Gaussian Process-Based Mortality Monitoring using Multivariate Cumulative Sum Procedures

Karim Barigou   Stéphane Loisel   Yahia Salhi

AFRIC 2023

karim.barigou@act.ulaval.ca

www.karimbarigou.com

July 25, 2023

UNIVERSITÉ LAVAL

Faculté des sciences et de génie
École d'actuariat

Outline

Monitoring insurance processes

Monitoring mortality rates is crucial for the risk management of life insurance.

Challenges:

- **Quickest detection:** In a rapidly changing environment, actuarial assumptions should be monitored quickly and efficiently.
  → Real-time sequential detection
- **Correlation:** Mortality data often exhibit interdependencies between different age groups or cohorts.
  → Gaussian Process (GP) regression
- **Multivariate detection:** Univariate detection methods ignore the complex dependence structure, limiting their effectiveness.
  → MCUSUM algorithm

Monitoring insurance processes

Monitoring mortality rates is crucial for the risk management of life insurance.

Challenges:

- **Quickest detection:** In a rapidly changing environment, actuarial assumptions should be monitored quickly and efficiently.
  → Real-time sequential detection
- **Correlation:** Mortality data often exhibit interdependencies between different age groups or cohorts.
  → Gaussian Process (GP) regression
- **Multivariate detection:** Univariate detection methods ignore the complex dependence structure, limiting their effectiveness.
  → MCUSUM algorithm

This presentation

- Proposed approach:
  - Forecasting: Mortality forecasting based on GP regression.
  - **MCUSUM monitoring:** Tracks differences between predicted and observed mortality rates, enabling real-time change detection.
- Which change?
  - Change of level by tracking mortality rates.
  - Change of trend by tracking mortality improvements.
- Empirical analysis for France, Japan, Canada, and the USA.
- The MCUSUM shows quicker detection to univariate alternatives that ignore dependence.

Gaussian process for mortality forecasting

- Training set: $(\boldsymbol{x}^i, y^i)$ $(i = 1, \ldots, n)$.
  - In our case: $\boldsymbol{x}^i = (x^i_{\text{age}}, x^i_{\text{year}})$ and $y^i = \log(D^i/E^i)$.
  - Age: M age-groups, e.g. $z_1 = [50; 55)$;$z_2 = [55; 60)$;. . .;$z_M = [85; 90)$.
  - T years: [1980,2020].
- Gaussian process:

$$f(\boldsymbol{x}) \sim \mathcal{N}\left(m(\boldsymbol{x}), k(\boldsymbol{x}, \boldsymbol{x})\right),$$

  where $k(\boldsymbol{x}, \boldsymbol{x})$ is the covariance matrix.
- Completely characterized by mean function $m(\boldsymbol{x})$ and covariance/kernel function $k(\boldsymbol{x}, \boldsymbol{x}')$.
- Key reference: Ludkovski et al. (2018).

Gaussian process for mortality forecasting

- GP posterior distribution is multivariate normal.
- Therefore, predicted log death rates are multivariate normal, i.e.

$$\begin{aligned} \boldsymbol{y}^{\boldsymbol{t}} &:= \log\left(\boldsymbol{\mu}_{\boldsymbol{t}}\right) \\ &\sim \mathcal{N}(\boldsymbol{m}_{\boldsymbol{t}}, \boldsymbol{\Sigma}_t) \end{aligned}$$

  for any prediction year $t = T+1, \ldots, T+N$.
- Monitoring central log death rates $\rightarrow$ Monitoring the mean of a multivariate normal process.

Gaussian process for mortality forecasting

- GP posterior distribution is multivariate normal.
- Therefore, predicted log death rates are multivariate normal, i.e.

$$\begin{aligned} \boldsymbol{y}^{\boldsymbol{t}} &:= \log\left(\boldsymbol{\mu}_{\boldsymbol{t}}\right) \\ &\sim \mathcal{N}(\boldsymbol{m}_{\boldsymbol{t}}, \boldsymbol{\Sigma}_t) \end{aligned}$$

  for any prediction year $t = T + 1, \ldots, T + N$.
- Monitoring central log death rates $\rightarrow$ Monitoring the mean of a multivariate normal process.

The MCUSUM for multivariate normal

- Sequential procedure for detecting a change in a process based on likelihood ratios.
- Observing a sequence $\boldsymbol{y} = (\boldsymbol{y}^t)_{t \geq 1}$ with unknown change point τ:

$$\begin{aligned} \boldsymbol{y}^t &\sim F_1, \quad 1 \leq t \leq \tau, \\ &\sim F_2, \quad \tau + 1 \leq t. \end{aligned}$$

- CUSUM algorithm signals change when:

$$S_t = \max\left(S_{t-1} + \log \frac{f_2\left(\boldsymbol{y}^{\boldsymbol{t}}\right)}{f_1\left(\boldsymbol{y}^{\boldsymbol{t}}\right)}, 0\right) > L,$$

where f_1 and f_2 are the density functions, and L is a fixed threshold.

The MCUSUM for multivariate normal (continued)

- For mortality monitoring, log death rates follow
  - In-control process: $\mathcal{N}(\boldsymbol{\mu_1}, \boldsymbol{\Sigma})$.
  - *Out-of-control process*: $\mathcal{N}(\boldsymbol{\mu_2}, \boldsymbol{\Sigma})$.
- The MCUSUM is

$$S_t = \max\Big(S_{t-1} + (\boldsymbol{\mu_2} - \boldsymbol{\mu_1})' \boldsymbol{\Sigma}^{-1} \left(\boldsymbol{y^t} - \boldsymbol{\mu_1}\right) - \frac{1}{2} (\boldsymbol{\mu_2} - \boldsymbol{\mu_1})' \boldsymbol{\Sigma}^{-1} (\boldsymbol{\mu_2} - \boldsymbol{\mu_1}), 0\Big).$$

What type of change?

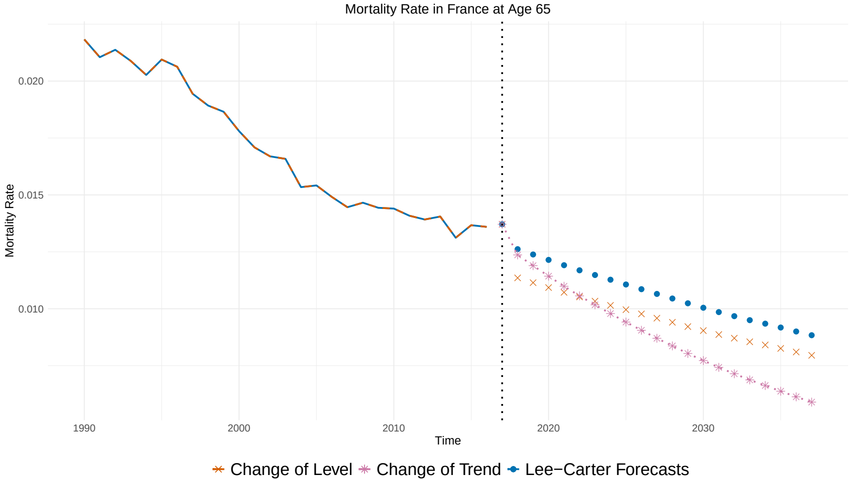

Figure: Mortality rate at age 65 in France with Lee-Carter forecasts, change of level and change of trend with a change point in 2017.

UNIVERSITÉ LAVAL

Change of level Detection

- The change-point model for level change detection can be expressed as

$$\mathbb{E}\left[\log(\boldsymbol{\mu_t})\right] \sim \begin{cases} \boldsymbol{m_t} & \text{for } i = 1, \ldots, \tau \\ \overline{\boldsymbol{m_t}} & \text{for } i = \tau + 1, \ldots \end{cases}$$

where e.g. $\overline{\boldsymbol{m_t}} = \boldsymbol{m_t} + \log(\alpha)\mathbf{1}$ with $\alpha = 0.9$ (longevity risk).

- The generalized MCUSUM is defined by:

$$S_t = \max\left(S_{t-1} + \left(\overline{\boldsymbol{m_t}} - \boldsymbol{m_t}\right)' \boldsymbol{\Sigma}_t^{-1}\left(\boldsymbol{y}^t - \boldsymbol{m_t}\right) - \frac{1}{2}\left(\overline{\boldsymbol{m_t}} - \boldsymbol{m_t}\right)' \boldsymbol{\Sigma}_t^{-1}\left(\overline{\boldsymbol{m}}_t - \boldsymbol{m_t}\right), 0\right),$$

where

1. $\boldsymbol{y}^t$ is the vector of observed log death rates.
2. $\boldsymbol{m_t}$ and $\boldsymbol{\Sigma}_t$ are the mean and covariance from GP-based forecasts

Change of level Detection

- The change-point model for level change detection can be expressed as

$$\mathbb{E}\left[\log(\boldsymbol{\mu_t})\right] \sim \begin{cases} \boldsymbol{m_t} & \text{for } i = 1, \ldots, \tau \\ \overline{\boldsymbol{m_t}} & \text{for } i = \tau + 1, \ldots \end{cases}$$

where e.g. $\overline{\boldsymbol{m_t}} = \boldsymbol{m_t} + \log(\alpha)\mathbf{1}$ with $\alpha = 0.9$ (longevity risk).

- The generalized MCUSUM is defined by:

$$S_t = \max\Big(S_{t-1} + (\overline{\boldsymbol{m_t}} - \boldsymbol{m_t})' \boldsymbol{\Sigma}_t^{-1} \left(\boldsymbol{y^t} - \boldsymbol{m_t}\right) \\ - \frac{1}{2}\left(\overline{\boldsymbol{m_t}} - \boldsymbol{m_t}\right)' \boldsymbol{\Sigma}_t^{-1} \left(\overline{\boldsymbol{m}}_t - \boldsymbol{m_t}\right), 0\Big),$$

where

1. $\boldsymbol{y^t}$ is the vector of observed log death rates.
2. $\boldsymbol{m_t}$ and $\boldsymbol{\Sigma}_t$ are the mean and covariance from GP-based forecasts.

Change of trend detection

- The change-point model for trend change detection can then be expressed as

$$\mathbb{E}\left[\Delta \log(\boldsymbol{\mu_t})\right] \sim \begin{cases} \boldsymbol{m_t^I} & \text{for } i = 1, \ldots, \tau \\ \overline{\boldsymbol{m_t^I}} & \text{for } i = \tau + 1, \ldots \end{cases}$$

  where

  1. **Mortality improvements:** $\Delta \log(\boldsymbol{\mu_t}) = \log(\boldsymbol{\mu_t}) - \log(\boldsymbol{\mu_{t-1}})$
  2. **Trend change:** $\overline{\boldsymbol{m}}_t^I = \log(\exp(\boldsymbol{m}_t^I) - \alpha)$

- How to fix the threshold L?

$$\mathbb{P}\left[\max_{1 \leq i \leq T} S_i \geq L \mid \text{no change}\right] = \alpha,$$

  determined by simulations.

Change of trend detection

- The change-point model for trend change detection can then be expressed as

$$\mathbb{E}\left[\Delta \log(\boldsymbol{\mu_t})\right] \sim \begin{cases} \boldsymbol{m_t^I} & \text{for } i = 1, \ldots, \tau \\ \overline{\boldsymbol{m_t^I}} & \text{for } i = \tau + 1, \ldots \end{cases}$$

where

1. **Mortality improvements:** $\Delta \log(\boldsymbol{\mu_t}) = \log(\boldsymbol{\mu_t}) - \log(\boldsymbol{\mu_{t-1}})$
2. **Trend change:** $\overline{\boldsymbol{m}}_t^I = \log(\exp(\boldsymbol{m}_t^I) - \alpha)$

- How to fix the threshold L?

$$\mathbb{P}\left[\max_{1 \leq i \leq T} S_i \geq L \mid \text{no change}\right] = \alpha,$$

determined by simulations.

Empirical analysis

- Countries: France, Canada, USA and Japan.
- Ages: 50-89 by 5-year age tranches.
- Years:
  1. Estimation: 1991-2010.
  2. Detection: 2011-2020.
- Detection types:
  1. +/- 5% in the level change.
  2. +/- 1% in the trend change.
- False alarm probability: 1%.

Empirical analysis: change of level

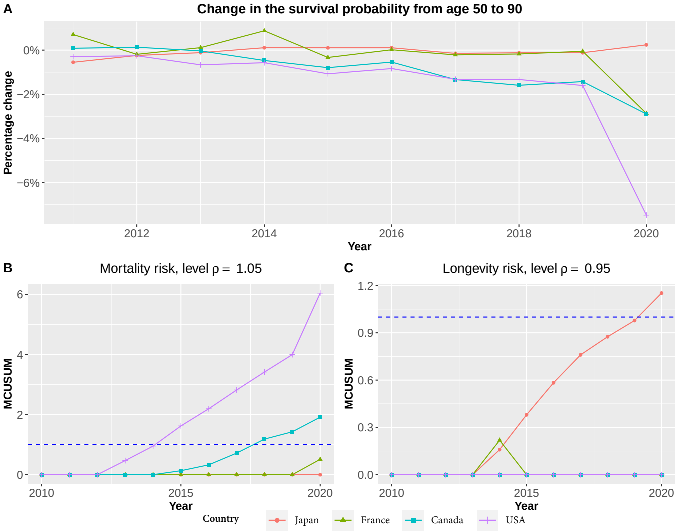

UNIVERSITÉ LAVAL

Empirical analysis: change of trend

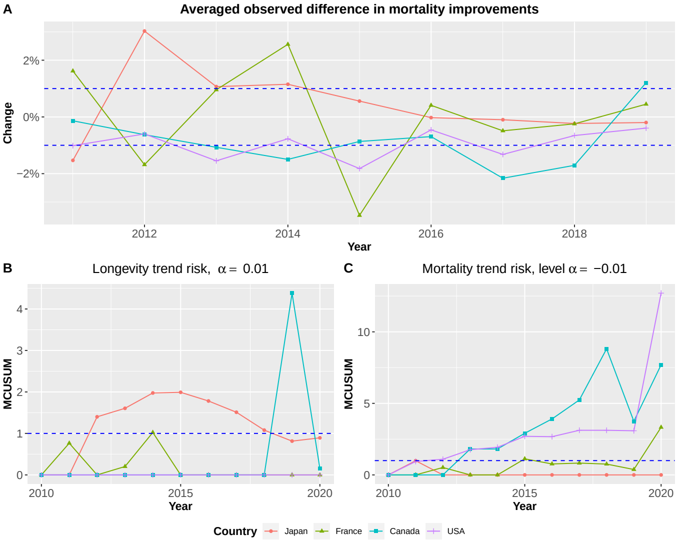

MCUSUM vs univariate CUSUM charts

What is the **added value** of the MCUSUM?

- Standard age-period-cohort models assume perfect correlation, e.g. for the Lee-Carter model:

$$\log(\mu_{x,t}) = \alpha_x + \beta_x \kappa_t,$$

  ⇒ The sum of log death rates across age is a comonotonic sum driven by κ_t.

- What is the **loss in detection performance** assuming comonotonicity between age classes?
  ⇒ **Detection performance:** Average Run Length (ARL) for a given false alarm probability.

UNIVERSITÉ LAVAL

MCUSUM vs univariate CUSUM charts

What is the **added value** of the MCUSUM?

- Standard age-period-cohort models assume perfect correlation, e.g. for the Lee-Carter model:

$$\log(\mu_{x,t}) = \alpha_x + \beta_x \kappa_t,$$

  $\Rightarrow$ The sum of log death rates across age is a comonotonic sum driven by κ_t.

- What is the **loss in detection performance** assuming comonotonicity between age classes?
  $\Rightarrow$ **Detection performance:** Average Run Length (ARL) for a given false alarm probability.

MCUSUM vs univariate CUSUM charts

What is the **added value** of the MCUSUM?

- Standard age-period-cohort models assume perfect correlation, e.g. for the Lee-Carter model:

$$\log(\mu_{x,t}) = \alpha_x + \beta_x \kappa_t,$$

  $\Rightarrow$ The sum of log death rates across age is a comonotonic sum driven by κ_t.

- What is the **loss in detection performance** assuming comonotonicity between age classes?
  $\Rightarrow$ **Detection performance:** Average Run Length (ARL) for a given false alarm probability.

MCUSUM vs univariate CUSUM charts

The comonomotonic CUSUM is defined as

$$S_t^c = \max\left(S_{t-1}^c + (\overline{\mu}_t - \mu_t)\frac{(s_t - \mu_t)}{\sigma_t} - \frac{1}{2}\frac{(\overline{\mu}_t - \mu_t)^2}{\sigma_t}, 0\right),$$

with

$$\mu_t = \sum_{x=1}^{M} m_{i,t} \qquad \sigma_t = \sum_{x=1}^{M} \sigma_{i,t}$$

$$\overline{\mu}_t = \mu_t + M\log(\alpha) \qquad s_t = \sum_{x=1}^{M} y_{i,t}$$

with $m_{i,t}$ and $\sigma_{i,t}$, the mean and standard deviations of the i-th component of the log death rates vector $\boldsymbol{y_t} = (y_{1,t}, \dots, y_{M,t})$.

Comparison of the MCUSUM and C-CUSUM charts

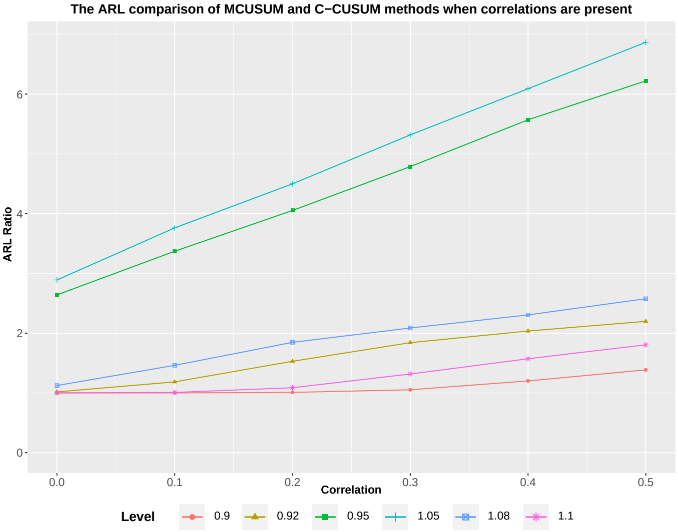

Conclusion

- GP-based mortality forecasts combined with the MCUSUM detection rule provide several benefits:
  1. Capture the dependence between age classes.
  2. **Efficient real-time multivariate monitoring** for e.g.
     - ★ Change of level.
     - ★ Change of trend.
  3. Detection of longevity risk in Japan and mortality risk in USA and Canada over the 10-year period 2011-2020.
  4. **Outperformance compared to univariate control charts** that ignore the dependence structure.

Thank you for your attention! Any questions?

UNIVERSITÉ LAVAL

Conclusion

- GP-based mortality forecasts combined with the MCUSUM detection rule provide several benefits:
  1. Capture the dependence between age classes.
  2. **Efficient real-time multivariate monitoring** for e.g.
     - ★ Change of level.
     - ★ Change of trend.
  3. Detection of longevity risk in Japan and mortality risk in USA and Canada over the 10-year period 2011-2020.
  4. **Outperformance compared to univariate control charts** that ignore the dependence structure.

Thank you for your attention! Any questions?

UNIVERSITÉ LAVAL

Extra slides *just in case*

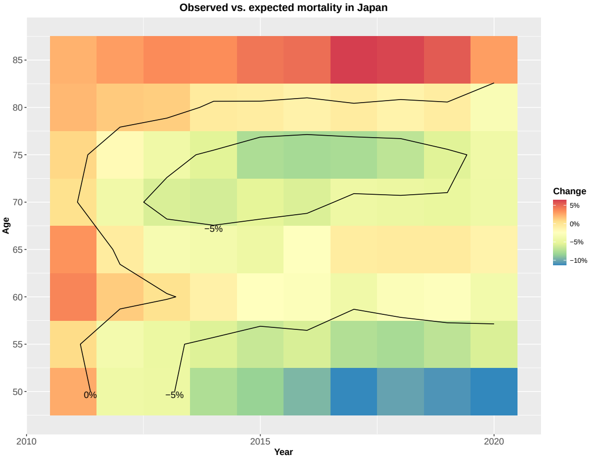

Figure: Percentage change between observed and GP-predicted death rates by age tranches for Japanese males.

Extra slides *just in case*

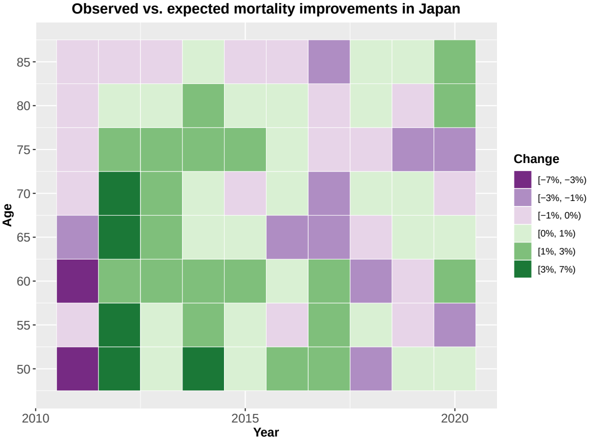

Figure: Difference between observed and GP-predicted mortality improvement rates by age tranches for Japanese Males.

Extra slides *just in case*

	[50,55)	[55,60)	[60,65)	[65,70)	...	[85,90)
[50; 55)	1	ρ	$\rho/2$	0	0	0
[55; 60)	ρ	1	ρ	$\rho/2$	0	0
[60,65)	$\rho/2$	ρ	1	ρ	$\rho/2$	0
[65,70)	0	$\rho/2$	ρ	1	ρ	$\rho/2$
...	0	0	$\rho/2$	ρ	1	ρ
[85; 90).	0	0	0	$\rho/2$	ρ	1

Table: Correlation matrix between age tranches used for the simulation study.

Extra slides *just in case*

	50	55	60	65	70	75	80	85
50	1.00	0.35	0.04	0.04	0.03	0.03	0.03	0.04
55	0.35	1.00	0.36	0.03	0.04	0.04	0.03	0.03
60	0.04	0.36	1.00	0.36	0.04	0.05	0.04	0.03
65	0.04	0.03	0.36	1.00	0.37	0.04	0.04	0.03
70	0.03	0.04	0.04	0.37	1.00	0.36	0.03	0.04
75	0.03	0.04	0.05	0.04	0.36	1.00	0.36	0.04
80	0.03	0.03	0.04	0.04	0.03	0.36	1.00	0.35
85	0.04	0.03	0.03	0.03	0.04	0.04	0.35	1.00

Color scale: 0, 0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9, 1

Figure: Estimated correlation matrix for Japanese male death rates in 2011.

Ludkovski, M., Risk, J. & Zail, H. (2018), 'Gaussian process models for mortality rates and improvement factors', *ASTIN Bulletin: The Journal of the IAA* **48**(3), 1307–1347.