9th WS of the COST Action TU1402, 29-30 May 2017, Dublin, Ireland

matteo pozzi

exploring infrastructure systems : free and constrained sensing optimization

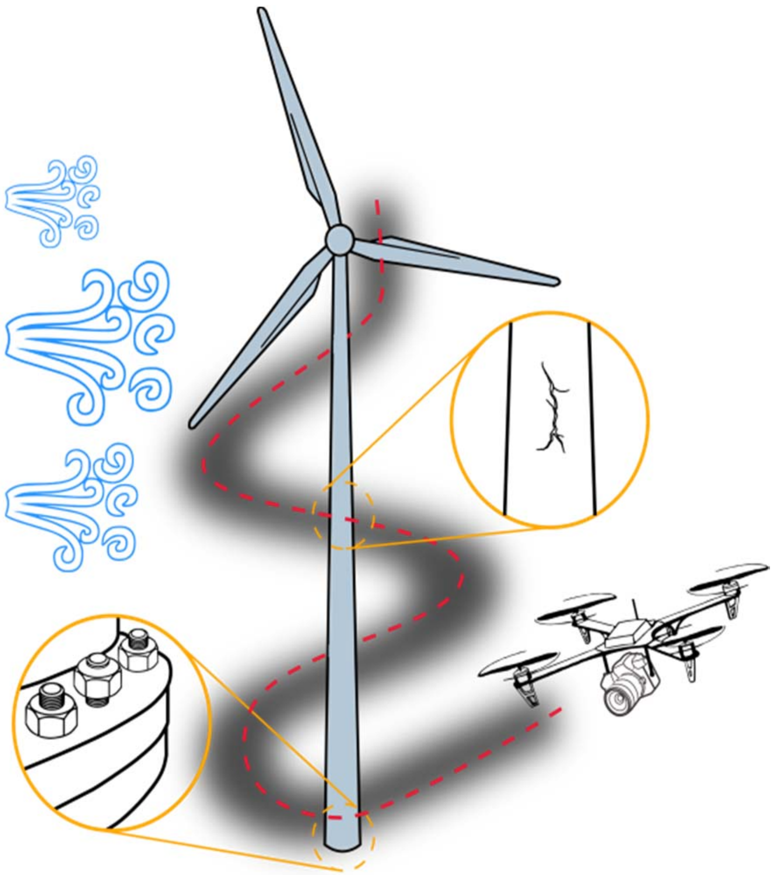

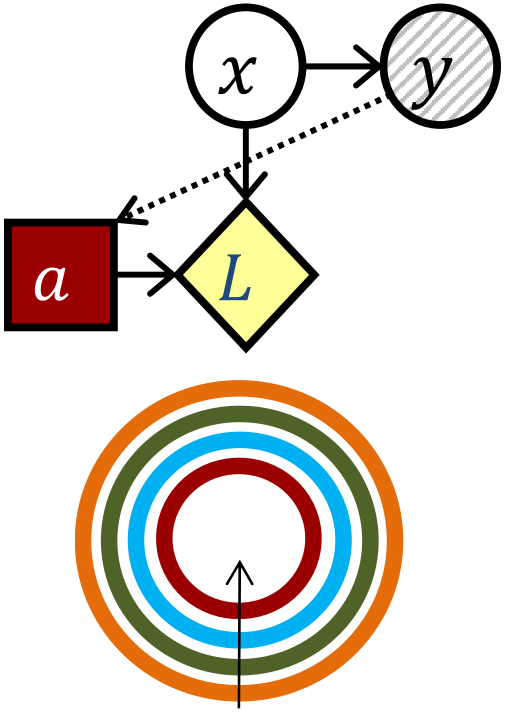

Carnegie Mellon

acknowledgements

Milad Memarzadeh

Carl Malings

Shawn Li

Metro21
Carnegie Mellon University

Carnegie Mellon University
Scott Institute
for Energy Innovation

PITA
Pennsylvania
Infrastructure
Technology Alliance

NSF

motivation: crumbling infrastructure systems

Roads: $101 billion in wasted time and fuel annually, $170 billion needed to improve conditions and performance.

Bridges: 600,000 bridges, average age 42 years, (1/9: structurally deficient). $20.5 billion needed, $12.8 billion currently spent.

[annual estimates by FHWA]

Energy: "America relies on an aging electrical grid and pipeline distribution systems [...] increasing number of failures"

www.infrastructurereportcard.org

Cyber-physical systems: How to integrate sensors and robotic inspectors in adaptive maintenance strategies for interconnected systems?

RedZone ROBOTICS

the value of information

VoI is metric based on Bayesian analysis and utility theory.

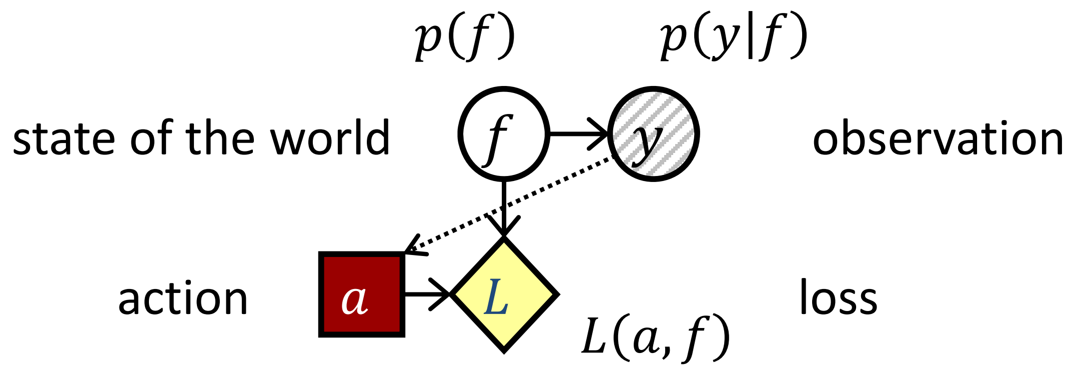

Exp. loss w/o Y $\quad L(\emptyset) = \min_a \mathbb{E}_F \, [L(a,f)]$

Expect loss observing Y $\quad L(Y) = \mathbb{E}_Y \min_a \mathbb{E}_{F|y}[L(a,f)]$

Value of information

$$VoI(Y) = L(\emptyset) - L(Y) \geq 0$$

[inference $\rightarrow p(f|y)$ *]*

[optimization]

[integration using all possible measures: $p(y)$ *]*

information gathering: $\quad Y^* = \text{argmin}\, L(Y) = \text{argmin}\, \mathbb{E}_Y \min_a \mathbb{E}_{F|y}[L(a,f)]$

[optimization]

why VoI is relevant: applications

spatial models

It can be used for assessing the maximum allowable investment for obtaining a piece of information.

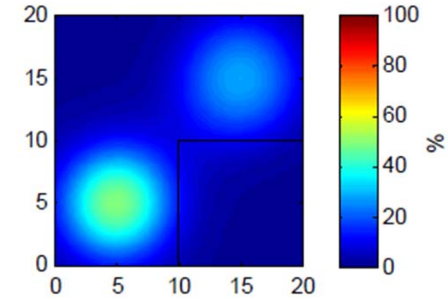

It can be used for comparing exploitative and explorative actions.

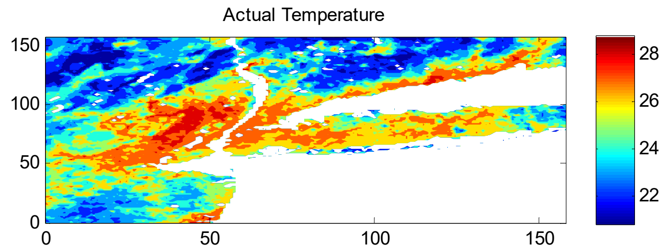

It can be used for giving priorities among observations that can be collected.

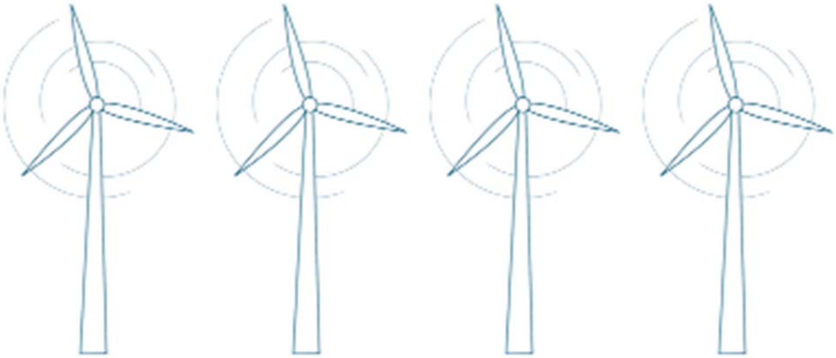

temporal model

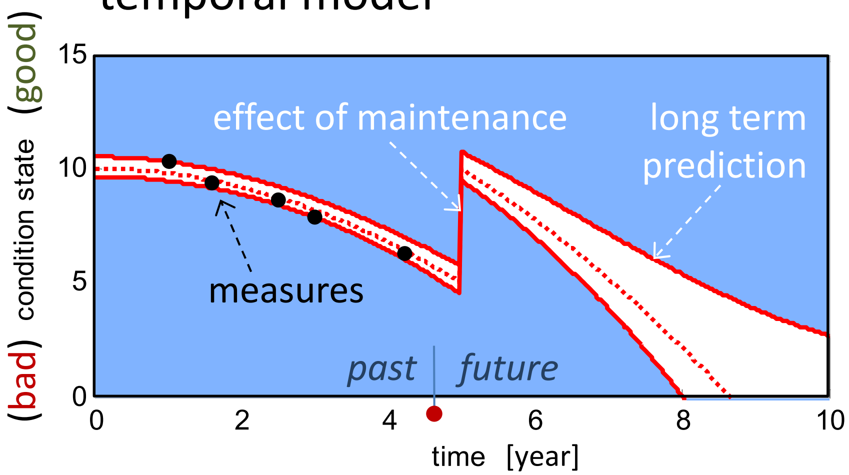

application to Urban Heat Risk

urban heat island effect

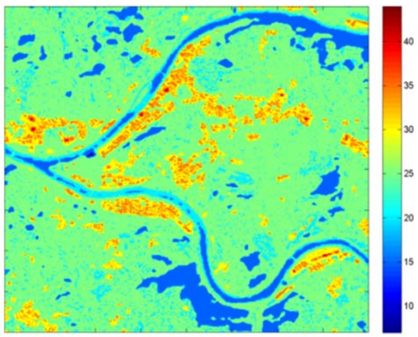

urban population density

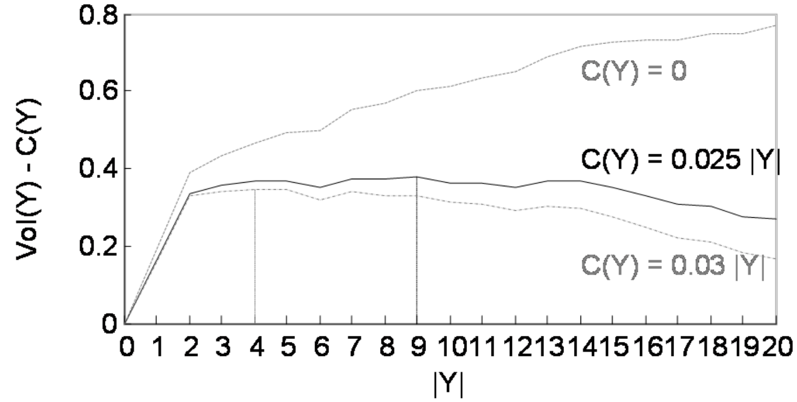

uncertainty in temperature prediction

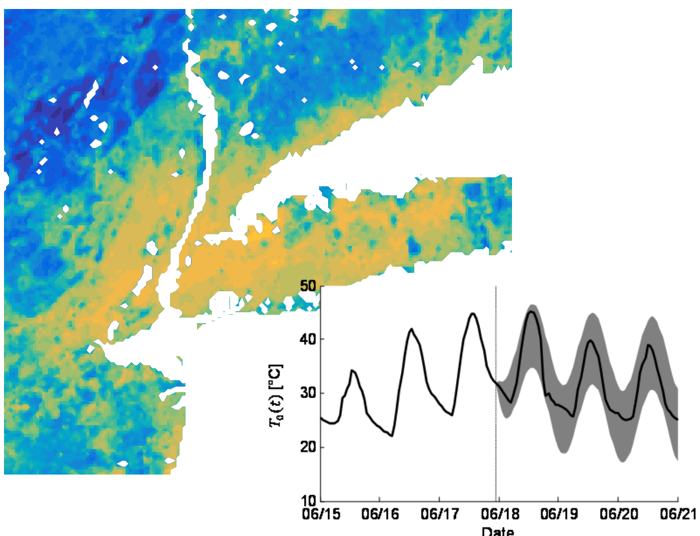

PRINCETON UNIVERSITY

Local Urban Weather Model (PUCM-WRF)

↓

Probabilistic Urban Temperature Model

sensor placement

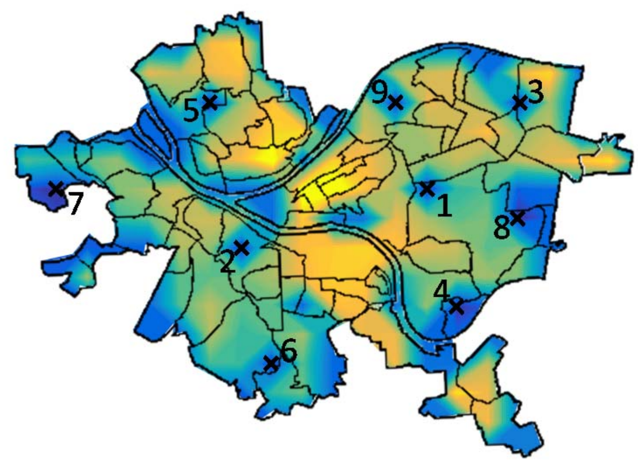

tasks: modelling temperature, optimizing data collection.

tools: Gaussian Processes, greedy optimization of Value of Information.

adaptive measurement scheduling

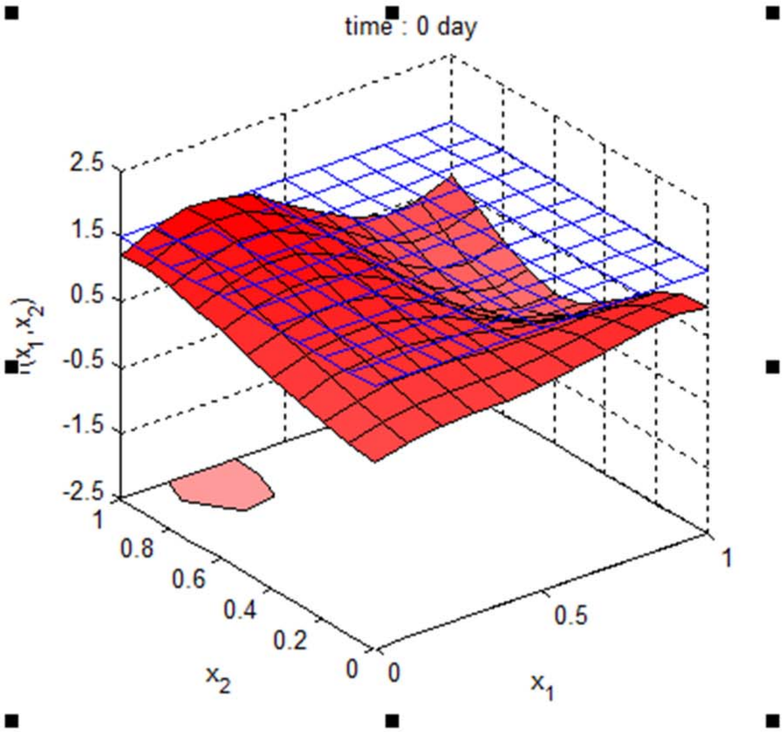

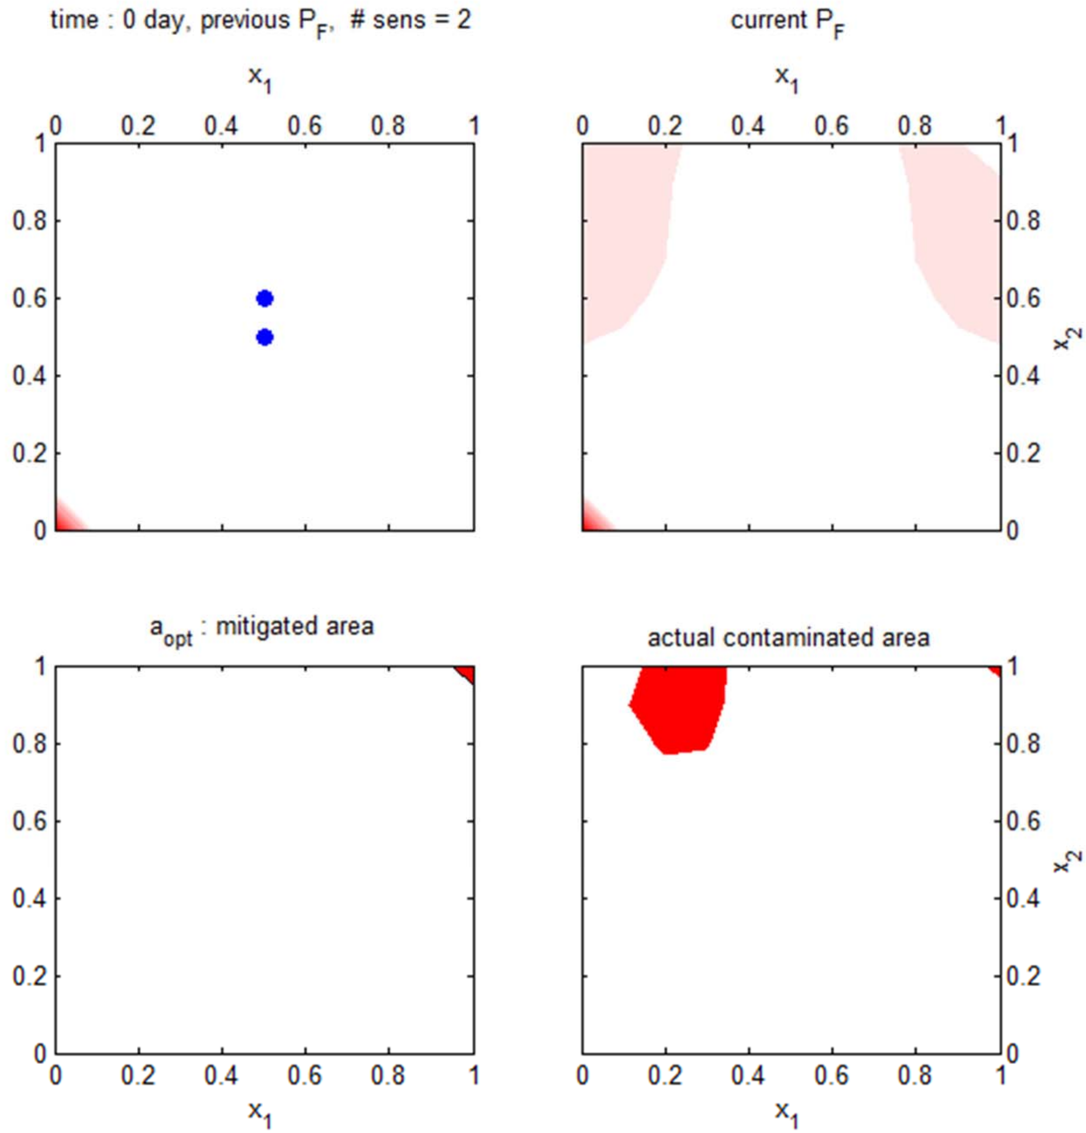

Contamination diffusion:
Optimal sensor placement

It is related to Uncertainty, Expected value, Correlation with other locations and with future values .

VoI in civil engineering research

Economics:

Howard Raiffa and R. Schlaifer, 1961
Ron Howard, 1966

Computer Science:

Andreas Krause

CE:

Michael Faber
Daniel Straub
Samer Madanat
Armen Der Kiureghian
Sebastian Thöns
Daniele Zonta
James Goulet
…

simplest maintenance problem

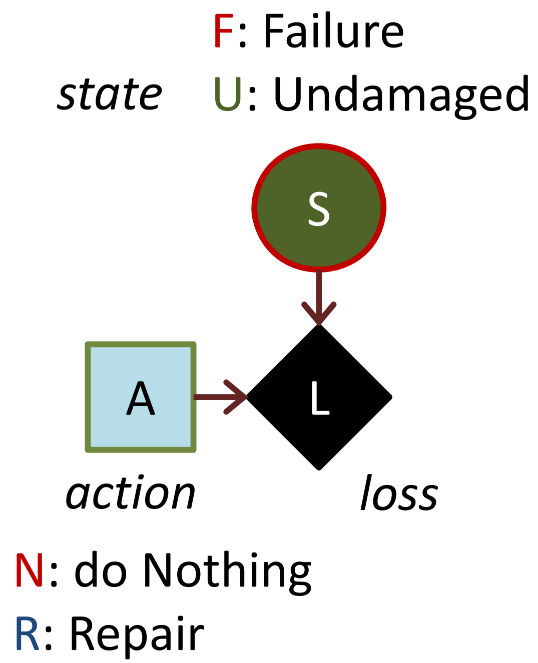

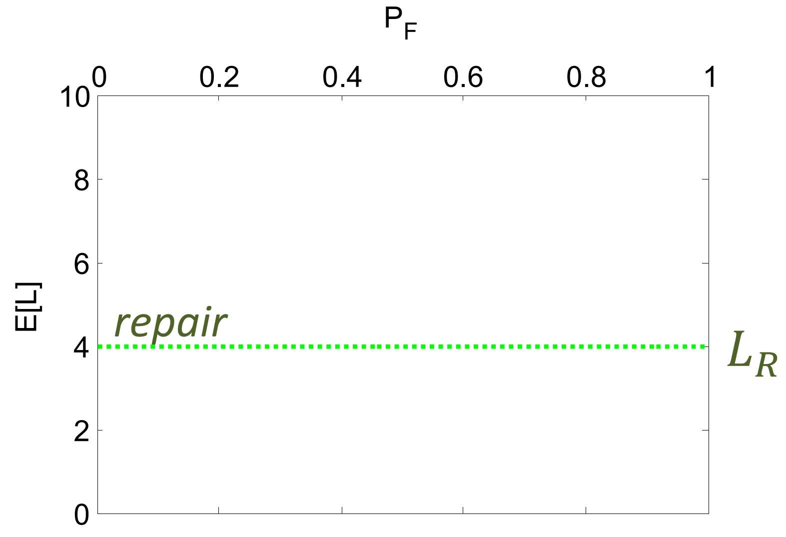

$L(S,A)$ (S) (A)

	U	F
N	0	L_F
R	L_R	L_R

agent's loss matrix

simplest maintenance problem

state
F: Failure
U: Undamaged

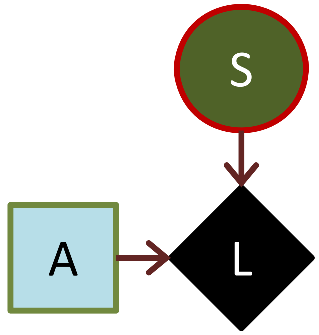

action *loss*

N: do Nothing
R: Repair

$L(S,A)$

(A) \ (S)	U	F
N	0	L_F
R	L_R	L_R

agent's loss matrix

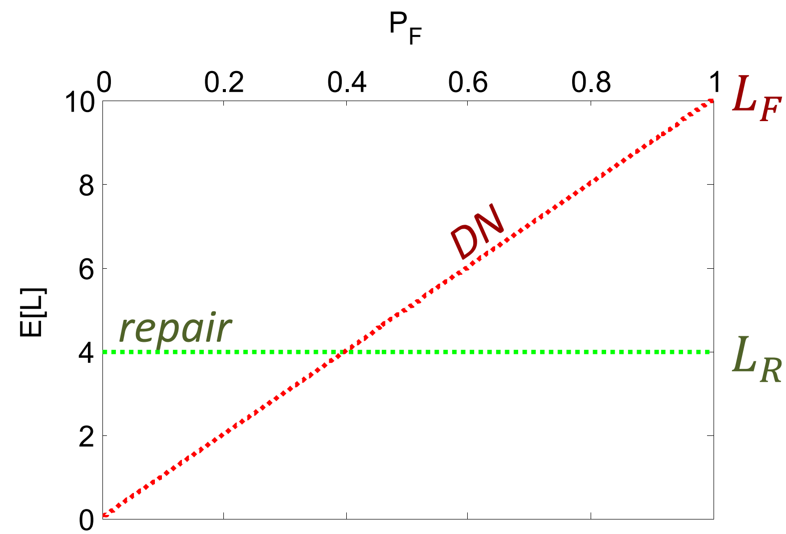

simplest maintenance problem

state F: Failure
U: Undamaged

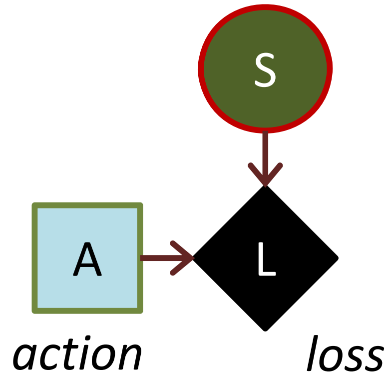

N: do Nothing
R: Repair

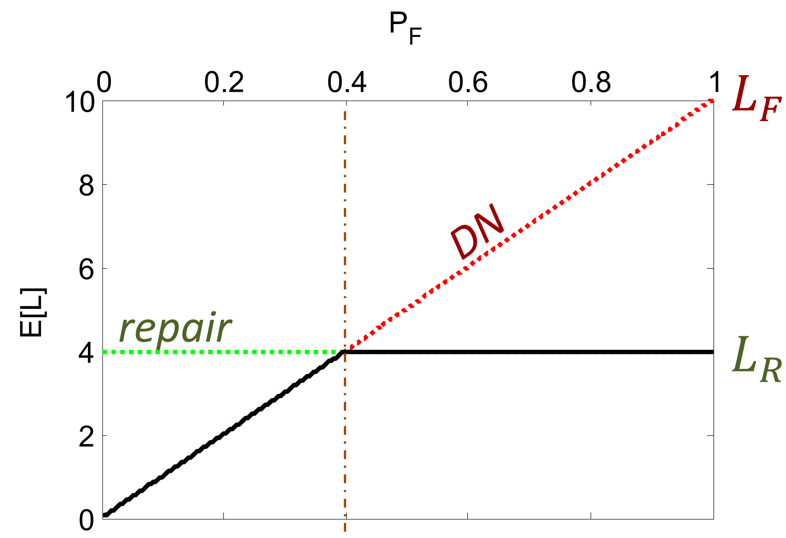

opt. threshold: $P_L = L_R / L_F$

DN *Repair*

$L(S,A)$ (S)

(A)		U	F
	N	0	L_F
	R	L_R	L_R

agent's loss matrix

simplest maintenance problem

state F: Failure
U: Undamaged

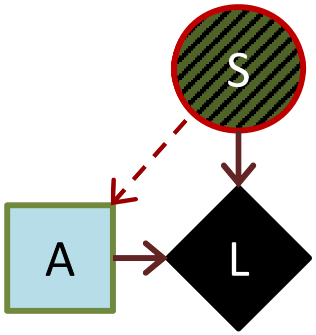

action *loss*

N: do Nothing
R: Repair

$L(S,A)$

(A) \ (S)	U	F
N	0	L_F
R	L_R	L_R

agent's loss matrix

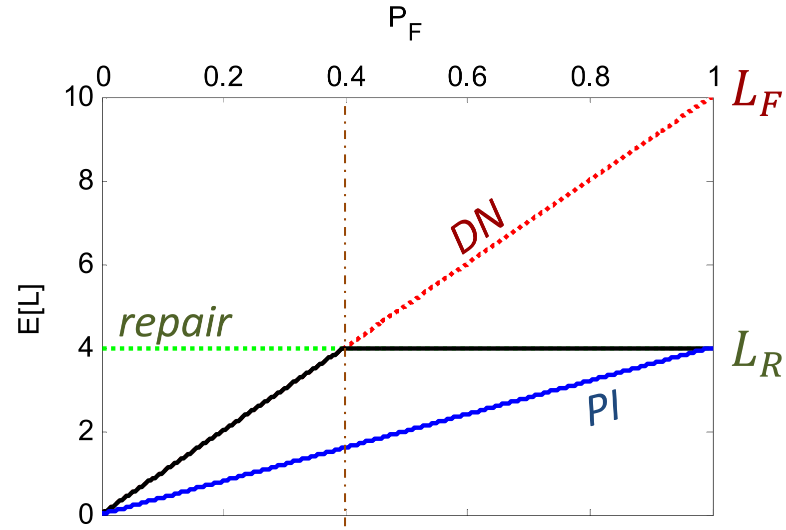

opt. threshold: $P_L = L_R / L_F$

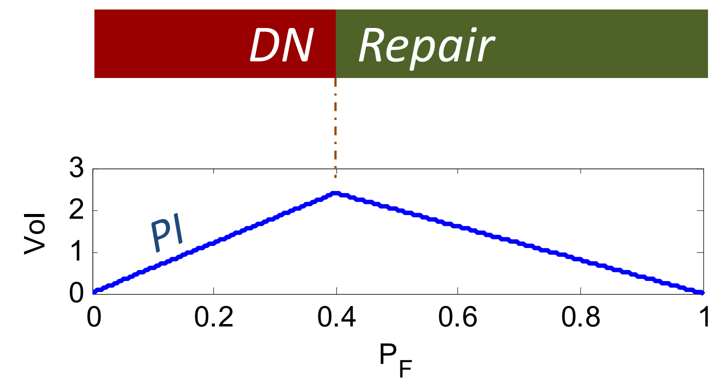

decision tree and Markov process

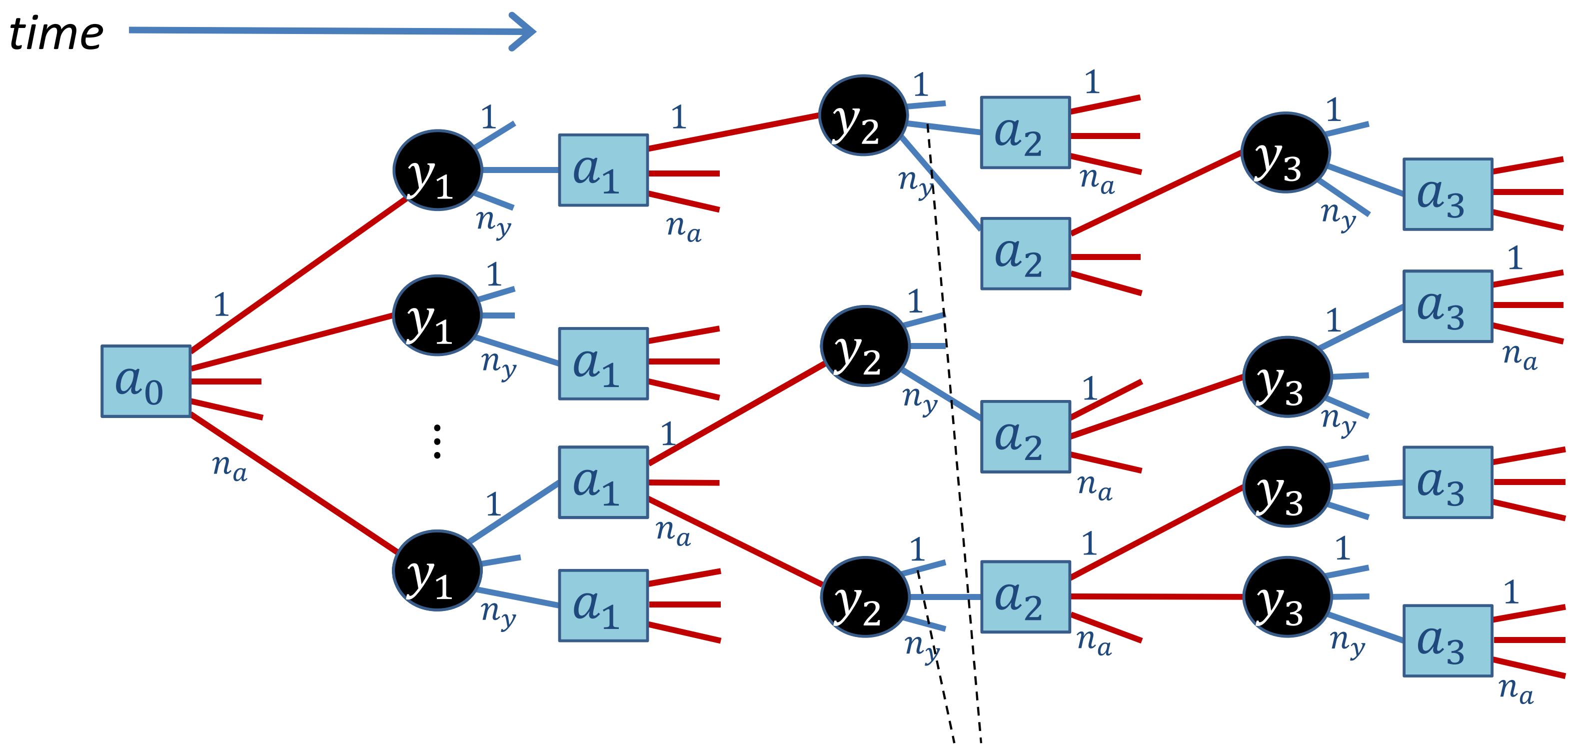

# of leaves on a tree grows exponentially with number of steps (depth growth).

if same sufficient statistics: then same optimal action

Bellman's equation:
complexity can grow linearly with number of steps.

sequential decision making

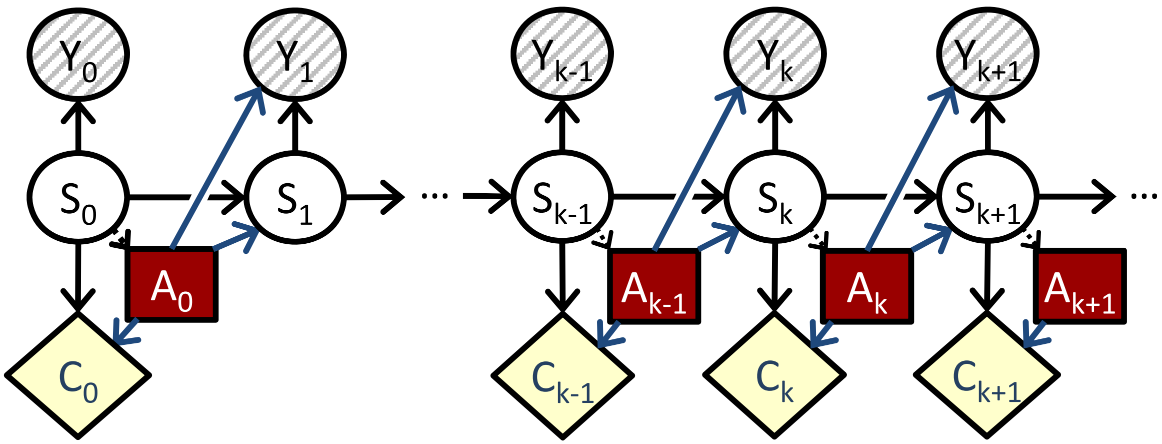

observations

state

~~one-stage decision, as before~~

multi-stage temporal process

partially observable MDP

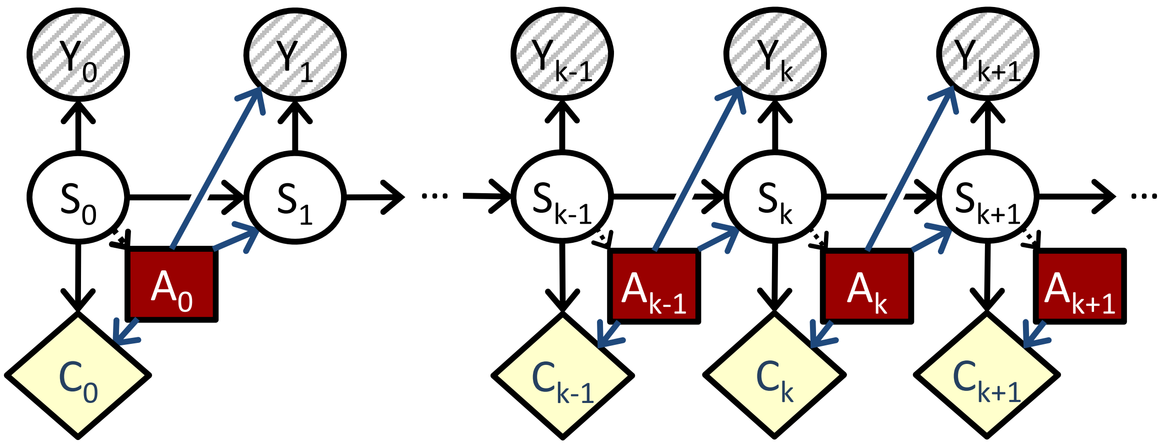

observations

state

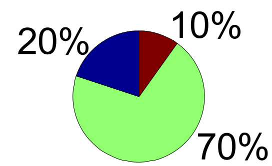

belief b: probability of current state

goal: minimize expected discounted sum of long-term costs

example: 3 states: undamaged, damaged, collapsed
3 actions: do nothing, inspect, repair.
4 observations

transition probability: "how the system evolves, depending on actions"

emission probability: "how physical state is related to observations"

belief: $b_k = P\left[S_k = i \middle| y_{1,\dots,k}\right]$

partially observable MDP

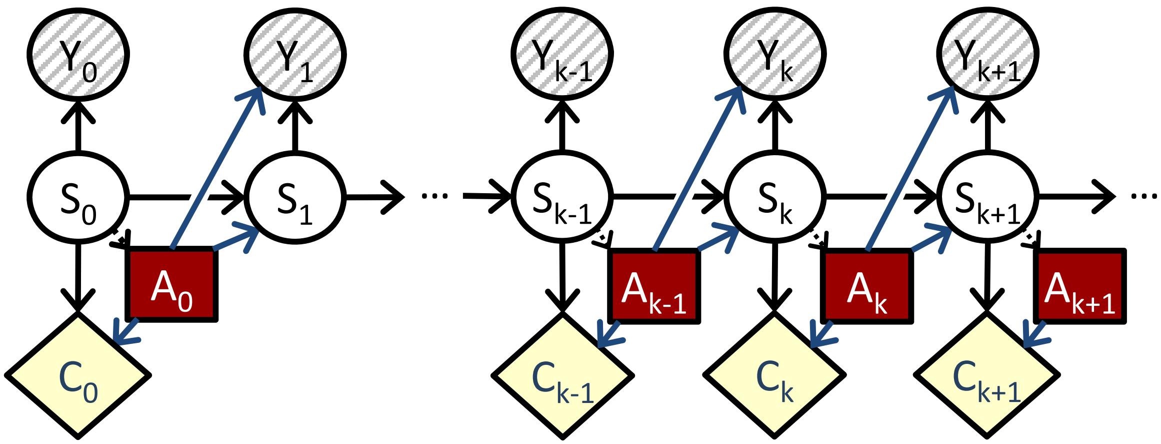

observations

state

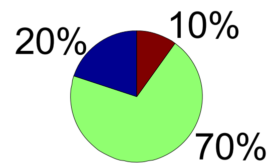

belief b: probability of current state

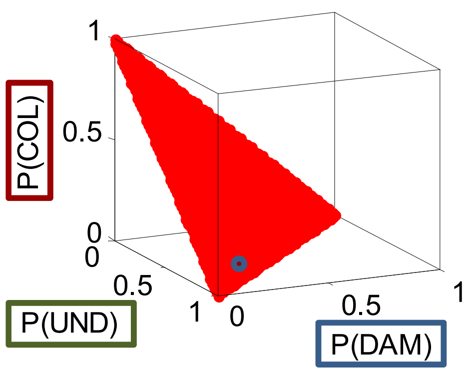

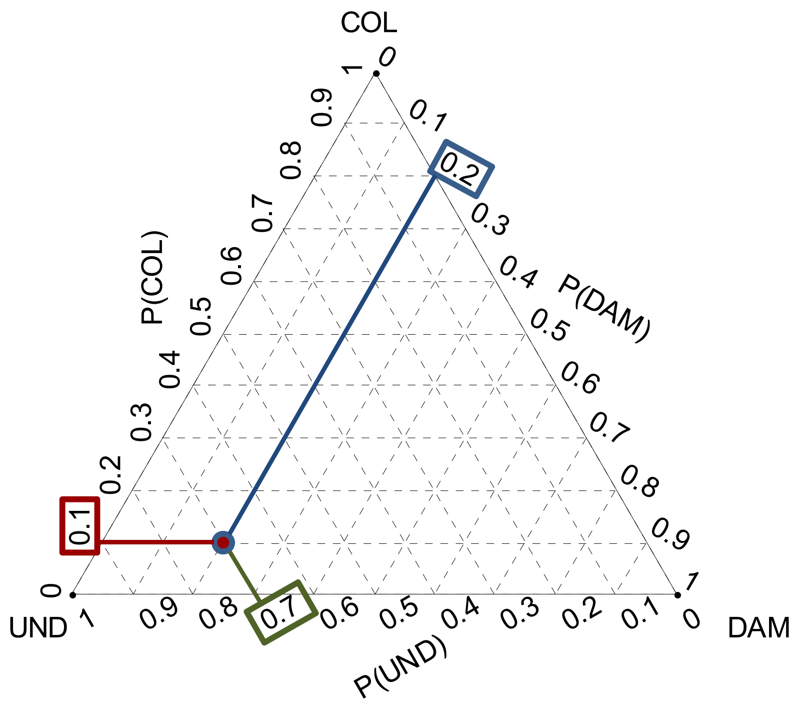

partially observable MDP

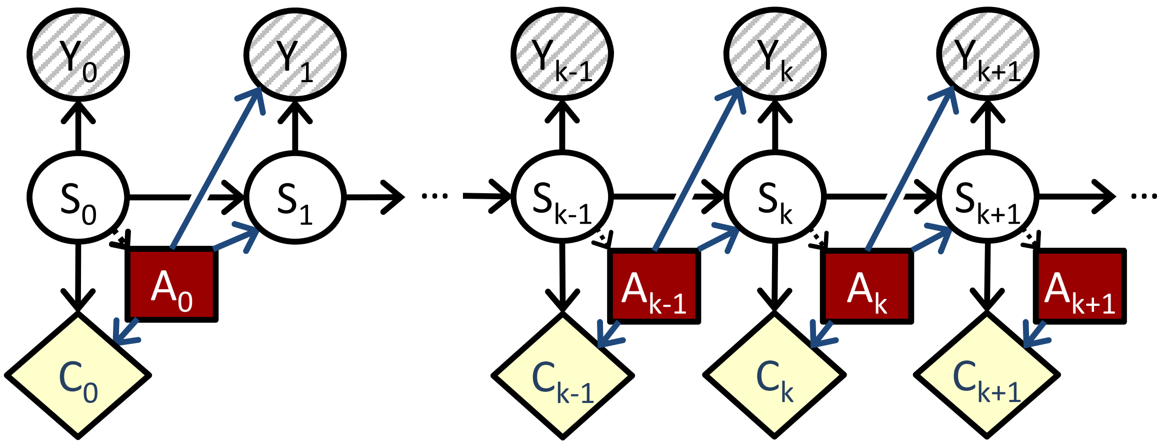

observations

state

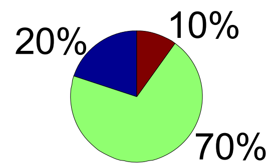

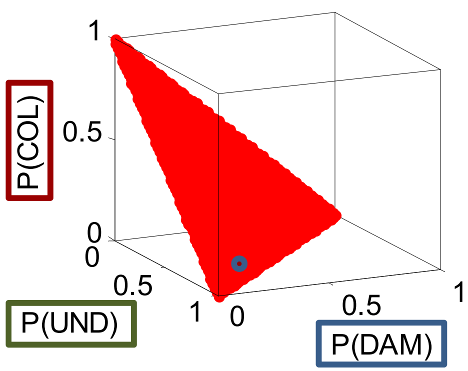

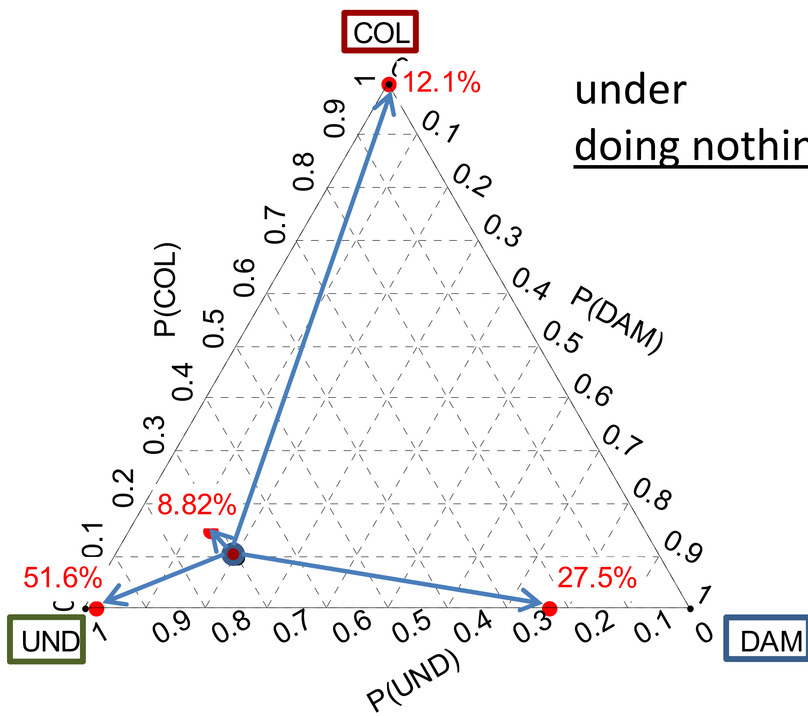

under
<u>doing nothing</u>

belief b:
probability of
current state

partially observable MDP

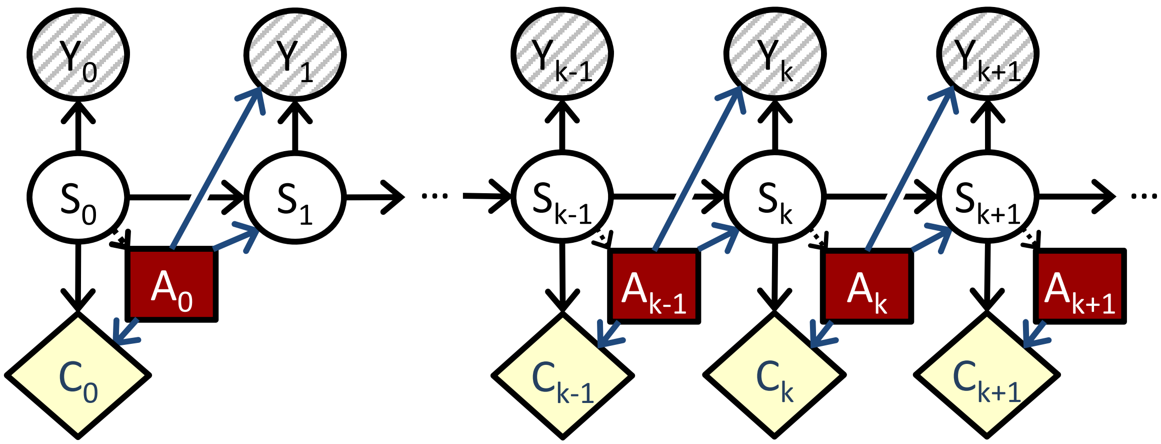

observations

state

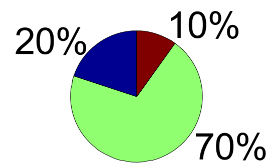

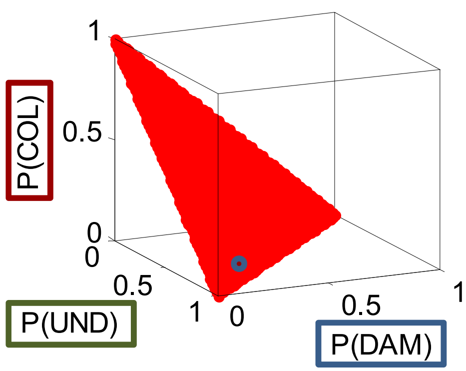

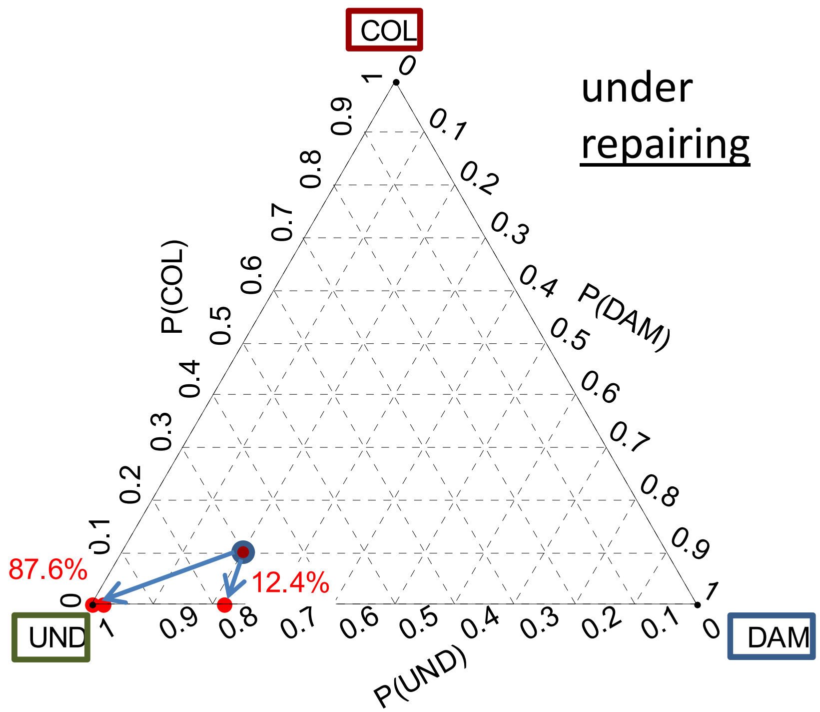

under

<u>repairing</u>

belief b:
probability of
current state

partially observable MDP

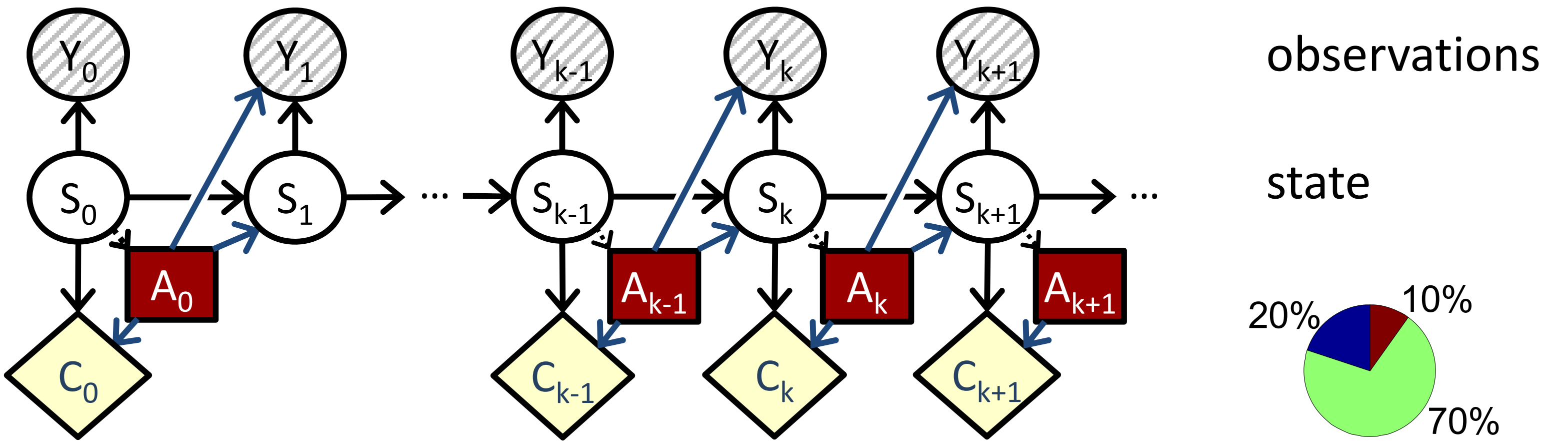

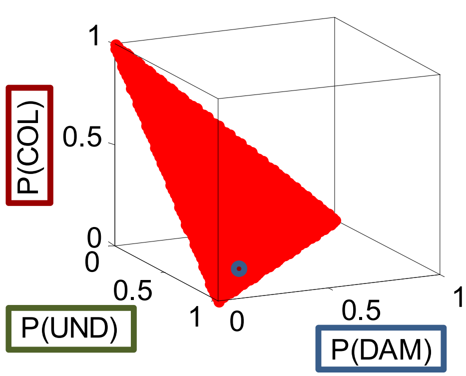

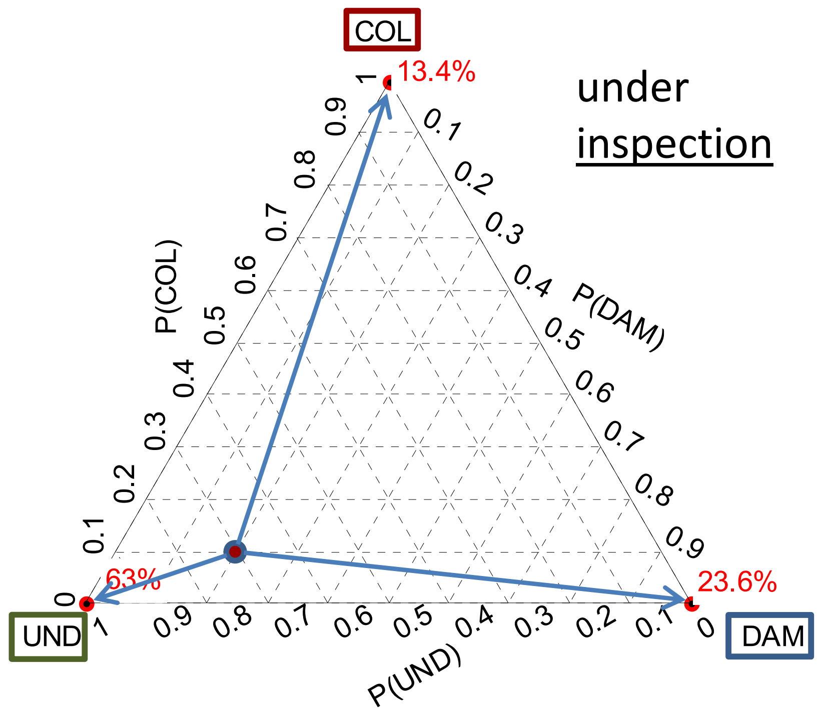

under <u>inspection</u>

belief b: probability of current state

partially observable MDP

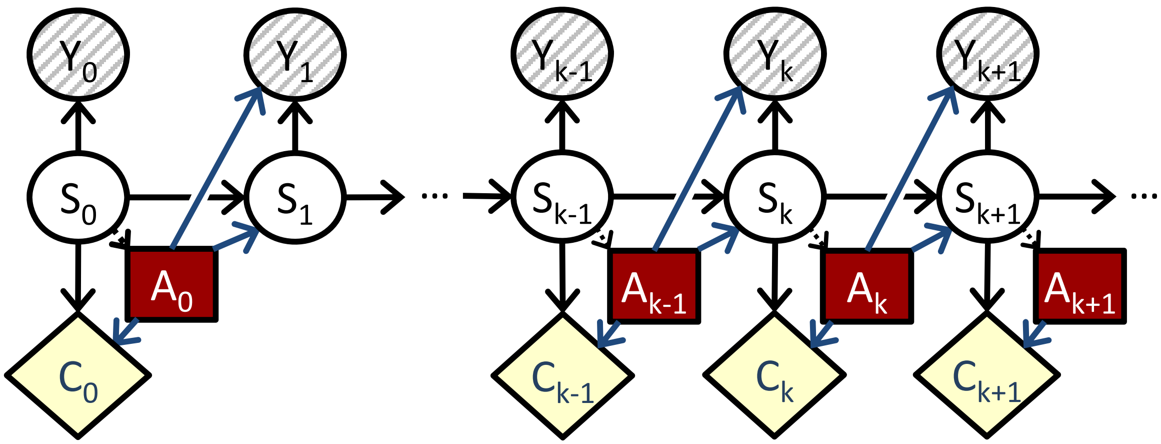

observations

state

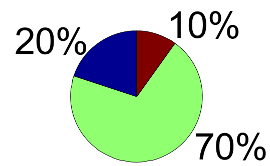

belief b: probability of current state

solving Bellman's equation: <u>optimal policy</u>

value (cumulative cost)

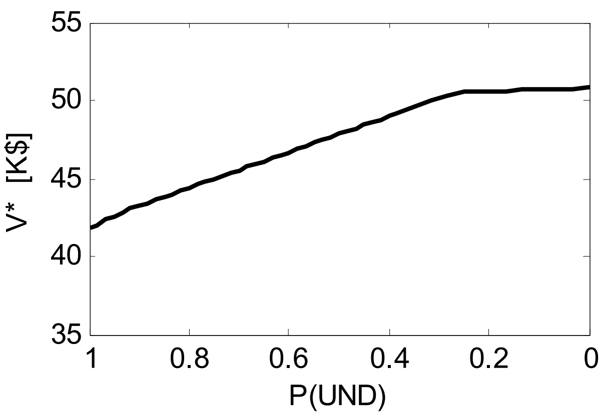

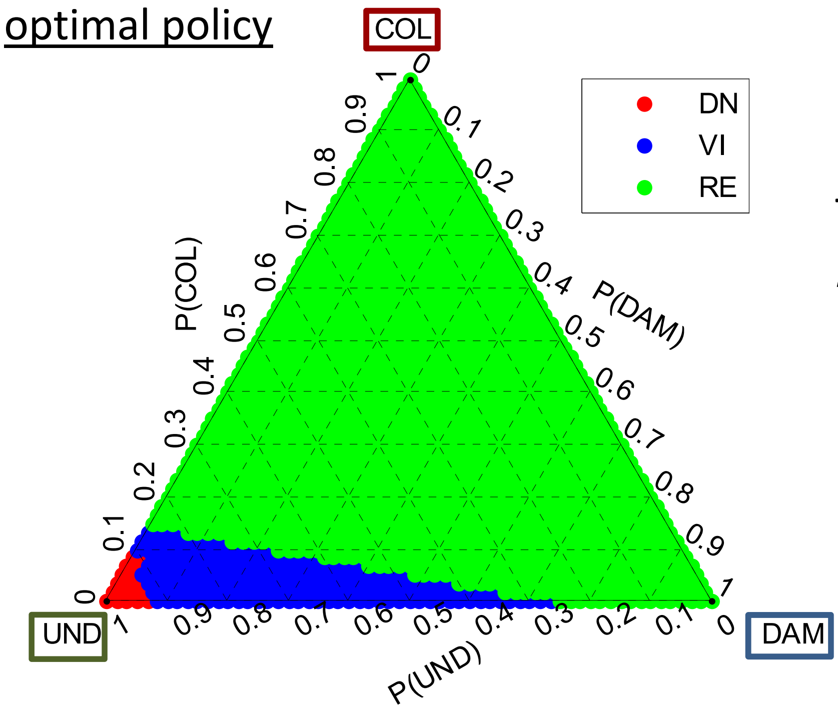

behavior: policy π as a function of the belief $b \longrightarrow A$: $A = \pi(b)$

optimal policy: it minimize the cumulative cost

partially observable MDP

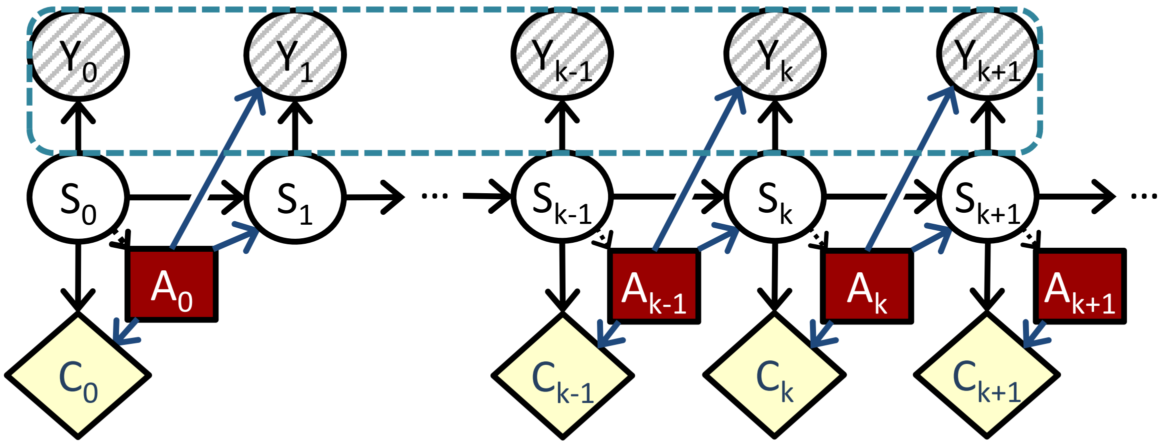

observations

state

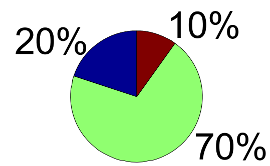

belief b: probability of current state

solving Bellman's equation: optimal policy

with more reliable observations

value (cumulative cost)

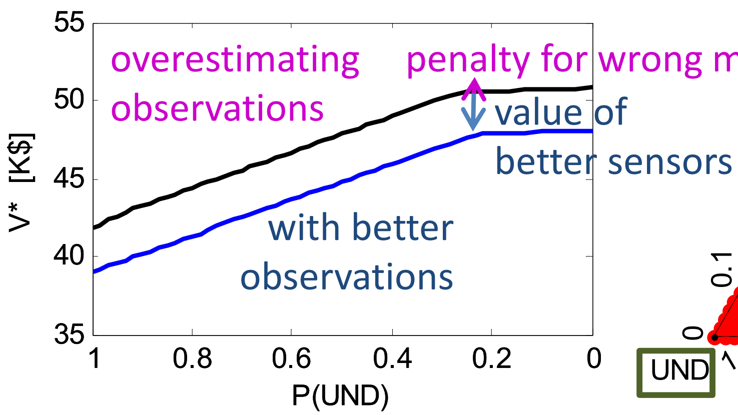

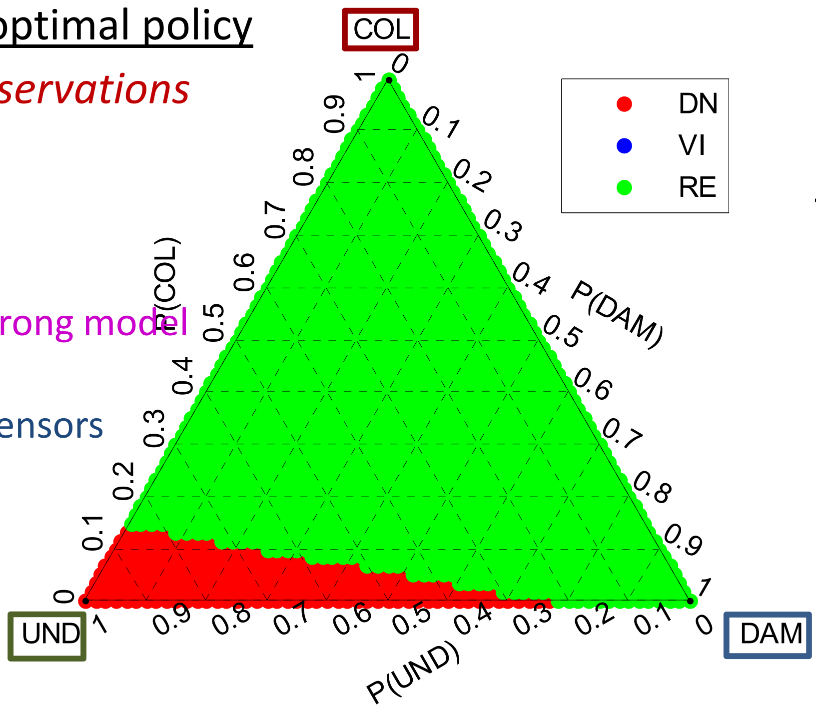

behavior: policy π as a function of the belief

$b \longrightarrow A$: $A = \pi(b)$

optimal policy:
it minimize the cumulative cost

solving POMDP: SARSOP

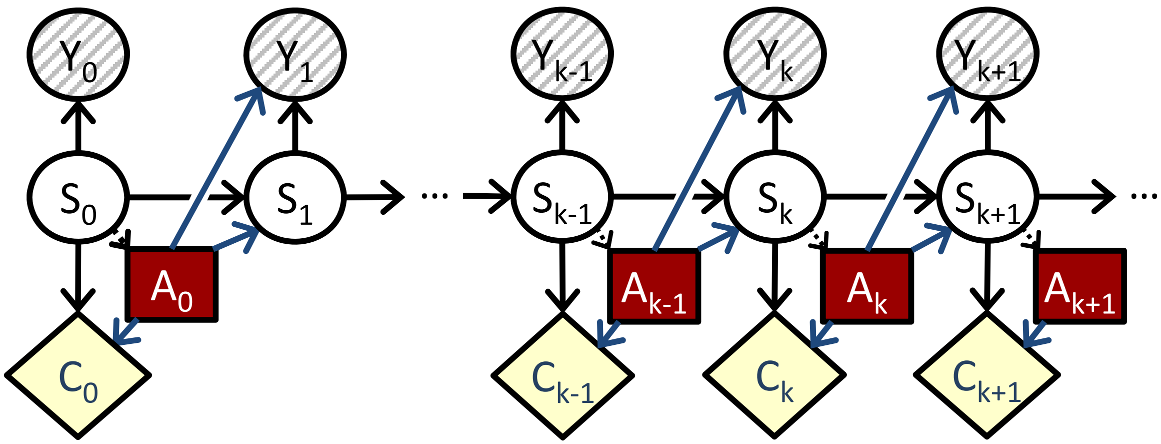

observations

state

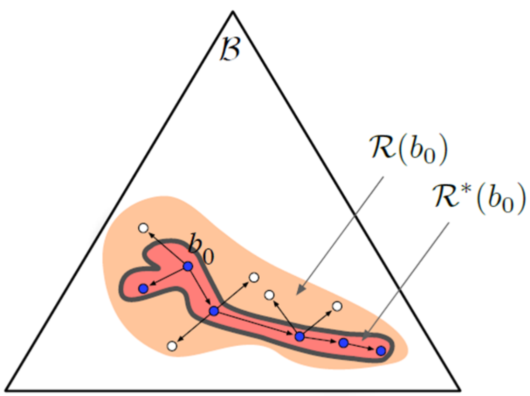

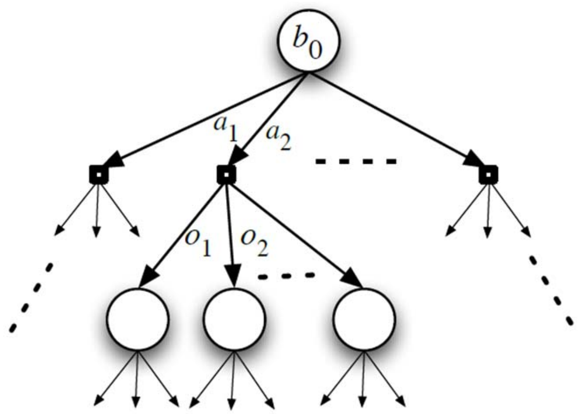

SARSOP software

successive approximations of the Reachable Space under Optimal Policies.

http://bigbird.comp.nus.edu.sg/pmwiki/farm/appl/

[pictures taken from:]

H Kurniawati, D Hsu, WS Lee, (2008), "SARSOP: Efficient Point-Based POMDP Planning by Approximating Optimally Reachable Belief Spaces." *Robotics: Science and Systems*

general setting for parametric analysis

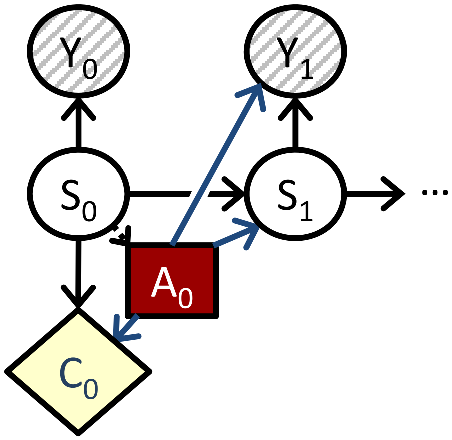

3 states:

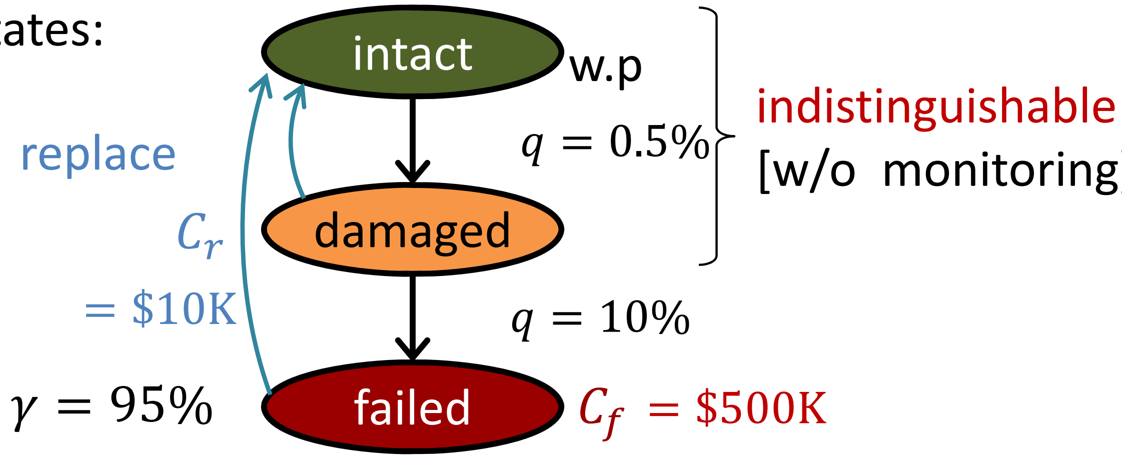

V_{OL}

$$VoI = V_{OL} - V_{CL}$$

V_{CL}

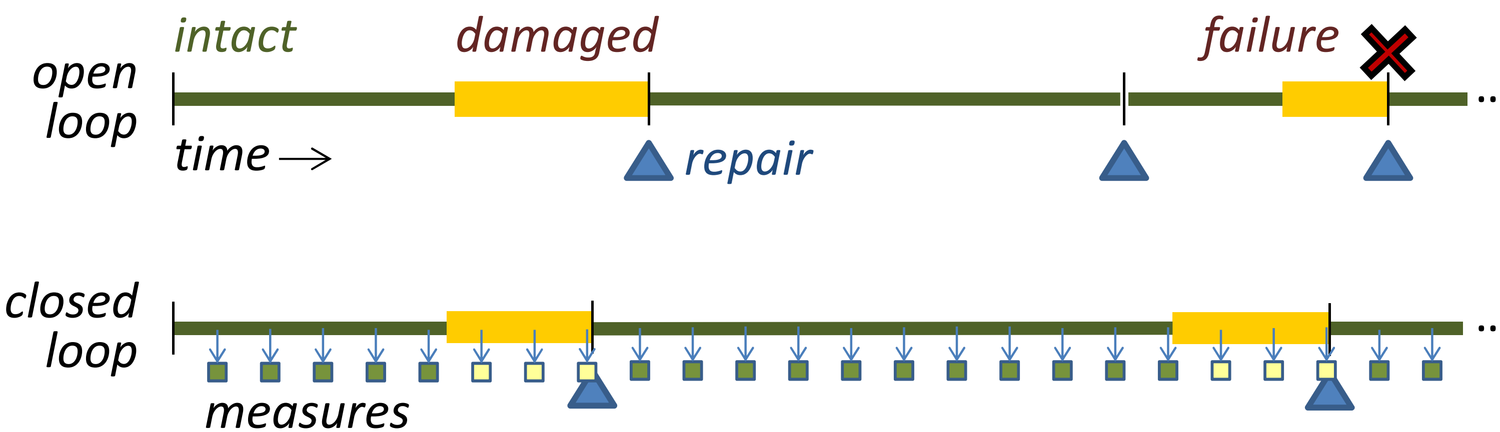

parameters: measure accuracy, measure availability, failure time predictability, repair cost, reaction time, discount factor.

sensitivity of VoI to availability

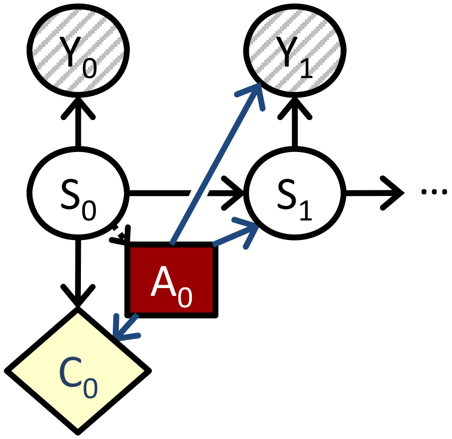

availability: "how frequently you measure"

VoI is monotonically increasing with measure availability.

In this setting, VoI grows faster for low availability and slower for high availability [submodularity]

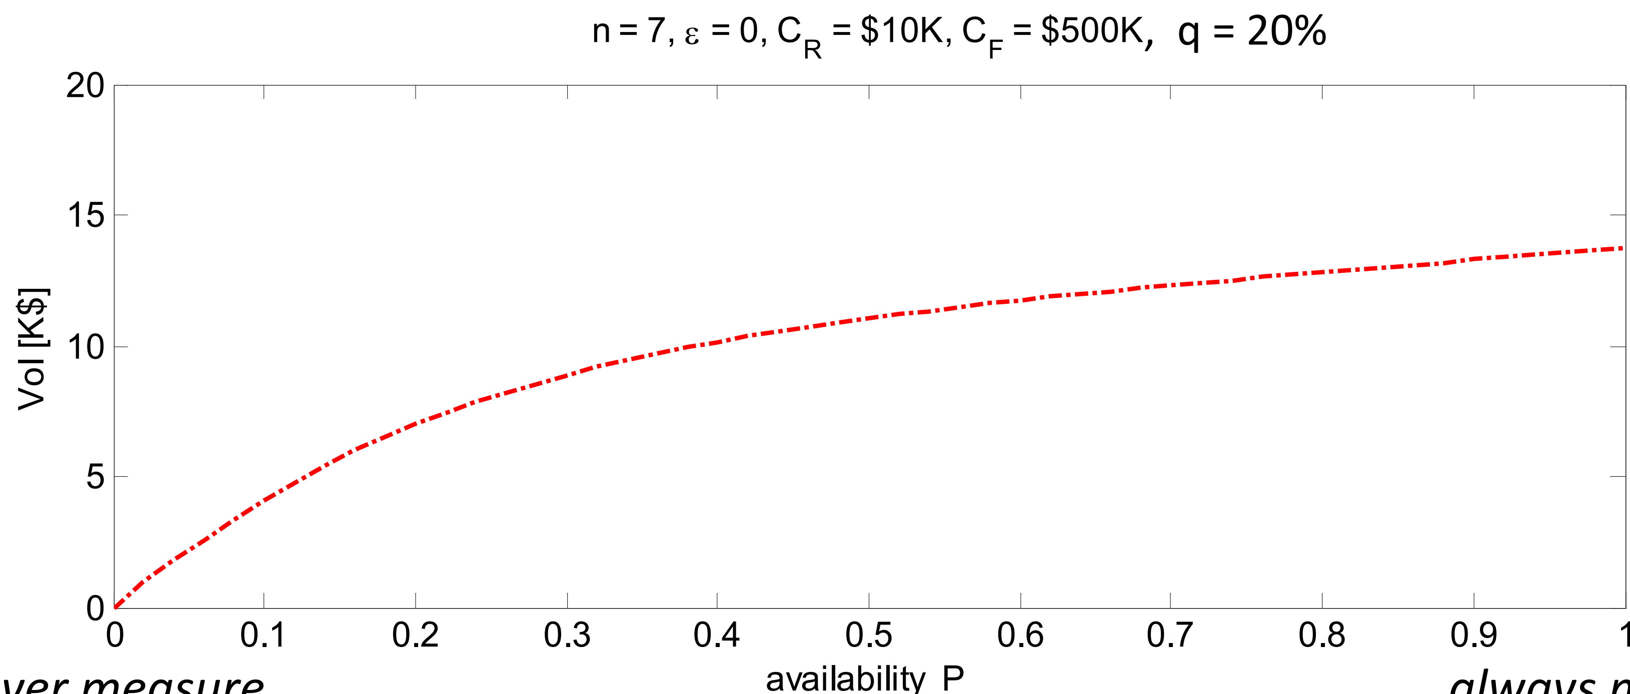

never measure

always measure

sensitivity of VoI to availability

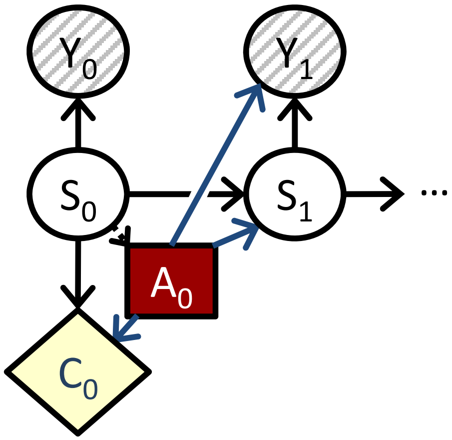

availability: "how frequently you measure"

VoI is monotonically increasing with measure availability.

In this setting, VoI grows faster for low availability and slower for high availability [submodularity].

e.g., if degradation is slow, the additional benefit of monitoring more than 40% of the time is negligible.

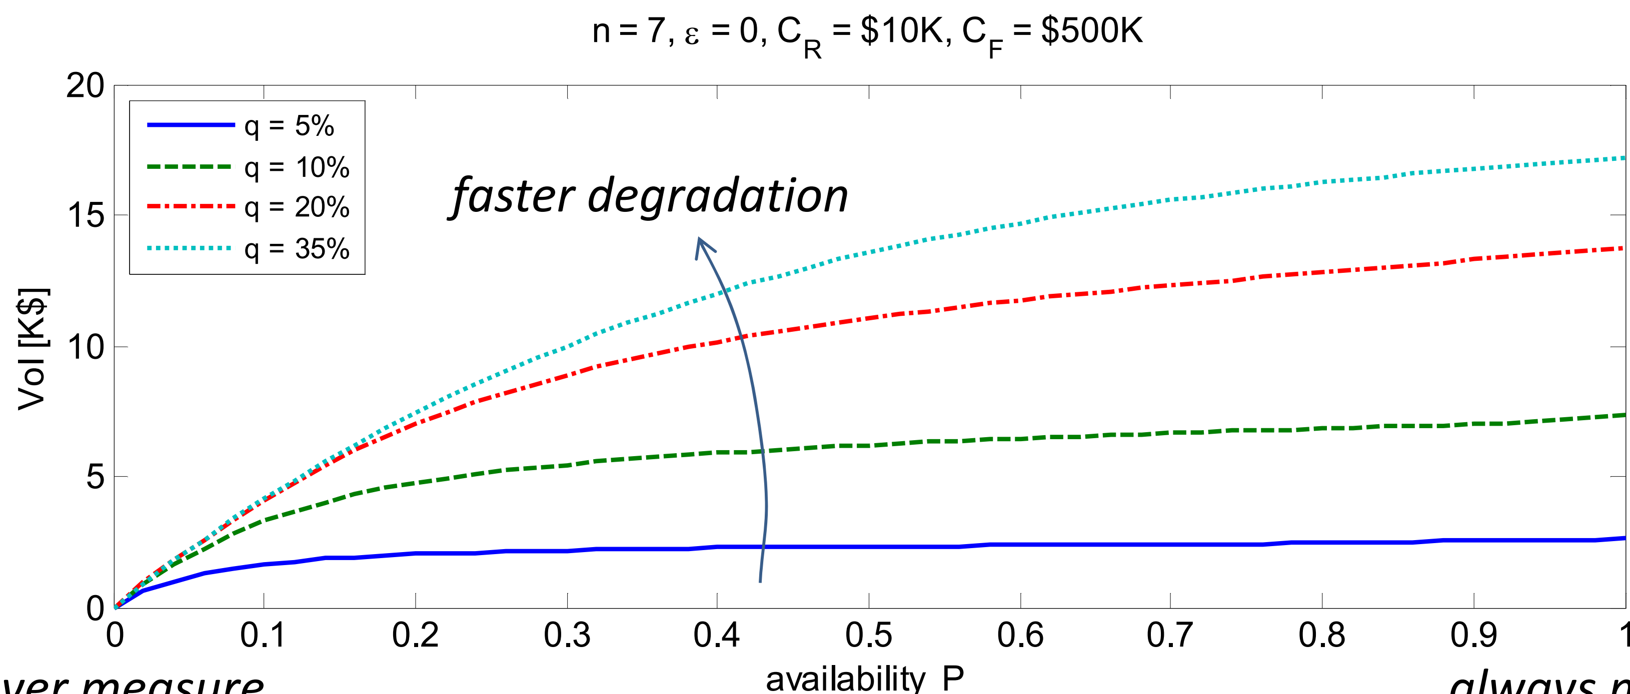

never measure *always measure*

sensitivity of VoI to inaccuracy

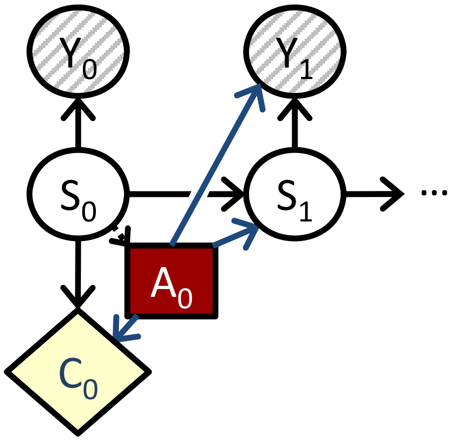

inaccuracy: "probability of incorrect detection"

VoI is monotonically decreasing with inaccuracy.

For zero inaccuracy, state observation is perfect.

Too inaccurate measures are useless.

This graph allows for comparing pair availability/inaccuracy.

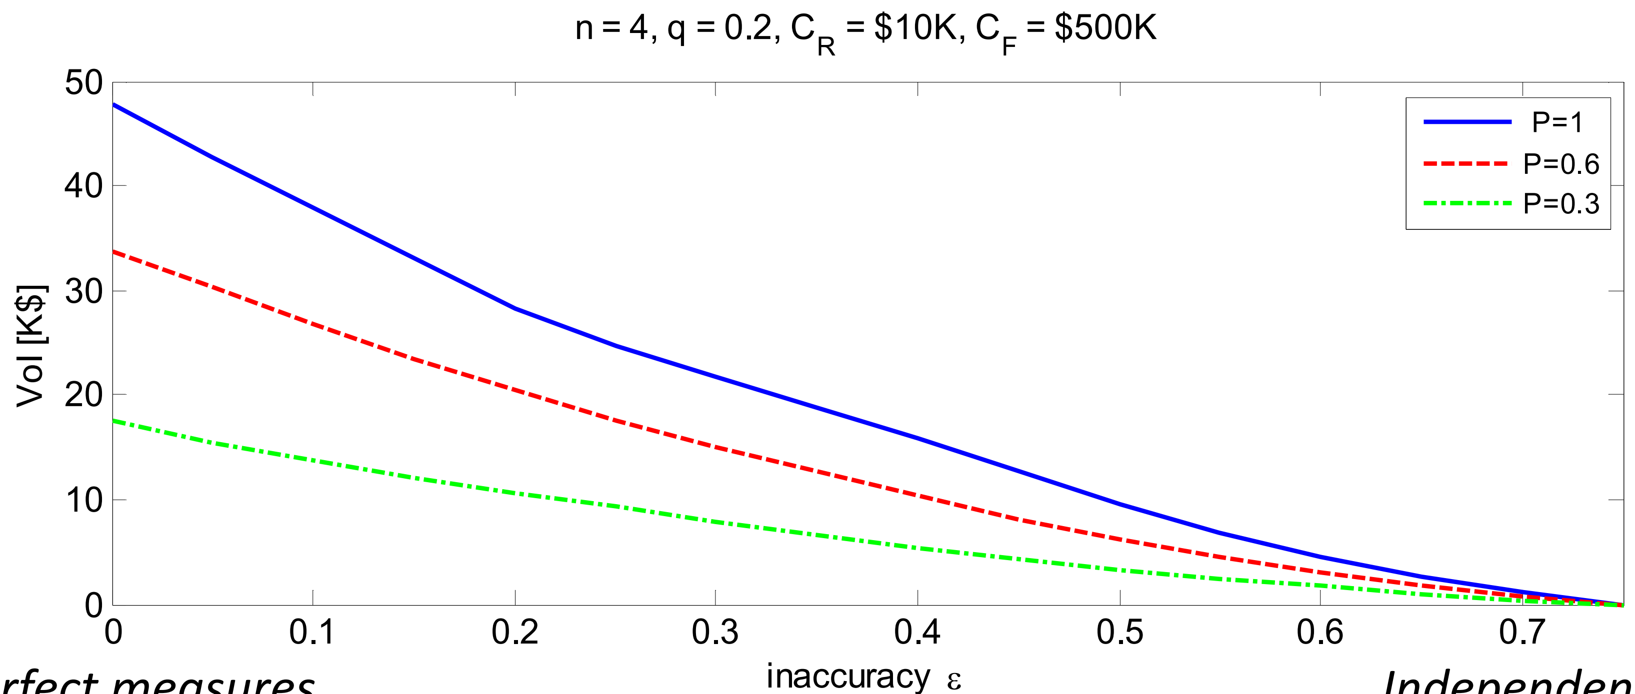

perfect measures *Independent measures*

sensitivity of VoI to unpredictability

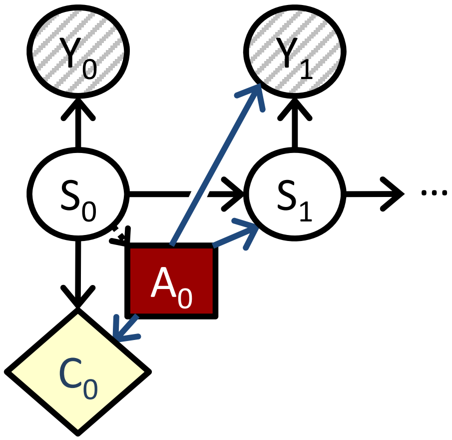

unpredictability: "expected prior error in guessing the time of failure"

Here VoI is monotonically increasing with unpredictability.

If the degradation process can be well predicted even without the monitoring support, the VoI is low.

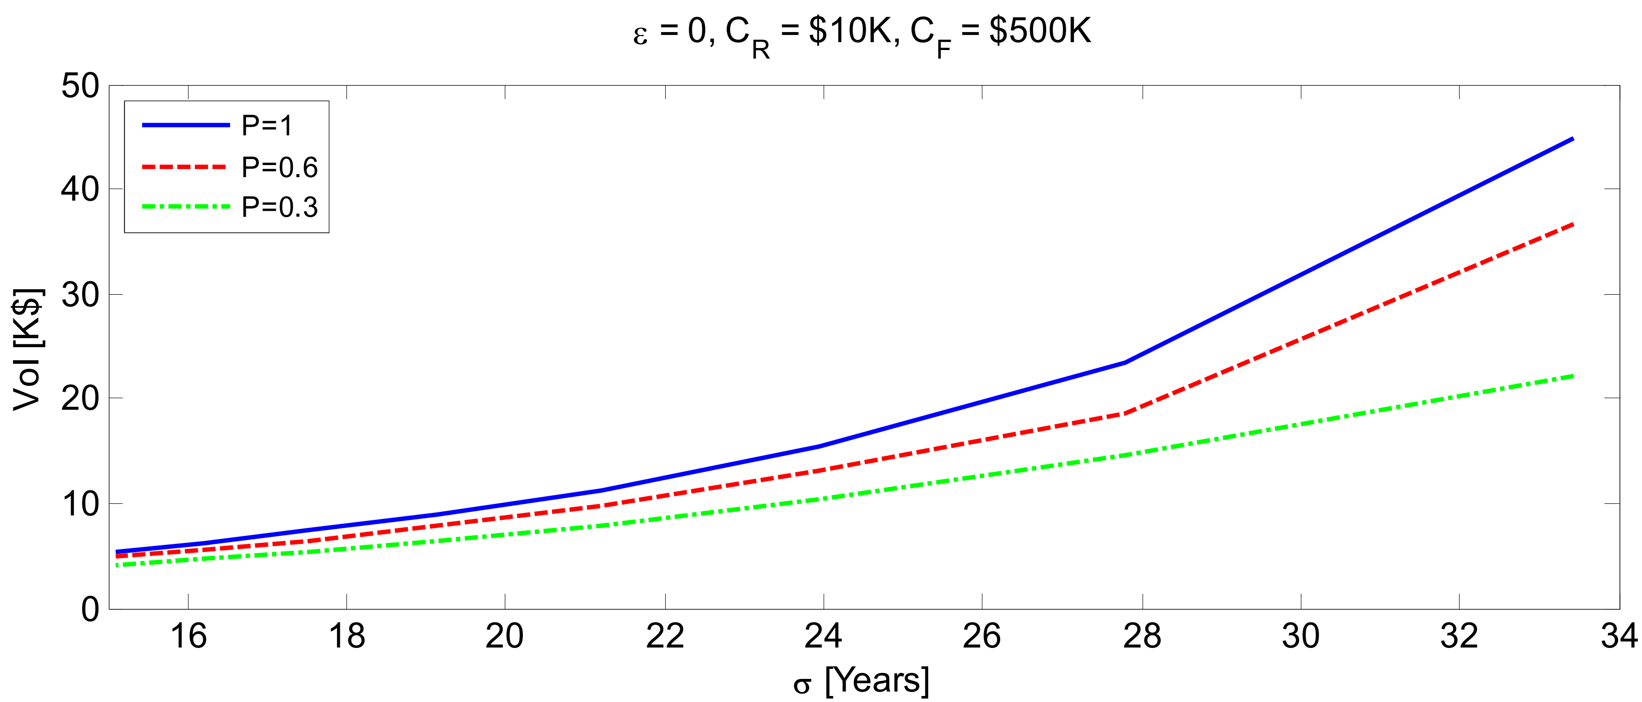

sensitivity of VoI to reaction time

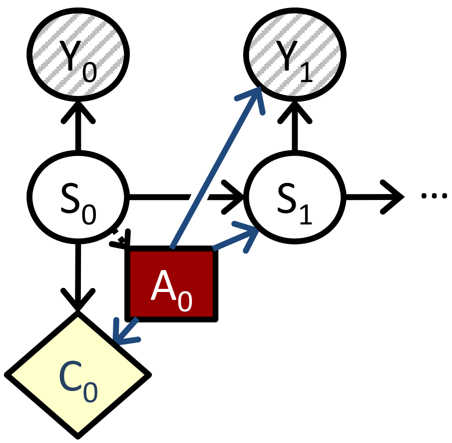

reaction time: "how many steps are needed for implementing a repair"

VoI is monotonically decreasing with reaction time.

Higher reaction times pose stronger constraints in using the information the sensors provides.

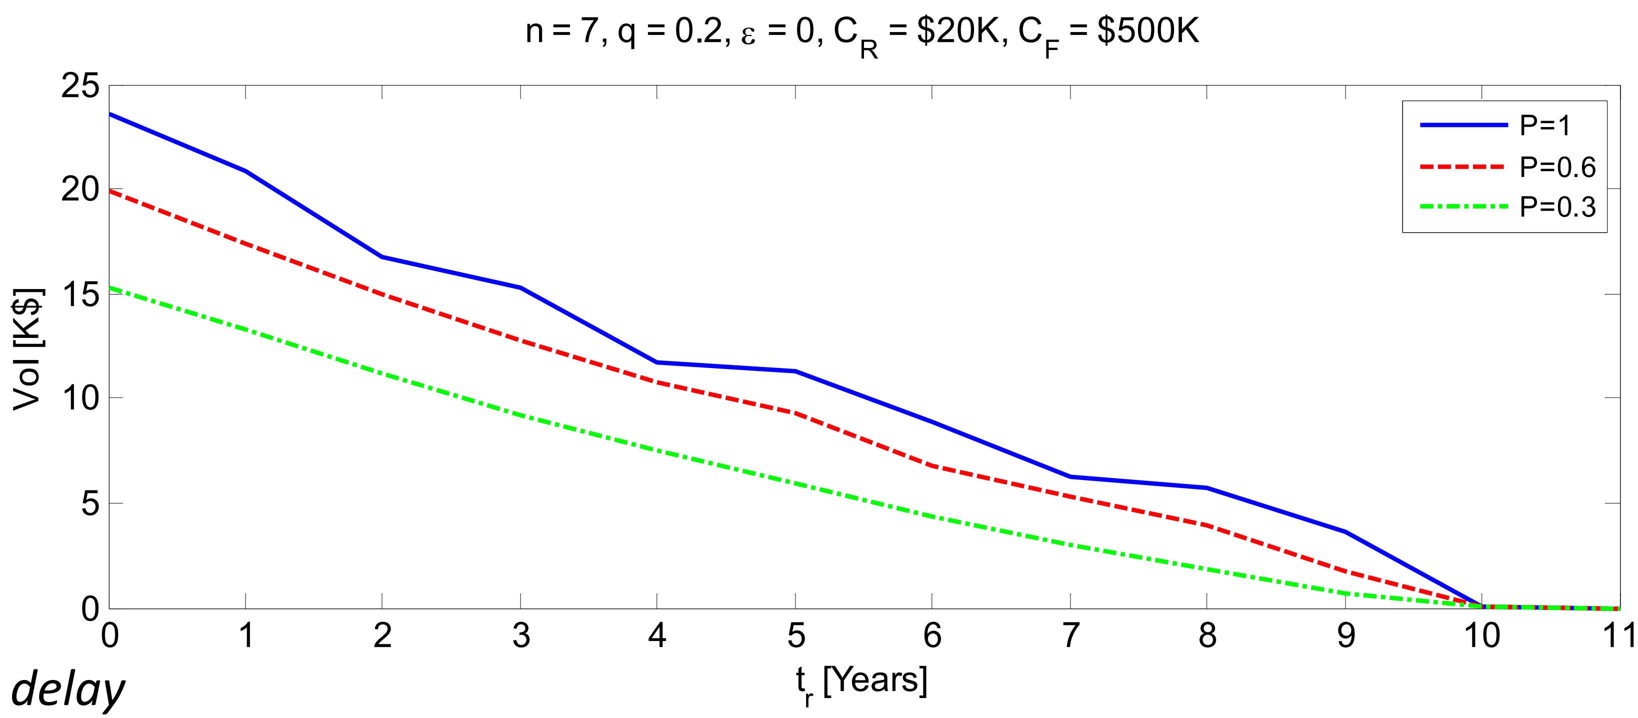

no delay

sensitivity of VoI to repair cost

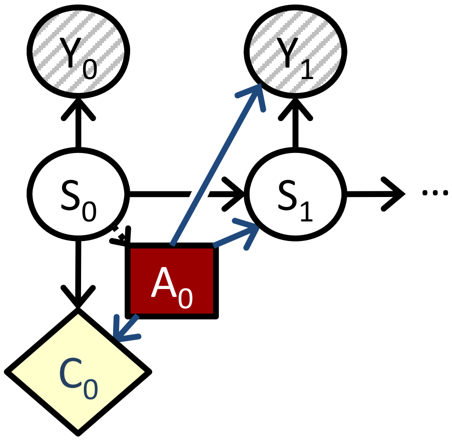

VoI is monotonically decreasing with reaction time.

It cannot be monotonic:

it needs to be zero when the cost of repair is zero and when it is infinite.

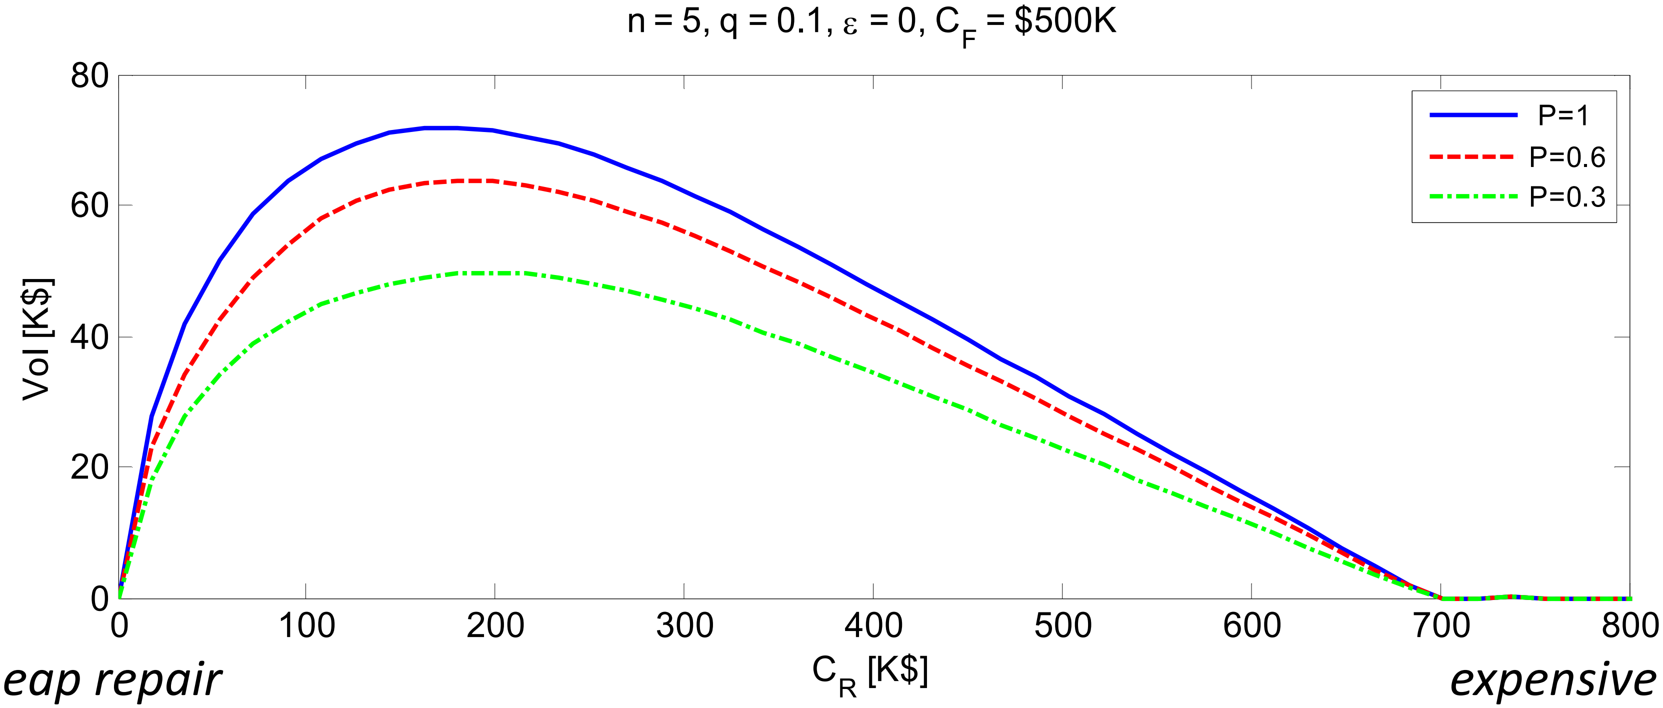

cheap repair *expensive repair*

sensitivity of VoI to discount factor

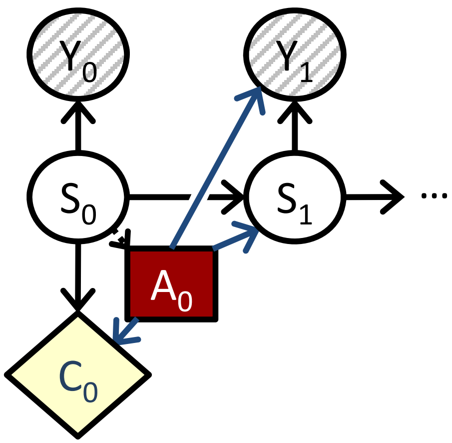

discount factor: "value of a dollar at next step"

VoI is monotonically increasing with discount factor. [?]

As initial state is intact, the VoI is negligible when γ<60%.

For factor going to one, the value goes to infinite.

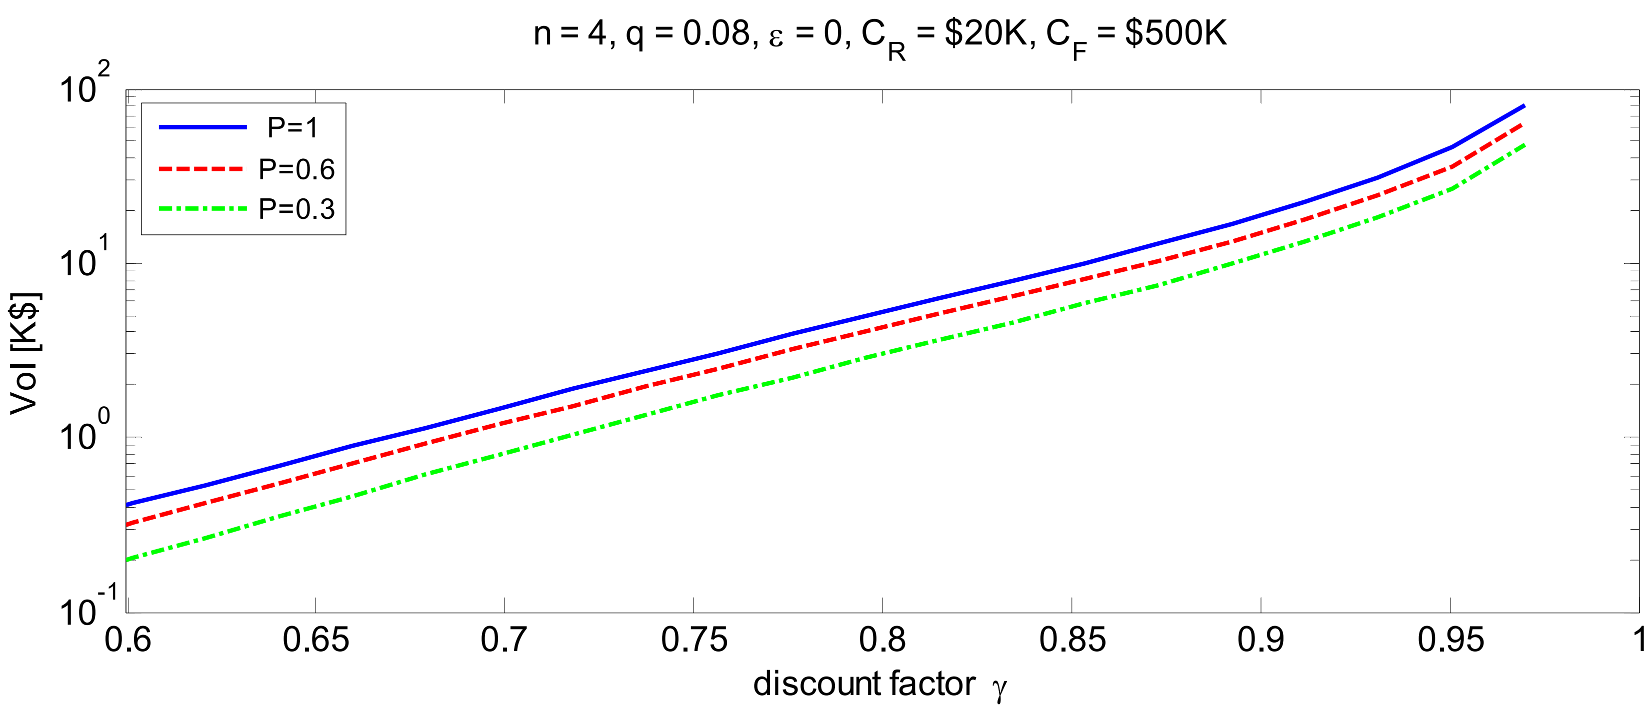

how to allocate inspectors across components?

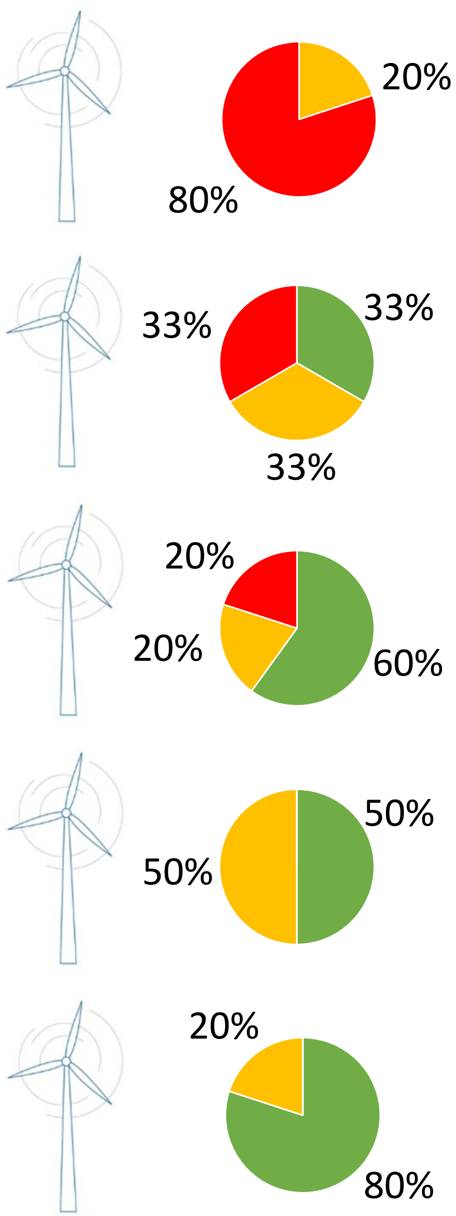

prob. of failure	entropy	VoI (pessimistic)
80%	0.5	0
33%	1.1	2.66
20%	0.95	4.79
0	0.69	1.52
0	0.5	0.897

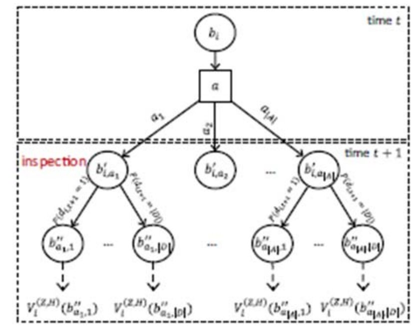

VoI incorporates risk and uncertainty in a consistent way.

it is the best metric for minimizing the long term cost.

a system made by parallel POMDPs

POMDP
turbine 1

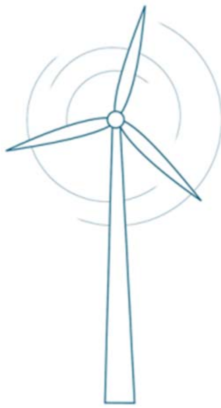

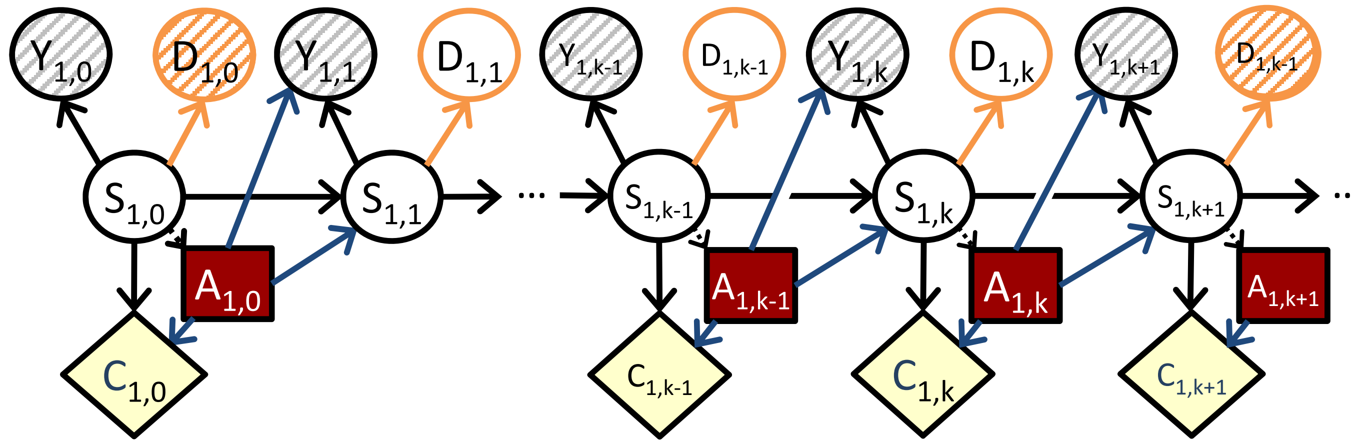

independent

POMDP
turbine n

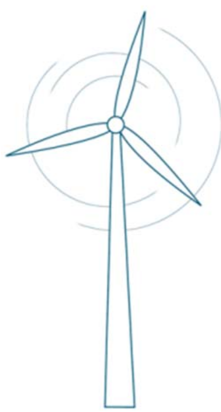

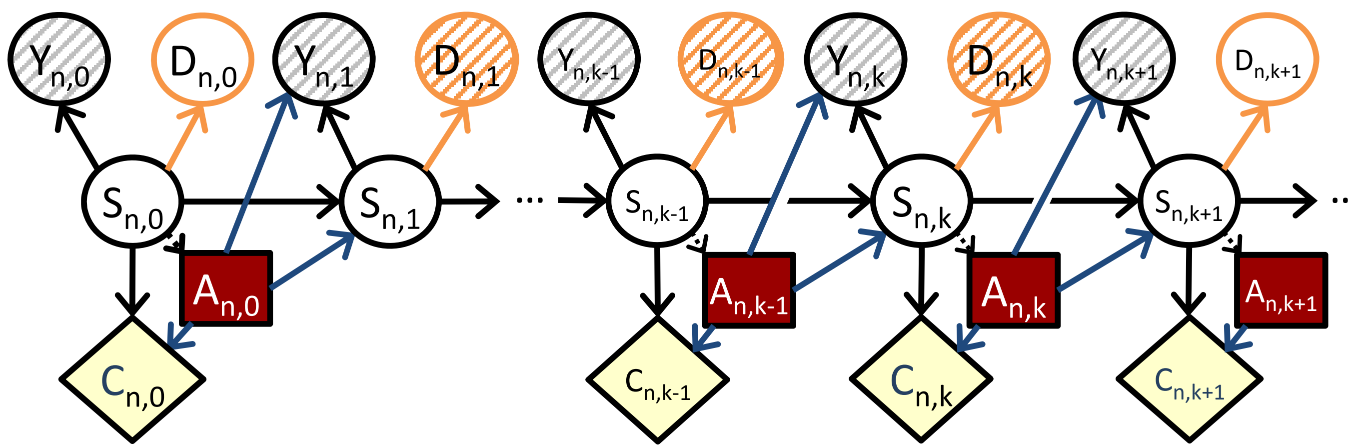

multi-state example: with Inspectors

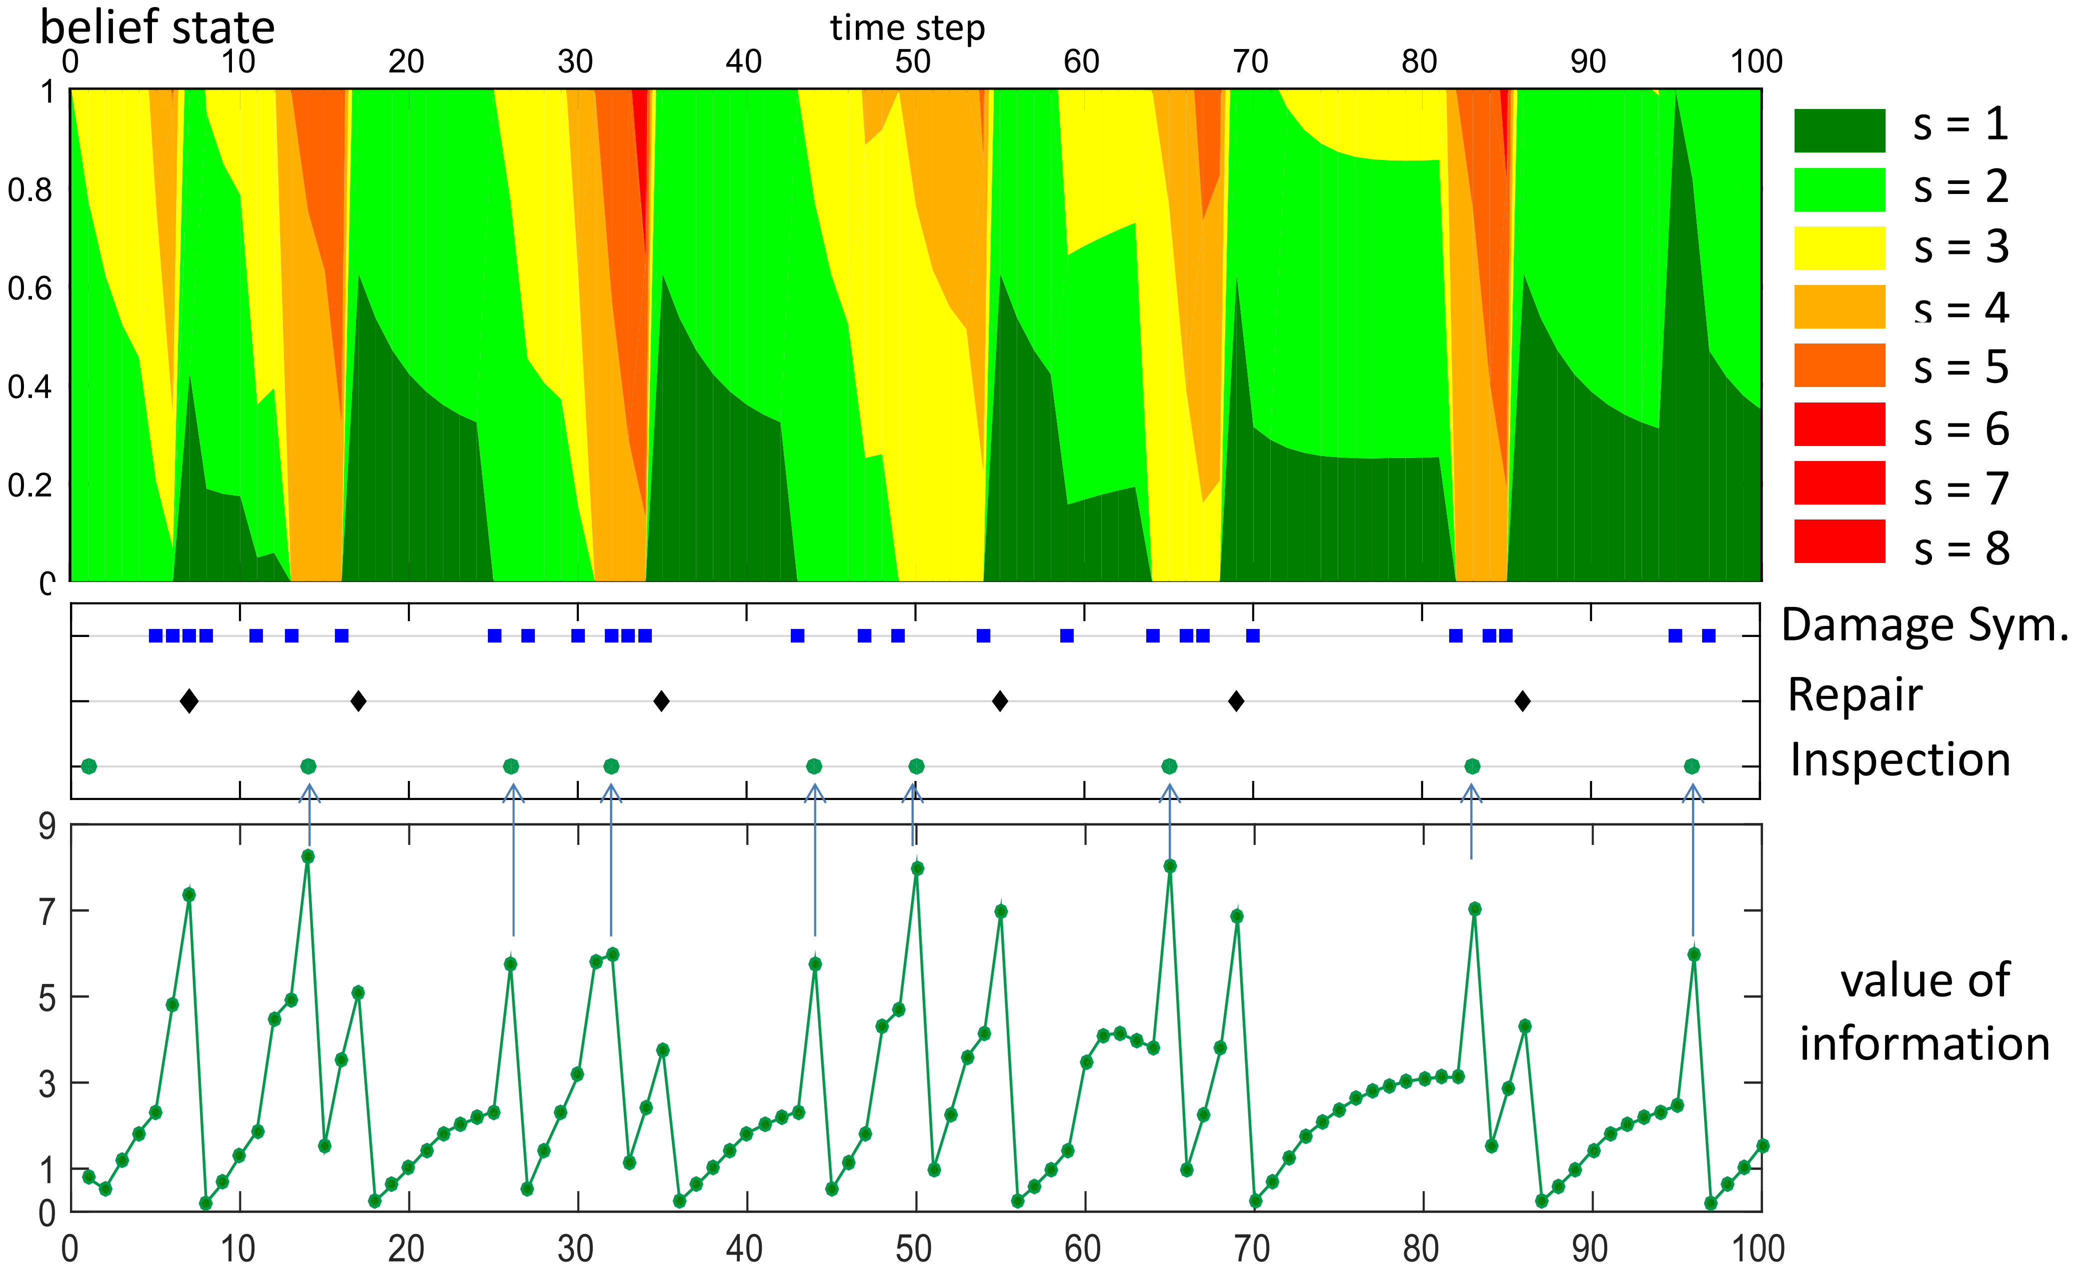

multi-state example: with Inspectors

impact of inspectors in the integrated management

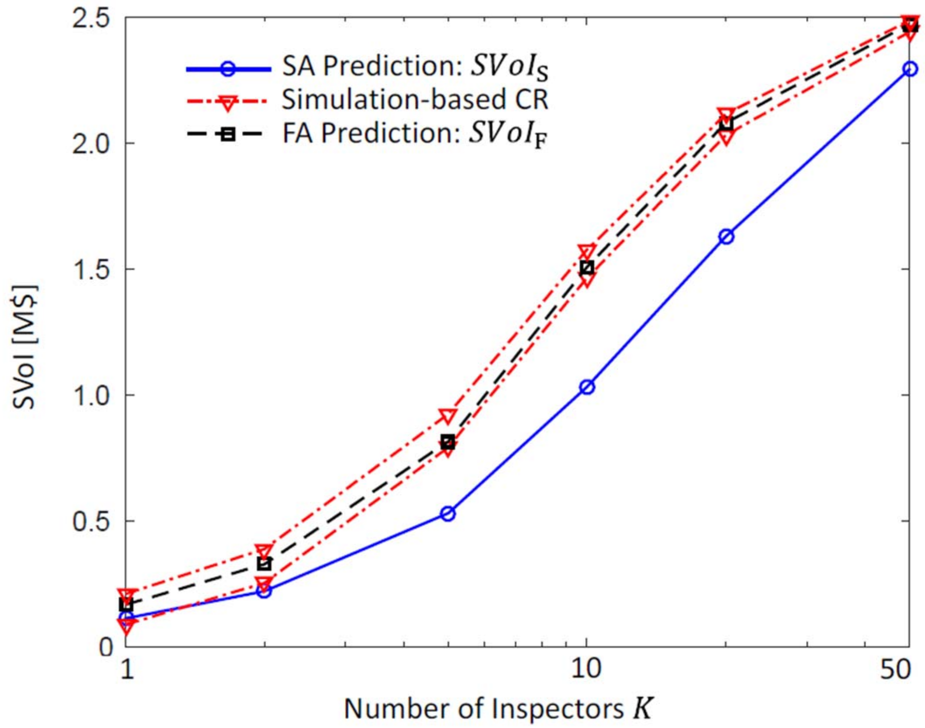

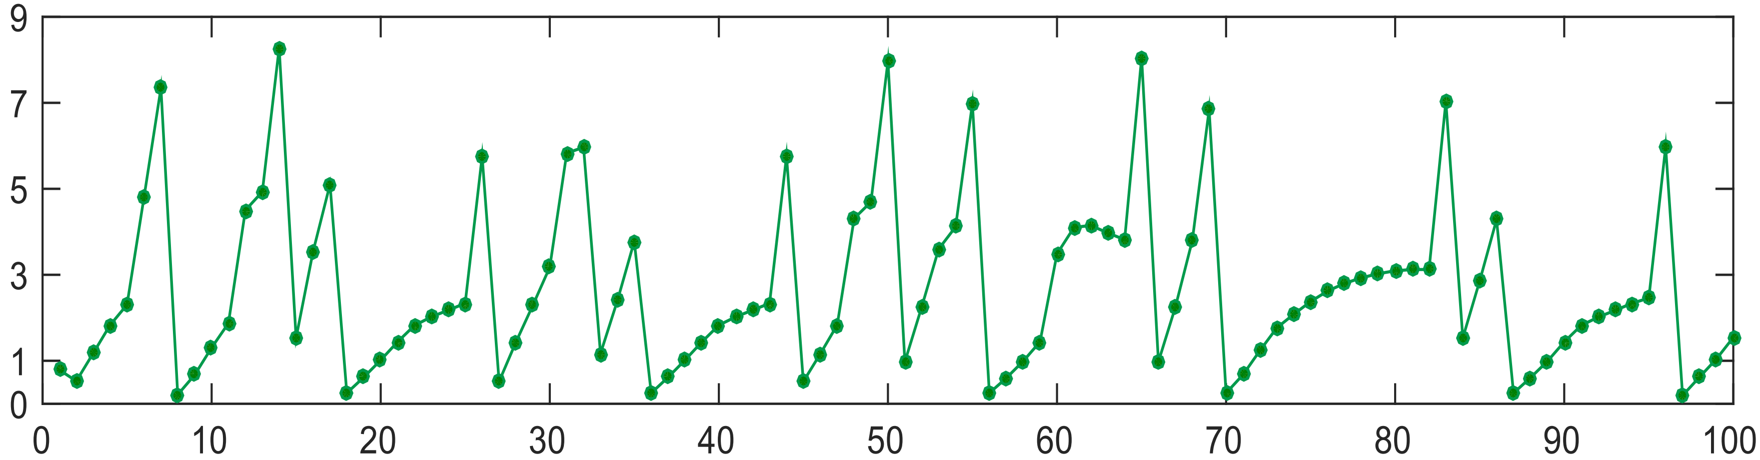

value of information

summary on VoI in sequential decision making

The Value of Information is a concept related to pre-posterior analysis, i.e. to decision theory and Bayesian analysis.

In sequential decision making, VoI can be computed by differentiating the Values in POMDPs.

VoI can be used for evaluating the impact of long-term monitoring.

Also, for inspection scheduling but, at system level, approximating the interaction between current and future observations, for avoiding curse of dimensionality.

To include model uncertainty into the VoI analysis is a challenging task still to be covered…

Memarzadeh, M., Pozzi, M. "Integrated inspection scheduling and maintenance planning for infrastructure systems," Computer-Aided Civil and Infrastructure Engineering (Wiley) DOI: 10.1111/mice.12178 (2015).

Memarzadeh, M., Pozzi, M. "Value of Information in Sequential Decision Making: component inspection, permanent monitoring and system-level scheduling," submitted to Reliability Engineering & System Safety.

Malings, C., Pozzi, M. "Value of Information for Spatially Distributed Systems: application to sensor placement," submitted to Reliability Engineering & System Safety.

information avoidance

VoI is guarantee to be non-negative.

Suppose society (say a building code) assigns a policy, unwelcomed by the agent:

An agent forced to repair,
may find convenient to buy info
to avoid unnecessary repairing.

But now consider another case:

This agent can avoid repairing,
but information can expose her
to the risk of repairing.
AVOID (FREE) INFORMATION ?

simplest maintenance problem: imperfect info

state F: Failure
U: Undamaged

sensor outcome

$P_{FA} = P(\mathrm{A}|\mathrm{U})$

$P_{FS} = P(\mathrm{S}|\mathrm{F})$

S : Silence
A : Alarm

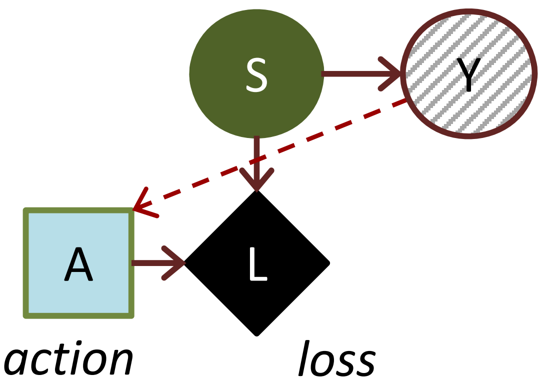

N: do Nothing
R: Repair

$L(S,A)$

		(S)	
		U	F
(A)	N	0	L_F
	R	L_R	L_R

agent's loss matrix

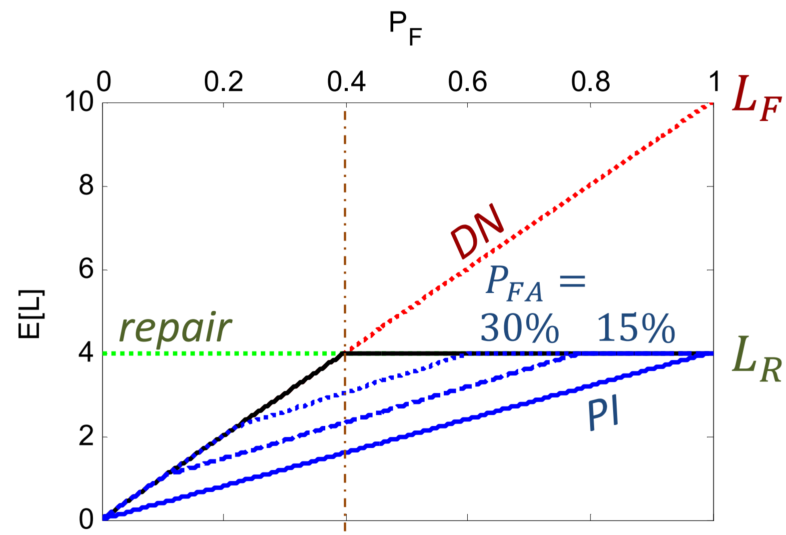

opt. threshold: $P_L = L_R/L_F$

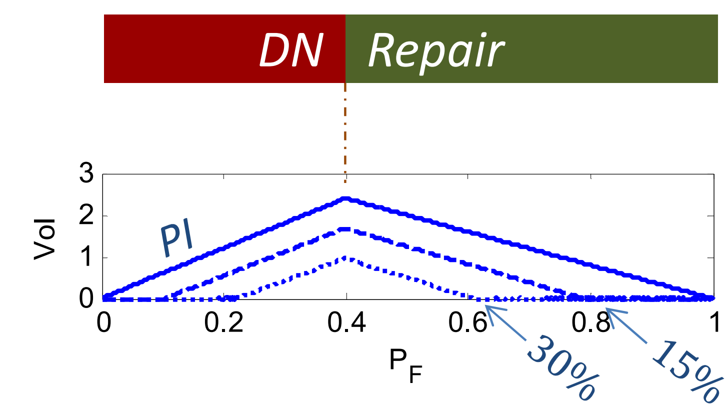

simplest maintenance problem

state F: Failure
U: Undamaged

sensor outcome $P_{FA} = P(A|U)$
$P_{FS} = P(S|F)$

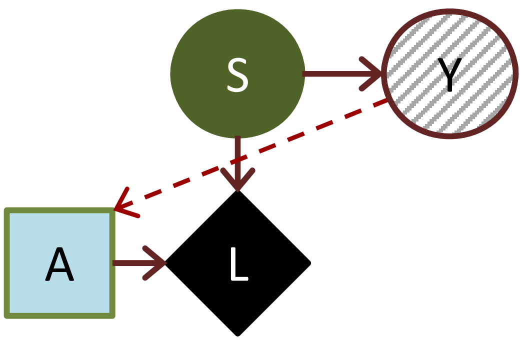

action *loss*

S : Silence
A : Alarm

N: do Nothing
R: Repair

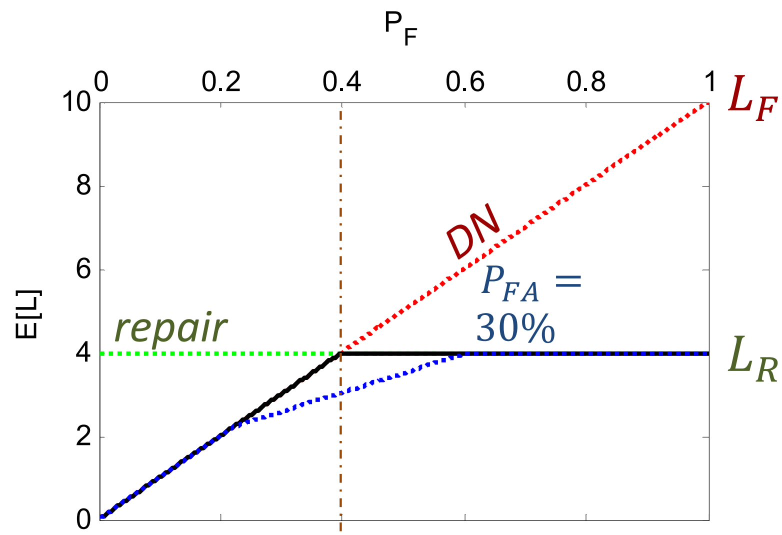

opt. threshold: $P_L = L_R / L_F$

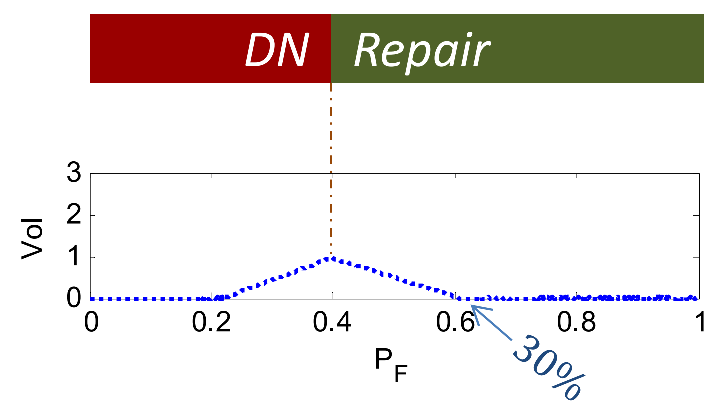

$L(S, A)$ (S)

		U	F
(A)	N	0	L_F
	R	L_R	L_R

agent's loss matrix

simplest maintenance problem

state F: Failure
U: Undamaged

sensor outcome $P_{FA} = P(A|U)$
$P_{FS} = P(S|F)$

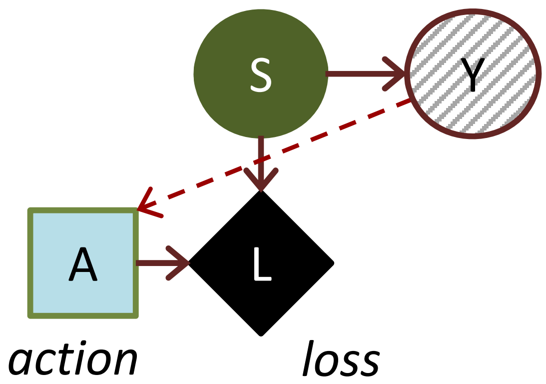

S : Silence
A : Alarm

N: do Nothing
R: Repair

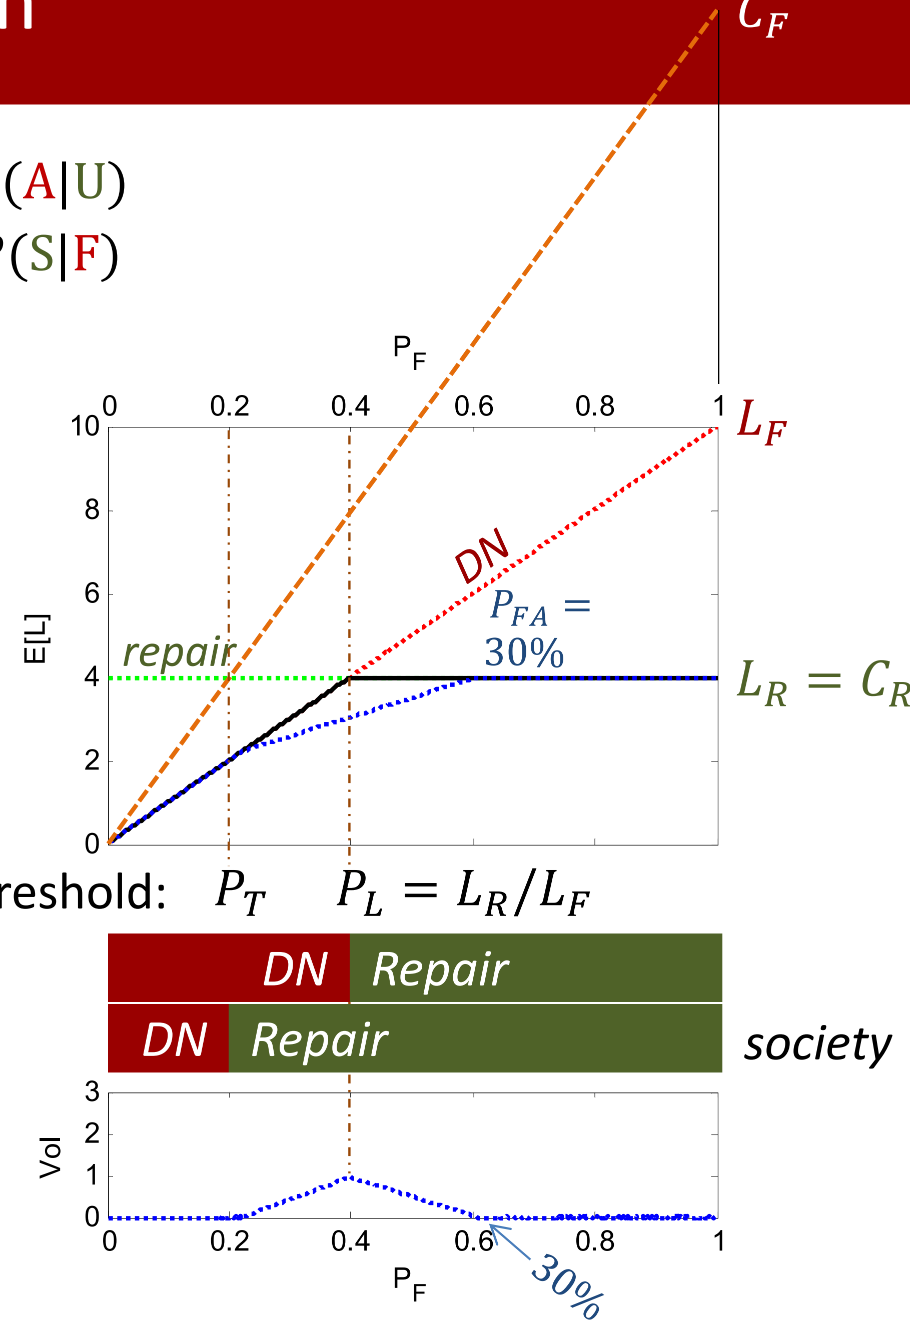

$L(S,A)$

(A) \ (S)	U	F
N	0	L_F
R	L_R	L_R

agent's loss matrix

$C(S,A)$

(A) \ (S)	U	F
N	0	C_F
R	C_R	C_R

society's cost matrix

simplest maintenance problem

state F: Failure
U: Undamaged

sensor outcome $P_{FA} = P(A|U)$
$P_{FS} = P(S|F)$

S : Silence
A : Alarm

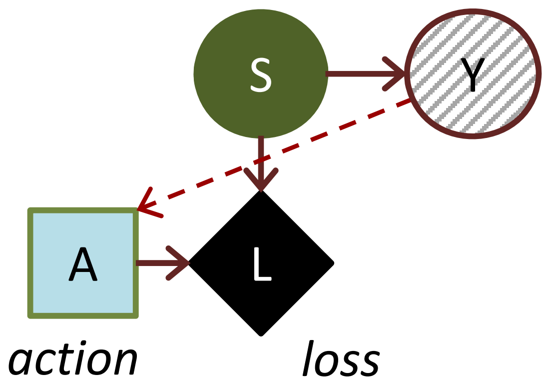

N: do Nothing
R: Repair

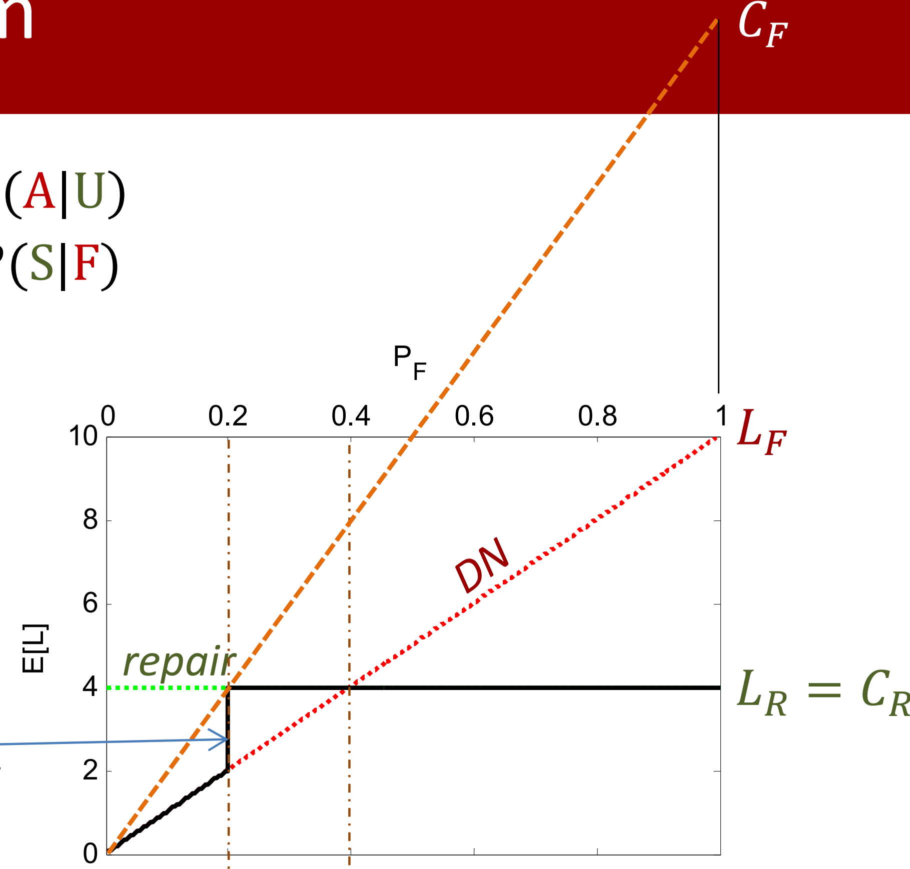

discontinuity in expected cost

opt. threshold: P_T $P_L = L_R / L_F$

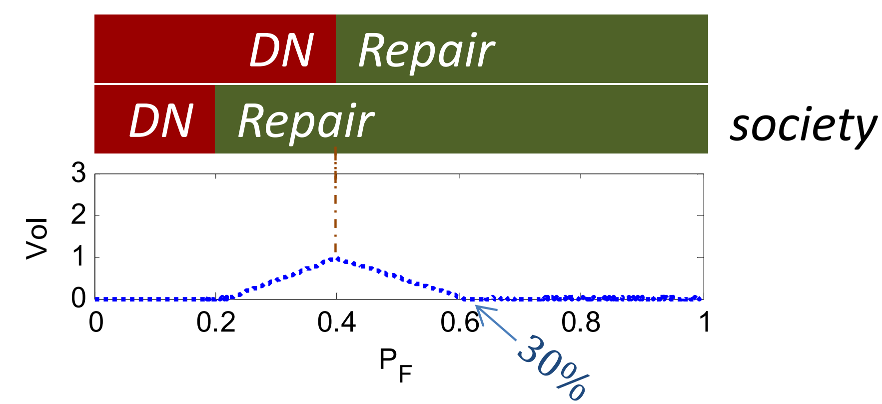

$L(S,A)$

(A) \ (S)	U	F
N	0	L_F
R	L_R	L_R

agent's loss matrix

$C(S,A)$

(A) \ (S)	U	F
N	0	C_F
R	C_R	C_R

society's cost matrix

simplest maintenance problem

state F: Failure, U: Undamaged

sensor outcome $P_{FA} = P(A|U)$, $P_{FS} = P(S|F)$

S : Silence
A : Alarm

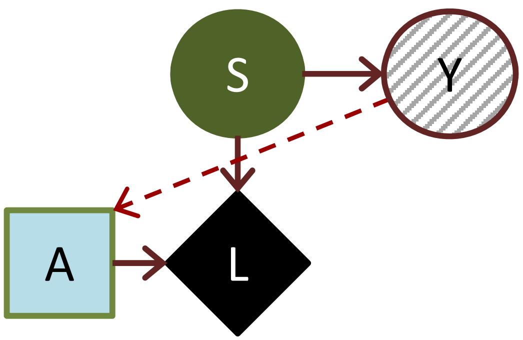

N: do Nothing
R: Repair

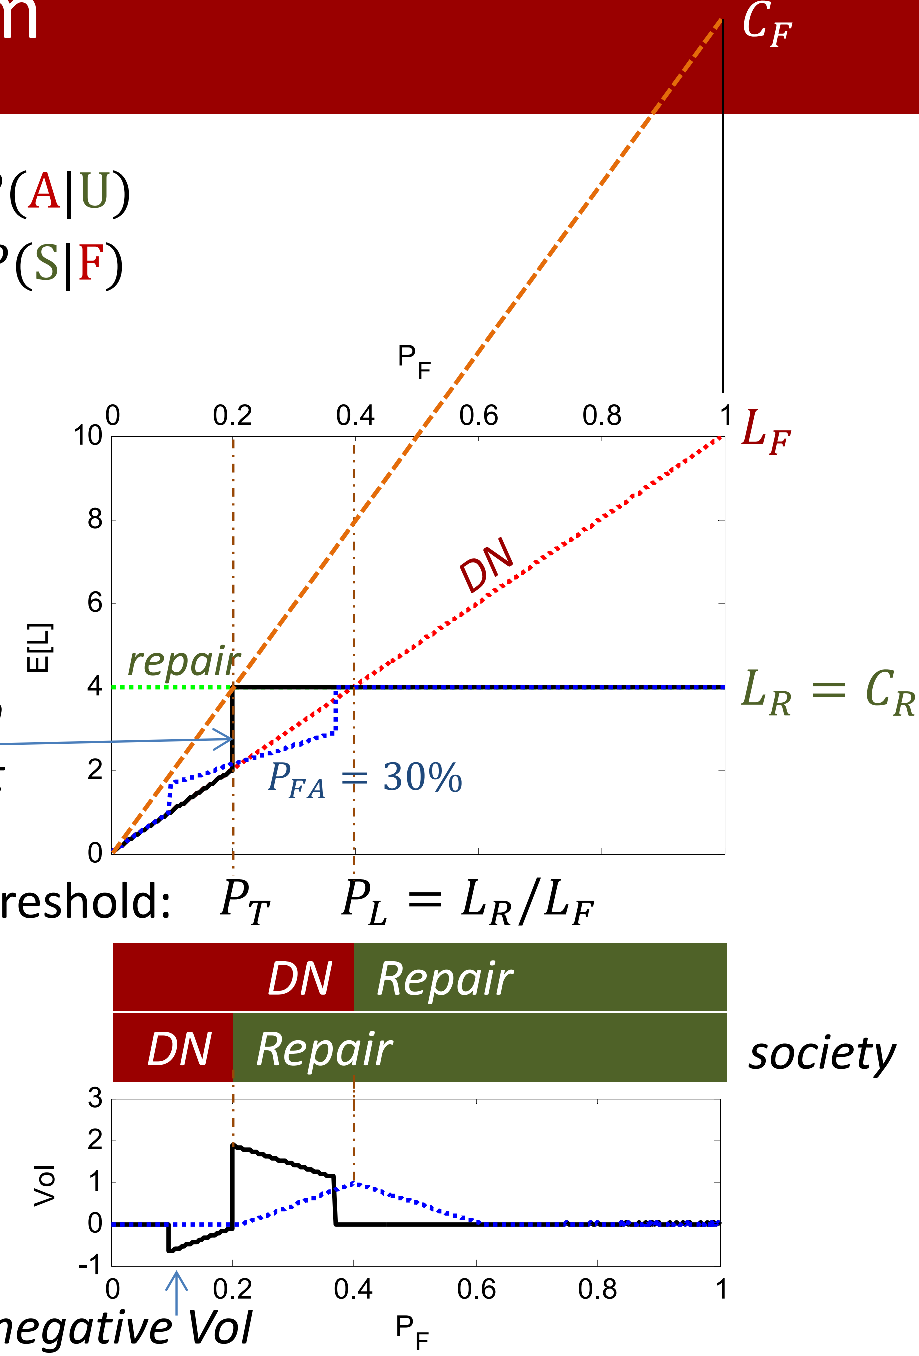

$L(S,A)$

(A) \ (S)	U	F
N	0	L_F
R	L_R	L_R

agent's loss matrix

$C(S,A)$

(A) \ (S)	U	F
N	0	C_F
R	C_R	C_R

society's cost matrix

simplest maintenance problem

state F: Failure
U: Undamaged

sensor outcome $P_{FA} = P(\text{A}|\text{U})$
$P_{FS} = P(\text{S}|\text{F})$

S : Silence
A : Alarm

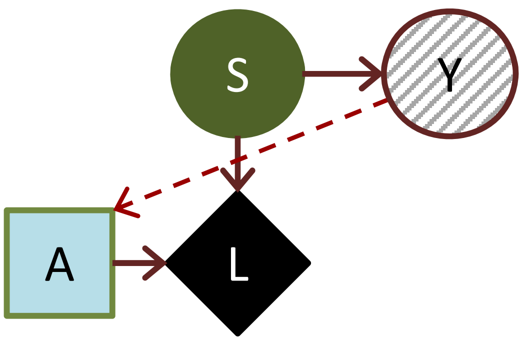

N: do Nothing
R: Repair

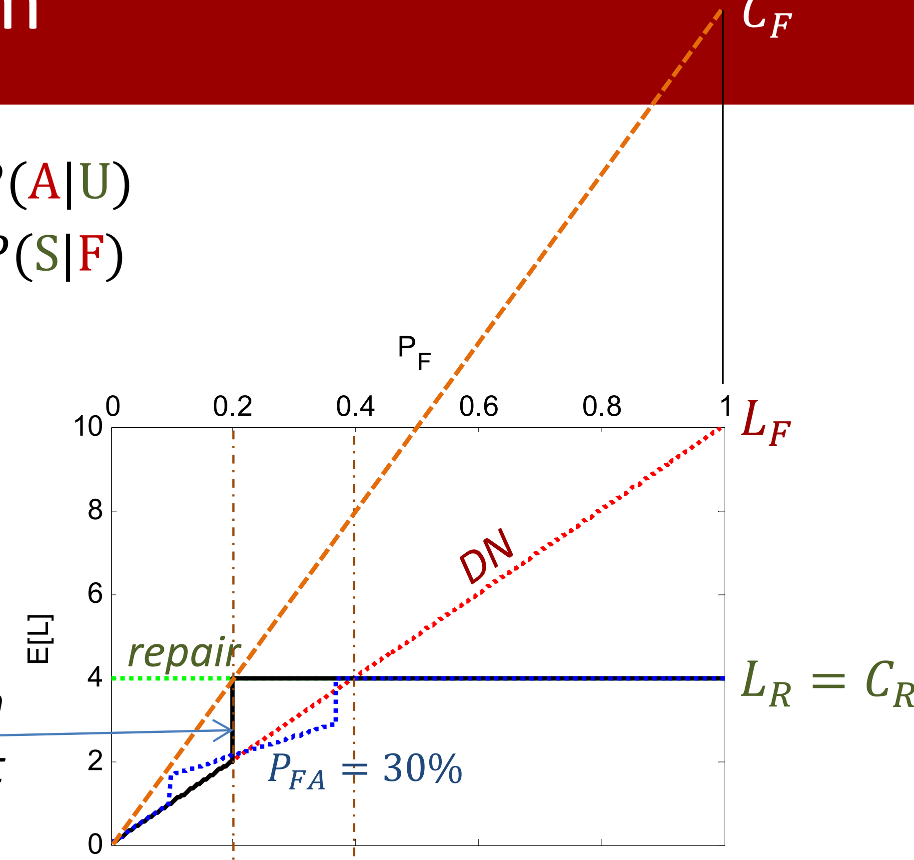

$L(S,A)$

		(S)	
		U	F
(A)	N	0	L_F
	R	L_R	L_R

agent's loss matrix

$C(S,A)$

		(S)	
		U	F
(A)	N	0	C_F
	R	C_R	C_R

society's cost matrix

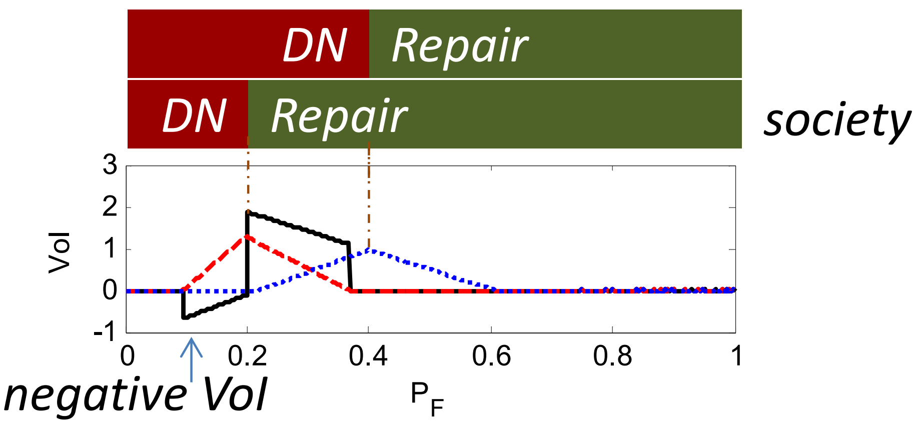

summary on VoI under external constraint

The constraint is effective in forcing agents to take decisions consistent with society's will.

But it has unwanted second-order effects on information avoidance.

How to solve this?

Codes can require to collect data, prescribing to evaluate VoI according to a given formula: buy if its cost is below that threshold.

Society could remove the constraint, and instead introduce incentives for aligning agents' preferences with societal ones.

optimal learning for infrastructure systems

thanks for your attention!

Matteo Pozzi
email: mpozzi@cmu.edu
website: http://faculty.ce.cmu.edu/pozzi/