A Risk-Theoretical Approach to $\mathcal{H}_2$-Optimal Control under Covert Attacks

Matias I. Müller,
Jezdimir Milošević, Henrik Sandberg and Cristian R. Rojas
(e-mails: `{mimr2,jezdimir,hsan,crro}@kth.se`)

Department of Automatic Control,
KTH Royal Institute of Technology,
Stockholm, Sweden

57th IEEE Conference on Decision and Control
December 18th, 2018

Introduction: Control Systems are Vulnerable

Introduction: Control Systems are Vulnerable

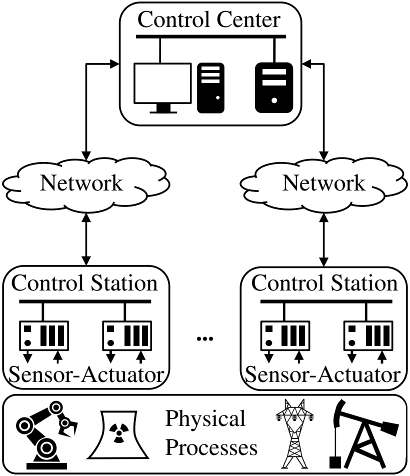

Introduction: Control Systems are Vulnerable

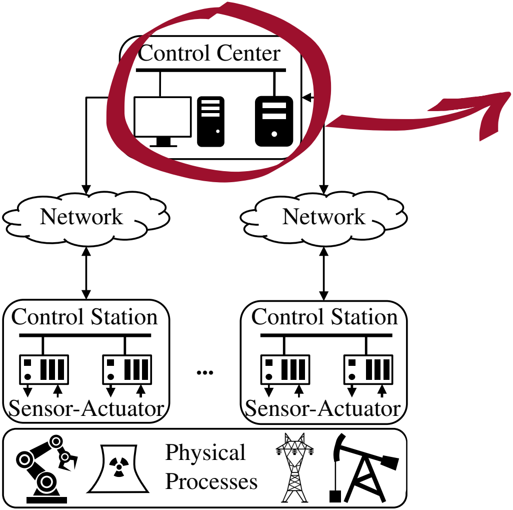

Inadequate protection against malware

Introduction: Control Systems are Vulnerable

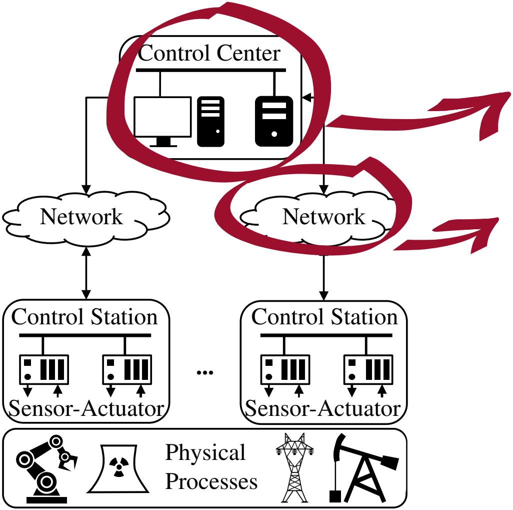

Inadequate protection against malware

Communication links are not encrypted

Control System under Attack

Control System under Attack

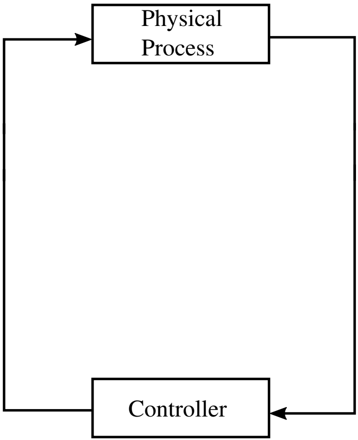

Control System under Attack

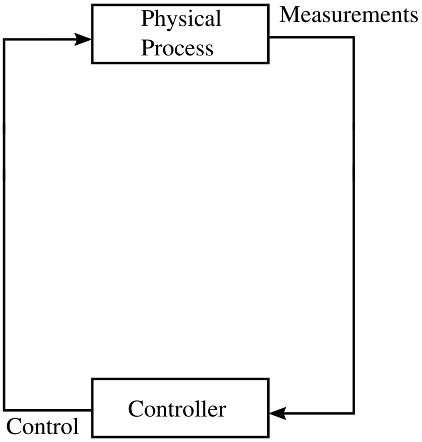

Control System under Attack

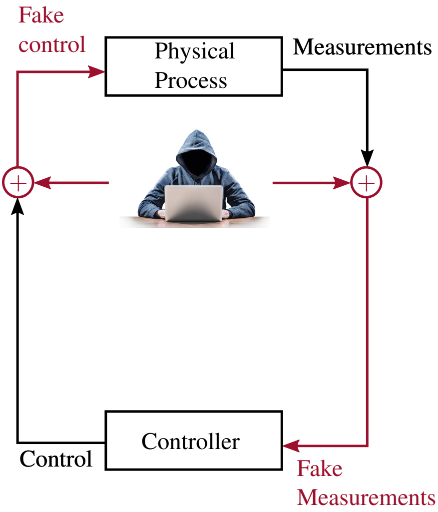

Control System under Attack

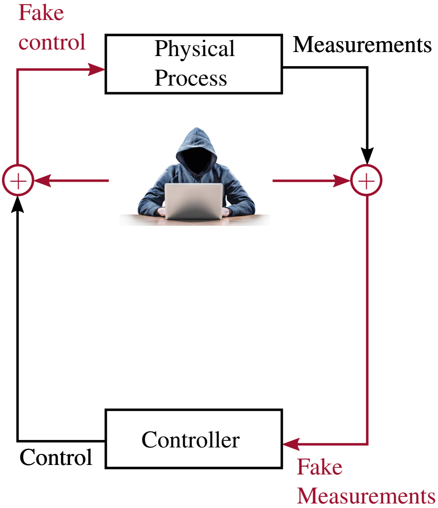

- The attacker corrupts measurements/control signals.

Control System under Attack

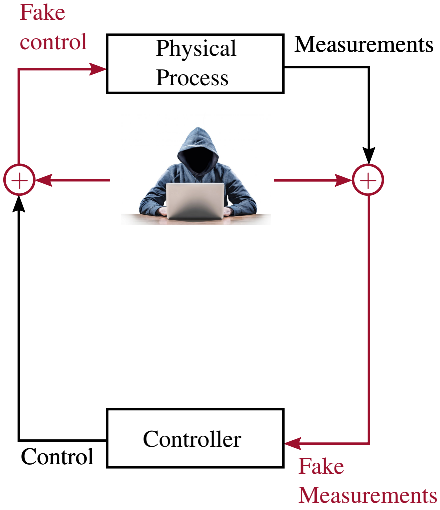

- The attacker corrupts measurements/control signals.
- It tries to remain stealthy.

Control System under Attack

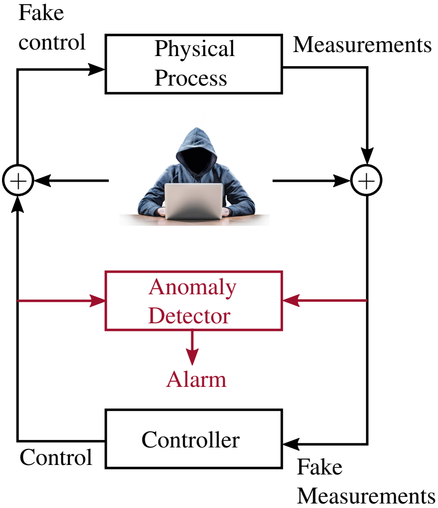

- The attacker corrupts measurements/control signals.
- It tries to remain stealthy.

✓ Presence of attack is revealed through alarm triggering.

Control System under Attack

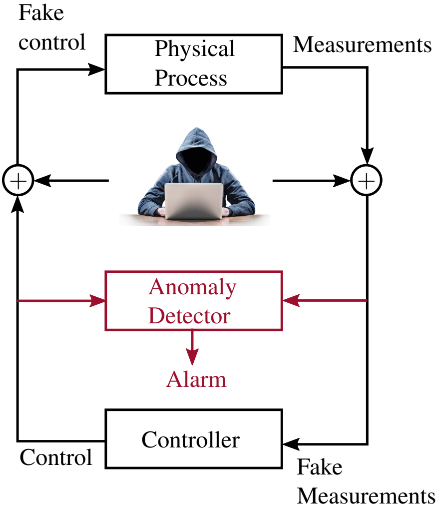

- The attacker corrupts measurements/control signals.
- It tries to remain stealthy.
- ✓ Presence of attack is revealed through alarm triggering.

Our concern: Control design under attack

Undetectable Attacks

R. S. Smith, "Covert misappropriation of networked control systems: Presenting a feedback structure," IEEE Control Systems, vol. 35, no. 1, pp. 82-92, Feb 2015.

$$\mu = (I - \Theta_\gamma \Pi_u)^{-1} \Theta_{\text{ref}} \gamma_{\text{ref}}, \tag{9}$$

$$\gamma = (I - \Theta_\gamma \Pi_u)^{-1} \Pi_u \Theta_{\text{ref}} \gamma_{\text{ref}}. \tag{10}$$

The covert controller is designed for reference tracking, with γ_{ref} as the reference input.

To see the consequences of the nested closed-loop systems in Figure 2, replace (3) by (5) and (2) by (6) and rearrange to get

$$y - y_m = (I - \Theta_\gamma \Pi_u)^{-1} \Theta_{\text{ref}} \Pi_u \gamma_{\text{ref}} - n \tag{11}$$

and

$$y_m = S P_u C_{\text{ref}} y_{\text{ref}} + S P_w w + S n + S(P_u - \Pi_u)\mu. \tag{12}$$

Note that in (11), the relative offset of the actual plant output y from that measured—and controlled—by the nominal controller y_m is the output of the covert agent's γ_{ref}-reference tracking controller. This gives the covert agent the ability to drive the actual plant output to a desired offset with respect to its nominal controlled value y_{ref}.

To examine the nominal controller's ability to detect the actions of the covert agent, the nominal controller's measurement y_m in the nominal case [given by (4)] is compared to the covert misappropriation strategy case [given by (12)]. The only difference between these two appears to the nominal controller as an output disturbance w_{covert}, given by

$$w_{\text{covert}} = S(P_u - \Pi_u)(I - \Theta_\gamma \Pi_u)^{-1} \Theta_{\text{ref}} \gamma_{\text{ref}}. \tag{13}$$

If the covert agent has perfect knowledge of the plant's input response, $\Pi_u = P_u$, then $w_{\text{covert}} = 0$ and the covert misappropriation is undetectable. This case was studied in [11] as a particular case of a slightly more general parameterization of the covert agent. It is important to note that the covert agent needs no knowledge of the nominal controller to execute an undetectable misappropriation strategy.

Specifying the covert agent's actions via the feedback structure in Figure 2 ensures that the covert controller's plant input signal μ is appropriate for the plant. Any feedback or actuation limitations imposed by the plant are taken into account in the design of the $\Pi_u - \Theta_\gamma$ feedback loop within the covert agent and will limit the range of γ_{ref} offset values that the covert controller can effectively command. These limitations make no difference to the extent to which the convert controller's actions can be detected.

It is more realistic to consider that the covert agent's knowledge of the plant is not perfect. In this case, define the covert agent's model error Δ via

covert agent's point of view, band-limiting the frequency content of γ_{ref} to those frequencies where the network control system operates well will make the covert actions harder to detect.

2) The size of the covert agent's model error Δ. The higher the quality of the covert agent's knowledge of the plant, the harder it will be to detect covert actions.
3) The covert agent's reference to actuation transfer function $(I - \Theta_\gamma \Pi_u)^{-1} \Theta_{\text{ref}}$. This is a function of the design of the covert agent and can be used to further hide the covert action. For example, by designing the bandwidth of $(I - \Theta_\gamma \Pi_u)^{-1} \Theta_{\text{ref}}$ to be lower than that of T the frequency components of μ will be in the range where S is small, reducing the size of w_{covert}.
4) The size of the covert offset command γ_{ref}.

Even if the covert agent's knowledge of the plant is not perfect, the nominal controller still sees, and responds to, the actual measurement noise and the actual plant disturbances. Furthermore, the dynamics of the controlled plant appear unchanged from the nominal case. The effect on the measured plant output y_m of any nominal controller control signal u_c is the same whether or not the covert controller is operating. These features hinder the nominal controller's ability to detect covert actions through probing signals, such as watermarks, or signal analysis, such as noise or disturbance statistics characterization.

A LINEAR FLOW CANAL CONTROL EXAMPLE

Nominal Model and Operation

To illustrate the action of the covert agent with an incorrect plant model, an irrigation canal example, originally described in [12], is studied. The geographical separation in this application explains the need for a networked control system. The security of similar applications has been studied in [4] and [5]. The irrigation system is illustrated in Figure 3. A reservoir at a fixed height feeds a flow canal through a controlled sluice gate. The outlet flow of the reservoir is proportional to the gate height u_1. The water flows through a narrow sloping canal to a second sluice gate with controlled height u_2 and from there into a second canal. The second canal ends in a spillway. The water heights at the ends of each canal are the measured variables and the outputs of interest in the control problem.

This system can be modeled by two partial differential equations, known as the Saint-Venant equations. The simplified model used here can be found in [12]

Undetectable Attacks

R. S. Smith, "Covert misappropriation of networked control systems: Presenting a feedback structure," IEEE Control Systems, vol. 35, no. 1, pp. 82-92, Feb 2015.

$$\mu = (I - \Theta_\gamma \Pi_u)^{-1} \Theta_{\text{ref}} \gamma_{\text{ref}}, \quad (9)$$

$$\gamma = (I - \Theta_\gamma \Pi_u)^{-1} \Pi_u \Theta_{\text{ref}} \gamma_{\text{ref}}. \quad (10)$$

The covert controller is designed for reference tracking, with γ_{ref} as the reference input.

To see the consequences of the nested closed-loop systems in Figure 2, replace (3) by (5) and (2) by (6) and rearrange to get

$$y - y_m = (I - \Theta_\gamma \Pi_u)^{-1} \Theta_{\text{ref}} \Pi_u \gamma_{\text{ref}} - n \quad (11)$$

and

$$y_m = SP_u C_{\text{ref}} y_{\text{ref}} + SP_w w + Sn + S(P_u - \Pi_u)\mu. \quad (12)$$

Note that in (11), the relative offset of the actual plant output y from that measured—and controlled—by the nominal controller y_m is the output of the covert agent's γ_{ref}-reference tracking controller. This gives the covert agent the ability to drive the actual plant output to a desired offset with respect to its nominal controlled value y_{ref}.

To examine the nominal controller's ability to detect the actions of the covert agent, the nominal controller's measurement y_m in the nominal case [given by (4)] is compared to the covert misappropriation strategy case [given by (12)]. The only difference between these two appears to the nominal controller as an output disturbance w_{covert}, given by

$$w_{\text{covert}} = S(P_u - \Pi_u)(I - \Theta_\gamma \Pi_u)^{-1} \Theta_{\text{ref}} \gamma_{\text{ref}}. \quad (13)$$

If the covert agent has perfect knowledge of the plant's input response, $\Pi_u = P_u$, then $w_{\text{covert}} = 0$ and the covert misappropriation is undetectable. This case was studied in [11] as a particular case of a slightly more general parameterization of the covert agent. It is important to note that the covert agent needs no knowledge of the nominal controller to execute an undetectable misappropriation strategy.

Specifying the covert agent's actions via the feedback structure in Figure 2 ensures that the covert controller's plant input signal μ is appropriate for the plant. Any feedback or actuation limitations imposed by the plant are taken into account in the design of the $\Pi_u - \Theta_\gamma$ feedback loop within the covert agent and will limit the range of γ_{ref} offset values that the covert controller can effectively command. These limitations make no difference to the extent to which the convert controller's actions can be detected.

It is more realistic to consider that the covert agent's knowledge of the plant is not perfect. In this case, define the covert agent's model error Δ via

covert agent's point of view, band-limiting the frequency content of γ_{ref} to those frequencies where the network control system operates well will make the covert actions harder to detect.

2) The size of the covert agent's model error Δ. The higher the quality of the covert agent's knowledge of the plant, the harder it will be to detect covert actions.
3) The covert agent's reference to actuation transfer function $(I - \Theta_\gamma \Pi_u)^{-1} \Theta_{\text{ref}}$. This is a function of the design of the covert agent and can be used to further hide the covert action. For example, by designing the bandwidth of $(I - \Theta_\gamma \Pi_u)^{-1} \Theta_{\text{ref}}$ to be lower than that of T the frequency components of μ will be in the range where S is small, reducing the size of w_{covert}.
4) The size of the covert offset command γ_{ref}.

Even if the covert agent's knowledge of the plant is not perfect, the nominal controller still sees, and responds to, the actual measurement noise and the actual plant disturbances. Furthermore, the dynamics of the controlled plant appear unchanged from the nominal case. The effect on the measured plant output y_m of any nominal controller control signal u_c is the same whether or not the covert controller is operating. These features hinder the nominal controller's ability to detect covert actions through probing signals, such as watermarks, or signal analysis, such as noise or disturbance statistics characterization.

A LINEAR FLOW CANAL CONTROL EXAMPLE

Nominal Model and Operation

To illustrate the action of the covert agent with an incorrect plant model, an irrigation canal example, originally described in [12], is studied. The geographical separation in this application explains the need for a networked control system. The security of similar applications has been studied in [4] and [5]. The irrigation system is illustrated in Figure 3. A reservoir at a fixed height feeds a flow canal through a controlled sluice gate. The outlet flow of the reservoir is proportional to the gate height u_1. The water flows through a narrow sloping canal to a second sluice gate with controlled height u_2 and from there into a second canal. The second canal ends in a spillway. The water heights at the ends of each canal are the measured variables and the outputs of interest in the control problem.

This system can be modeled by two partial differential equations, known as the Saint-Venant equations. The simplified model used here can be found in [12]

FEBRUARY 2015 « IEEE CONTROL SYSTEMS MAGAZINE 85

"If the covert agent has perfect knowledge of the plant's input response then the covert misappropriation is undetectable [to the controller]."

Undetectable Attacks

R. S. Smith, "Covert misappropriation of networked control systems: Presenting a feedback structure," IEEE Control Systems, vol. 35, no. 1, pp. 82-92, Feb 2015.

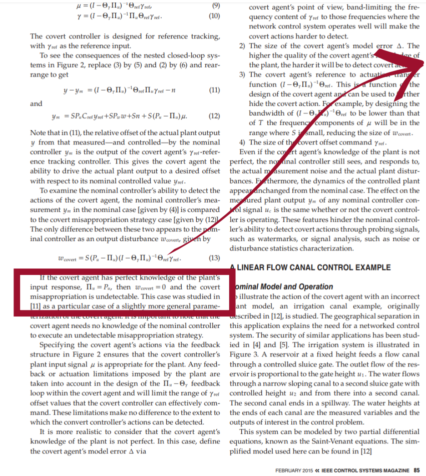

$$\mu = (I-\Theta_\gamma\Pi_u)^{-1}\Theta_{\text{ref}}\gamma_{\text{ref}}, \quad (9)$$
$$\gamma = (I-\Theta_\gamma\Pi_u)^{-1}\Pi_u\Theta_{\text{ref}}\gamma_{\text{ref}}. \quad (10)$$

The covert controller is designed for reference tracking, with γ_{ref} as the reference input.

To see the consequences of the nested closed-loop systems in Figure 2, replace (3) by (5) and (2) by (6) and rearrange to get

$$y - y_m = (I-\Theta_\gamma\Pi_u)^{-1}\Theta_{\text{ref}}\Pi_u\gamma_{\text{ref}} - n \quad (11)$$

and

$$y_m = SP_uC_{\text{ref}}y_{\text{ref}} + SP_w w + Sn + S(P_u - \Pi_u)\mu. \quad (12)$$

Note that in (11), the relative offset of the actual plant output y from that measured—and controlled—by the nominal controller y_m is the output of the covert agent's γ_{ref}-reference tracking controller. This gives the covert agent the ability to drive the actual plant output to a desired offset with respect to its nominal controlled value y_{ref}.

To examine the nominal controller's ability to detect the actions of the covert agent, the nominal controller's measurement y_m in the nominal case [given by (4)] is compared to the covert misappropriation strategy case [given by (12)]. The only difference between these two appears to the nominal controller as an output disturbance w_{covert}, given by

$$w_{\text{covert}} = S(P_u - \Pi_u)(I-\Theta_\gamma\Pi_u)^{-1}\Theta_{\text{ref}}\gamma_{\text{ref}}. \quad (13)$$

If the covert agent has perfect knowledge of the plant's input response, $\Pi_u = P_u$, then $w_{\text{covert}} = 0$ and the covert misappropriation is undetectable. This case was studied in [11] as a particular case of a slightly more general parameterization of the covert agent. It is important to note that the covert agent needs no knowledge of the nominal controller to execute an undetectable misappropriation strategy.

Specifying the covert agent's actions via the feedback structure in Figure 2 ensures that the covert controller's plant input signal μ is appropriate for the plant. Any feedback or actuation limitations imposed by the plant are taken into account in the design of the $\Pi_u - \Theta_\gamma$ feedback loop within the covert agent and will limit the range of γ_{ref} offset values that the covert controller can effectively command. These limitations make no difference to the extent to which the convert controller's actions can be detected.

It is more realistic to consider that the covert agent's knowledge of the plant is not perfect. In this case, define the covert agent's model error Δ via

covert agent's point of view, band-limiting the frequency content of γ_{ref} to those frequencies where the network control system operates well will make the covert actions harder to detect.

2) The size of the covert agent's model error Δ. The higher the quality of the covert agent's knowledge of the plant, the harder it will be to detect covert actions.
3) The covert agent's reference to actuation transfer function $(I-\Theta_\gamma\Pi_u)^{-1}\Theta_{\text{ref}}$. This is a function of the design of the covert agent and can be used to further hide the covert action. For example, by designing the bandwidth of $(I-\Theta_\gamma\Pi_u)^{-1}\Theta_{\text{ref}}$ to be lower than that of T the frequency components of μ will be in the range where S is small, reducing the size of w_{covert}.
4) The size of the covert offset command γ_{ref}.

Even if the covert agent's knowledge of the plant is not perfect, the nominal controller still sees, and responds to, the actual measurement noise and the actual plant disturbances. Furthermore, the dynamics of the controlled plant appear unchanged from the nominal case. The effect on the measured plant output y_m of any nominal controller control signal u_c is the same whether or not the covert controller is operating. These features hinder the nominal controller's ability to detect covert actions through probing signals, such as watermarks, or signal analysis, such as noise or disturbance statistics characterization.

A LINEAR FLOW CANAL CONTROL EXAMPLE

Nominal Model and Operation

To illustrate the action of the covert agent with an incorrect plant model, an irrigation canal example, originally described in [12], is studied. The geographical separation in this application explains the need for a networked control system. The security of similar applications has been studied in [4] and [5]. The irrigation system is illustrated in Figure 3. A reservoir at a fixed height feeds a flow canal through a controlled sluice gate. The outlet flow of the reservoir is proportional to the gate height u_1. The water flows through a narrow sloping canal to a second sluice gate with controlled height u_2 and from there into a second canal. The second canal ends in a spillway. The water heights at the ends of each canal are the measured variables and the outputs of interest in the control problem.

This system can be modeled by two partial differential equations, known as the Saint-Venant equations. The simplified model used here can be found in [12]

FEBRUARY 2015 « IEEE CONTROL SYSTEMS MAGAZINE 85

"If the covert agent has perfect knowledge of the plant's input response then the covert misappropriation is undetectable [to the controller]."

The controller cannot compensate for these attacks either!

In This Work

Assuming full model knowledge is quite conservative...

In This Work

Assuming full model knowledge is quite conservative...

- Inaccurate/Outdated model

In This Work

Assuming full model knowledge is quite conservative...

- Inaccurate/Outdated model
- Fictitious uncertainty (multiplicative watermarking, Teixeira'18)

In This Work

Assuming full model knowledge is quite conservative...

- Inaccurate/Outdated model
- Fictitious uncertainty (multiplicative watermarking, Teixeira'18)

Condition: covert attack + partial knowledge of the attacker:

In This Work

Assuming full model knowledge is quite conservative...

- Inaccurate/Outdated model
- Fictitious uncertainty (multiplicative watermarking, Teixeira'18)

Condition: covert attack + partial knowledge of the attacker:

> How to design a controller that performs well in most of the feasible attacker scenarios?

In This Work

Assuming full model knowledge is quite conservative...

- Inaccurate/Outdated model
- Fictitious uncertainty (multiplicative watermarking, Teixeira'18)

Condition: covert attack + partial knowledge of the attacker:

How to design a controller that performs well
in most of the feasible attacker scenarios?

Modeling the lack of knowledge of the attacker as *uncertainty*

In This Work

Assuming full model knowledge is quite conservative...

- Inaccurate/Outdated model
- Fictitious uncertainty (multiplicative watermarking, Teixeira'18)

Condition: covert attack + partial knowledge of the attacker:

How to design a controller that performs well
in most of the feasible attacker scenarios?

Modeling the lack of knowledge of the attacker as *uncertainty*

Defender's
confidence

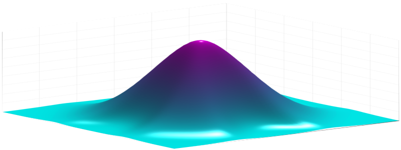

Models the attacker might potentially pocess

Model Setup

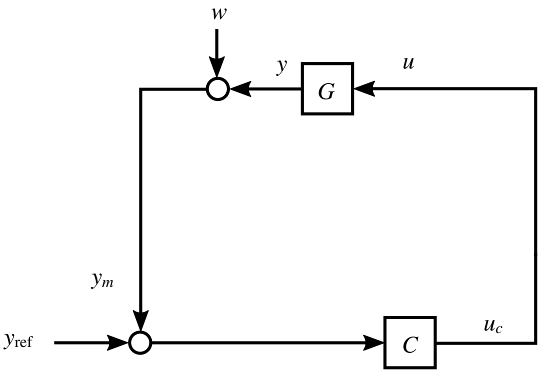

Model Setup

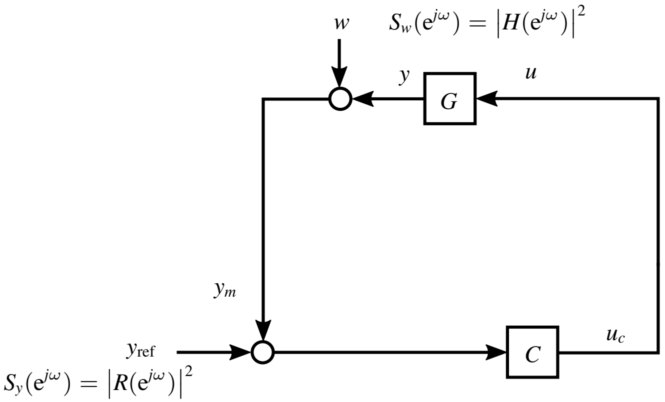

w
$S_w(e^{j\omega}) = \left|H(e^{j\omega})\right|^2$
y
G
u
y_m
y_{ref}
$S_y(e^{j\omega}) = \left|R(e^{j\omega})\right|^2$
C
u_c

Model Setup

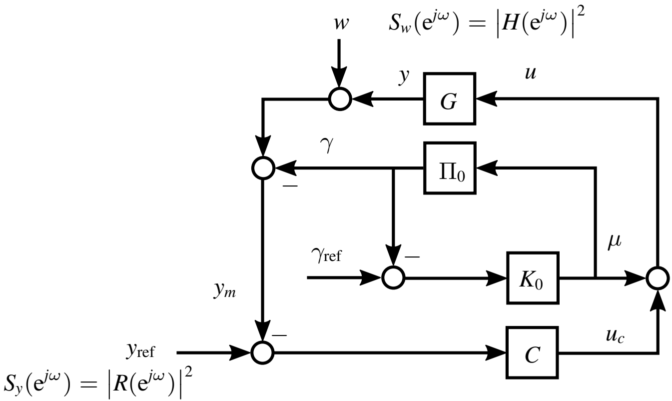
w
$S_w(e^{j\omega}) = \left|H(e^{j\omega})\right|^2$
y
G
u
γ
Π_0
$-$
μ
γ_{ref}
$-$
K_0
y_m
$-$
y_{ref}
C
u_c
$S_y(e^{j\omega}) = \left|R(e^{j\omega})\right|^2$

Model Setup

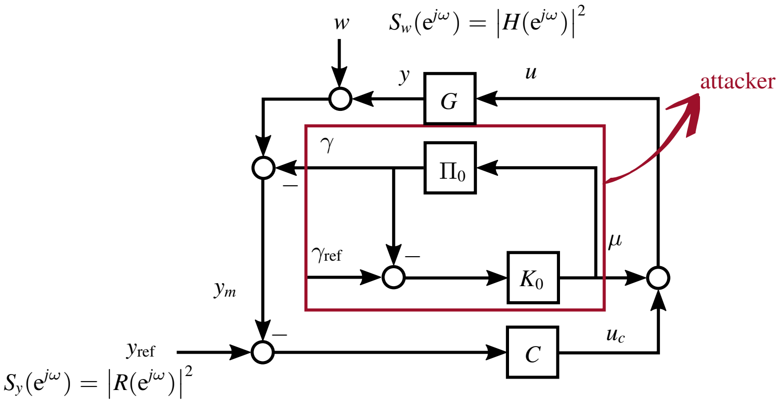
w
$S_w(e^{j\omega}) = |H(e^{j\omega})|^2$
y
G
u
attacker
γ
Π_0
$-$
γ_{ref}
$-$
μ
K_0
y_m
$-$
u_c
y_{ref}
C
$S_y(e^{j\omega}) = |R(e^{j\omega})|^2$

Model Setup

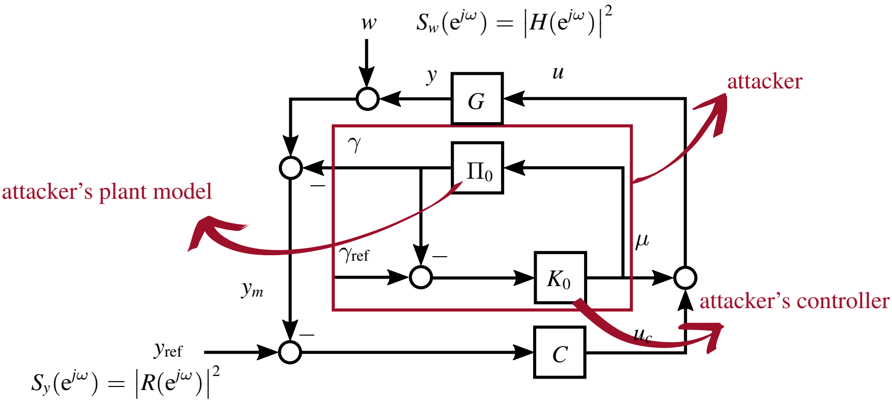

w
$S_w(e^{j\omega}) = \left|H(e^{j\omega})\right|^2$
y
G
u
attacker
γ
Π_0
$-$
attacker's plant model
μ
γ_{ref}
$-$
K_0
y_m
attacker's controller
$-$
u_c
y_{ref}
C
$S_y(e^{j\omega}) = \left|R(e^{j\omega})\right|^2$

Model Setup

$$S_w(e^{j\omega}) = \left|H(e^{j\omega})\right|^2$$

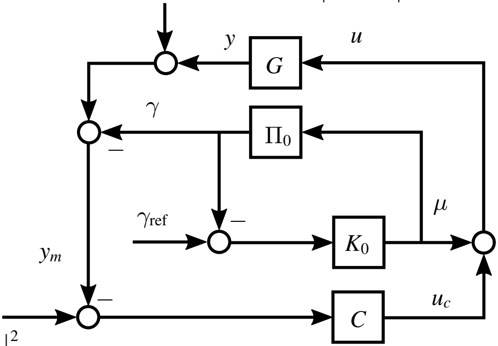

$$S_y(e^{j\omega}) = \left|R(e^{j\omega})\right|^2$$

$$\boldsymbol{J}_C := \|y_{\mathrm{ref}} - y\|_2^2$$

Model Setup

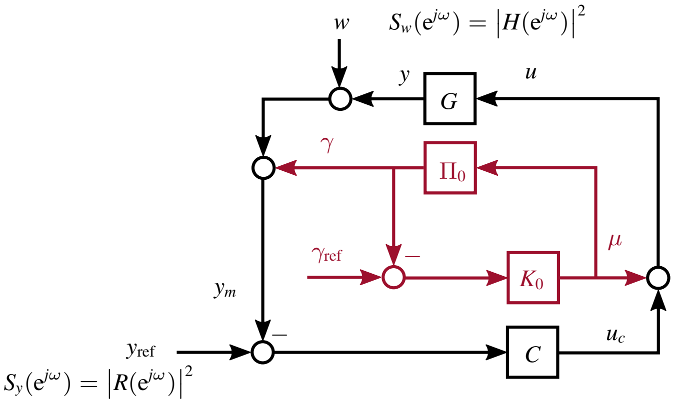

$$J_C := \left\| \left(1 - \frac{[G - \Pi_0]\, C}{1 + GC} \right) \frac{GSK_0}{1 + K_0 \Pi_0} \right\|_2^2$$

Model Setup

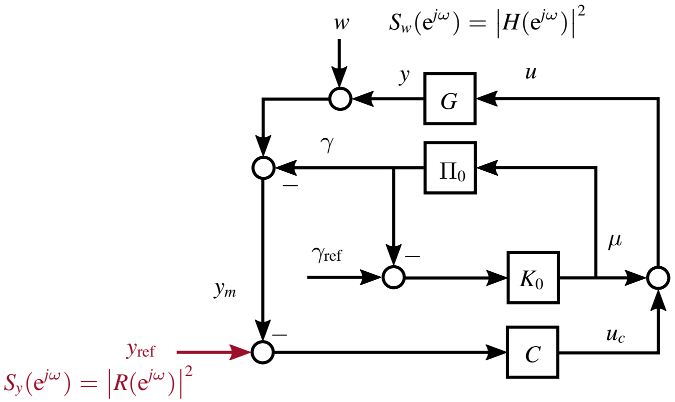

$$J_C := \left\| \left(1 - \frac{[G - \Pi_0]\, C}{1 + GC}\right) \frac{GSK_0}{1 + K_0 \Pi_0} \right\|_2^2 + \left\| \left(1 - \frac{GC}{1 + GC}\right) R \right\|_2^2$$

Model Setup

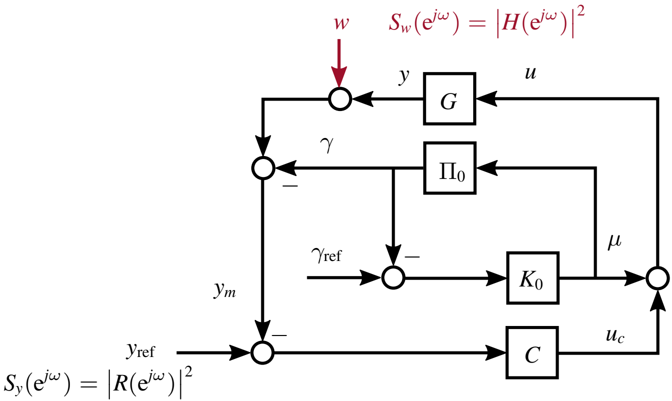

$$J_C := \left\| \left(1 - \frac{[G - \Pi_0]\, C}{1 + GC}\right) \frac{GSK_0}{1 + K_0 \Pi_0} \right\|_2^2 + \left\| \left(1 - \frac{GC}{1 + GC}\right) R \right\|_2^2 + \left\| \frac{HGC}{1 + GC} \right\|_2^2$$

Model Setup

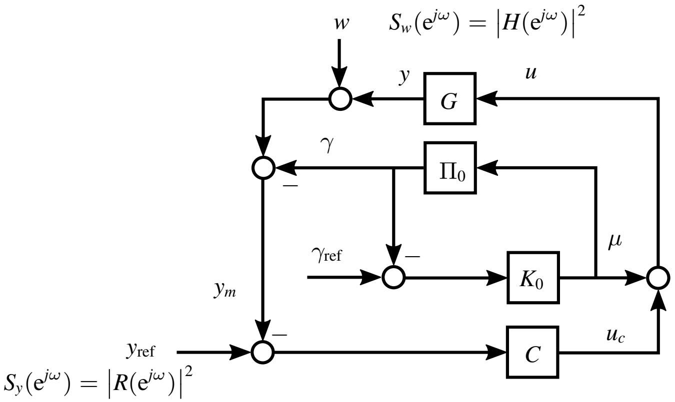

$$J_C := \left\| \left(1 - \frac{[G-\Pi_0]\,C}{1+GC}\right) \frac{GSK_0}{1+K_0\Pi_0} \right\|_2^2 + \left\| \left(1 - \frac{GC}{1+GC}\right) R \right\|_2^2 + \left\| \frac{HGC}{1+GC} \right\|_2^2$$

Model Setup

$$S_w(\mathrm{e}^{j\omega}) = \left|H(\mathrm{e}^{j\omega})\right|^2$$

$\theta \sim p(\theta)$
$\theta \in \Theta$, compact

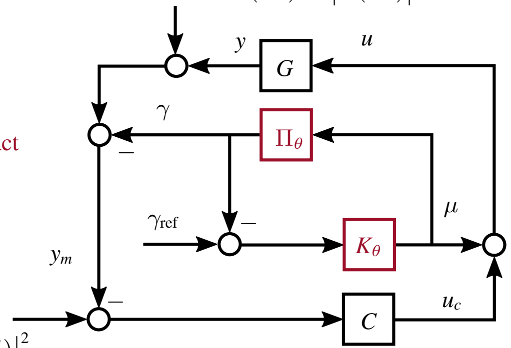

$$S_y(\mathrm{e}^{j\omega}) = \left|R(\mathrm{e}^{j\omega})\right|^2$$

$$\boldsymbol{J_C} := \left\| \left(1 - \frac{[G - \Pi_0]\, C}{1 + GC}\right) \frac{GSK_0}{1 + K_0 \Pi_0} \right\|_2^2 + \left\| \left(1 - \frac{GC}{1 + GC}\right) R \right\|_2^2 + \left\| \frac{HGC}{1 + GC} \right\|_2^2$$

Model Setup

$$S_w(e^{j\omega}) = \left|H(e^{j\omega})\right|^2$$

$\theta \sim p(\theta)$
$\theta \in \Theta$, compact

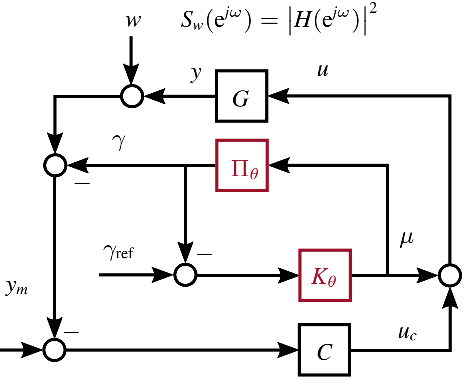

$$S_y(e^{j\omega}) = \left|R(e^{j\omega})\right|^2$$

RANDOM COST

$$J_C(\theta) := \left\| \left(1 - \frac{[G - \Pi_\theta]\, C}{1 + GC}\right) \frac{GSK_\theta}{1 + K_\theta \Pi_\theta} \right\|_2^2 + \left\| \left(1 - \frac{GC}{1+GC}\right) R \right\|_2^2 + \left\| \frac{HGC}{1+GC} \right\|_2^2$$

Problem Formulation

Problem 1: H2RCA ($\mathcal{H}_2$-optimal Risk control under Covert Attacks)

$$\min_{C \in \mathcal{H}_2} \mathcal{R}(\boldsymbol{J}_C)$$

$$\boldsymbol{J}_C(\theta) = \left\| \left(1 - \frac{[G - \Pi_\theta]\, C}{1 + GC}\right) \frac{GSK_\theta}{1 + K_\theta \Pi_\theta} \right\|_2^2 + \left\| \left(1 - \frac{GC}{1 + GC}\right) R \right\|_2^2 + \left\| \frac{HGC}{1 + GC} \right\|_2^2$$

Problem Formulation

Problem 1: H2RCA ($\mathcal{H}_2$-optimal Risk control under Covert Attacks)

$$\min_{C\in\mathcal{H}_2} \mathcal{R}(\boldsymbol{J}_C)$$

$$\boldsymbol{J}_C(\theta) = \left\| \left(1 - \frac{[G-\Pi_\theta]\,C}{1+GC}\right)\frac{GSK_\theta}{1+K_\theta\Pi_\theta} \right\|_2^2 + \left\| \left(1 - \frac{GC}{1+GC}\right) R \right\|_2^2 + \left\| \frac{HGC}{1+GC} \right\|_2^2$$

Contribution:

1. How to choose $\mathcal{R}()$

Problem Formulation

Problem 1: H2RCA ($\mathcal{H}_2$-optimal Risk control under Covert Attacks)

$$\min_{C\in\mathcal{H}_2} \mathcal{R}(\boldsymbol{J}_C)$$

$$\boldsymbol{J}_C(\theta) = \left\| \left(1 - \frac{[G-\Pi_\theta]\,C}{1+GC}\right)\frac{GSK_\theta}{1+K_\theta\Pi_\theta}\right\|_2^2 + \left\| \left(1 - \frac{GC}{1+GC}\right) R \right\|_2^2 + \left\| \frac{HGC}{1+GC} \right\|_2^2$$

Contribution:

1. How to choose $\mathcal{R}()$
2. Risk theoretic framework for attack resilient controller design

Problem Formulation

Problem 1: H2RCA ($\mathcal{H}_2$-optimal Risk control under Covert Attacks)

$$\min_{C \in \mathcal{H}_2} \mathcal{R}(\boldsymbol{J}_C)$$

$$\boldsymbol{J}_C(\theta) = \left\| \left(1 - \frac{[G - \Pi_\theta]\, C}{1 + GC}\right) \frac{GSK_\theta}{1 + K_\theta \Pi_\theta} \right\|_2^2 + \left\| \left(1 - \frac{GC}{1 + GC}\right) R \right\|_2^2 + \left\| \frac{HGC}{1 + GC} \right\|_2^2$$

Contribution:

1. How to choose $\mathcal{R}()$
2. Risk theoretic framework for attack resilient controller design
3. Comparison between different measures of risk

Risk Theory

Risk Theory

$\theta \sim p(\theta)$

Random
variable

Risk Theory

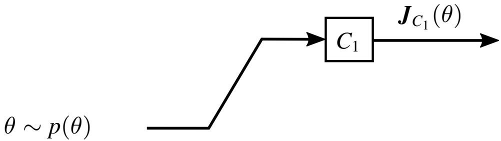

Random variable

Decision

Risk Theory

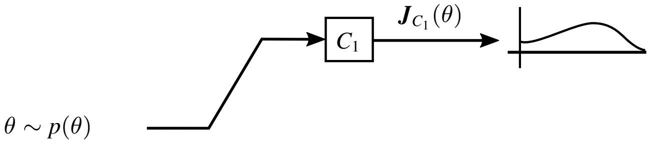

Random variable

Decision

Cost (pdf)

Risk Theory

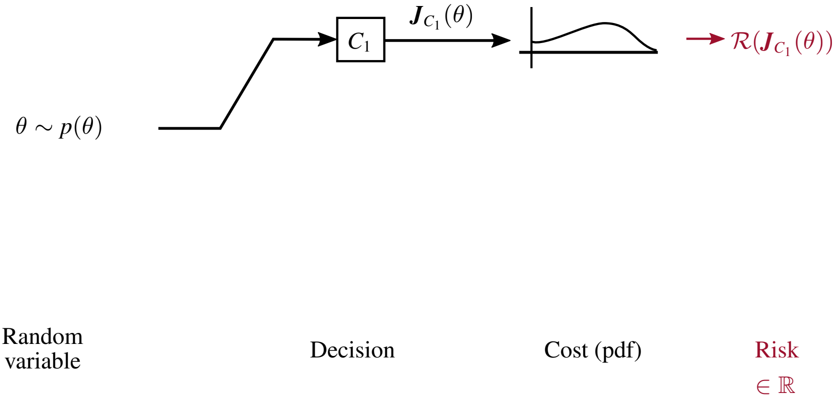
$\theta \sim p(\theta)$
C_1
$\boldsymbol{J}_{C_1}(\theta)$
$\mathcal{R}(\boldsymbol{J}_{C_1}(\theta))$
Random variable
Decision
Cost (pdf)
Risk
$\in \mathbb{R}$

Risk Theory

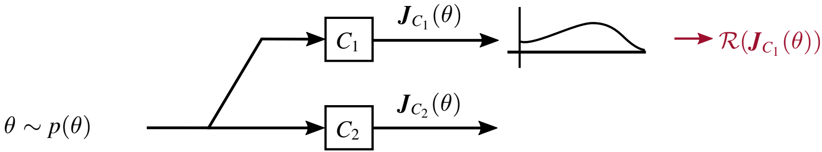

Random variable

Decision

Cost (pdf)

Risk

$\in \mathbb{R}$

Risk Theory

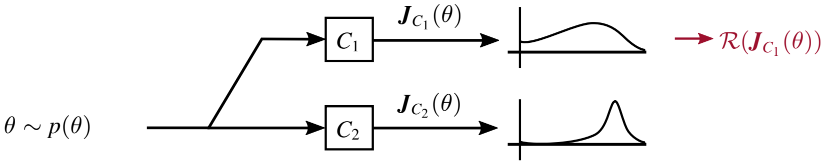

Random variable

Decision

Cost (pdf)

Risk

$\in \mathbb{R}$

Risk Theory

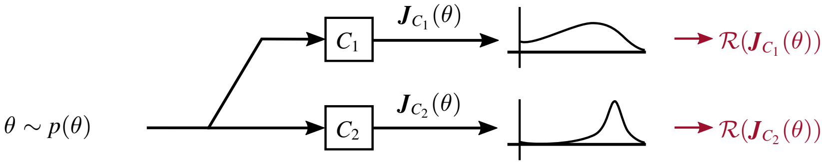

Random variable

Decision

Cost (pdf)

Risk

$\in \mathbb{R}$

Risk Theory

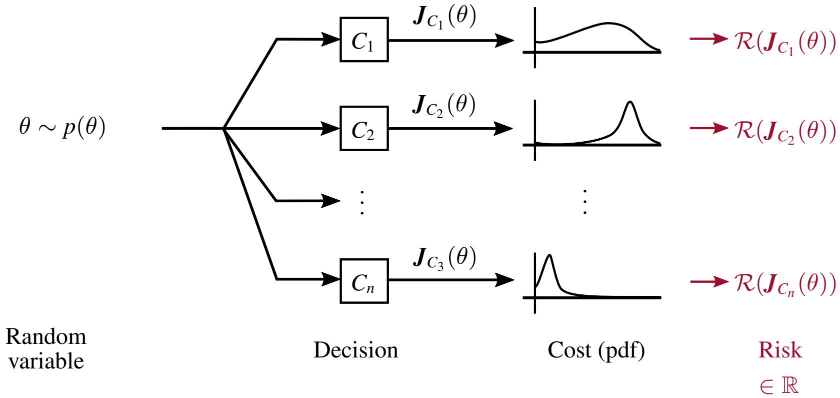
$\theta \sim p(\theta)$
C_1
C_2
C_n
$\boldsymbol{J}_{C_1}(\theta)$
$\boldsymbol{J}_{C_2}(\theta)$
$\boldsymbol{J}_{C_3}(\theta)$
$\mathcal{R}(\boldsymbol{J}_{C_1}(\theta))$
$\mathcal{R}(\boldsymbol{J}_{C_2}(\theta))$
$\mathcal{R}(\boldsymbol{J}_{C_n}(\theta))$
Random variable
Decision
Cost (pdf)
Risk
$\in \mathbb{R}$

Risk Theory

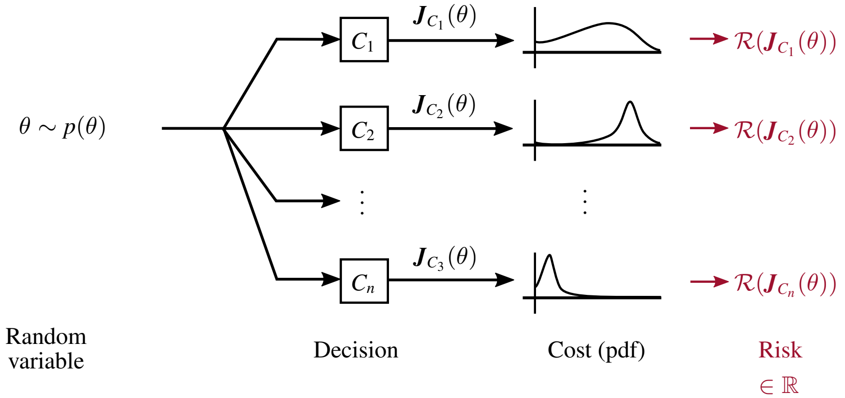

Problem 1: (H2RCA) $\quad \min_{C \in \mathcal{H}_2} \mathcal{R}(\boldsymbol{J}_C)$

Risk Theory

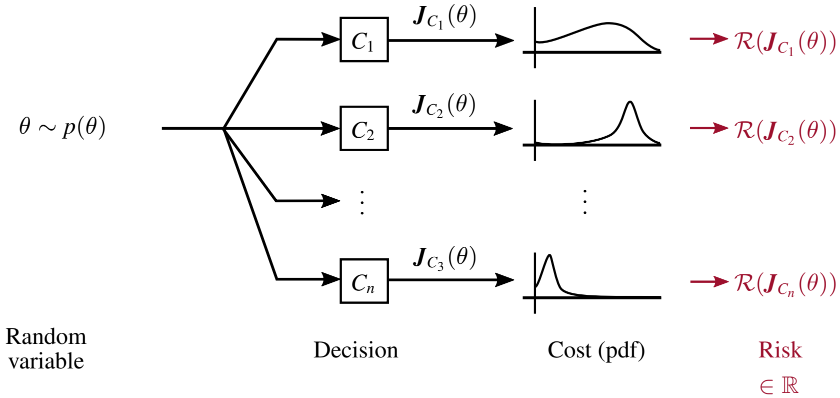

Problem 1: (H2RCA) $\min_{C \in \mathcal{H}_2} \mathcal{R}(J_C)$ $\mathcal{R} = ?$

Conditional Value-at-Risk

Common choices of $\mathcal{R}()$ are $\mathbb{E}\{\}$,

Conditional Value-at-Risk

Common choices of $\mathcal{R}()$ are $\mathbb{E}\left\{\right\}$, worst-case (robust),

Conditional Value-at-Risk

Common choices of $\mathcal{R}()$ are $\mathbb{E}\{\}$, worst-case (robust), nominal.

Conditional Value-at-Risk

Common choices of $\mathcal{R}()$ are $\mathbb{E}\{\}$, worst-case (robust), nominal.

Definition (Conditional value-at-risk)

For $\alpha \in [0, 1]$:

$$\mathrm{CVaR}_\alpha(Y) := \frac{1}{1-\alpha} \int_{y:\mathrm{P}\{Y \leq y\} \geq \alpha} y\, p(y) dy$$

Conditional Value-at-Risk

Common choices of $\mathcal{R}()$ are $\mathbb{E}\{\}$, worst-case (robust), nominal.

Definition (Conditional value-at-risk)

For $\alpha \in [0, 1]$:

$$\mathrm{CVaR}_\alpha(Y) := \frac{1}{1-\alpha} \int_{y:\mathrm{P}\{Y \leq y\} \geq \alpha} y\, p(y) dy$$

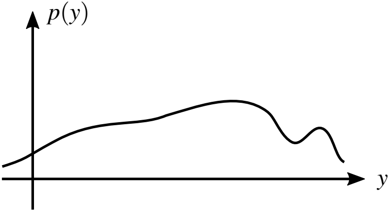

Conditional Value-at-Risk

Common choices of $\mathcal{R}()$ are $\mathbb{E}\{\}$, worst-case (robust), nominal.

Definition (Conditional value-at-risk)

For $\alpha \in [0, 1]$:

$$\mathrm{CVaR}_\alpha(Y) := \frac{1}{1-\alpha} \int_{y:\mathrm{P}\{Y \leq y\} \geq \alpha} y\, p(y) dy$$

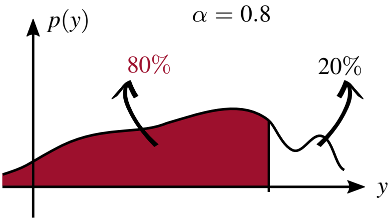

Conditional Value-at-Risk

Common choices of $\mathcal{R}()$ are $\mathbb{E}\{\}$, worst-case (robust), nominal.

Definition (Conditional value-at-risk)

For $\alpha \in [0, 1]$:

$$\mathrm{CVaR}_\alpha(Y) := \frac{1}{1-\alpha} \int_{y:\mathrm{P}\{Y \leq y\} \geq \alpha} y\, p(y) dy$$

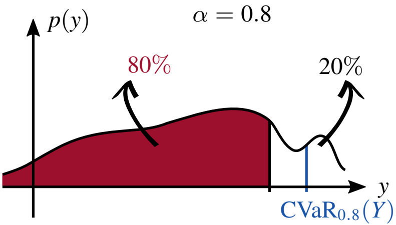

Conditional Value-at-Risk

Common choices of $\mathcal{R}()$ are $\mathbb{E}\{\}$, worst-case (robust), nominal.

Definition (Conditional value-at-risk)

For $\alpha \in [0, 1]$:

$$\mathrm{CVaR}_\alpha(Y) := \frac{1}{1-\alpha} \int_{y:\mathrm{P}\{Y \leq y\} \geq \alpha} y\, p(y) dy = \mathbb{E}\{Y | Y \geq \mathrm{VaR}_\alpha(Y)\}$$

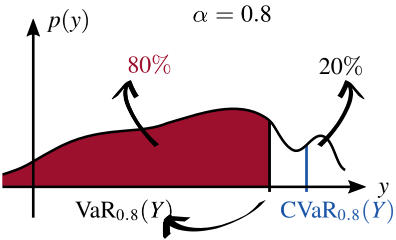

Conditional Value-at-Risk

Common choices of $\mathcal{R}()$ are $\mathbb{E}\{\}$, worst-case (robust), nominal.

Definition (Conditional value-at-risk)

For $\alpha \in [0, 1]$:

$$\mathrm{CVaR}_\alpha(Y) := \frac{1}{1-\alpha} \int_{y:\mathrm{P}\{Y \le y\} \ge \alpha} y\, p(y) dy = \mathbb{E}\{Y | Y \ge \mathrm{VaR}_\alpha(Y)\}$$

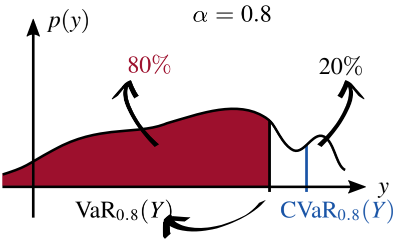

CVaR is convex $\rightarrow$ easy to optimize

An alternative expression for CVaR

An alternative expression for CVaR

Remark

By considering the dual problem (Rockafellar and Uryasev, 2000):

An alternative expression for CVaR

Remark

By considering the dual problem (Rockafellar and Uryasev, 2000):

$$\mathrm{CVaR}_\alpha(\boldsymbol{J}_C) = \min_{\mu\in\mathbb{R}} \mu + \frac{1}{1-\alpha}\mathbb{E}\left\{[\boldsymbol{J}_C - \mu]_+\right\}$$

CVaR Control Design

CVaR Control Design

1. Reparametrize the cost function: Youla Parameter

CVaR Control Design

1. Reparametrize the cost function: Youla Parameter
Let $Q := \frac{C}{1+GC} \iff C = \frac{Q}{1-CQ}$. Then

CVaR Control Design

1. Reparametrize the cost function: Youla Parameter
Let $Q := \frac{C}{1+GC} \iff C = \frac{Q}{1-CQ}$. Then

$$J_C(\theta) = \left\| \left(1 - \frac{[G - \Pi_\theta]\, C}{1+GC}\right) \frac{GSK_\theta}{1 + K_\theta \Pi_\theta} \right\|_2^2 + \left\| \left(1 - \frac{GC}{1+GC}\right) R \right\|_2^2 + \left\| \frac{HGC}{1+GC} \right\|_2^2$$

CVaR Control Design

1. Reparametrize the cost function: Youla Parameter
   Let $Q := \frac{C}{1+GC} \iff C = \frac{Q}{1-CQ}$. Then

$$J_C(\theta) = \left\| \left(1 - \frac{[G - \Pi_\theta]\, C}{1 + GC}\right) \frac{GSK_\theta}{1 + K_\theta \Pi_\theta} \right\|_2^2 + \left\| \left(1 - \frac{GC}{1 + GC}\right) R \right\|_2^2 + \left\| \frac{HGC}{1 + GC} \right\|_2^2$$

CVaR Control Design

1. Reparametrize the cost function: Youla Parameter
Let $Q := \frac{C}{1+GC} \iff C = \frac{Q}{1-CQ}$. Then

$$J_C(\theta) = \left\| \left(1 - \frac{[G-\Pi_\theta]\, C}{1+GC}\right) \frac{GSK_\theta}{1+K_\theta\Pi_\theta} \right\|_2^2 + \left\| \left(1 - \frac{GC}{1+GC}\right) R \right\|_2^2 + \left\| \frac{HGC}{1+GC} \right\|_2^2$$

$$= \left\| (1 - [G-\Pi_\theta]\, Q) \frac{GSK_\theta}{1+K_\theta\Pi_\theta} \right\|_2^2 + \|(1-GQ)R\|_2^2 + \|HGQ\|_2^2$$

CVaR Control Design

1. Reparametrize the cost function: Youla Parameter
Let $Q := \frac{C}{1+GC} \iff C = \frac{Q}{1-CQ}$. Then

$$\boldsymbol{J}_C(\theta) = \left\| \left(1 - \frac{[G-\Pi_\theta]\, C}{1+GC}\right) \frac{GSK_\theta}{1+K_\theta\Pi_\theta} \right\|_2^2 + \left\| \left(1 - \frac{GC}{1+GC}\right) R \right\|_2^2 + \left\| \frac{HGC}{1+GC} \right\|_2^2$$

$$= \left\| (1 - [G-\Pi_\theta]\, Q) \frac{GSK_\theta}{1+K_\theta\Pi_\theta} \right\|_2^2 + \|(1-GQ)R\|_2^2 + \|HGQ\|_2^2$$

$$=: \boldsymbol{V}_Q(\theta)$$

CVaR Control Design

1. Reparametrize the cost function: Youla Parameter
Let $Q := \frac{C}{1+GC} \iff C = \frac{Q}{1-CQ}$. Then

$$J_C(\theta) = \left\| \left(1 - \frac{[G-\Pi_\theta]\, C}{1+GC}\right) \frac{GSK_\theta}{1+K_\theta\Pi_\theta} \right\|_2^2 + \left\| \left(1 - \frac{GC}{1+GC}\right) R \right\|_2^2 + \left\| \frac{HGC}{1+GC} \right\|_2^2$$

$$= \left\| (1 - [G-\Pi_\theta]\, Q) \frac{GSK_\theta}{1+K_\theta\Pi_\theta} \right\|_2^2 + \|(1-GQ)R\|_2^2 + \|HGQ\|_2^2$$

$$=: \boldsymbol{V}_Q(\theta)$$

H2RCA is then equivalent to

$$\min_{Q\in\mathcal{H}_2} \mathrm{CVaR}_\alpha(\boldsymbol{V}_Q)$$

CVaR Control Design

❷ Approximate the feasible set $\mathcal{H}_2$ in

$$\min_{Q \in \mathcal{H}_2} \mathrm{CVaR}_\alpha(V_Q)$$

CVaR Control Design

2. Approximate the feasible set $\mathcal{H}_2$ in

$$\min_{Q \in \mathcal{H}_2} \mathrm{CVaR}_\alpha(\boldsymbol{V}_Q) \approx \min_{Q \in \mathcal{Q}_L} \mathrm{CVaR}_\alpha(\boldsymbol{V}_Q)$$

CVaR Control Design

2. Approximate the feasible set $\mathcal{H}_2$ in

$$\min_{Q \in \mathcal{H}_2} \mathrm{CVaR}_\alpha(\boldsymbol{V}_Q) \approx \min_{Q \in \mathcal{Q}_L} \mathrm{CVaR}_\alpha(\boldsymbol{V}_Q)$$

with $\mathcal{Q}_L = \{Q \colon Q(z) = \sum_{k=0}^{L} x_k z^{-k}, x_0, \ldots, x_L \in \mathbb{R}\}$.

CVaR Control Design

2. Approximate the feasible set $\mathcal{H}_2$ in

$$\min_{Q \in \mathcal{H}_2} \mathrm{CVaR}_\alpha(\boldsymbol{V}_Q) \approx \min_{Q \in \mathcal{Q}_L} \mathrm{CVaR}_\alpha(\boldsymbol{V}_Q)$$

with $\mathcal{Q}_L = \{Q \colon Q(z) = \sum_{k=0}^{L} x_k z^{-k}, x_0, \ldots, x_L \in \mathbb{R}\}$.

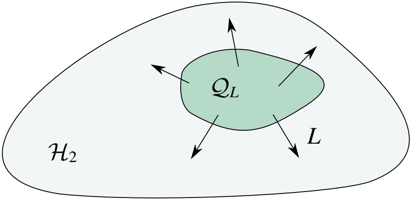

$\lim_{L\to\infty} \mathcal{Q}_L = \mathcal{H}_2$

CVaR Control Design

3. We approximate $\mathrm{CVaR}_{\alpha}(V_Q(\theta))$:

CVaR Control Design

3. We approximate $\mathrm{CVaR}_\alpha(\boldsymbol{V}_Q(\theta))$:

$$\mathrm{CVaR}_\alpha(\boldsymbol{V}_Q) = \min_{\mu\in\mathbb{R}} \mu + \frac{1}{1-\alpha}\mathbb{E}\left\{[\boldsymbol{V}_Q - \mu]_+\right\}$$

CVaR Control Design

3. We approximate $\mathrm{CVaR}_\alpha(\boldsymbol{V}_Q(\theta))$:

$$\mathrm{CVaR}_\alpha(\boldsymbol{V}_Q) = \min_{\mu \in \mathbb{R}} \mu + \frac{1}{1-\alpha} \underbrace{\mathbb{E}\left\{[\boldsymbol{V}_Q - \mu]_+\right\}}_{\text{hard to compute!}}$$

CVaR Control Design

3. We approximate $\mathrm{CVaR}_\alpha(\boldsymbol{V}_Q(\theta))$:

$$\mathrm{CVaR}_\alpha(\boldsymbol{V}_Q) = \min_{\mu \in \mathbb{R}} \mu + \frac{1}{1-\alpha} \underbrace{\mathbb{E}\left\{[\boldsymbol{V}_Q - \mu]_+\right\}}_{\text{hard to compute!}}$$

$\{\theta_i\}_{i=1}^N$: N iid samples from $p(\theta)$

CVaR Control Design

3. We approximate $\mathrm{CVaR}_\alpha(\boldsymbol{V}_Q(\theta))$:

$$\mathrm{CVaR}_\alpha(\boldsymbol{V}_Q) = \min_{\mu\in\mathbb{R}} \mu + \frac{1}{1-\alpha}\underbrace{\mathbb{E}\left\{[\boldsymbol{V}_Q-\mu]_+\right\}}_{\text{hard to compute!}}$$

$$\approx \min_{\mu\in\mathbb{R}} \mu + \frac{1}{1-\alpha}\frac{1}{N}\sum_{i=1}^{N}[\boldsymbol{V}_Q(\theta_i)-\mu]_+$$

$\{\theta_i\}_{i=1}^N$: N iid samples from $p(\theta)$

CVaR Control Design

3. We approximate $\mathrm{CVaR}_\alpha(\boldsymbol{V}_Q(\theta))$:

$$\mathrm{CVaR}_\alpha(\boldsymbol{V}_Q) = \min_{\mu\in\mathbb{R}} \mu + \frac{1}{1-\alpha}\underbrace{\mathbb{E}\left\{[\boldsymbol{V}_Q-\mu]_+\right\}}_{\text{hard to compute!}}$$

$$\approx \min_{\mu\in\mathbb{R}} \mu + \frac{1}{1-\alpha}\frac{1}{N}\sum_{i=1}^{N}[\boldsymbol{V}_Q(\theta_i)-\mu]_+$$

$$=: \overline{\mathrm{CVaR}}_\alpha(\{\boldsymbol{V}_Q(\theta_i)\}_{i=1}^N)$$

$\{\theta_i\}_{i=1}^N$: N iid samples from $p(\theta)$

CVaR Control Design

3. We approximate $\mathrm{CVaR}_\alpha(\boldsymbol{V}_Q(\theta))$:

$$\mathrm{CVaR}_\alpha(\boldsymbol{V}_Q) = \min_{\mu\in\mathbb{R}} \mu + \frac{1}{1-\alpha}\underbrace{\mathbb{E}\left\{[\boldsymbol{V}_Q-\mu]_+\right\}}_{\text{hard to compute!}}$$

$$\approx \min_{\mu\in\mathbb{R}} \mu + \frac{1}{1-\alpha}\frac{1}{N}\sum_{i=1}^{N}[\boldsymbol{V}_Q(\theta_i)-\mu]_+$$

$$=: \overline{\mathrm{CVaR}}_\alpha(\{\boldsymbol{V}_Q(\theta_i)\}_{i=1}^N)$$

$\{\theta_i\}_{i=1}^N$: N iid samples from $p(\theta)$

Problem 2:

$$\min_{Q\in\mathcal{H}_2} \mathrm{CVaR}_\alpha(\boldsymbol{V}_Q) \approx \min_{Q\in\mathcal{Q}_L} \overline{\mathrm{CVaR}}_\alpha(\{\boldsymbol{V}_Q(\theta_i)\}_{i=1}^N)$$

CVaR Control Design

Lemma (Convergence of cost functions)

Let N = # iid samples from $p(\theta)$, and L = length of the FIR filter:

CVaR Control Design

Lemma (Convergence of cost functions)

Let N = # iid samples from $p(\theta)$, and L = length of the FIR filter:

$$\lim_{N,L\to\infty} \min_{Q\in\mathcal{Q}_L} \overline{\mathrm{CVaR}_\alpha}(\{\boldsymbol{V}_Q(\theta_i)\}_{i=1}^N) = \min_{Q\in\mathcal{H}_2} \mathrm{CVaR}_\alpha(\boldsymbol{V}_Q)$$

Main Result

Theorem (QCLP)

Let N = # iid samples from $p(\theta)$, and L = length of the FIR filter.

Main Result

Theorem (QCLP)

Let N = # iid samples from $p(\theta)$, and L = length of the FIR filter. Then
$Q^\star := \arg\min_{Q \in \mathcal{Q}_L} \overline{\mathrm{CVaR}_\alpha}(\{V_Q(\theta_i)\}_{i=1}^N) = \sum_{k=0}^{L} x_k^\star z^{-k}$,

Main Result

Theorem (QCLP)

Let N = # iid samples from $p(\theta)$, and L = length of the FIR filter. Then
$Q^\star := \arg\min_{Q\in\mathcal{Q}_L} \overline{\mathrm{CVaR}_\alpha(}\{V_Q(\theta_i)\}_{i=1}^N) = \sum_{k=0}^L x_k^\star z^{-k}$,
$\boldsymbol{x}^\star := \begin{bmatrix} x_0^\star & x_1^\star & \dots & x_L^\star \end{bmatrix}$,

Main Result

Theorem (QCLP)

Let N = # iid samples from $p(\theta)$, and L = length of the FIR filter. Then $Q^\star := \arg\min_{Q \in \mathcal{Q}_L} \mathrm{CVaR}_\alpha(\{V_Q(\theta_i)\}_{i=1}^N) = \sum_{k=0}^{L} x_k^\star z^{-k}$, $\boldsymbol{x}^\star := \begin{bmatrix} x_0^\star & x_1^\star & \dots & x_L^\star \end{bmatrix}$, *where*

$$\begin{bmatrix} \boldsymbol{x}^\star & \mu^\star & \boldsymbol{t}^\star \end{bmatrix}^\top := \underset{\begin{bmatrix} \boldsymbol{x} & \mu & \boldsymbol{t} \end{bmatrix}^\top \in \mathbb{R}^{L+N+2}}{\arg\min} \ \mu + \frac{1}{N(1-\alpha)} \mathbf{1}_N^\top \boldsymbol{t}$$

subject to

$$t_i \geq k(\theta_i) + \boldsymbol{x}^\top \boldsymbol{M}(\theta_i)\boldsymbol{x} - 2\boldsymbol{c}^\top(\theta_i)\boldsymbol{x} - \mu,$$

$$t_i \geq 0, \quad i = 1, \dots, N$$

Main Result

Theorem (QCLP)

Let N = # iid samples from $p(\theta)$, and L = length of the FIR filter. Then $Q^\star := \arg\min_{Q \in \mathcal{Q}_L} \mathrm{CVaR}_\alpha(\{\boldsymbol{V}_Q(\theta_i)\}_{i=1}^N) = \sum_{k=0}^L x_k^\star z^{-k}$, $\boldsymbol{x}^\star := \begin{bmatrix} x_0^\star & x_1^\star & \dots & x_L^\star \end{bmatrix}$, *where*

$$\begin{bmatrix} \boldsymbol{x}^\star & \mu^\star & \boldsymbol{t}^\star \end{bmatrix}^\top := \underset{\begin{bmatrix} \boldsymbol{x} & \mu & \boldsymbol{t} \end{bmatrix}^\top \in \mathbb{R}^{L+N+2}}{\arg\min} \mu + \frac{1}{N(1-\alpha)} \mathbf{1}_N^\top \boldsymbol{t}$$

subject to

$$t_i \geq k(\theta_i) + \boldsymbol{x}^\top \boldsymbol{M}(\theta_i)\boldsymbol{x} - 2\boldsymbol{c}^\top(\theta_i)\boldsymbol{x} - \mu,$$

$$t_i \geq 0, \quad i = 1, \dots, N$$

H2RCA ≈ Problem 2 $\rightarrow$ QCLP (easy to solve)

Example

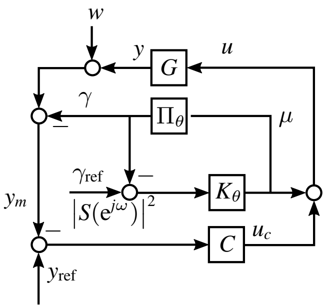

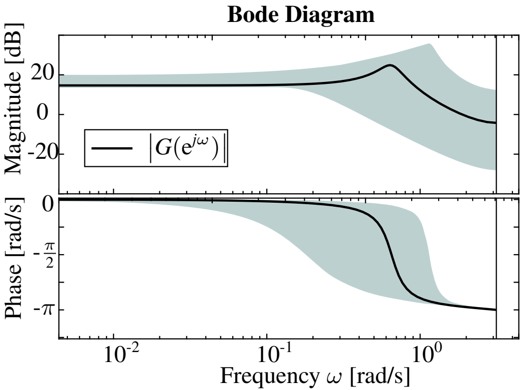

$$K_\theta := \arg\min_{K \in \mathcal{H}_2} \left\| \left(1 - \frac{\Pi_\theta K}{1 + \Pi_\theta K}\right) S \right\|_2^2$$

Example

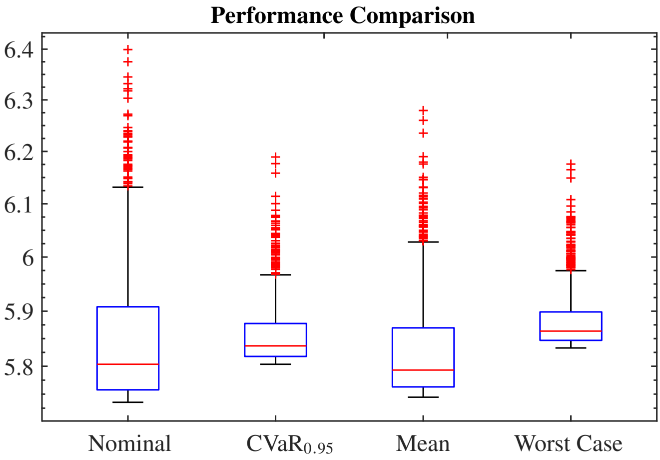

Performance Comparison
6.4
6.3
6.2
6.1
6
5.9
5.8
Nominal
CVaR$_{0.95}$
Mean
Worst Case

Conclusions

- Link between Control Design and Financial Theory of Risk

Conclusions

- Link between Control Design and Financial Theory of Risk
- When $\mathcal{R} = \mathrm{CVaR}_\alpha$, a QLCP approximates the solution

Conclusions

- Link between Control Design and Financial Theory of Risk
- When $\mathcal{R} = \mathrm{CVaR}_\alpha$, a QLCP approximates the solution
- Better control performance by using $p(\theta)$

A Risk-Theoretical Approach to $\mathcal{H}_2$-Optimal Control under Covert Attacks

Matias I. Müller,
Jezdimir Milošević, Henrik Sandberg and Cristian R. Rojas
(e-mails: `{mimr2,jezdimir,hsan,crro}@kth.se`)

Department of Automatic Control,
KTH Royal Institute of Technology,
Stockholm, Sweden

57th IEEE Conference on Decision and Control
December 18th, 2018

CVaR Properties

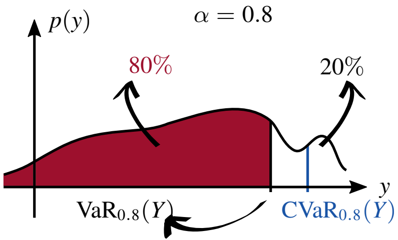

CVaR Properties

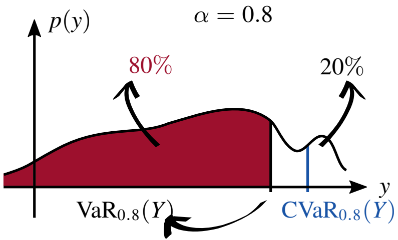

$\mathrm{CVaR}_0(Y) = \mathbb{E}\{Y\}$

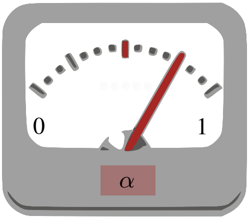

$\mathrm{CVaR}_1(Y) = \max\{\text{support } Y\}$