Mixed-Strategy Chance Constrained Optimal Control

June 18, 2013
American Control Conference

Masahiro Ono
J. (Bob) Balaram

Jet Propulsion Laboratory
California Institute of Technology

Yoshiaki Kuwata

Space Exploration Technologies Corp.

Government sponsorship acknowledged.

Mixed Strategy in Game Theory

- A player can be better off by making a decision stochastically

Payoff function		Player 1	
		Head	Tail
Player 2	Head	1, -1	-1, 1
	Tail	-1, 1	1, -1

(Both players maximizes the payoff)

- If Player 1 chooses Head deterministically
  - Player 2 chooses Tail; Player 1's payoff : -1
- If Player 1 chooses Tail deterministically
  - Player 2 chooses Head; Player 1's payoff : -1
- If Player 1 chooses Head and Tail with 50% chances
  - Player 1's payoff : 0

Mixed Strategy in Optimal Control

Three available control options

Control	Cost	Risk
u_1	40	0.5%
u_2	30	1.0%
u_3	10	1.5%

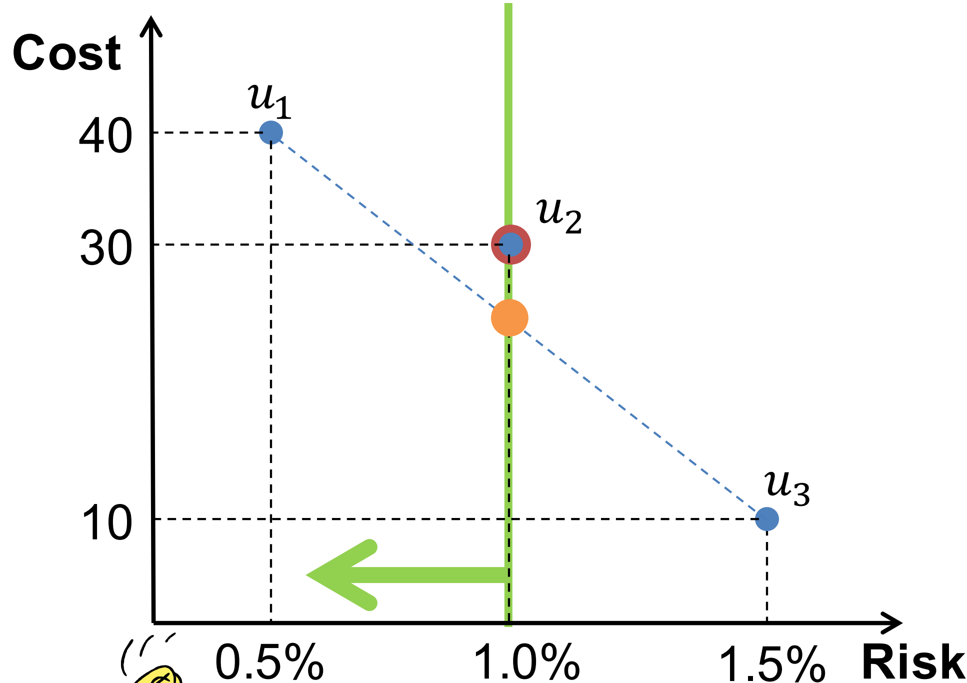

Problem:
- Minimize expected cost
- Risk $\leq$ 1%

Optimal control input: u_2

= pure control strategy

Optimal *mixed control strategy*:

- Choose u_1 with 50%
- Choose u_3 with 50%

Expected cost: 25, Risk: 1.0%

Problem Formulation

- u: *control strategy*
  - Control sequence for MPC
  - Control policy for DP
- Feasible control set $\mathcal{U}$
- w: uncertain parameter
  - With known prob. distribution
- Cost function: $f(u, w)$
- Constraint: $g(u, w) \succ 0$

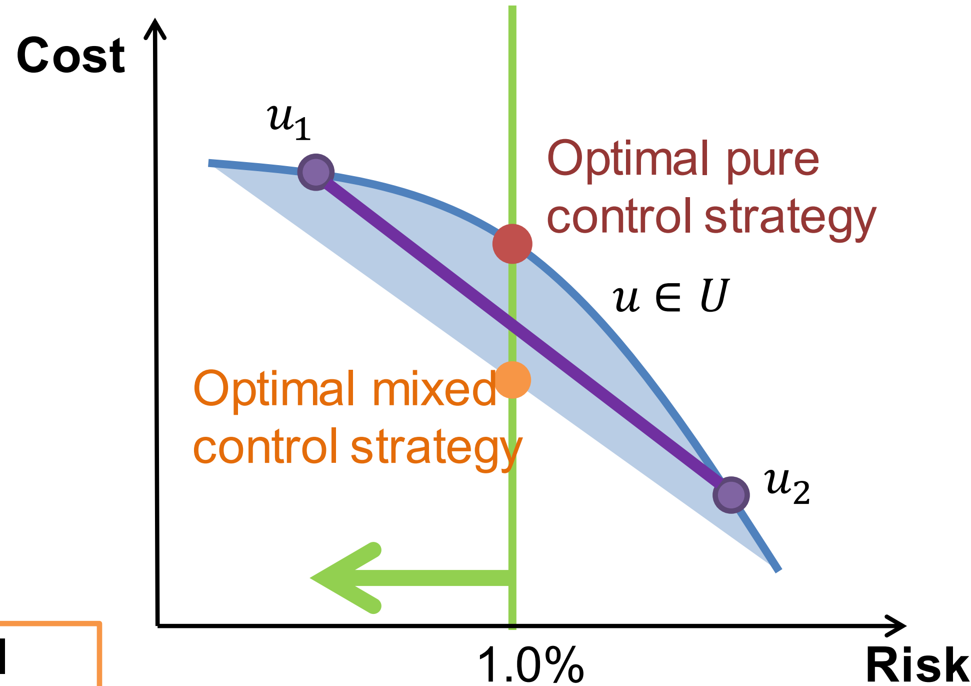

Chance-constrained optimal control problem

$$\min_{u \in \mathcal{U}} \mathbb{E}[f(u, w)]$$
$$\text{s.t. } \Pr[g(u, w) \succ 0]$$

Cost and Risk of Mixed Control Strategy

Pure control strategy u_1

- Expected cost : $c_1 \coloneqq \mathbb{E}[f(u_1, w)]$
- Risk: $r_1 \coloneqq \Pr[g(u_1, w) \succ 0]$

Pure control strategy u_2

- Expected cost : $c_2 \coloneqq \mathbb{E}[f(u_1, w)]$
- Risk: $r_2 \coloneqq \Pr[g(u_1, w) \succ 0]$

Mixed control strategy

- Choose u_1 with probability p_1
- Choose u_2 with probability p_2
- Expected cost : $p_1 c_1 + p_2 c_2$
- Risk: $p_1 r_1 + p_2 r_2$

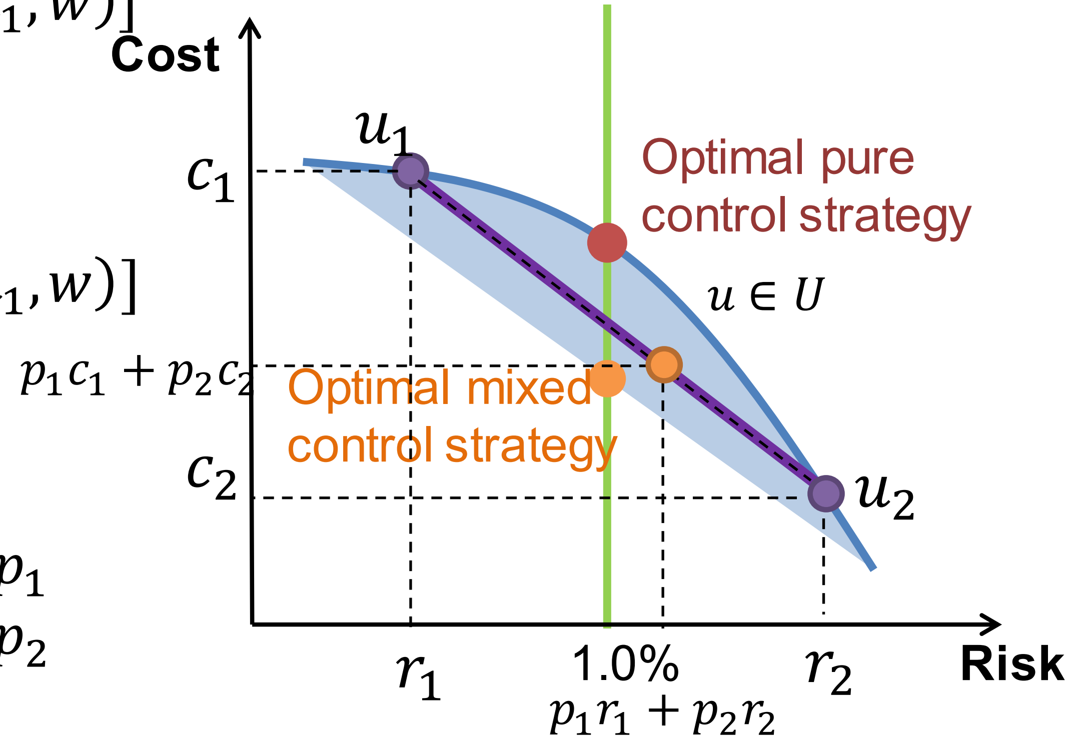

Cost and Risk of Mixed Control Strategy

Theorem 1

The set of the cost and risk of all mixed control strategies is the convex hull of that of all pure control strategies.

Remark

A mixed control strategy can outperform pure control strategies if the optimization problem is nonconvex

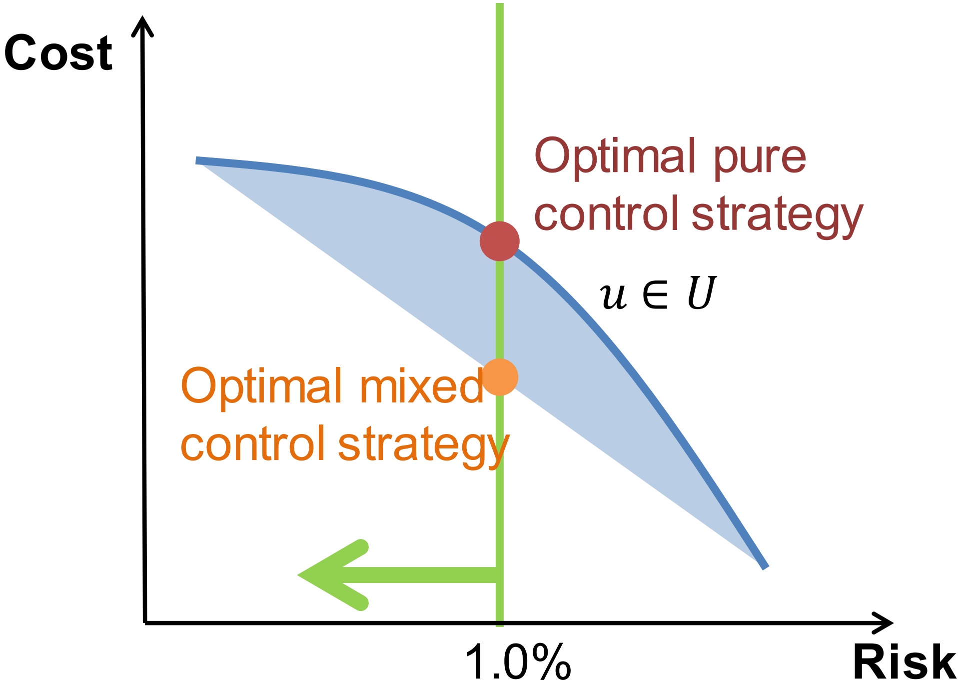

Related Work

- MDP typically considers mixed strategy
- In Optimization
  - Vajda and Greenberg
  - Mukherjee, 1980: Considered a special case with two dicision variables
- No prior work in the domain of chance-constrained optimal control

Strong Duality of Mixed Strategy CCOC

Theorem 2

Mixed-strategy CCOC problem does not have a duality gap

Theorem 3

The optimal mixed control strategy consists of up to two pure control strategies

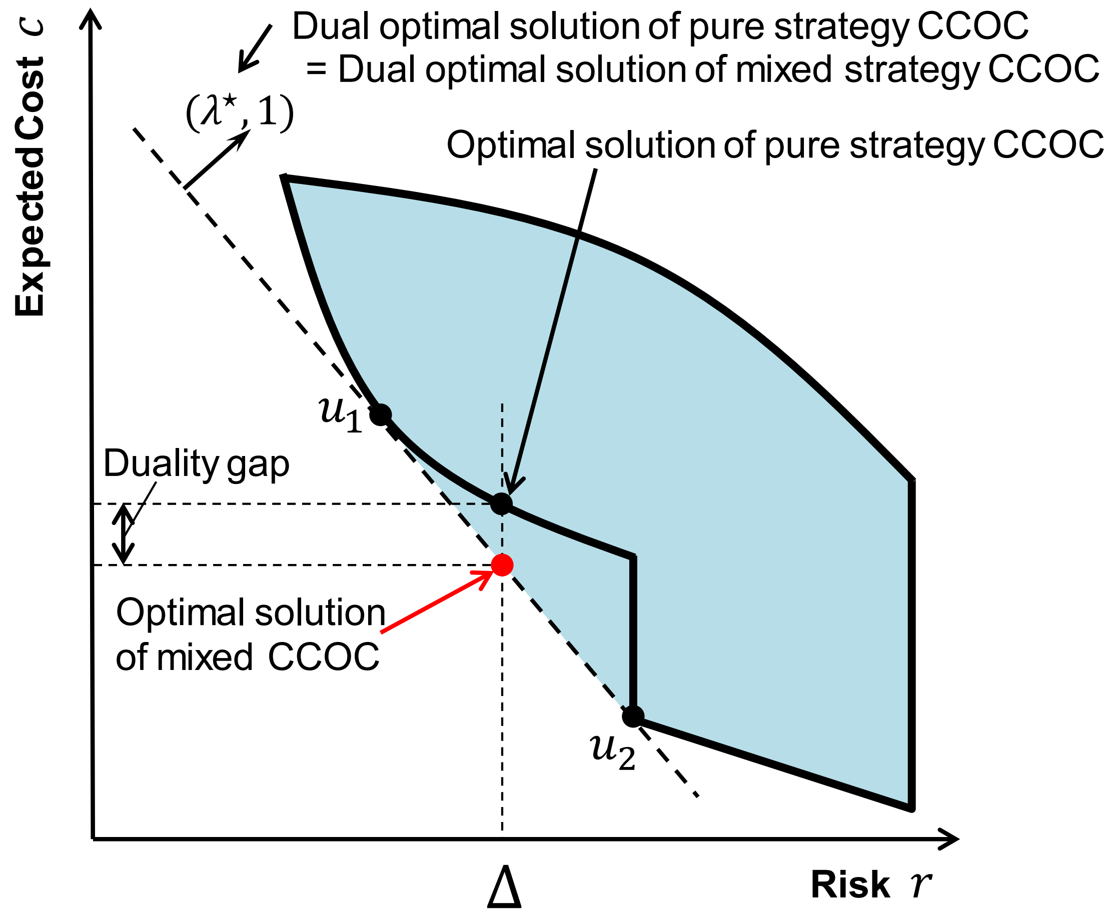

How to Compute an Optimal Mixed Control Strategy?

Idea:

1. Solve the dual of pure strategy CCOC
2. Obtain u_1 and u_2
3. $p_1 = \frac{r_2 - \Delta}{r_2 - r_1}$, $p_2 = \frac{\Delta - r_1}{r_2 - r_1}$
4. **Optimal mixed control strategy**: choose u_1 with the probability of p_1, choose u_2 with the probability of p_2

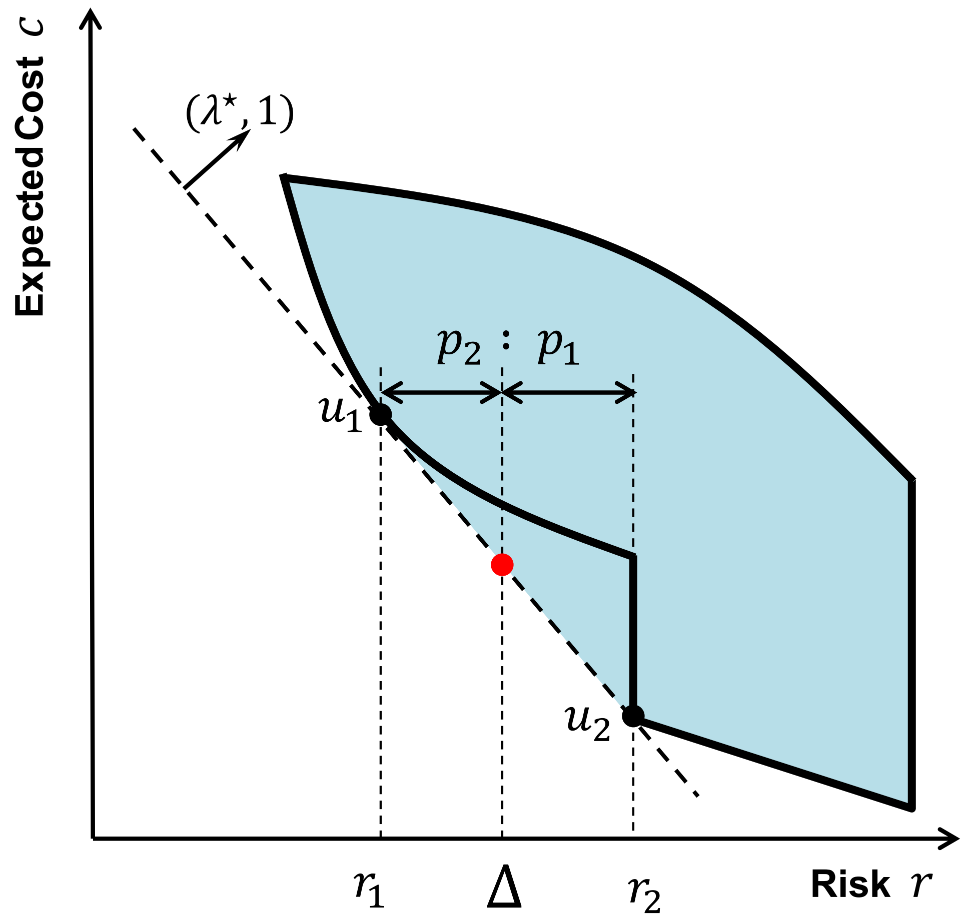

How to Compute an Optimal Mixed Control Strategy?

Idea:

1. **Solve the dual of pure strategy CCOC**
2. Obtain u_1 and u_2
3. $p_1 = \frac{r_2 - \Delta}{r_2 - r_1}$, $p_2 = \frac{\Delta - r_1}{r_2 - r_1}$
4. **Optimal mixed control strategy**: choose u_1 with the probability of p_1, choose u_2 with the probability of p_2

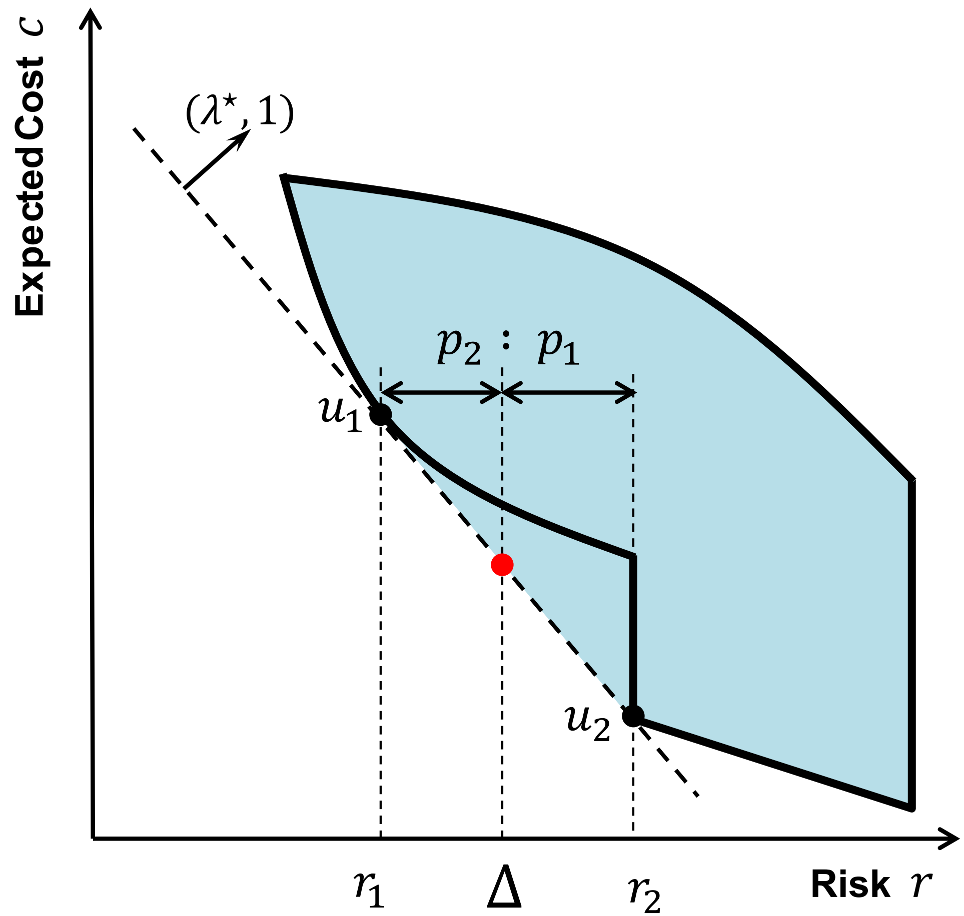

Solving the Dual of CCOC

Dual Problem

$$\max_{\lambda} \min_{u \in \mathcal{U}} \mathbb{E}[f(u,w)] + \lambda(\Pr[g(u,w) \succ 0] - \Delta)$$

$$\|$$

$$q(\lambda)$$

Primal Problem

$$\min_{u \in \mathcal{U}} \mathbb{E}[f(u,w)]$$

s.t. $\Pr[g(u,w) \succ 0] \leq \Delta$

Solving the Dual of CCOC

Dual Problem

$$\max_{\lambda} \min_{u \in \mathcal{U}} \mathbb{E}[f(u,w)] + \lambda(\Pr[g(u,w) \succ 0] - \Delta)$$

$$||$$

$$q(\lambda)$$

- $q(\lambda)$ is always concave
- Optimality condition: $0 \in \partial q(\lambda)$
  - $\partial q(\lambda)$: subgradient
- $\Pr[g(u,w) \succ 0] - \Delta \in \partial q(\lambda)$
- Therefore, dual optimization is reduce to a zero-finding problem over $\partial q(\lambda)$

Solving the Dual of CCOC

Dual Problem

$$\max_{\lambda} \min_{u \in \mathcal{U}} \mathbb{E}[f(u,w)] + \lambda(\Pr[g(u,w) \succ 0] - \Delta)$$

$$||$$

$$q(\lambda)$$

- $q(\lambda)$ is always concave
- Optimality condition: $0 \in \partial q(\lambda)$
  - $\partial q(\lambda)$: subgradient
- $\Pr[g(u,w) \succ 0] - \Delta \in \partial q(\lambda)$
- Therefore, dual optimization is reduce to a zero-finding problem over $\partial q(\lambda)$

Solving the Dual of CCOC

Dual Problem

$$\max_{\lambda} \min_{u \in \mathcal{U}} \mathbb{E}[f(u,w)] + \lambda(\Pr[g(u,w) \succ 0] - \Delta)$$

$$||$$

$$q(\lambda)$$

- For a given λ, solve the unconstrained optimal control problem to find the optimal control $u^*(\lambda)$
- Evaluate $\Pr[g(u^*(\lambda), w) \succ 0] - \Delta$
- If it is not zero, adjust λ.

Solving the Dual of CCOC

Dual Problem

$$\max_{\lambda} \underbrace{\min_{u \in \mathcal{U}} \mathbb{E}[f(u,w)] + \lambda(\Pr[g(u,w) \succ 0] - \Delta)}_{=\, q(\lambda)}$$

- Intuition
  - λ is a sensitivity to risk
  - Smaller λ : greater $\Pr[g(u^*(\lambda),w) \succ 0]$
  - Greater λ : smaller $\Pr[g(u^*(\lambda),w) \succ 0]$
- Dual optimization = find λ that results in exactly Δ of risk.

CCDP: Dual Reformulation

Dual Problem

$$\max_{\lambda} \; q(\lambda)$$

$$\Pr[g(u^*(\lambda), w) \succ 0] - \Delta$$

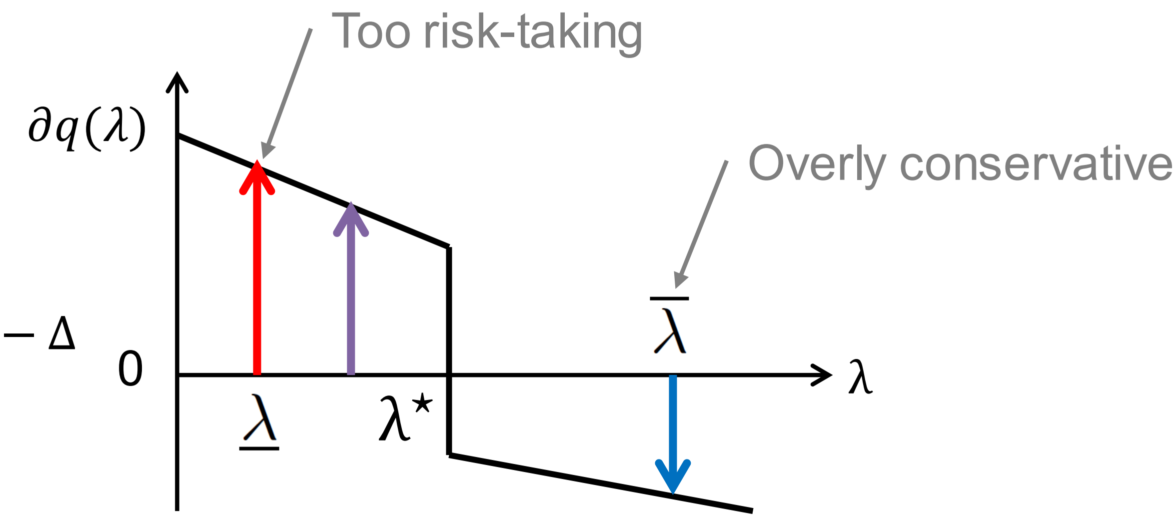

Dual Problem

$$\max_{\lambda} \; q(\lambda)$$

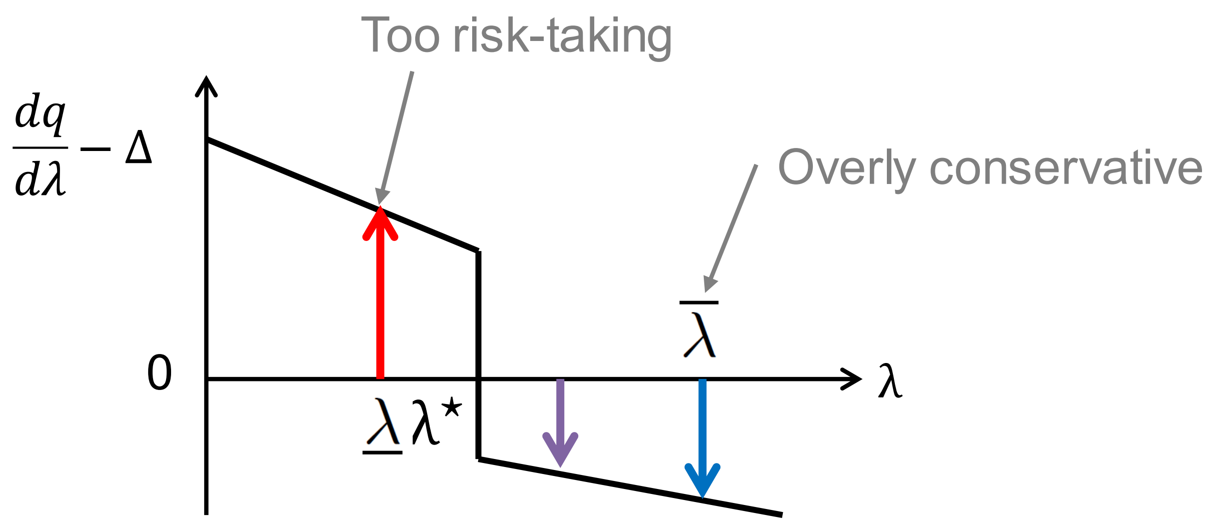

CCDP: Problem Formulation

Dual Problem

$$\max_{\lambda} \; q(\lambda)$$

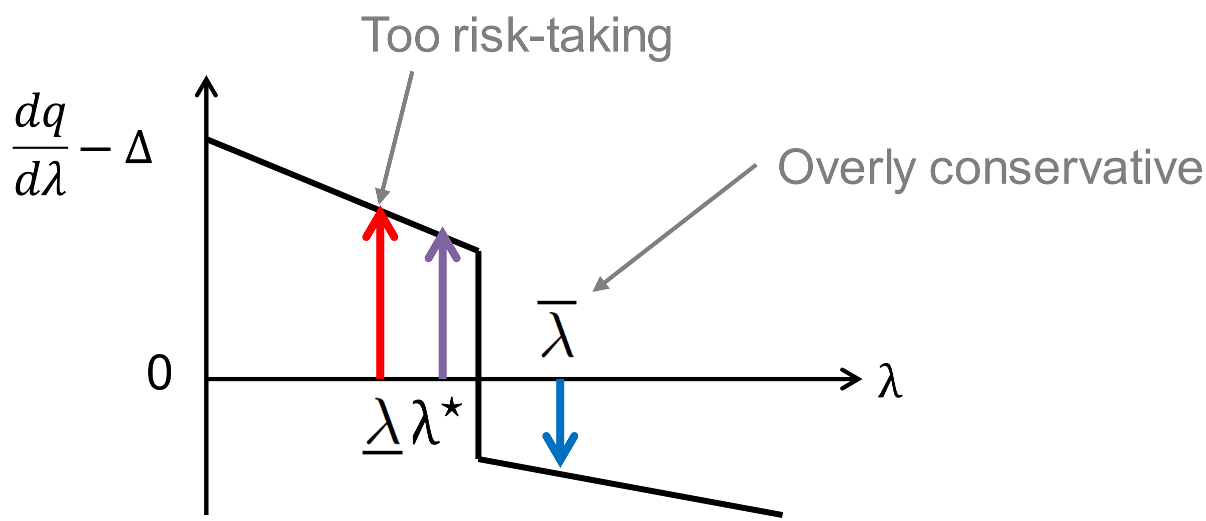

CCDP: Problem Formulation

Dual Problem

$$\max_{\lambda} \quad q(\lambda)$$

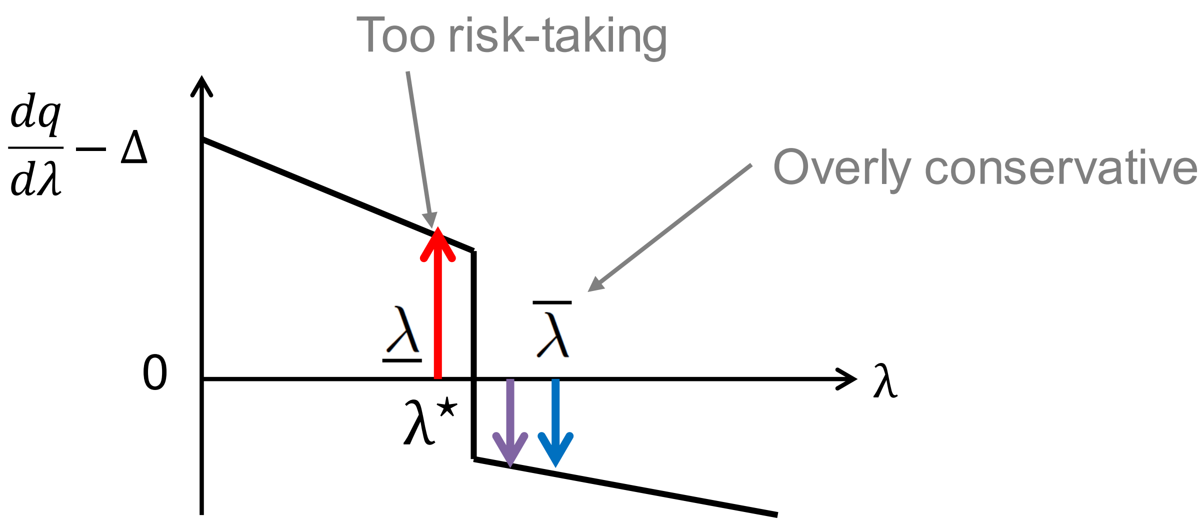

CCDP: Problem Formulation

Dual Problem

$$\max_{\lambda} \; q(\lambda)$$

Interpretation: dual variable = *penalty*

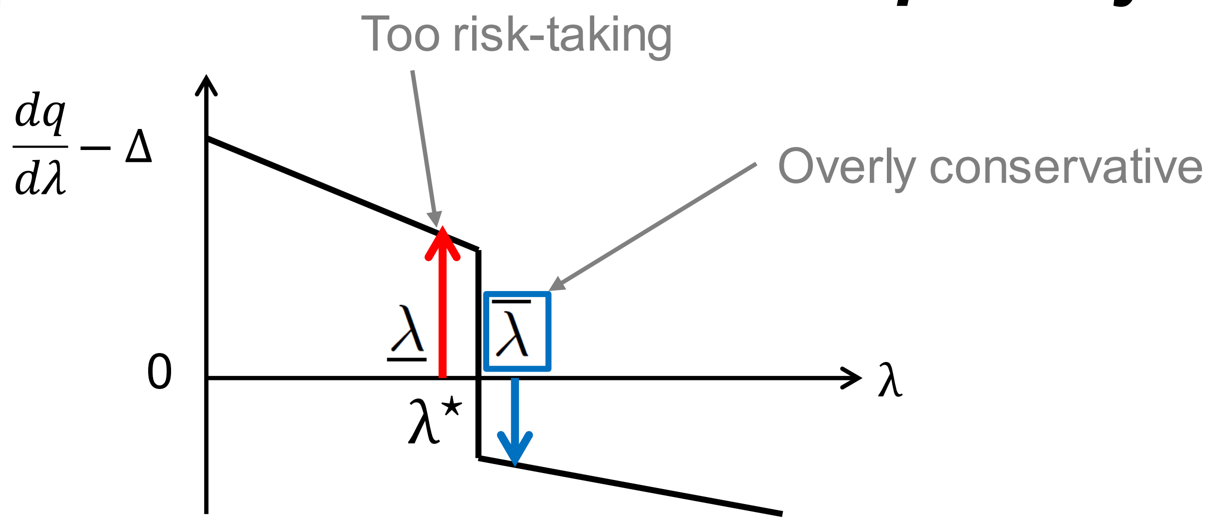

How to Compute an Optimal Mixed Control Strategy?

Idea:

1. Solve the dual of pure strategy CCOC
2. **Obtain u_1 and u_2**
3. $p_1 = \frac{r_2 - \Delta}{r_2 - r_1}$, $p_2 = \frac{\Delta - r_1}{r_2 - r_1}$
4. **Optimal mixed control strategy**: choose u_1 with the probability of p_1, choose u_2 with the probability of p_2

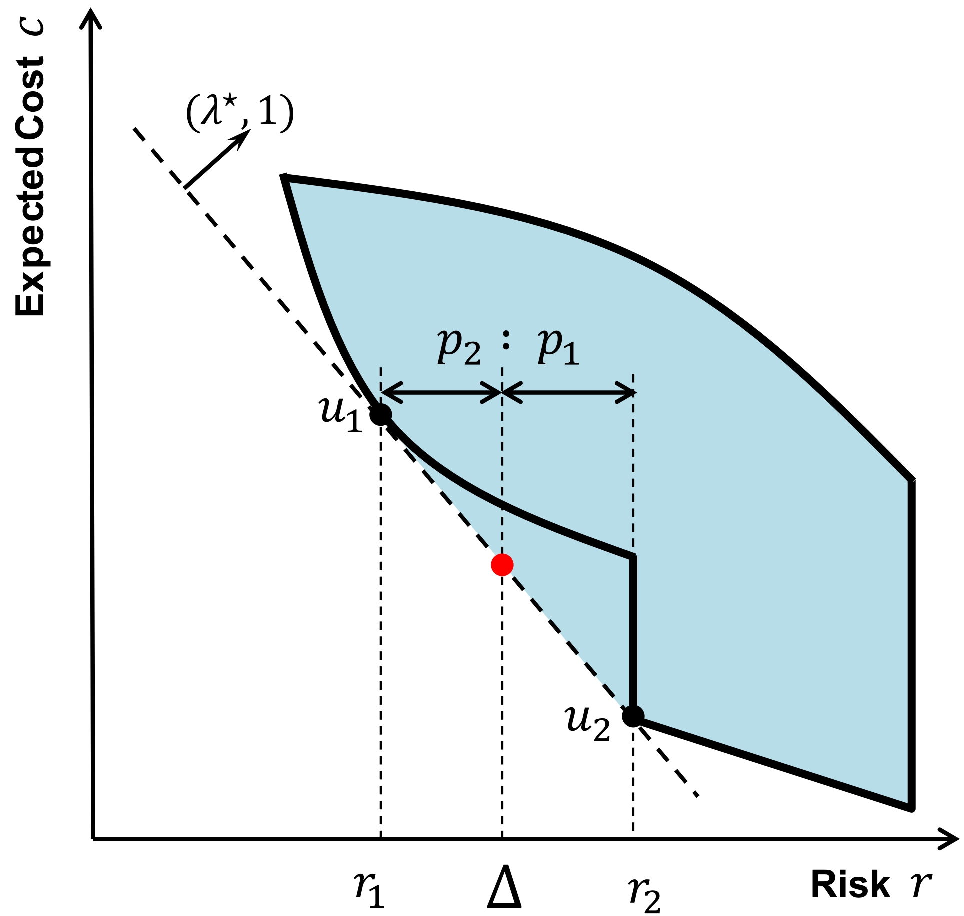

CCDP: Problem Formulation

Dual Problem

$$\max_{\lambda} \; q(\lambda)$$

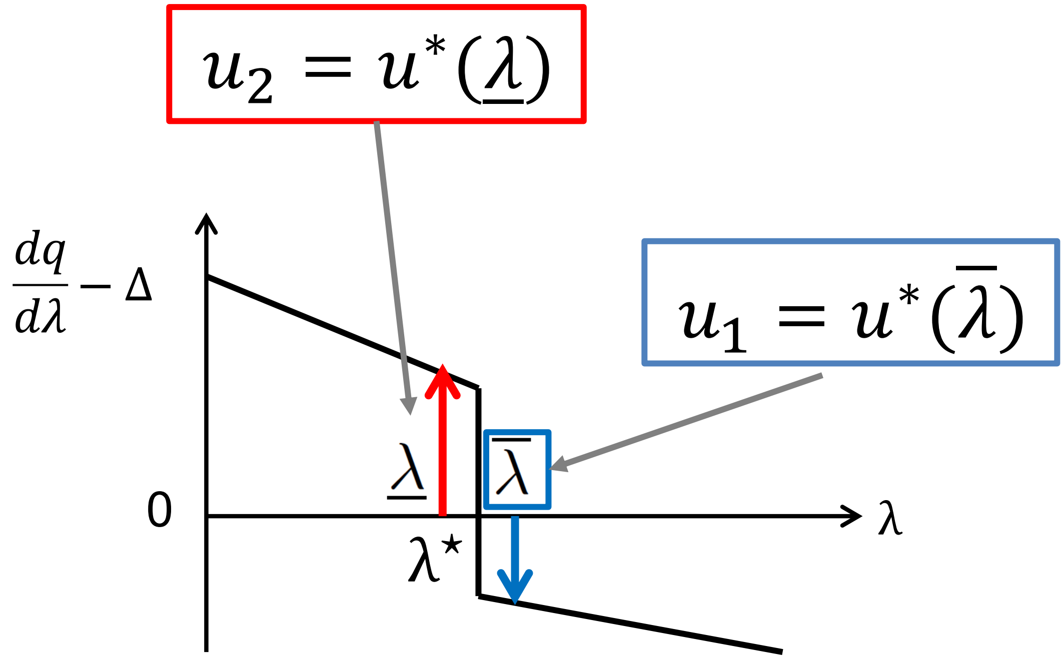

How to Compute an Optimal Mixed Control Strategy?

Idea:

1. Solve the dual of pure strategy CCOC
2. Obtain u_1 and u_2
3. $p_1 = \frac{r_2 - \Delta}{r_2 - r_1}$, $p_2 = \frac{\Delta - r_1}{r_2 - r_1}$
4. **Optimal mixed control strategy**: choose u_1 with the probability of p_1, choose u_2 with the probability of p_2

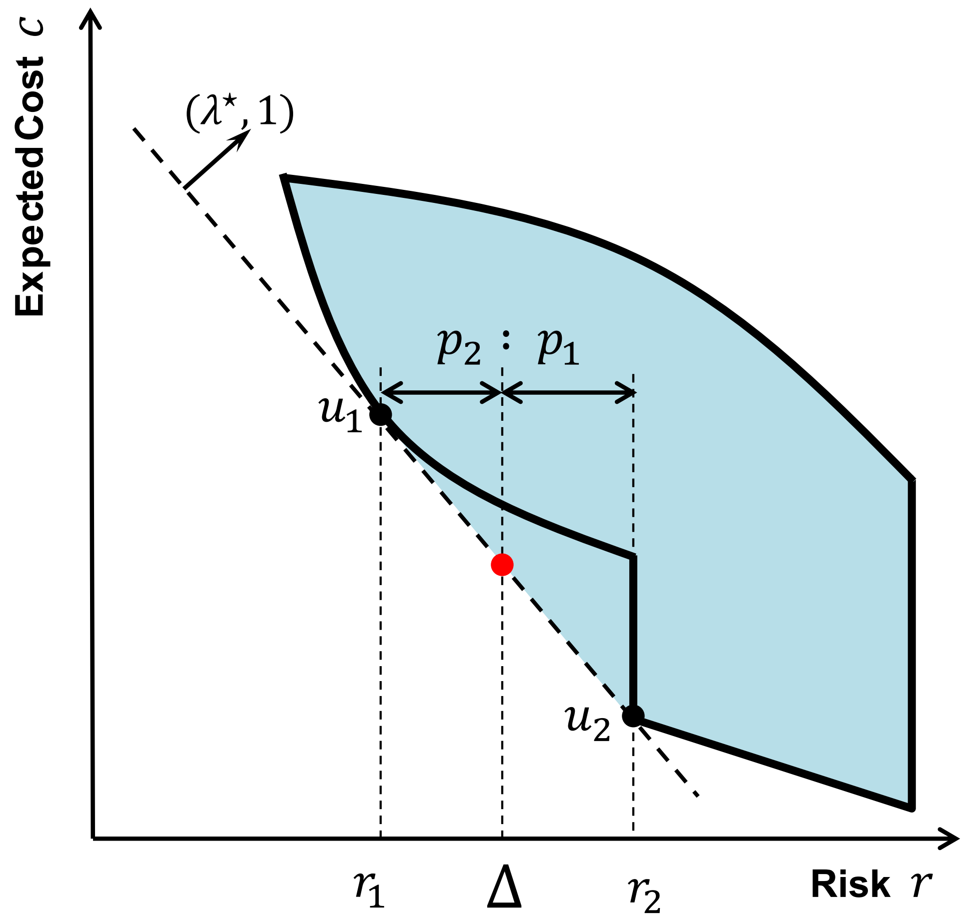

Implementation

- Can be used with any chance-constrained optimal control problem, as long as its dual can be solved
- Implemented on chance-constrained dynamic programming (CCDP)*

*Ono, Kuwata, Baralam, "Joint chance-constrained dynamic programming," CDC-12

Application to Path Planning

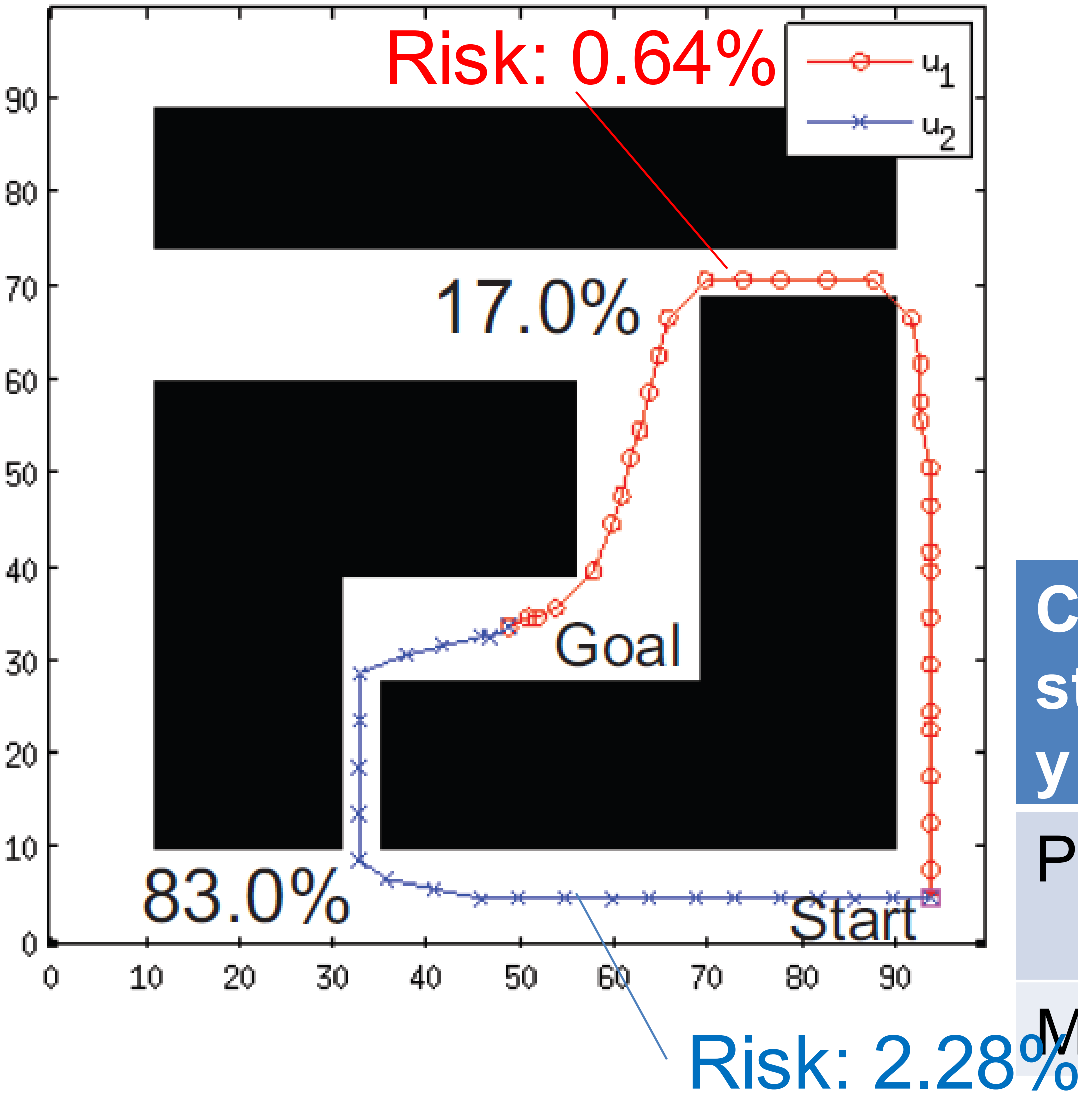

Problem:
- Minimize path length
- Risk $\leq$ 2%

- $x_{k+1} = x_k + u_k + w_k$
- 100x100 discrete state space

Control strategy	Expected path length	Risk
Pure	130.8	0.64 %
Mixed	104.2	2.0%

Application to Mars Entry, Descent, and Landing

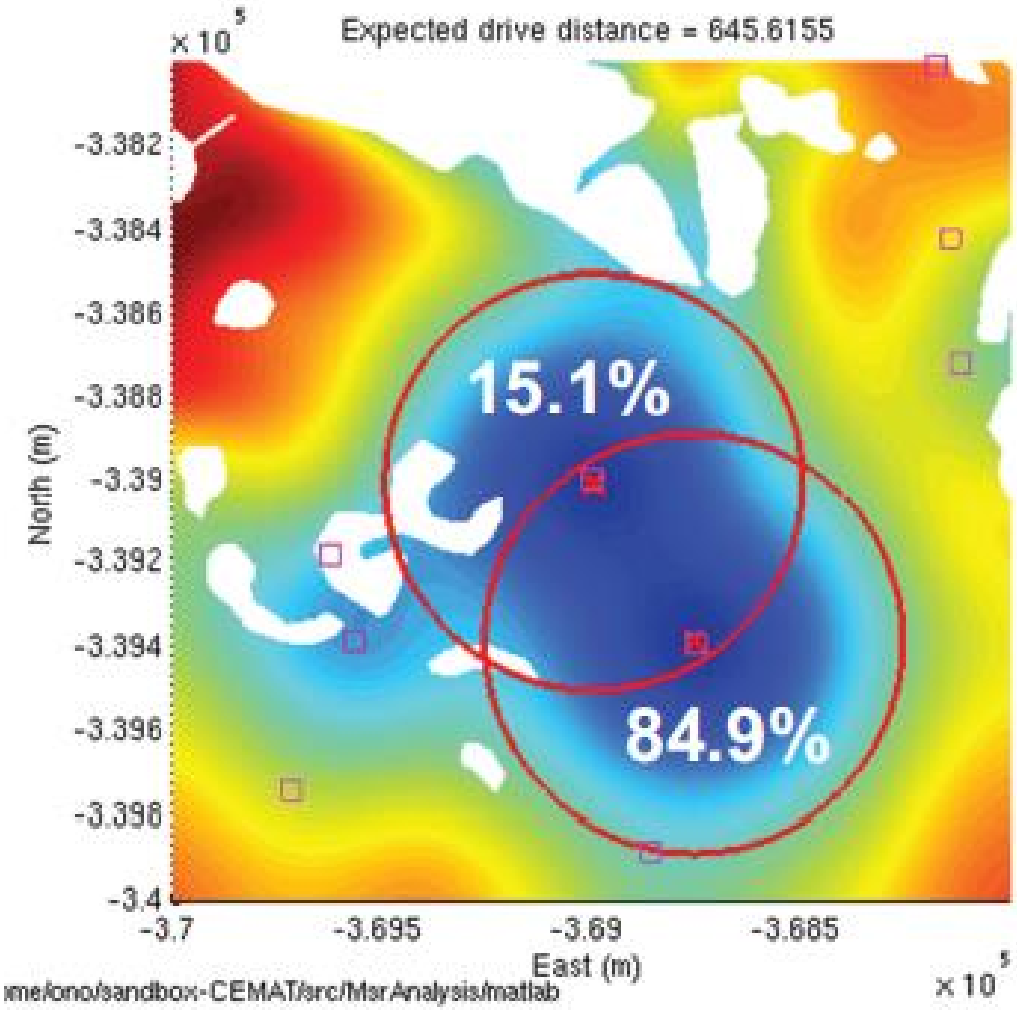

Problem:

- Minimize driving distance to visit two science targets after landing
- Risk ≤ 0.1%
- 2000x2000 discrete state space
- Used terrain data at E. Margaritifer on Mars

Control strategy	Expected cost	Risk
Pure	645.49	0.016 %
Mixed	644.81	0.1%

Conclusions

- Characterized a mixed-strategy chance-constrained optimal control problem using the MC/MC framework
  - It is a convexification of a pure-strategy chance-constrained optimal control problem
  - Hence, there is no duality gap
- Developed an algorithm that obtains an optimal mixed control strategy
  - Build upon the dual solution to a pure-strategy chance-constrained optimal control problem
- Demonstrated the proposed algorithm on path planning and Mars EDL scenarios