Integrated learning and decision making

Stella Kapodistria

joint work with Collin Drent, Paulo Serra

Department of Mathematics, TU/e

My goal: Given a trove of data, be able to choose when to perform action A vs action B (at costs c_A vs c_B)

Data

Estimation

Statistical model

Optimization

Cost optimal decision

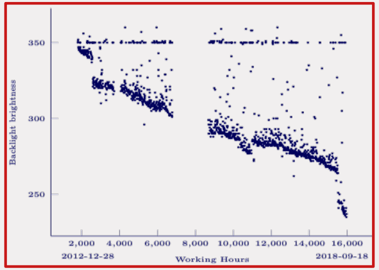

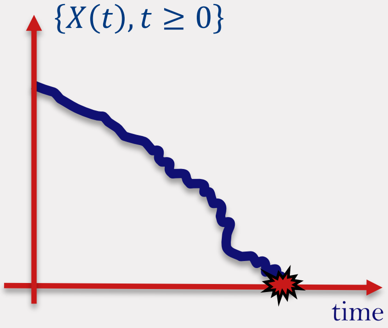

If $X(t) \leq \tau_t$, then do PM.
If $X(t) \leq 0$, then do CM.

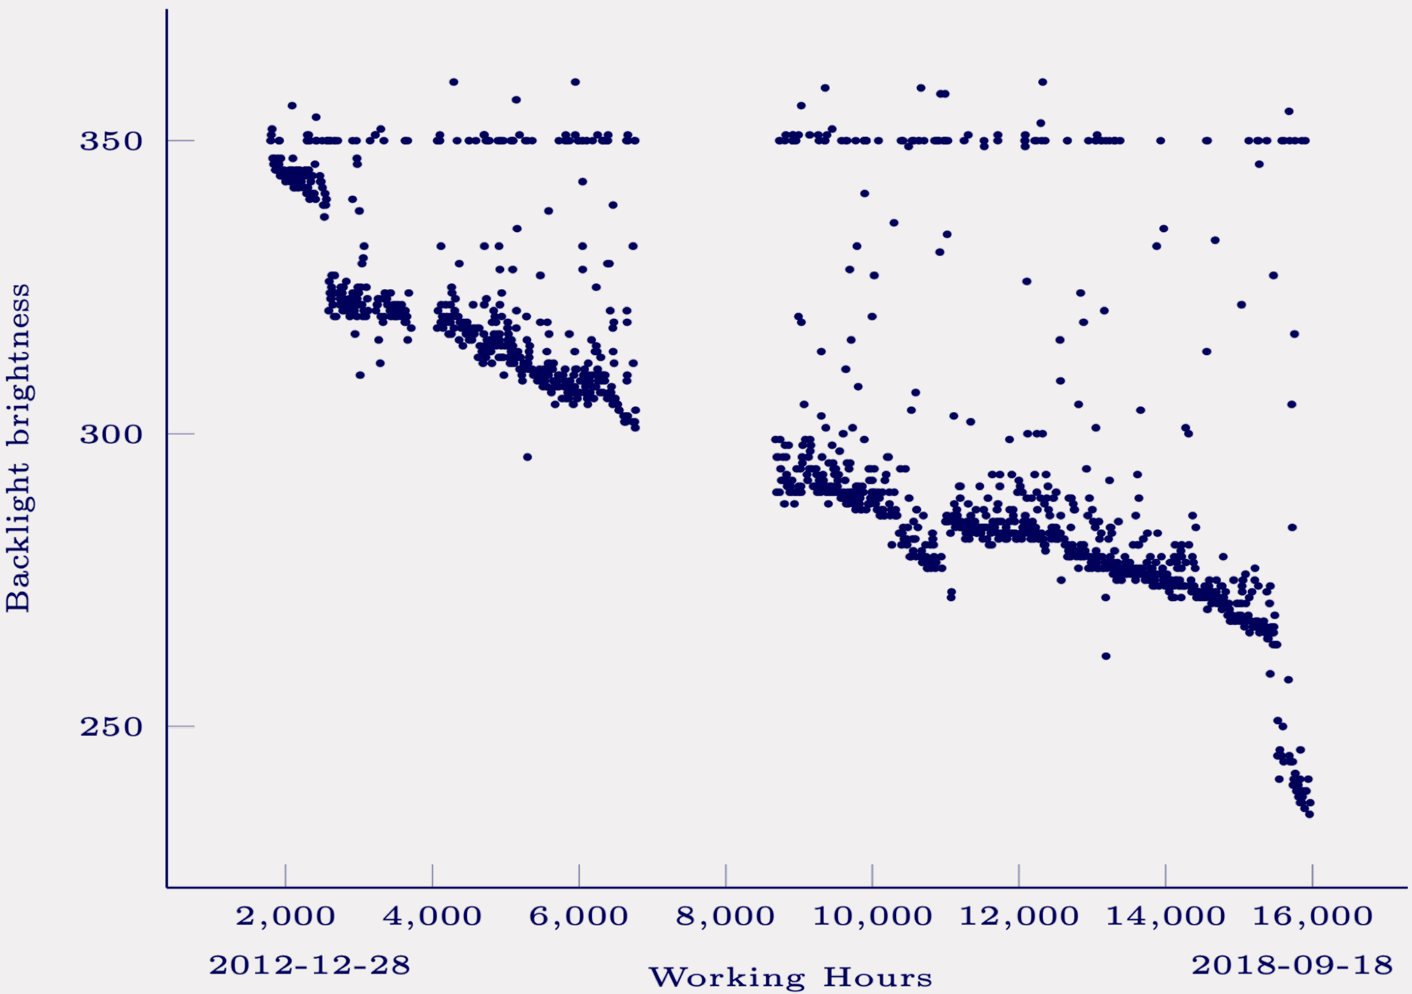
350
300
250
Backlight brightness
2,000
4,000
6,000
8,000
10,000
12,000
14,000
16,000
2012-12-28
Working Hours
2018-09-18

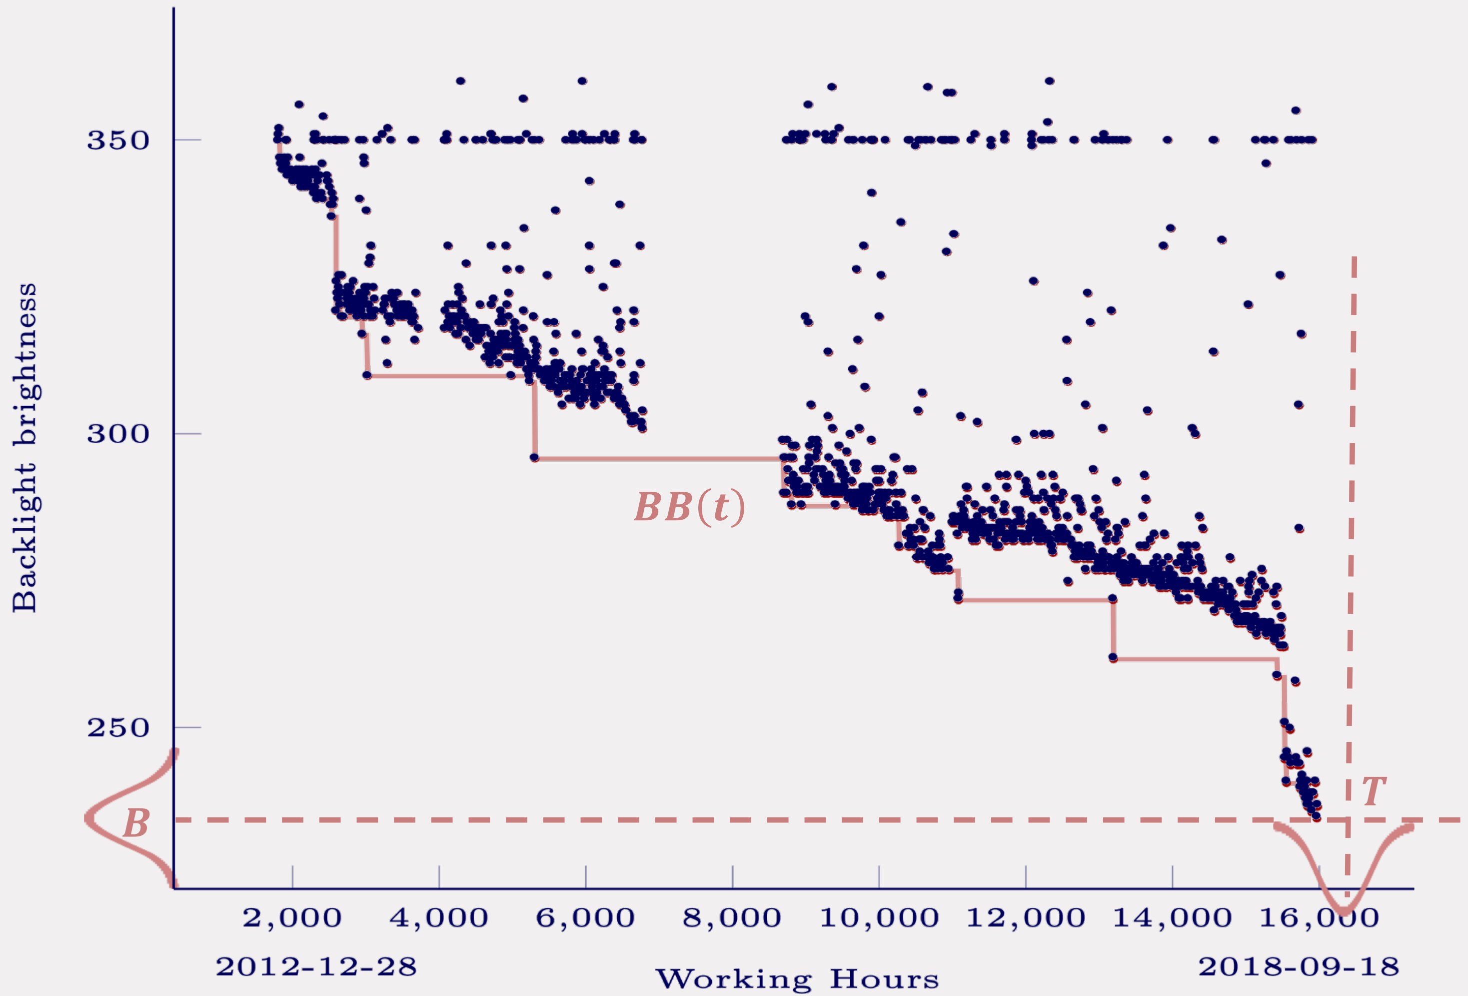

350
300
250
Backlight brightness
BB(t)
B
T
2,000
4,000
6,000
8,000
10,000
12,000
14,000
16,000
2012-12-28
2018-09-18
Working Hours

Upon failure, it can be replaced with an as-good-as new iid component at cost c_c.
Before failure, the component can be preventively replaced at cost $c_p < c_c$.

Data = Lifetimes
Age replacement policy

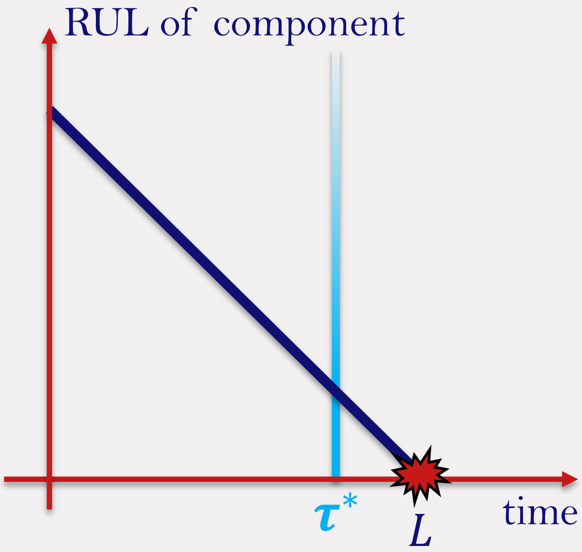

Data = Condition
Condition replacement policy

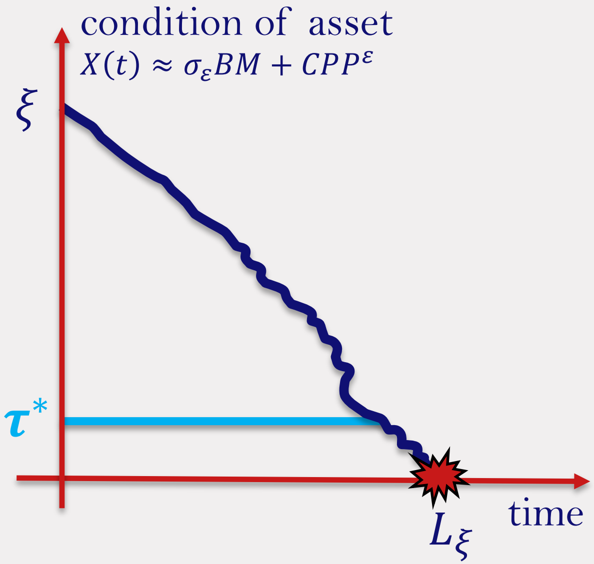

Condition prescriptive maintenance policy for a heterogeneous population

Traditional condition replacement policy

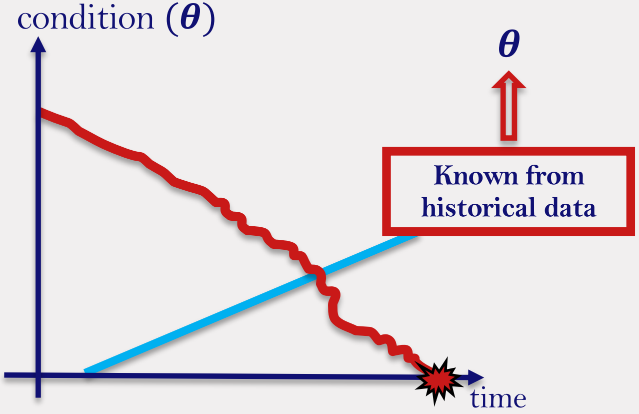

Integrated condition replacement policy

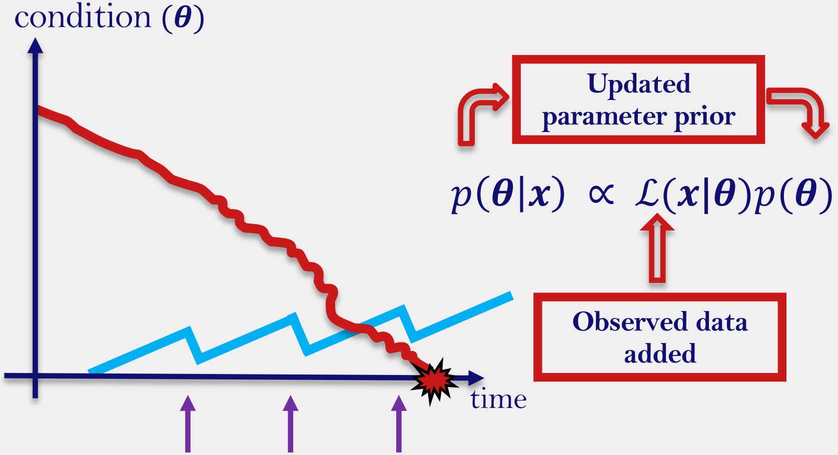

Condition maintenance policy: continuous time

$\{X(t) := \alpha + \beta t + \sigma\varepsilon(t), t \geq 0\}$ denotes the condition, with $\varepsilon(t)$ a standard Brownian motion
Failure is defined as the first passage time of $X(t)$ to $\boldsymbol{\xi}$
Preventive replacement cost $= c_p = c < c_c = c + z =$ Corrective replacement cost

Question:

What is the optimal preventive maintenance (replacement) policy?

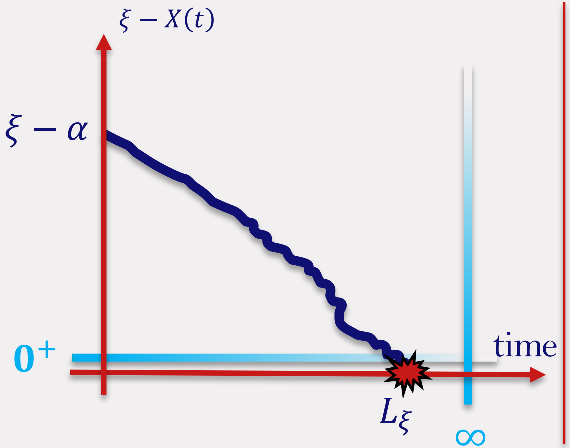

Optimal policy (no information): $\boldsymbol{\tau}^*$

$$\boldsymbol{\tau}^* = \arg\min_\tau \frac{\mathbb{E}[\text{cost in } L_\xi \wedge \tau]}{\mathbb{E}[L_\xi \wedge \tau]} = \infty$$

Optimal policy (full information @ $\mathrm{d}t \to 0^+$): $\mathbf{0}^+$

Optimal policy (full information @ $\mathrm{d}t > 0$) ?

Condition prescriptive policy: discrete time

$\{X(t_n) := \alpha + \beta t_n + \sigma\varepsilon(t_n), n \geq 0\}$ denotes the condition, with $t_n - t_{n-1} = \Delta$

Failure is defined as the first passage time of $X(t_n)$ to $\boldsymbol{\xi}$

Preventive replacement cost $= c_p = c < c_c = c + z =$ Corrective replacement cost

Question:

What is the optimal preventive maintenance (replacement) policy?

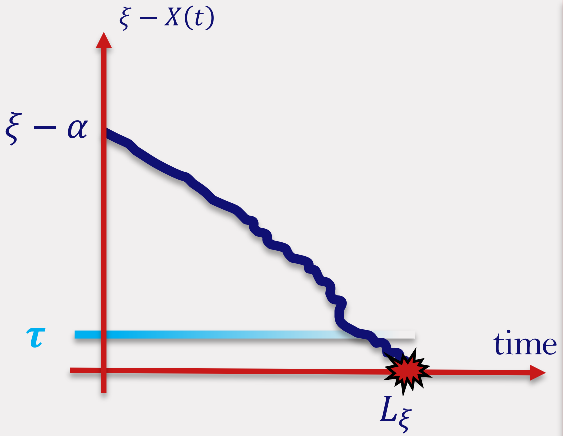

Optimal Policy Theorem [K, Serra 2018]: if $c < (c+z)e^{-r\,\Delta}$, then there exists an optimal stationary policy $\boldsymbol{\tau}$.

$V_\tau :=$ total expected discounted cost for the infinite horizon problem

$$V_\tau = \mathbb{E}[e^{-rL_\tau}](c + V_\tau) + \mathbb{E}\left[e^{-rL_\tau}\mathbb{I}_{\{\xi - X(L_\tau) \leq 0\}}\right]z$$

$$\Rightarrow V_\tau = \frac{\mathbb{E}[e^{-rL_\tau}]c + \mathbb{E}\left[e^{-rL_\tau}\mathbb{I}_{\{\xi - X(L_\tau) \leq 0\}}\right]z}{1 - \mathbb{E}[e^{-rL_\tau}]}$$

Condition prescriptive policy: discrete time

$\{X(t_n) := \alpha + \beta t_n + \sigma\varepsilon(t_n), n \geq 0\}$ denotes the condition, with $t_n - t_{n-1} = \Delta$
Failure is defined as the first passage time of $X(t_n)$ to $\boldsymbol{\xi}$
Preventive replacement cost $= c_p = c < c_c = c + z =$ Corrective replacement cost

Question:

What is the optimal preventive maintenance (replacement) policy?

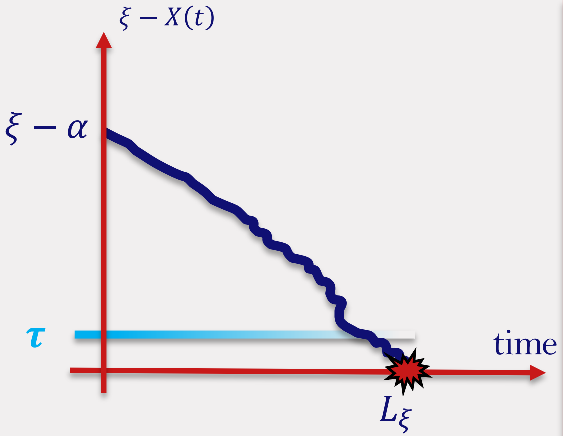

Optimal Policy Theorem [K, Serra 2018]: if $c < (c+z)e^{-r\Delta}$, then there exists an optimal stationary policy $\boldsymbol{\tau}$.

$V_\tau :=$ total expected discounted cost for the infinite horizon problem

$$\lim_{r\to 0} r\, V_\tau = \frac{c + \mathbb{P}[X(L_\tau) \geq \xi]z}{\mathbb{E}[L_\tau]}$$

$\lim_{r\to 0} r\, V_\tau :=$ long-run rate of cost

Condition prescriptive policy: discrete time

$\{X(t_n) := \alpha + \beta t_n + \sigma\varepsilon(t_n), n \geq 0\}$ denotes the condition, with $t_n - t_{n-1} = \Delta$

Failure is defined as the first passage time of $X(t_n)$ to $\boldsymbol{\xi}$

Preventive replacement cost $= c_p = c < c_c = c + z =$ Corrective replacement cost

Question:

What is the optimal preventive maintenance (replacement) policy?

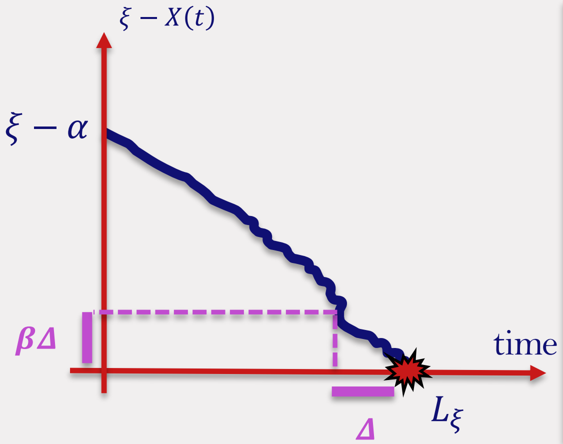

Optimal Policy Theorem [K, Serra 2018]: if $c < (c+z)e^{-r\Delta}$, then there exists an optimal stationary policy $\boldsymbol{\tau}$.

First order approximation: $\boldsymbol{\tau} \approx \beta\Delta$

Second order approximation: $\boldsymbol{\tau} \approx \beta\Delta + \zeta_{c,z}\sigma\sqrt{\Delta}$

Bayesian condition replacement policy

$\{X(t_n) := \alpha + \beta t_n + \sigma\varepsilon(t_n), n \geq 0\}$ denotes the condition, with $t_n - t_{n-1} = \Delta$, $\alpha, \beta \sim$Normal prior
Failure is defined as the first passage time of $X(t)$ to $\boldsymbol{\xi}$
Preventive replacement cost $= c_p = c < c_c = c + z =$ Corrective replacement cost

Question:

What is the optimal preventive maintenance (replacement) policy?

Objective: determine the (cost) optimal preventive maintenance policy

$$V(\xi - x, n; \alpha_n, \beta_n) \equiv V(\xi - x, n)$$

Approach: Bellman equations

Results: optimal policy = update-dependent control limit policy

$$V(\xi - x, n) = \begin{cases} c_c + e^{-r\,\Delta}\mathbb{E}_{\alpha_n,\beta_n}[V(\xi - \alpha_n - \beta_n\delta - \sigma\varepsilon_\delta, n+1)], & \text{if } \xi - x \leq 0 \\ \min\{c_p + e^{-r\,\Delta}\mathbb{E}_{\alpha_n,\beta_n}[V(\xi - \alpha_n - \beta_n\delta - \sigma\varepsilon_\delta, 0\}, n+1)], e^{-r\,\Delta}\mathbb{E}_{\alpha_n,\beta_n}[V(\xi - x - \beta_n\delta - \sigma\varepsilon_\delta, n+1)]\}, & \text{if } \xi - x > 0 \end{cases}$$

Bayesian condition replacement policy

Design matrix

$$\boldsymbol{X}_n = \boldsymbol{T}_n\boldsymbol{\theta} + \boldsymbol{\varepsilon}_n, \boldsymbol{T}_n = \begin{bmatrix} 1 & T_1 \\ \vdots & \vdots \\ 1 & T_n \end{bmatrix}, \boldsymbol{\theta} = \begin{bmatrix} \alpha \\ \beta \end{bmatrix}, \boldsymbol{\varepsilon}_n \sim N(0, \boldsymbol{\Sigma}_n)$$

$$T_1 = \Delta, T_{n+1} = (T_n + \delta)\mathbb{I}_{\{X_n \le \tau_n\}} + \delta\mathbb{I}_{\{X_n > \tau_n\}}$$

Posterior distribution

$$\Pi(\alpha, \beta | L_1, \dots, L_n) = \text{Normal}\left((\boldsymbol{I}_2 + \boldsymbol{T}'_n\boldsymbol{\Sigma}_n^{-1}\boldsymbol{T}_n)^{-1}\boldsymbol{T}'_n\boldsymbol{\Sigma}_n^{-1}\boldsymbol{T}_n, (\boldsymbol{I}_2 + \boldsymbol{T}'_n\boldsymbol{\Sigma}_n^{-1}\boldsymbol{T}_n)^{-1}\right)$$

Note that

$$\mathbb{E}\left\|\widehat{\boldsymbol{\theta}}_{\boldsymbol{n}} - \boldsymbol{\theta}\right\|^2 \lesssim \frac{\sigma^2}{n} + \frac{\|\boldsymbol{\theta}\|^2}{n^2}$$

(root-n consistent estimator)

TU/e

Bayesian condition replacement policy

In the model with known parameters, second order approximation

$$\boldsymbol{\tau} \approx \beta\Delta + \zeta_{c,z}\sigma\sqrt{\Delta}$$

So we create a regret function that converges to the second order approximation

$$\mathfrak{C}(\boldsymbol{\tau}; x_n, \alpha_n, \beta_n) = c\mathbb{I}_{\{0<\xi-X_{n+1},\boldsymbol{\tau}<\xi-x_n\}} + (c+z)\mathbb{I}_{\{0\geq\xi-X_{n+1},\boldsymbol{\tau}<\xi-x_n\}}\xi + (2c+z)\mathbb{I}_{\{0\geq\xi-X_{n+1},\boldsymbol{\tau}\geq\xi-x_n\}}$$

We compute the $\boldsymbol{\tau}$ that minimizes the conditional expectation $\mathbb{E}[\mathfrak{C}(\boldsymbol{\tau}; x_n, \alpha_n, \beta_n)]$. That produces a sequence of $\boldsymbol{\tau}$'s, say $\{\tau_n\}$.

Theorem [K, Serra 2018]:

$$\tau_n = \beta_n\Delta + \Phi^{-1}\left(\Phi\left(\xi - \alpha_n - \beta_n\Delta \Big/ \sigma\sqrt{\Delta}\right) - c\big/ c+z\right)\sigma\sqrt{\Delta}$$

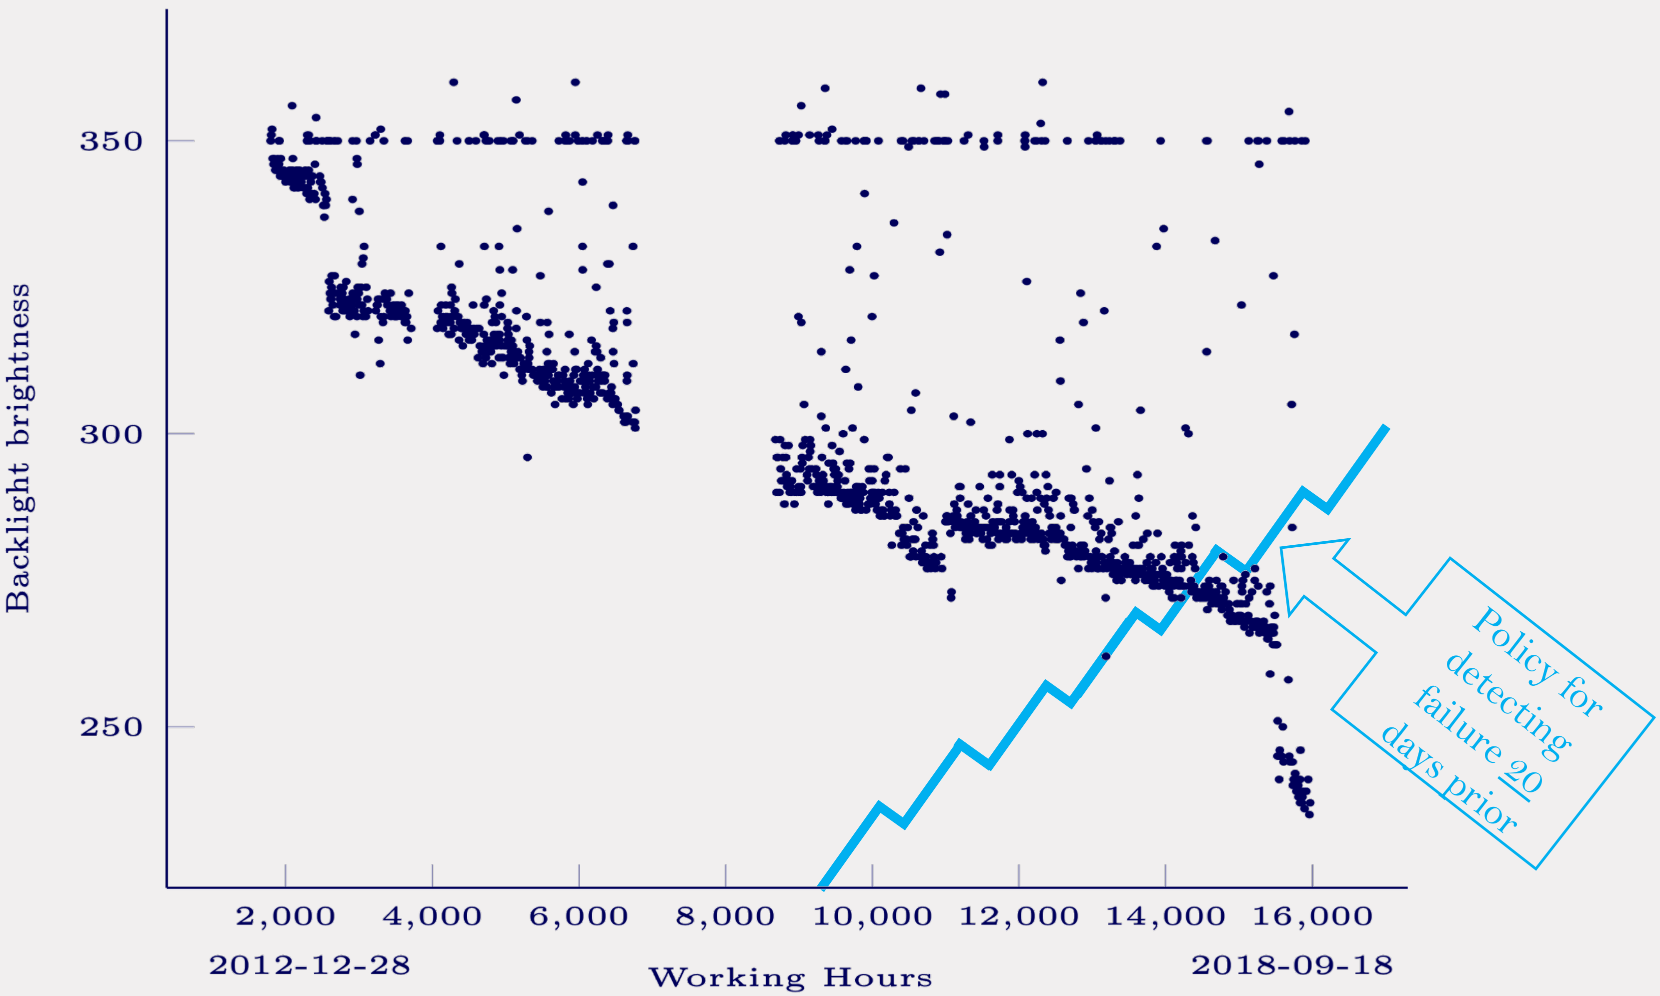

Backlight brightness
350
300
250
2,000
4,000
6,000
8,000
10,000
12,000
14,000
16,000
2012-12-28
2018-09-18
Working Hours
Policy for detecting failure 20 days prior

Cost savings in comparison to the oracle

(2%, 5%)

Bayes cost gap < Naïve $\boldsymbol{\theta}$ updating cost gap < Historical $\boldsymbol{\theta}$ cost gap

(10%, 25%)

“Many small jumps”

“Few large jumps”

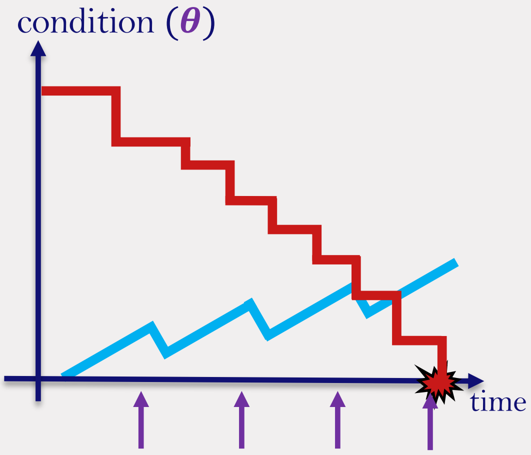

My goal: Given a trove of data, be able to choose when to perform action A vs action B (at costs c_A vs c_B)

Integrated learning and decision making

Stella Kapodistria
joint work with Collin Drent, Paulo Serra

Department of Mathematics, TU/e