \mathbf{c}_T(\boldsymbol{\xi}_{[T-1]})^\top \mathbf{x}_T \right) \right)$$

  - Different risk measures ρ_t can be applied in each stage
- Coherence of ρ is mostly expected [Artzner et al., 2007]

Time consistency

- Many different definitions
- Need to distinguish between time consistency of the risk measure and time consistency of the model
- **TC1** [Carpentier, et al., 2012] The sequence of dynamic optimization problems is dynamically consistent if the optimal strategies obtained when solving the original problem remain optimal for all subsequent problems.
- **TC2** [Shapiro, 2009] At each state of the system, optimality of a decision policy should not involve states which cannot happen in the future.
- Risk-neutral stochastic programs are time consistent
- In general, time consistency for risk-averse stochastic programs does not hold true

Nested CVaR risk measure

- Consider sequence of random costs $\mathbf{Z} = (Z_1, \ldots, Z_T)$
- Nested CVaR risk measure is given by:

$$\rho^n(\mathbf{Z}) = \text{CVaR}_\alpha\left[\cdot|\mathcal{F}_1\right] \circ \cdots \circ \text{CVaR}_\alpha\left[\cdot|\mathcal{F}_{T-1}\right]\left(\sum_{t=1}^{T} Z_t\right)$$

- The interpretation is not straightforward
  - can be viewed as the cost we would be willing to pay at the first stage instead of incurring the sequence of random costs $Z_1, \ldots, Z_T$
  - cf. Ruszczyński [2010]

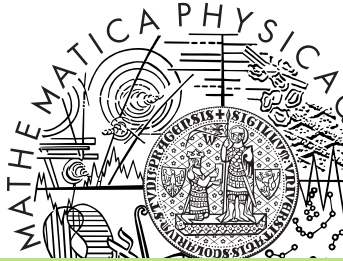

Nested CVaR model

- Risk measures are usually combined with expectation to get efficient solutions
- Given risk coefficients λ_t and random loss variable Z we define:

$$\rho_{t,\boldsymbol{\xi}_{[t-1]}}[Z] = (1-\lambda_t)\,\mathbb{E}\left[Z|\boldsymbol{\xi}_{[t-1]}\right] + \lambda_t\,\mathrm{CVaR}_{\alpha_t}\left[Z|\boldsymbol{\xi}_{[t-1]}\right]$$

- Nested model can be written:

$$\min_{\mathbf{A}_1\mathbf{x}_1=\mathbf{b}_1,\mathbf{x}_1\geq 0} \mathbf{c}_1^\top\mathbf{x}_1 + \rho_{2,\boldsymbol{\xi}_{[1]}}\left[\min_{\mathbf{B}_2\mathbf{x}_1+\mathbf{A}_2\mathbf{x}_2=\mathbf{b}_2,\mathbf{x}_2\geq 0} \mathbf{c}_2^\top\mathbf{x}_2 + \cdots\right.$$
$$\left.\cdots + \rho_{T,\boldsymbol{\xi}_{[T-1]}}\left[\min_{\mathbf{B}_T\mathbf{x}_{T-1}+\mathbf{A}_T\mathbf{x}_T=\mathbf{b}_T,\mathbf{x}_T\geq 0} \mathbf{c}_T^\top\mathbf{x}_T\right]\right]$$

- Time consistent w.r.t. [TC1] and [TC2]

Nested CVaR model

- Allows to develop dynamic programming equations, using:

$$\text{CVaR}_\alpha [Z] = \min_u \left[u + \frac{1}{\alpha} \mathbb{E} [Z - u]_+ \right]$$

- Denote $Q_t(\mathbf{x}_{t-1}, \boldsymbol{\xi}_{[t]})$, $t = 2, \ldots, T$ as the optimal value of:

$$\begin{aligned} Q_t(\mathbf{x}_{t-1}, \boldsymbol{\xi}_{[t]}) = \min_{\mathbf{x}_t, u_t} \; & \mathbf{c}_t^\top \mathbf{x}_t + \lambda_{t+1} u_t + \mathcal{Q}_{t+1}(\mathbf{x}_t, u_t, \boldsymbol{\xi}_{[t]}) \\ \text{s.t.} \; & \mathbf{A}_t \mathbf{x}_t = \mathbf{b}_t - \mathbf{B}_t \mathbf{x}_{t-1} \\ & \mathbf{x}_t \geq 0, \end{aligned}$$

- Recourse function $\mathcal{Q}_{t+1}(\mathbf{x}_t, u_t, \boldsymbol{\xi}_{[t]})$ is given by ($\mathcal{Q}_{T+1}(\cdot) \equiv 0$):

$$\mathbb{E}_{\mathbb{P}_{t+1}[\cdot|\boldsymbol{\xi}_{[t]}]} \left[(1 - \lambda_{t+1}) Q_{t+1}(\mathbf{x}_t, \boldsymbol{\xi}_{[t+1]}) + \frac{\lambda_{t+1}}{\alpha_{t+1}} \left[Q_{t+1}(\mathbf{x}_t, \boldsymbol{\xi}_{[t+1]}) - u_t \right]_+ \right.$$

Multiperiod CVaR risk measure

- Based on the following risk measure:

$$\rho^m(\mathbf{Z}) = \sum_{t=2}^{T} \mu_t \mathbb{E}\left[\mathrm{CVaR}_{\alpha_t}\left[Z_t | \mathcal{F}_{t-1}\right]\right].$$

- The difference between this risk measure and the nested CVaR risk measure is that here we apply expectation instead of the nesting
- Easier interpretation
  - Averaging of the risks in future stages
- Polyhedral risk measure
  - Solution of a multi-stage stochastic linear program of a special form
  - Optimization of the original problem can be combined with the optimization problem which defines the risk measure

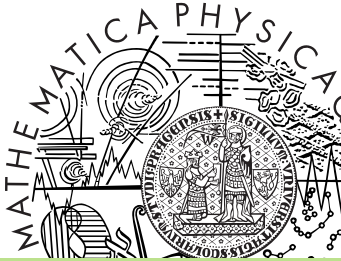

Multiperiod CVaR model

- Stochastic programming model:

$$\min_{\mathbf{x}_1,\ldots,\mathbf{x}_T} \ \mathbf{c}_1^\top \mathbf{x}_1 + \mu_2 \mathbb{E}\left[\rho_{2,\boldsymbol{\xi}_{[1]}}\left[\mathbf{c}_2^\top \mathbf{x}_2\right]\right] + \cdots + \mu_T \mathbb{E}\left[\rho_{T,\boldsymbol{\xi}_{[T-1]}}\left[\mathbf{c}_T^\top \mathbf{x}_T\right]\right]$$

$$\begin{aligned}
\text{s.t. } & \mathbf{A}_1\mathbf{x}_1 = \mathbf{b}_1 \\
& \mathbf{A}_2\mathbf{x}_2 = \mathbf{b}_2 - \mathbf{B}_2\mathbf{x}_1 \\
& \quad \vdots \\
& \mathbf{A}_T\mathbf{x}_T = \mathbf{b}_T - \mathbf{B}_T\mathbf{x}_{T-1} \\
& \mathbf{x}_t \geq 0, \ \mathbf{x}_t \in \mathcal{L}_p\left(\Omega, \mathcal{F}_t, \mathbb{P}\right) \quad t = 1, \ldots, T.
\end{aligned}$$

- Time consistent w.r.t. [TC1] and [TC2]
- Dynamic programming equations are developed, similarly to the nested case

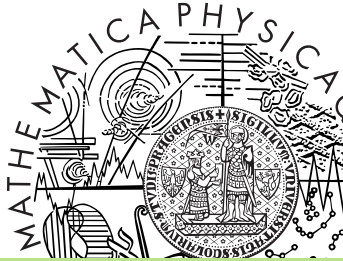

Sum of CVaR model

- The weighted sum of CVaR model is based on the following risk measure:

$$\rho^s(\mathbf{Z}) = \sum_{t=2}^{T} \mu_t \, \mathrm{CVaR}_{\alpha_t}[Z_t]$$

  with $\sum_{t=2}^{T} \mu_t = 1$, $\mu_t \geq 0 \, \forall t$.
- No nesting or averaging is present
- Easy interpretation - weighted sum of CVaR for all stages
- Related to multi-criteria optimization
  - "We want to hedge against risk in all stages separately"
- Dynamic programming equations show that all u_t are decided in the first stage
- Corresponding model is time consistent w.r.t. [TC1], but not w.r.t. [TC2]

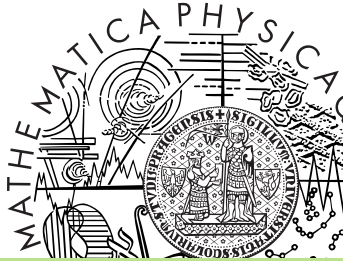

Asset allocation model

- At stage t we observe the price ratio between the new price and the old price $\mathbf{r}_t$
- $\mathbf{x}_t$ contains the optimal allocation (in USD, say)
- The total portfolio value is tracked as a multiple of the initial value
- Dynamic programming equations are very simple:

$$\begin{aligned} \min_{\mathbf{x}_t, u_t} \quad & -\mathbf{1}^\top \mathbf{x}_t + \lambda_{t+1} u_t + \mathcal{Q}_{t+1}(\mathbf{x}_t, u_t) \\ \text{s.t.} \quad & \mathbf{r}_t^\top \mathbf{x}_{t-1} - \mathbf{1}^\top \mathbf{x}_t = 0 \\ & \mathbf{x}_t \geq 0 \end{aligned}$$

- Transaction costs of $f_t \mathbf{1}^\top |\mathbf{x}_t - \mathbf{x}_{t-1}|$ can be included
- We solve problems up to 15 stages with 10^{24} scenarios, using SDDP with importance sampling

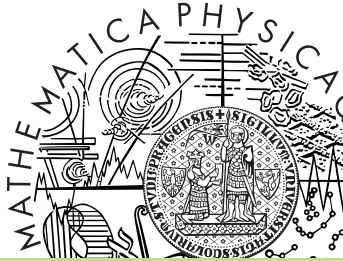

SDDP algorithm

- Starts with SAA of the problem - scenario tree, given or sampled
- Forward iteration
  - Samples $\xi^1, \ldots, \xi^J$ sample paths
  - Policy is evaluated using all the cuts collected so far
  - Value of the policy gives the upper bound
- Backward iteration
  - Subset of the scenarios from the forward iteration is chosen
  - For every chosen node the Benders' cut is calculated
    - Using all of its immediate descendants
  - Optimal value of the root problem gives the lower bound
- The bounds are compared and the process is repeated
- Relies on the stage-independence assumption
  - Cuts are valid for all nodes from given stage
  - Low memory requirements to store scenarios
  - Linear complexity w.r.t. number of stages
- CPLEX and COIN-OR used as solvers for the LPs

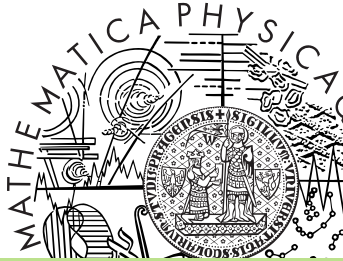

Upper bound overview

- Risk-neutral problems
  - The value of the current optimal policy can be estimated easily
  - Expectation at each node can be estimated by single chosen descendant
- Risk-averse problems
  - To estimate the CVaR value we need more descendants in practice
  - Leads to intractable estimators with exponential computational complexity, denoted by $U^{\mathbf{e}}$
- Current solution (to our knowledge)
  - Run the risk-neutral version of the same problem and determine the number of iterations needed to stop the algorithm, then run the same number of iterations on the risk-averse problem
  - Inner approximation scheme proposed by Philpott et al. [2013]
    - Works with different policy than the outer approximation
    - Probably the best alternative so far
    - Does not scale well with the dimension of $\mathbf{x}$

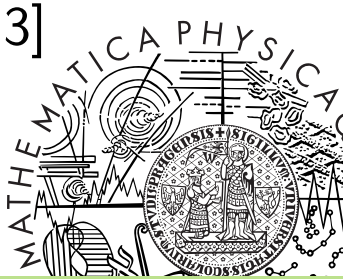

Upper bound enhancements

- State-of-the-art estimator runs with exponential complexity
- We need an estimator with linear complexity to build valid bounds
- Ideally it should be unbiased, or in practice, have small bias
- We start with the linear estimator from the risk-neutral case and include:
  - Importance sampling, with an additional assumption needed
  - Further enhancements to reduce bias and volatility

Assumption

Let $a_t(\mathbf{x}_{t-1}, \boldsymbol{\xi}_t)$ approximate the recourse value of our decisions $\mathbf{x}_{t-1}$ after the random parameters $\boldsymbol{\xi}_t$ have been observed, and let $a_t(\mathbf{x}_{t-1}, \boldsymbol{\xi}_t)$ be cheap to evaluate.

- For example in our portfolio model:
  $a_t(\mathbf{x}_{t-1}, \boldsymbol{\xi}_t) = -\boldsymbol{\xi}_t^\top \mathbf{x}_{t-1} = -\mathbf{r}_t^\top \mathbf{x}_{t-1}$

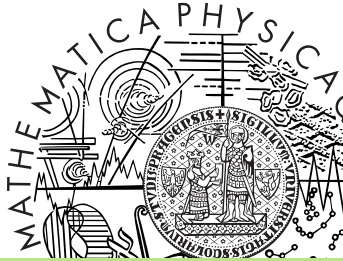

Importance sampling

- We start with standard pmf, all probabilities equal for D_t scenarios:

$$g_t(\boldsymbol{\xi}_t) = \frac{1}{D_t}\mathbb{I}\left[\boldsymbol{\xi}_t \in \left\{\boldsymbol{\xi}_t^1, \ldots, \boldsymbol{\xi}_t^{D_t}\right\}\right]$$

  - This is not required by SDDP and can be easily relaxed
- Denote $u_a = \mathrm{VaR}_\alpha\left[a_t(\mathbf{x}_{t-1}, \boldsymbol{\xi}_t)\right]$
- We change the measure to put more weight to the CVaR nodes:

$$h_t(\boldsymbol{\xi}_t|\mathbf{x}_{t-1}) = \begin{cases} \dfrac{\beta_t}{\alpha_t} g_t\, \mathbb{I}\left[\boldsymbol{\xi}_t \in \left\{\boldsymbol{\xi}_t^1, \ldots, \boldsymbol{\xi}_t^{D_t}\right\}\right], & \text{if } a_t(\mathbf{x}_{t-1}, \boldsymbol{\xi}_t) \geq u_a \\ \dfrac{1-\beta_t}{1-\alpha_t} g_t\, \mathbb{I}\left[\boldsymbol{\xi}_t \in \left\{\boldsymbol{\xi}_t^1, \ldots, \boldsymbol{\xi}_t^{D_t}\right\}\right], & \text{if } a_t(\mathbf{x}_{t-1}, \boldsymbol{\xi}_t) < u_a, \end{cases}$$

  - We select forward nodes according to this measure
  - $\mathbb{E}_{g_t}[Z] = \mathbb{E}_{h_t}\left[Z\,\frac{g_t}{h_t}\right]$
  - We start with $\beta_t = \frac{1}{2}$ and optimize its value later

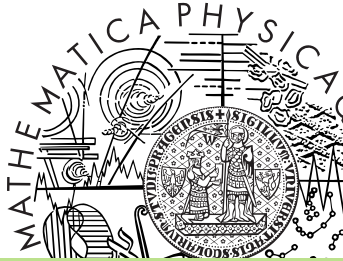

Linear estimators

- The nodes can be selected randomly from the standard i.i.d. measure or from the importance sampling measure
- For stages $t = 2, \dots, T$ is given by:

$$\hat{\mathfrak{v}}_t(\xi_{t-1}^{j_{t-1}}) = (1-\lambda_t)\left((\mathbf{c}_t^{j_t})^\top \mathbf{x}_t^{j_t} + \hat{\mathfrak{v}}_{t+1}(\xi_t^{j_t})\right) + \\ + \lambda_t u_{t-1}^{j_{t-1}} + \frac{\lambda_t}{\alpha_t}\left[(\mathbf{c}_t^{j_t})^\top \mathbf{x}_t^{j_t} + \hat{\mathfrak{v}}_{t+1}(\xi_t^{j_t}) - u_{t-1}^{j_{t-1}}\right]_+$$

- $\hat{\mathfrak{v}}_{T+1}(\xi_T^{j_T}) \equiv 0$
- Along a single path for scenario j the cost is estimated by:

$$\hat{\mathfrak{v}}(\xi^j) = \mathbf{c}_1^\top \mathbf{x}_1 + \hat{\mathfrak{v}}_2$$

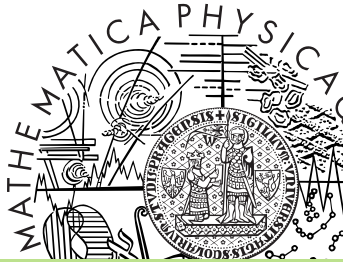

Linear estimators

- For scenarios selected via the original pmf we have the naive estimator

$$U^{\mathbf{n}} = \frac{1}{M} \sum_{j=1}^{M} \hat{\mathfrak{v}}(\boldsymbol{\xi}^j)$$

- With weights again defined via

$$w(\boldsymbol{\xi}^j) = \prod_{t=2}^{T} \frac{g_t(\xi_t)}{h_t(\xi_t | \mathbf{x}_{t-1})}$$

- For scenarios selected via the IS pmf we have the IS estimator

$$U^{\mathbf{i}} = \frac{1}{\sum_{j=1}^{M} w(\boldsymbol{\xi}^j)} \sum_{j=1}^{M} w(\boldsymbol{\xi}^j) \hat{\mathfrak{v}}(\boldsymbol{\xi}^j)$$

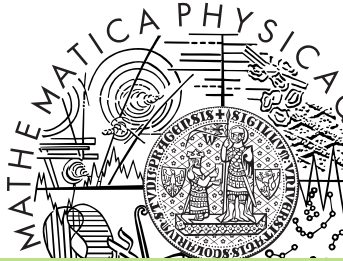

Linear estimators - validity

Function $U^{\mathbf{i}}$ provides an asymptotic upper bound estimator for the SAA version of the presented optimization problem.

Proposition

Assume finite optimal value, relatively complete recourse and interstage independence. Let φ denote the optimal value. Let $\boldsymbol{\xi}$ denote a sample path selected under the empirical distribution, and let $\hat{\mathfrak{v}}(\boldsymbol{\xi})$ be defined for that sample path. Then $\mathbb{E}_g\left[\hat{\mathfrak{v}}(\boldsymbol{\xi})\right] \geq \varphi$. Furthermore if $\boldsymbol{\xi}^j$, $j = 1, \ldots, M$, are i.i.d. and generated by the IS pmfs then $U^{\mathbf{i}} \to \mathbb{E}_g\left[\hat{\mathfrak{v}}(\boldsymbol{\xi})\right]$, w.p.1, as $M \to \infty$.

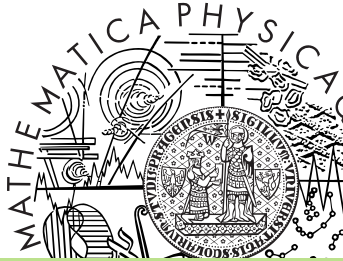

Upper bound enhancements

- Linear estimator still degrades for problems with 10 or 15 stages
- The reason for the bias of the estimator comes from poor estimates of CVaR
  - Once the cost estimate for stage t exceeds u_{t-1} the difference is multiplied by α_t^{-1}
  - When estimating stage $t-1$ costs in the nested model we sum stage $t-1$ costs and stage t estimate which means that we usually end up with costs greater than u_{t-2} so another multiplication occurs
  - This brings both bias and large variance

Assumption

For every stage $t = 2, \ldots, T$ and decision $\mathbf{x}_{t-1}$ the approximation function a_t satisfies:

$$Q_t \geq \mathrm{VaR}_{\alpha_t}[Q_t] \text{ if and only if } a_t \geq \mathrm{VaR}_{\alpha_t}[a_t].$$

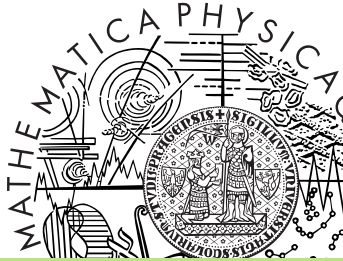

Improved estimator

- Provided that the equivalence assumption holds we can reduce the bias of the estimator
  - The positive part operator in the equation is used only in the case of CVaR node
- For stages $t = 2, \ldots, T$ we have

$$\hat{\mathfrak{v}}_t^{\mathbf{a}}(\xi_{t-1}^{j_{t-1}}) = (1-\lambda_t)\left((\mathbf{c}_t^{j_t})^\top \mathbf{x}_t^{j_t} + \hat{\mathfrak{v}}_{t+1}^{\mathbf{a}}(\xi_t^{j_t})\right) + \lambda_t u_{t-1}^{j_{t-1}} + \\ + \mathbb{I}[a_t > \mathrm{VaR}_{\alpha_t}[a_t]]\frac{\lambda_t}{\alpha_t}\left[(\mathbf{c}_t^{j_t})^\top \mathbf{x}_t^{j_t} + \hat{\mathfrak{v}}_{t+1}^{\mathbf{a}}(\xi_t^{j_t}) - u_{t-1}^{j_{t-1}}\right]_+$$

- $\hat{\mathfrak{v}}_{T+1}^{\mathbf{a}}(\xi_T^{j_T}) \equiv 0$
-

$$U^{\mathbf{a}} = \frac{1}{\sum_{j=1}^M w(\xi^j)} \sum_{j=1}^M w(\xi^j)\hat{\mathfrak{v}}^{\mathbf{a}}(\xi^j)$$

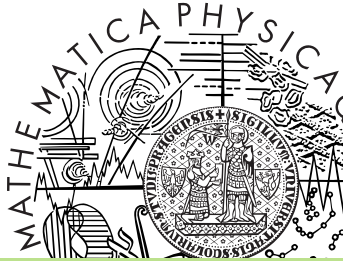

Improved estimator - validity

Function $U^{\mathbf{a}}$ provides an asymptotic upper bound estimator for the SAA version of the presented optimization problem.

Proposition

Assume finite optimal value, relatively complete recourse and interstage independence. Let φ denote the optimal value Let $\boldsymbol{\xi}$ denote a sample path selected under the empirical distribution and let the "perfect ordering" assumption hold. Then $\mathbb{E}_g\left[\hat{\mathfrak{v}}^{\mathbf{a}}(\boldsymbol{\xi})\right] \geq \varphi$. If $\boldsymbol{\xi}^j$, $j = 1, \ldots, M$, are i.i.d. and generated by the IS pmfs then $U^{\mathbf{a}} \to \mathbb{E}_g\left[\hat{\mathfrak{v}}^{\mathbf{a}}(\boldsymbol{\xi})\right]$, w.p.1, as $M \to \infty$. Furthermore, if the subproblems induce the same policy for both $\hat{\mathfrak{v}}(\boldsymbol{\xi})$ and $\hat{\mathfrak{v}}^{\mathbf{a}}(\boldsymbol{\xi})$ then $\mathbb{E}_g\left[\hat{\mathfrak{v}}(\boldsymbol{\xi})\right] \geq \mathbb{E}_g\left[\hat{\mathfrak{v}}^{\mathbf{a}}(\boldsymbol{\xi})\right]$.

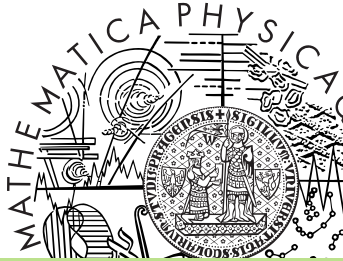

Improved estimator results

- Comparison with exponential estimator $U^{\mathbf{e}}$ from the literature:

T	$\underline{z}$	$U^{\mathbf{n}}$ (s.d.)	$U^{\mathbf{i}}$ (s.d.)	$U^{\mathbf{a}}$ (s.d.)	$U^{\mathbf{e}}$ (s.d.)
2	-0.9518	-0.9515 (0.0020)	-0.9517 (0.0012)	-0.9517 (0.0011)	-0.9518 (0.0019)
3	-1.8674	-1.8300 (0.0145)	-1.8285 (0.0108)	-1.8656 (0.0060)	-1.8013 (0.0302)
4	-2.7811	-2.4041 (0.1472)	-2.3931 (0.1128)	-2.7764 (0.0126)	-2.6027 (0.0883)
5	-3.6794	-3.4608 (0.1031)	-3.4963 (0.1008)	-3.6731 (0.0303)	-2.9031 (0.5207)
10	-7.6394	9.3×10^4 (1.4×10^4)	9.0×10^4 (8.7×10^4)	-7.5465 (0.2562)	1.5×10^7 (1.3×10^6)
15	-11.5188	NA	NA	-11.0148 (0.6658)	NA

- For $T = 2, \ldots, 5$ variance reduction of $U^{\mathbf{a}}$ relative to $U^{\mathbf{e}}$: 3 to 25 to 50 to 300.
- Computation time for $U^{\mathbf{n}}$ for $T = 5, 10, 15$: 8.7 sec. to 31.6 sec. to 67.4 sec.
- Computation time for $U^{\mathbf{a}}$ for $T = 5, 10, 15$: 6.8 sec. to 30.0 sec. to 66.5 sec.
- Can be extended to handle more complex models, for instance asset allocation with transaction costs

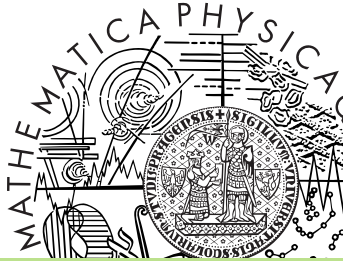

Variance reduction

- We consider a functional from our model:

$$\mathcal{Q}_\alpha[Z] = (1-\lambda)\,\mathbb{E}[Z] + \lambda\,\mathrm{CVaR}_\alpha[Z]$$

- We define:

$$Q^s = (1-\lambda)\,Z + \lambda\left(u_Z + \frac{1}{\alpha}[Z-u_Z]_+\right)$$

$$Q^i = \frac{g}{h}\left((1-\lambda)\,Z + \lambda\left(u_Z + \frac{1}{\alpha}[Z-u_Z]_+\right)\right)$$

- It clearly holds $\mathcal{Q} = \mathbb{E}_h\left[Q^i\right] = \mathbb{E}_g\left[Q^s\right]$
- Our aim is to find suitable parameter β for our importance sampling scheme, so that $\mathrm{var}_h\left[Q^i\right] < \mathrm{var}_g\left[Q^s\right]$

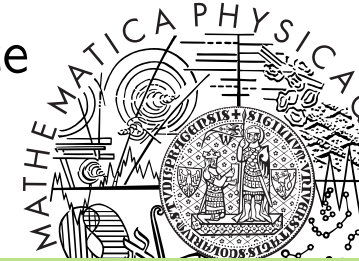

Example - normal distribution

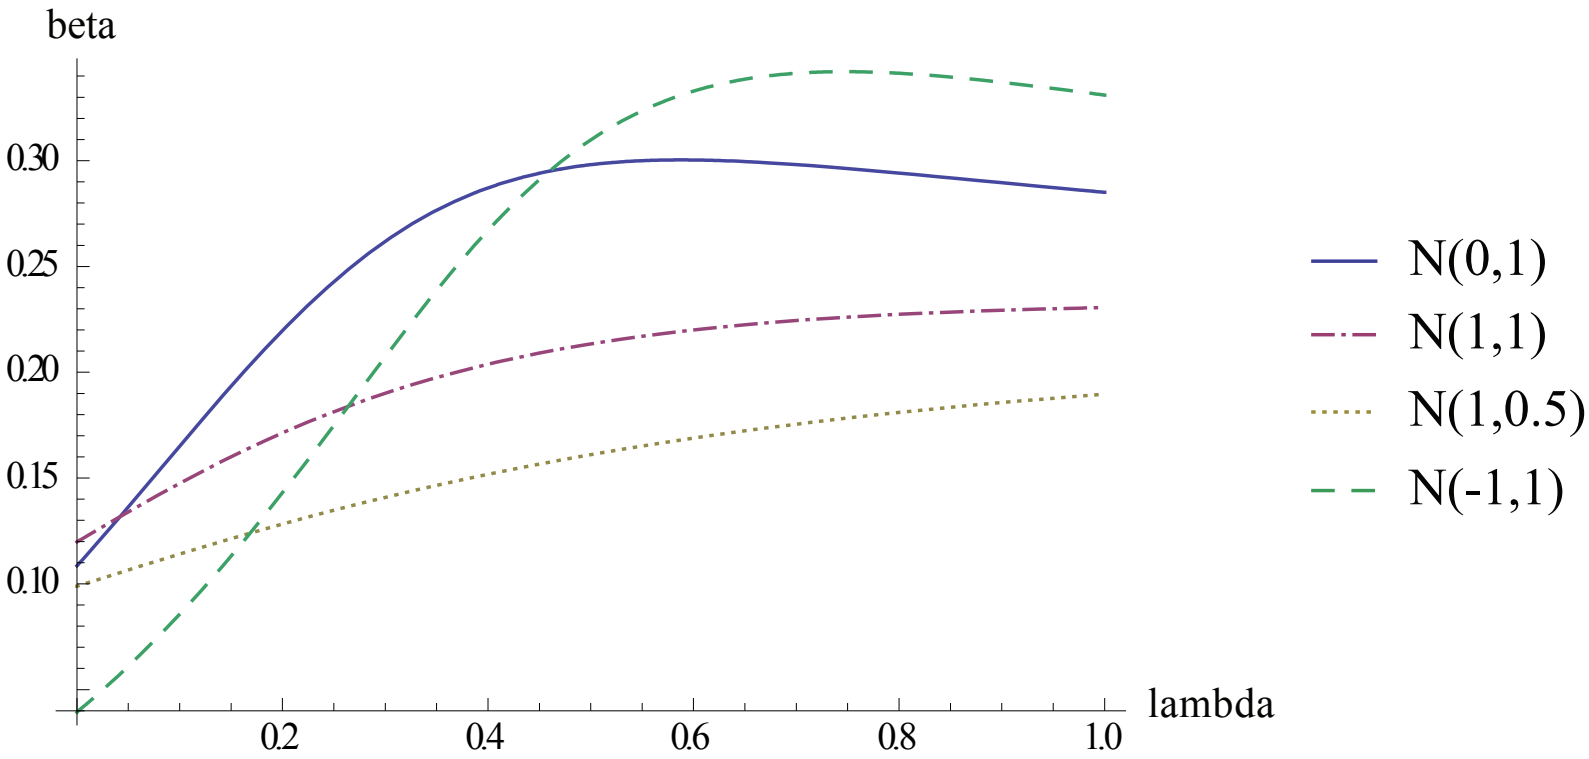

Other distributions

- We can also estimate the suitable β by sampling
  - We choose a mesh of possible values, e.g. $\mathcal{B} = \{0.01, 0.02, \ldots, 0.99\}$
  - For each of them, we sample prescribed number of scenarios, Z^j
  - We compute the mean and variance of the values Q^j given by Z^j
  - The lowest variance is selected as a suitable choice of β
- Even though we start with log-normal distribution, convolution and nested structure of our model brings complex transformations
- Different values of β should be selected for every stage, as the parameters of the distributions also vary
- We have values β estimated by the algorithm for general distributions on $100,000$ scenarios
- An alternative would be $\beta = 0.3$ which provides low variance for most of our charts

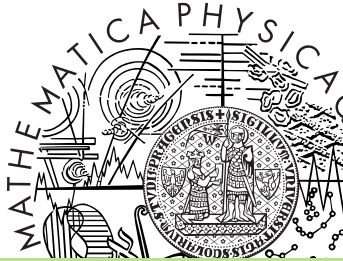

Optimal beta for the asset allocation model

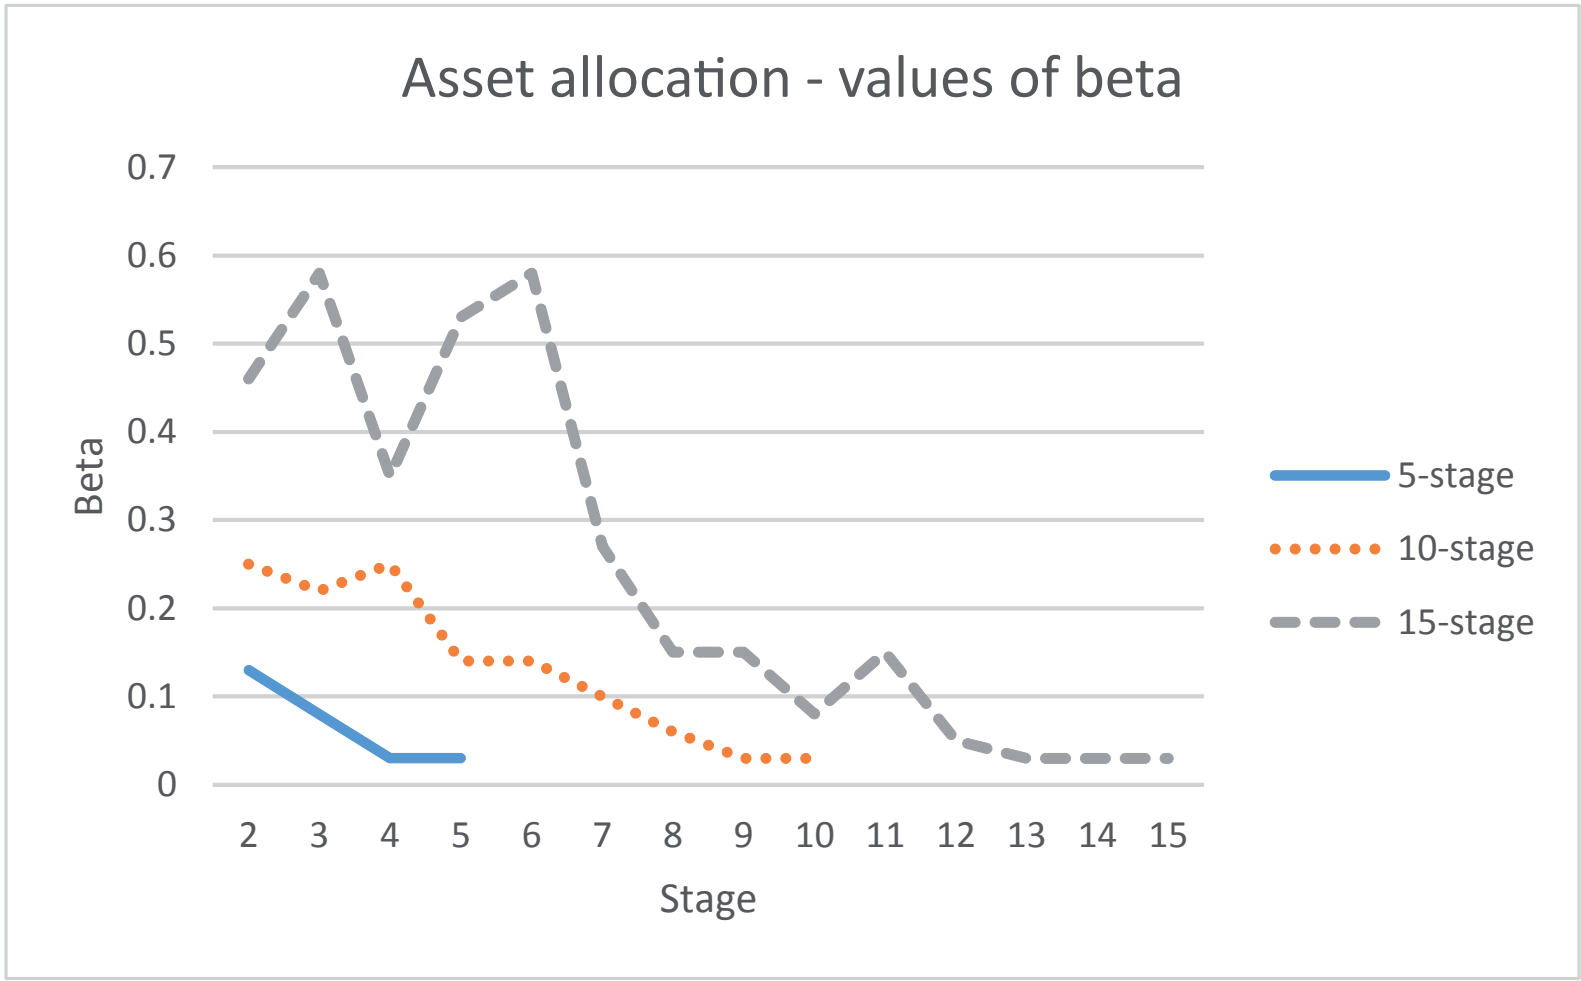

Results

- Standard Monte Carlo setup $\hat{\mathcal{Q}}^s$ ($\beta_t = \alpha_t = 0.05$)
- Improved estimator $\hat{\mathcal{Q}}^i$ with variable β_t based on our analysis
- Lower bound $\underline{z}$

T	total scenarios	$\underline{z}$	$\hat{\mathcal{Q}}^s$ (s.d.)	$\hat{\mathcal{Q}}^i$ (s.d.)
5	$6,250,000$	-3.5212	-3.5166 (0.0168)	-3.5170 (0.0111)
10	$\approx 10^{15}$	-7.3885	-7.2833 (0.2120)	-7.2838 (0.0303)
15	$\approx 10^{24}$	-10.4060	-10.1482 (0.8184)	-10.1245 (0.1355)

- For $T = 5,\ 10,\ 15$ we achieved roughly 35%, 85% and 85% reduction of standard deviation
- Negligible effect on computation times

Contamination for multistage risk-averse problems

- Captures various changes in the probability distribution
- Assume that it's possible to reformulate the stochastic program as:

$$\min_{\mathbf{x}\in\mathcal{X}} F(\mathbf{x},\mathbb{P}) := \min_{\mathbf{x}\in\mathcal{X}} \int_{\Omega} f(\mathbf{x},\boldsymbol{\xi})\mathbb{P}(\mathrm{d}\boldsymbol{\xi})$$

  - Simplest case of contamination, we obtain global bounds
  - It's possible to reformulate our CVaR model with auxiliary variable u in this way
- Let $\mathbb{Q}$ be another *fixed* probability distribution and define contaminated distributions

$$\mathbb{P}_k := (1-k)\mathbb{P} + k\mathbb{Q},\ k \in [0,1]$$

  - Suppose that the stochastic program has a solution $\varphi(k) := \min_{\mathbf{x}\in\mathcal{X}} F(\mathbf{x},\mathbb{P}_k)$ for all these distributions

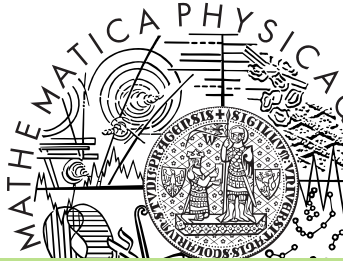

Contamination for multistage risk-averse problems

- Suppose nonempty, bounded set of optimal solutions $\mathcal{X}^*(\mathbb{P})$ of the initial stochastic program
- Then the directional derivative is given by:

$$\varphi'(0^+) = \min_{\mathbf{x} \in \mathcal{X}^*(\mathbb{P})} F(x, \mathbb{Q}) - \varphi(0)$$

- $\varphi(k)$ concave on $[0, 1]$
- The contamination bounds follow:

$$(1-k)\varphi(0) + k\varphi(1) \leq \varphi(k) \leq \varphi(0) + k\varphi'(0^+),\ k \in [0, 1]$$

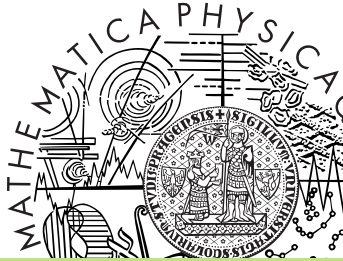

Contamination for multistage risk-averse problems

- For large-scale problems we cannot compute a precise solution
  - We apply SDDP to the sampled distributions $\hat{\mathbb{P}}$ and $\hat{\mathbb{Q}}$
  - We have deterministic lower bound $\underline{\varphi}$ for problems under $\hat{\mathbb{P}}$ and $\hat{\mathbb{Q}}$
  - We use our estimator to obtain upper bound $\overline{\varphi}$ under $\hat{\mathbb{P}}$
- Let $\tilde{\mathbf{x}}^*$ be the approximate solution of the initial problem
  - We compute upper estimate $\bar{F}(\tilde{\mathbf{x}}^*, \hat{\mathbb{Q}})$ of $F(\tilde{\mathbf{x}}^*, \hat{\mathbb{Q}})$ in the similar manner as we compute our improved upper bound
    - The solution of the initial problem is represented by the sets of cuts
    - We sample the scenarios from $\hat{\mathbb{Q}}$
    - We calculate the solution using the cuts from the initial problem and the scenario from the contaminating problem
    - This solution is used in the upper bound calculation
- The approximate contamination bounds are given by:

$$(1-k)\underline{\varphi}(\hat{\mathbb{P}}) + k\underline{\varphi}(\hat{\mathbb{Q}}) \leq \varphi(\hat{\mathbb{P}}_k) \leq (1-k)\overline{\varphi}(\hat{\mathbb{P}}) + k\bar{F}(\tilde{\mathbf{x}}^*, \hat{\mathbb{Q}})$$

Numerical results

- Monthly data from Prague Stock Exchange, January 2009 to February 2012
- Risk aversion coefficients set to $\lambda_t = 10\%$
- Contaminating distribution $\mathbb{Q}$ was obtained by increasing the variance by 20%
- 3 and 5 stage problems with $1,000$ descendants per node
- We have calculated the derivative values for 10 times and used their mean as well as empirical statistical upper bound

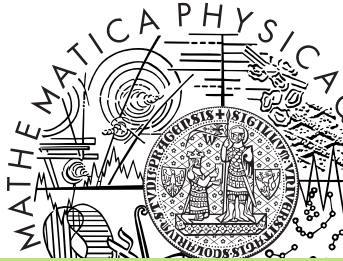

Results - 3 stages without transaction costs

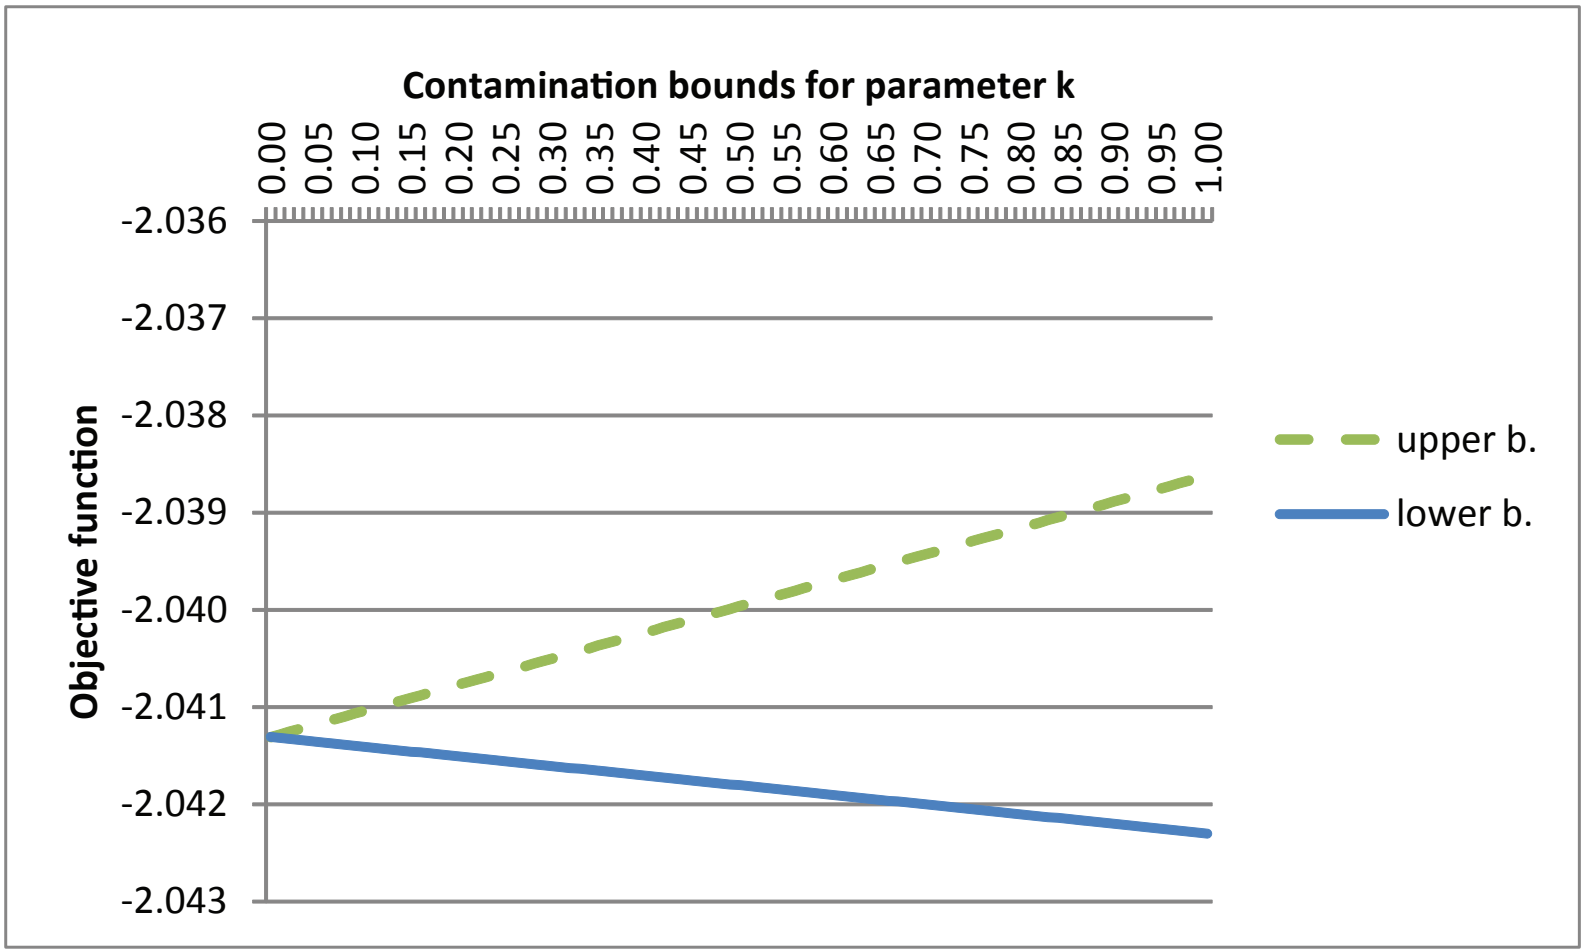

Results - 5 stages without transaction costs

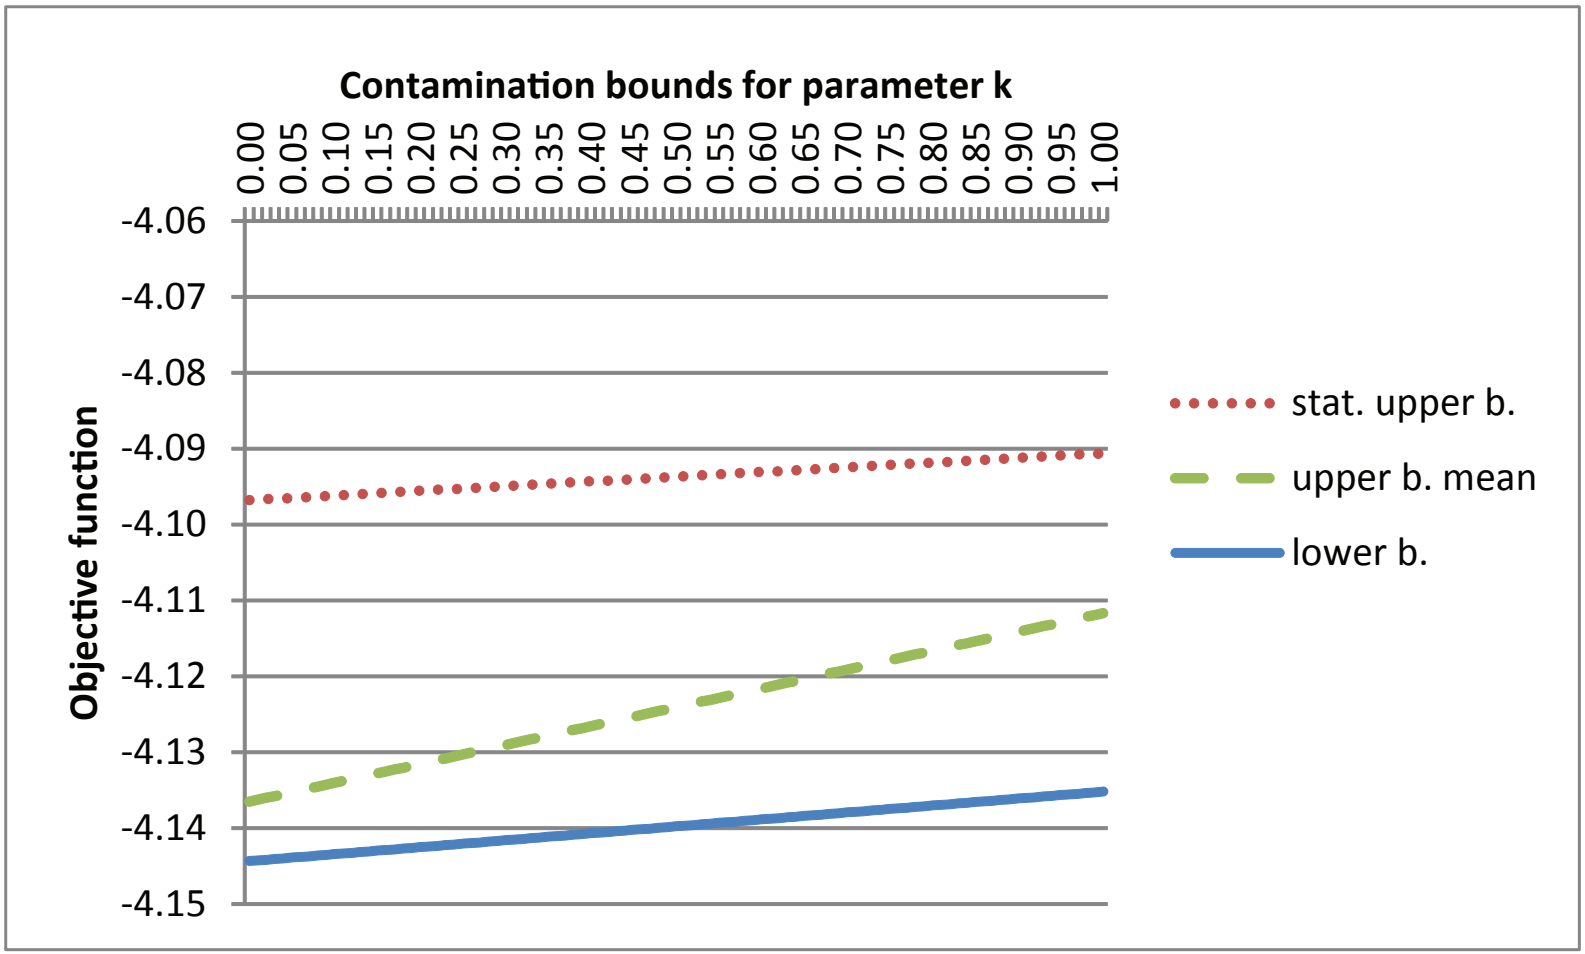

Model comparison

- Monthly data from Prague Stock Exchange, January 2009 to February 2012

asset	**mean**	**std. deviation**
AAA	1.0290	0.1235
CETV	0.9984	0.2469
ČEZ	0.9990	0.0647
ERSTE GROUP BANK	1.0172	0.1673
KOMERČNÍ BANKA	1.0110	0.1157
ORCO	1.0085	0.2200
PEGAS NONWOVENS	1.0221	0.0863
PHILIP MORRIS ČR	1.0213	0.0719
TELEFÓNICA C.R.	0.9993	0.0595
UNIPETROL	1.0079	0.0843
VIENNA INSURANCE GROUP	1.0074	0.1100

Model comparison

- Two different models, nested and multi-period CVaR risk measure
- Two different distributions:
  - $\mathbb{P}$ based on the input data
  - $\mathbb{Q}$ constructed from $\mathbb{P}$ by increasing the variance by 20% to test the stability
- We have repeated the sampling for 10 times
- The CVaR levels α_t were always set to 5%
- No transaction costs, $f_t = 0\%$
- Three-stage model with $1,000$ descendants per node, total of $1,000,000$ scenarios
- Best performing assets - AAA, PEGAS and PHILLIP MORRIS

Model comparison

- Both models relatively stable with respect to variance of the underlying distribution
- Nested model has more stable solutions and better diversification
- When increasing λ_t, solutions become more stable in both models

λ_t	model	distr.	AAA	PEGAS	PHILIP M.
0.1	nested	$\hat{\mathbb{P}}$	0.2388 (0.1133)	0.3893 (0.1109)	0.3720 (0.1011)
0.1	nested	$\hat{\mathbb{Q}}$	0.2718 (0.1600)	0.3582 (0.0902)	0.3700 (0.1565)
0.1	multiper.	$\hat{\mathbb{P}}$	0.6034 (0.3681)	0.2262 (0.2084)	0.1704 (0.2000)
0.1	multiper.	$\hat{\mathbb{Q}}$	0.6032 (0.3453)	0.1660 (0.1562)	0.2308 (0.2369)
0.2	nested	$\hat{\mathbb{P}}$	0.1774 (0.0681)	0.4132 (0.0774)	0.4032 (0.0907)
0.2	nested	$\hat{\mathbb{Q}}$	0.1730 (0.0541)	0.3471 (0.0566)	0.4545 (0.0429)
0.2	multiper.	$\hat{\mathbb{P}}$	0.3081 (0.1472)	0.2993 (0.1757)	0.3926 (0.0990)
0.2	multiper.	$\hat{\mathbb{Q}}$	0.3127 (0.1776)	0.3963 (0.0975)	0.2910 (0.1781)

Future research

- Statistical properties of the proposed upper bound estimators
- Approximation functions applicable in importance sampling schemes for various practical problems
- Application in hydroelectric scheduling under inflow uncertainty
- Develop analogous schemes and estimators for other risk measures
  - Spectral risk measures based on CVaR
- Scenario reduction techniques under stage-wise independence
- More general structures without the stage-wise independence assumption
  - Markov chains
  - Additive dependence models
- Implement parallel processing for SDDP

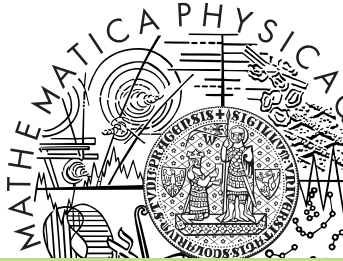

Conclusion

Thank you!

Václav Kozmík
vaclav@kozmik.cz

References

- ARTZNER, P., DELBAEN, F., EBER, J.-M., HEATH, D. and KU, H. (2007): *Coherent multiperiod risk adjusted values and Bellman's principle*, Annals of Operations Research vol. 152, pp. 2–22
- CARPENTIER, P., CHANCELIER, J. P., COHEN, G., DE LARA, M. and GIRARDEAU, P. (2012): *Dynamic consistency for stochastic optimal control problems*, Annals of Operations Research vol. 200, pp. 247–263
- DUPAČOVÁ, J. (1995): *Postoptimality for multistage stochastic linear programs*, Ann. Oper. Res. vol. 56, pp. 65–78
- PEREIRA, M. V. F. and PINTO, L. M. V. G. (1991): *Multi-stage stochastic optimization applied to energy planning*, Mathematical Programming vol. 52, pp. 359–375.

References

- PHILPOTT, A. B., DE MATOS, V. L., FINARDI, E. C. (2013): *On Solving Multistage Stochastic Programs with Coherent Risk Measures*. Oper. Res. vol. 61, pp. 957–970
- PHILPOTT, A. B. and GUAN, Z. (2008): *On the convergence of sampling-based methods for multi-stage stochastic linear programs*, Operations Research Letters vol. 36, pp. 450–455.
- ROCKAFELLAR, R. T. and URYASEV, S. (2002): *Conditional value at risk for general loss distributions*, Journal of Banking & Finance vol. 26, pp. 1443–1471
- RUSZCZYŃSKI, A. (2010): *Risk-averse dynamic programming for Markov decision processes*, Mathematical Programming, vol. 125, pp. 235–261.

References

- RUSZCZYŃSKI, A. and SHAPIRO, A. (2006): *Conditional risk mappings*, Mathematics of Operations Research vol. 31, pp. 544–561.
- SHAPIRO, A. (2011): *Analysis of stochastic dual dynamic programming method*, European Journal of Operational Research 209, pp. 63-72.
- SHAPIRO, A. (2009): *On a time consistency concept in risk averse multistage stochastic programming*, Operations Research Letters vol. 37, pp. 143–147

Publications

- DUPAČOVÁ, J. and KOZMÍK, V. (2014): *Stress Testing for Risk-Averse Stochastic Programs*
  - To appear in Acta Mathematica Universitatis Comenianae
- KOZMÍK, V. (2014): *On Variance Reduction of Mean-CVaR Monte Carlo Estimators*
  - Accepted in Computational Management Science, doi: 10.1007/s10287-014-0225-7
- DUPAČOVÁ, J. and KOZMÍK, V. (2014): *Structure of Risk-Averse Multistage Stochastic Programs*
  - Accepted in OR Spectrum, doi: 10.1007/s00291-014-0379-2
- KOZMÍK, V. and MORTON, D. (2014): *Evaluating policies in risk-averse multi-stage stochastic programming*
  - Accepted in Mathematical Programming Series A, doi: 10.1007/s10107-014-0787-8

Publications

- KOZMÍK, V. (2012): *Multistage risk-averse asset allocation with transaction costs*, Proceedings of 30^{th} International Conference Mathematical Methods in Economics, pp. 455–460.
- KOZMÍK, V. and KOPA, M. (2011): *Stability of Mean-Risk Models with Log-Normal Distribution of Returns*, Mathematical Methods in Economics 2011 Proceedings, pp. 363–368
- KOZMÍK, V. (2010): *Convergence of Approximate Solutions in Mean-Risk Models*, Proceedings of 47th EWGFM meeting pp. 81–88

Citations

- RUSZCZYŃSKI, A. (2014): *Dynamic risk-averse optimization*, Safety, Reliability, Risk and Life-Cycle Performance of Structures and Infrastructures, New York: CRC Press, ISBN 978-1138000865.
- HOMEM-DE-MELLO, T. and BAYRAKSAN, G. (2014): *Monte Carlo Sampling-Based Methods for Stochastic Optimization*, Surveys in Operations Research and Management Science, vol. 19, pp. 56–85.
- GUIGUES, V. (2014): *Convergence analysis of sampling-based decomposition methods for risk-averse multi-stage stochastic convex programs*, Preprint
- MACEIRA, M. E. P., et al. (2014): *Application of CVaR Risk Aversion Approach in the Expansion and Operation Planning and for Setting the Spot Price in the Brazilian Hydrothermal Interconnected System*. In: Power Systems Computation Conference, Wroclav.

Citations

- DUPAČOVÁ, J. (2013): *Ways and means with scenarios*, Bulletin of the Czech Econometric Society, vol. 20, pp. 112–123.
- RUSZCZYŃSKI, A. (2013): *Advances in Risk-Averse Optimization*, Tutorials in Operations Research: Theory Driven by Influential Application, Maryland: Informs, ISBN 978-0-9843378-4-2.
- WERNER, A. S. , PICHLER, A., MIDTHUN, K. T., HELLEMO, L. and TOMASGARD, A. (2013): *Risk Measures in Multi-Horizon Scenario Trees*, Handbook of Risk Management in Energy Production and Trading, New York: Springer, ISBN 978-1461490340.
- ZHANG, W. (2013): *Water Network Design and Management via Stochastic Programming*, PhD Thesis, The University of Arizona.