Cyber Risk Frequency Modelling Using Hawkes Processes : Calibration on Attack and Vulnerability Data

Yousra Cherkaoui-Zirnheld, CREST - AXA GIE

Joint work with : Alexandre Boumezoued, Milliman R&D

Caroline Hillaire, ENSAE IP Paris - CREST

6th European
Congress of Actuaries
www.eca2026.org

Agenda

- Context
- Cyber risk modelling using Hawkes processes with cyber vulnerabilities
- Cyber attacks and vulnerability databases
- Calibration results
- Response measures using the second phase of the Hawkes process
- Future research questions

Cyber risk

Examples of cumulative cyber incidents

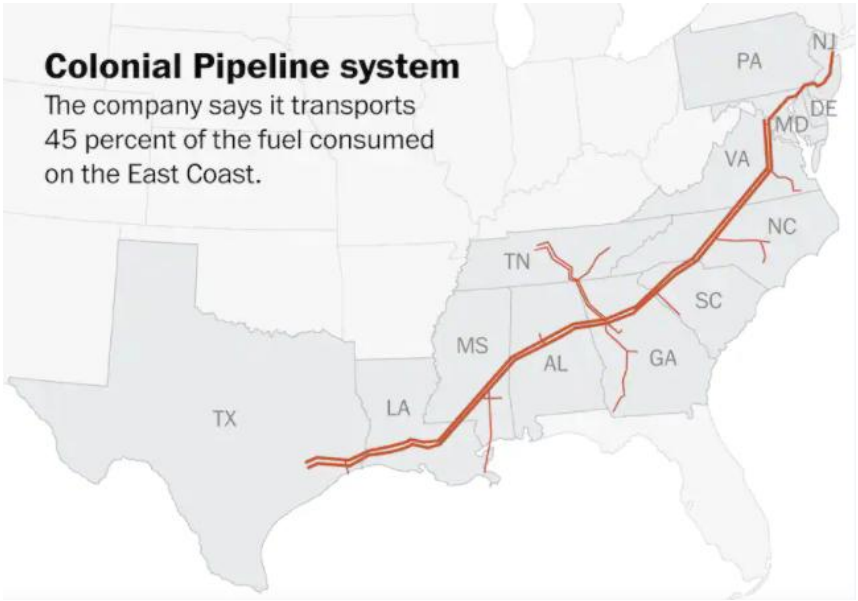

Wannacry/NotPetya (2017)

Cyber attack triggered by a vulnerability exploit, leading to global contagion.

Colonial Pipeline (2021)

Cyber attack on a critical infrastructure company, causing cascading costs across the supply chain.

CrowdStrike (2024) – single point of failure

IT incident from a single point of failure, disrupting systems worldwide.

CREST
CENTER FOR RESEARCH IN ECONOMICS AND STATISTICS

Cyber risk

Consequences and cyber insurance

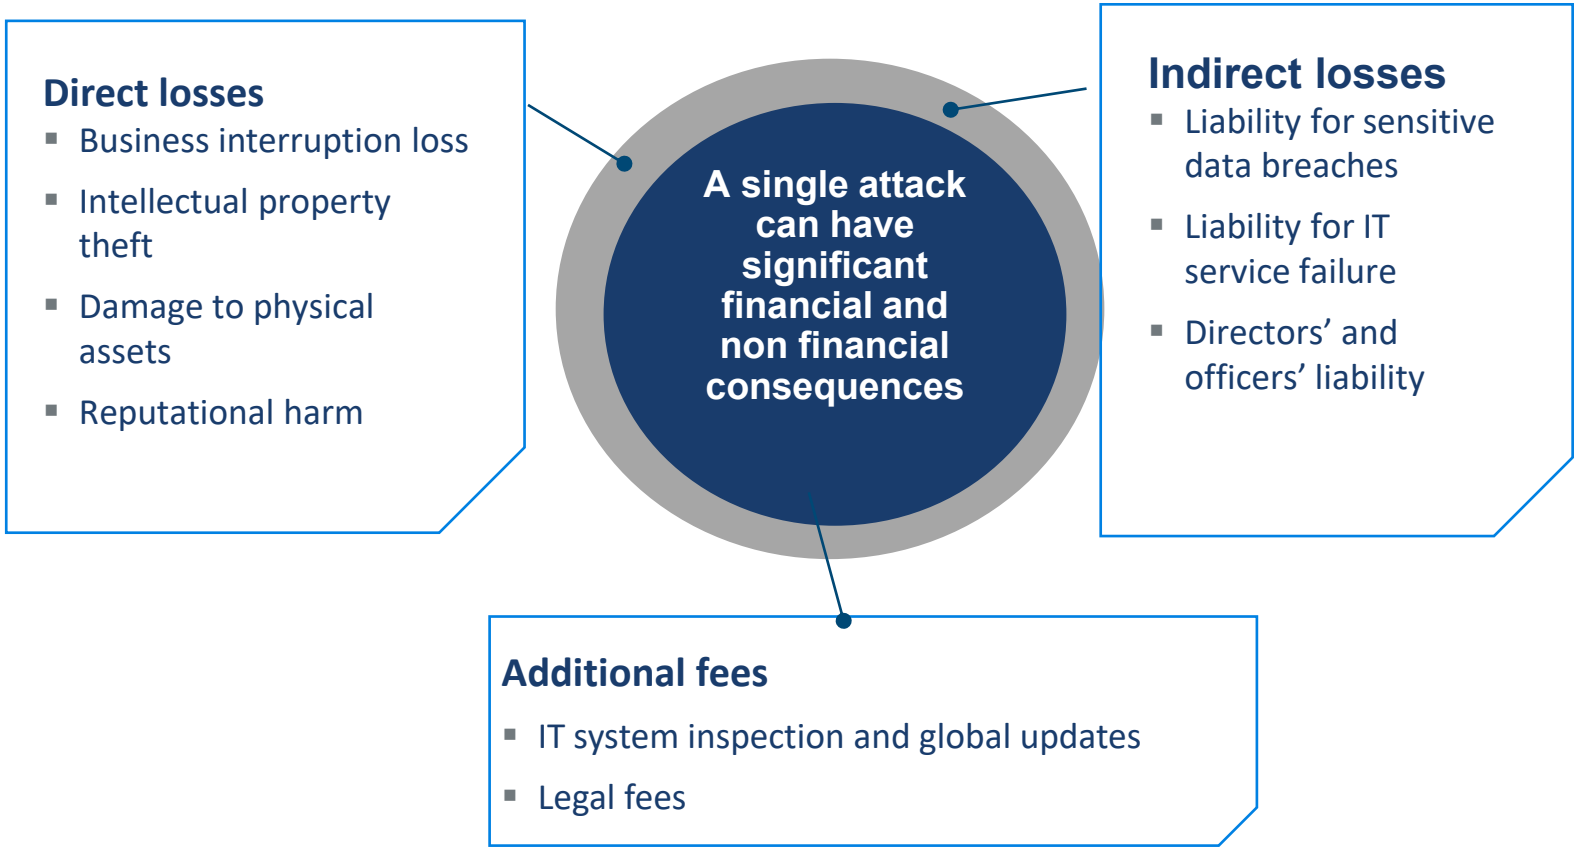

- Such **losses** can be **covered** by **cyber insurance** contracts.
- **Multiple policyholders** can be hit **simultaneously** by the **same cyber event**, creating an **accumulation** risk across the portfolio and potentially **overwhelming** the **insurer**'s assistance and claims teams.
- Insurers need **mathematical tools** to measure **cyber risk frequency accumulation**
- **Cyber attacks** may originate from **cyber vulnerabilities** exploit

CREST
CENTER FOR RESEARCH IN ECONOMICS AND STATISTICS

Cyber risk modeling

Hawkes processes

The **number of claims** up to time **t (Frequency component)**

Cumulative loss process of a cyber-insurance portfolio $L(t) := \sum_{i=1}^{N_t} L_i$ The cost of i^{th} claim (Severity component)

- A **Hawkes process** is a counting process $N_t = \sum_{i \geq 1} 1_{T_i \leq t}$ with intensity:

$$\lambda_t = \lambda_0 + \sum_{T_i < t} \phi(t - T_i)$$

Baseline intensity — Excitation kernel

- The sum $\sum_{T_i<t} \phi(t - T_i)$ represents **the impact of past events** and captures the **self-excitation** property. Hawkes processes allow the modeling of **intense arrivals** and **regime changes**.
- **Self-excitation**, causality between attacks
- Adapted to the following **cyber datasets (PRC, VCDB, Hackmageddon)**
- **Parametric** and **tractable**

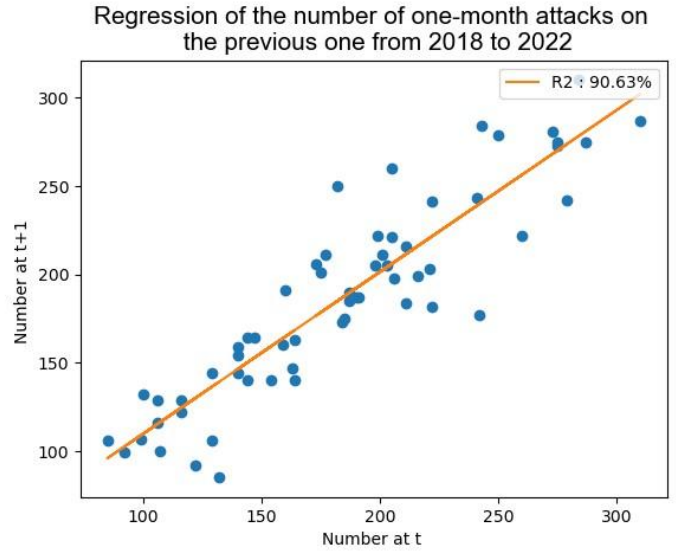

Hackmageddon database

Cyber risk modeling

Hawkes processes with external excitation

$\overline{N_t} = \sum_{k\geq 1} 1_{\overline{T_k}\leq t}$ vulnerabilities Poisson process with intensity $\boldsymbol{\rho}$

$$\lambda_t = \lambda_0 + \sum_{\overline{T_k}<t} \overline{\phi}\left(t - \overline{T_k}\right) + \sum_{T_i<t} \phi(t - T_i)$$

Baseline intensity

External excitation

Vulnerabilities dates

Internal excitation

Cyber attacks dates

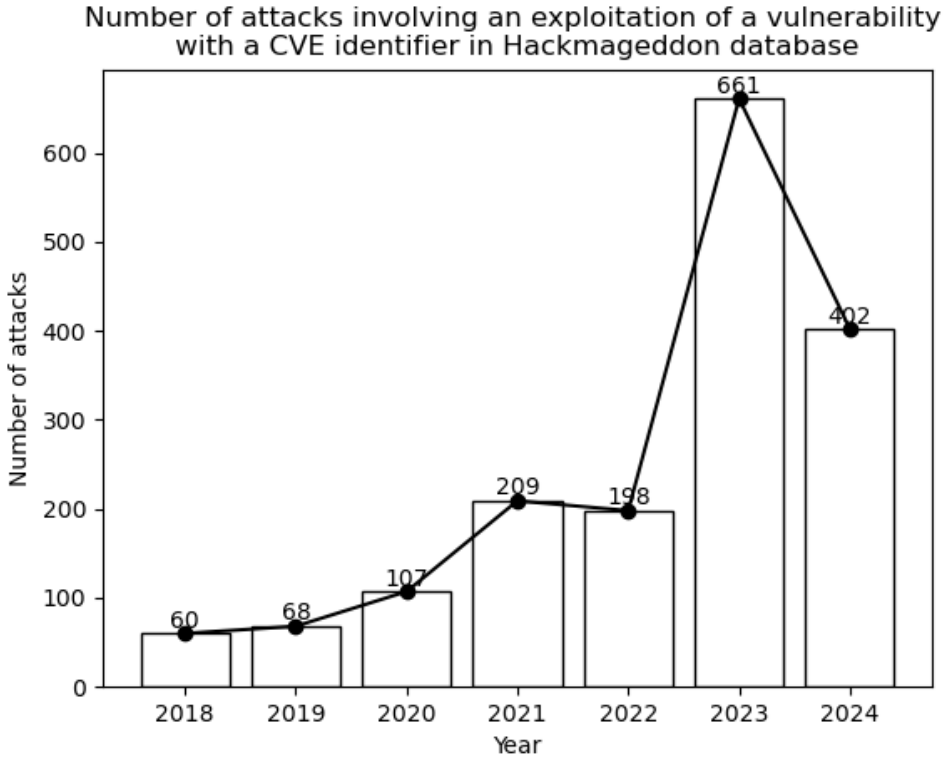

- *Dassios, A., & Zhao, H. (2011). A dynamic contagion process. Advances in applied probability, 43(3), 814-846.*
- Here the **external excitation** represents the regularly published **vulnerabilities** in public datatabases (**NVD database** for example)

Cyber risk modeling

Hawkes processes with external excitation

- **Two kernels :**
  1. **Exponential kernel** : instantaneous impact of the jump time on the intensity – Markovian intensity
  2. **Erlang kernel** : a delayed impact – Non markovian intensity

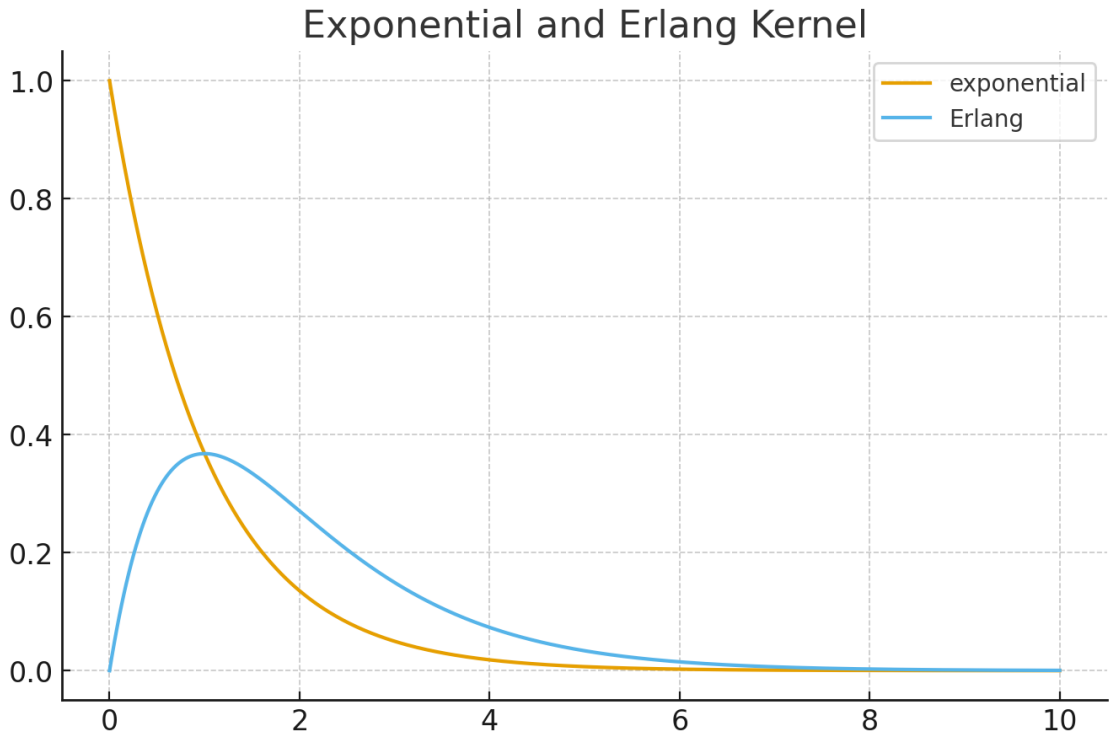

Cyber risk modeling

Hawkes processes with external excitation and stochastic marks

$$\lambda_t = \lambda_0 + \sum_{\overline{T_k}<t} \overline{\phi}\left(t - \overline{T_k}, \overline{Y_k}\right) + \sum_{T_i<t} \phi(t - T_i, Y_i)$$

Vulnerabilities marks

Cyber attacks marks

- Add **marks** into the modeling to account for **the uncertain impact** of vulnerabilities and **heterogeneity** in cyber attacks propagation
- An illustrative simulated example using the thinning algorithm **:**
  - External events (here cyber vulnerabilities) increase the intensity of the cyber attacks process leading to internal contagion
  - Heterogeneous impacts of vulnerabilities and cyber attacks

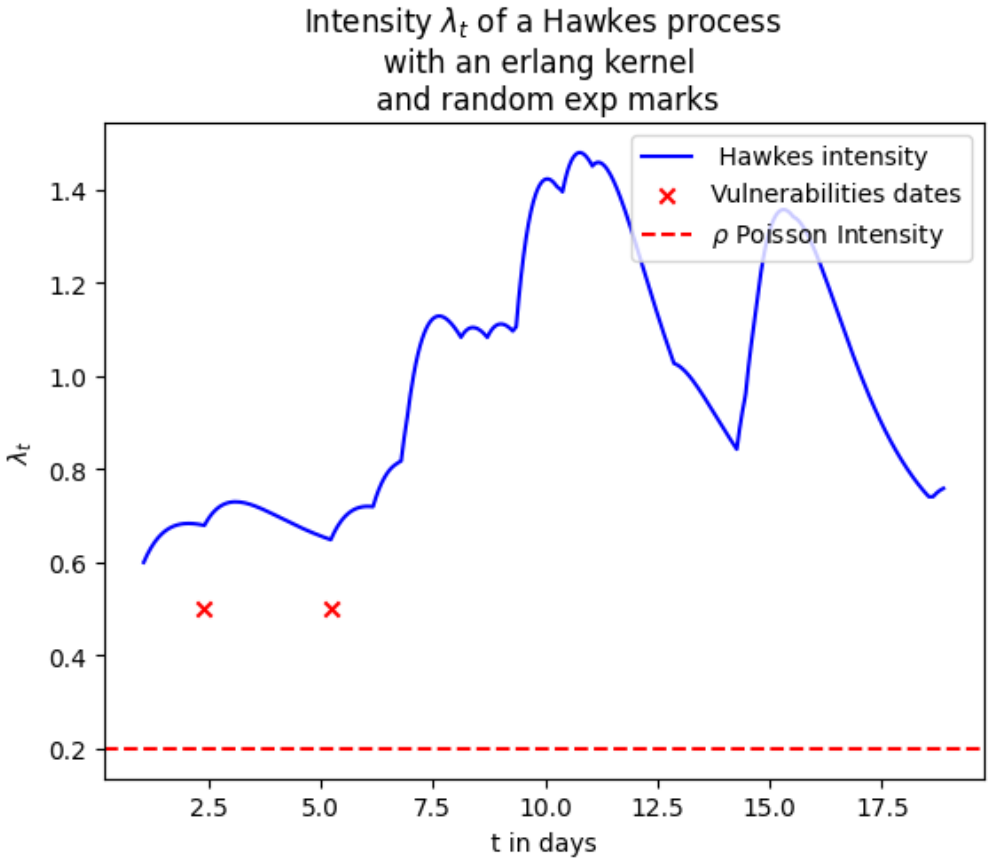

Calibration of a Hawkes Process with external excitation

Comparison of calibration methods on simulated data

Method	Data Requirements	Unicity of solution	Parameter estimation accuracy	$\| \phi \|$	$\frac{\rho \| \overline{\phi} \| + \lambda_0}{1 - \| \phi \|}$
MSE	**Number of attacks in given intervals**	**Multiple possible solutions**	**Misestimated**	**Misestimated**	**Accurate**
Likelihood	**Dates of attacks and vulnerabilities required**	**One solution**	**Accurate**	**Accurate**	**Accurate**

- $\| \boldsymbol{\phi} \|$ (the **endogeneity degree** of the system) represents the **average number** of attacks **an attack** will lead to.
- **The ergodicity ratio** ($\frac{\rho \| \overline{\phi} \| + \lambda_0}{1 - \| \phi \|}$) is preserved for the tested calibration techniques

Cyber databases

Vulnerability databases

- The NVD (National Vulnerability Database) is a US government database that lists computer vulnerabilities, while KEV (Known Exploited Vulnerabilities) lists vulnerabilities that hackers are actively exploited.

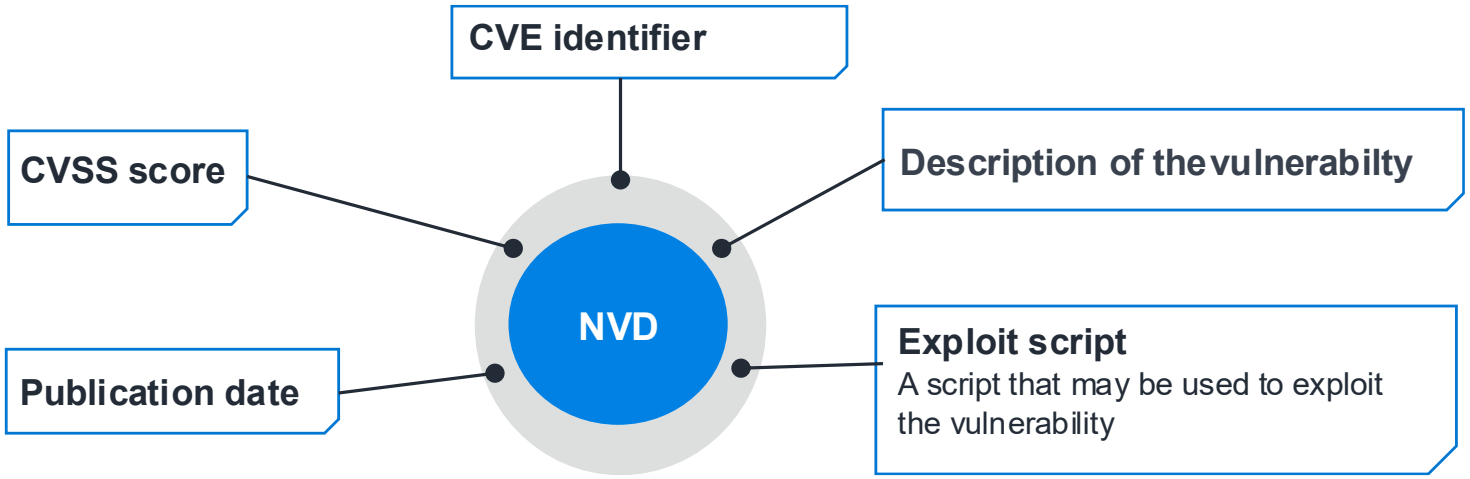

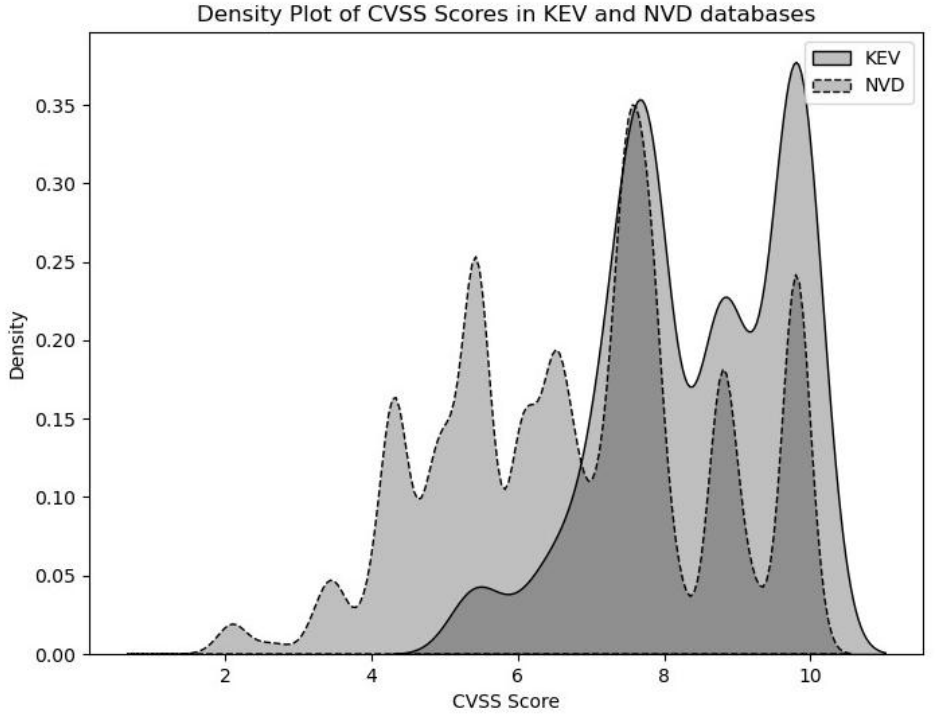

Cyber databases

Hackmageddon database

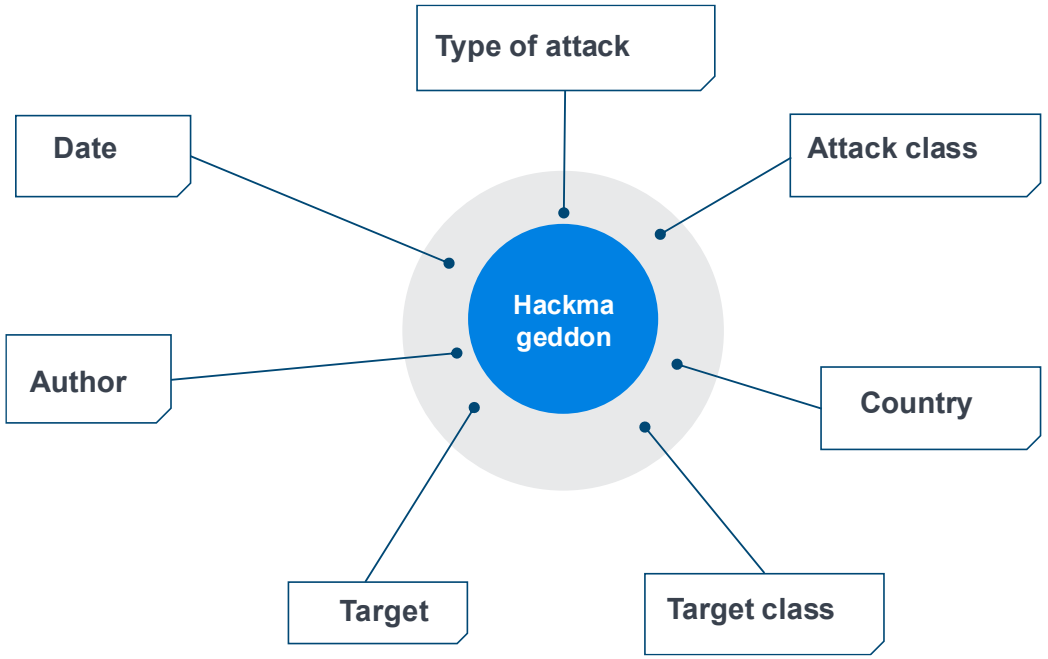

- **Hackmageddon** database :
  - Gathered by **Paolo Passeri** according to his **expert judgment**
  - **Different incidents** types, among these, incidents with known **CVE identifiers**, which we link to **vulnerability databases**

Different countries are represented in this database - The US is still the most represented

Calibration of the Hawkes process with external excitation

Calibration periods

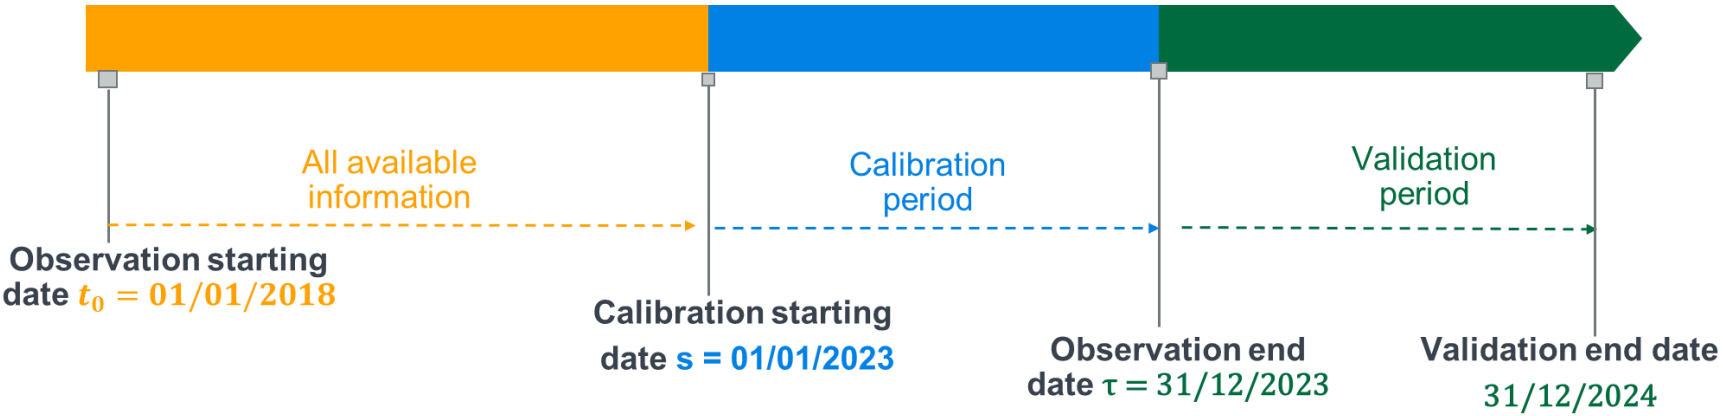

Year	Nb. of attacks
2018	1310
2019	1788
2020	2324
2021	2525
2022	3075
2023	4154
2024	4024

- Cyber risk is **rapidly evolving** :
  - **Calibration** is done on **2023** knowing **historical events** from **2018 to 2022**
  - **Validation** is done on **2024**

CREST
CENTER FOR RESEARCH IN ECONOMICS AND STATISTICS

Calibration of the Hawkes process with external excitation

Calibrated models

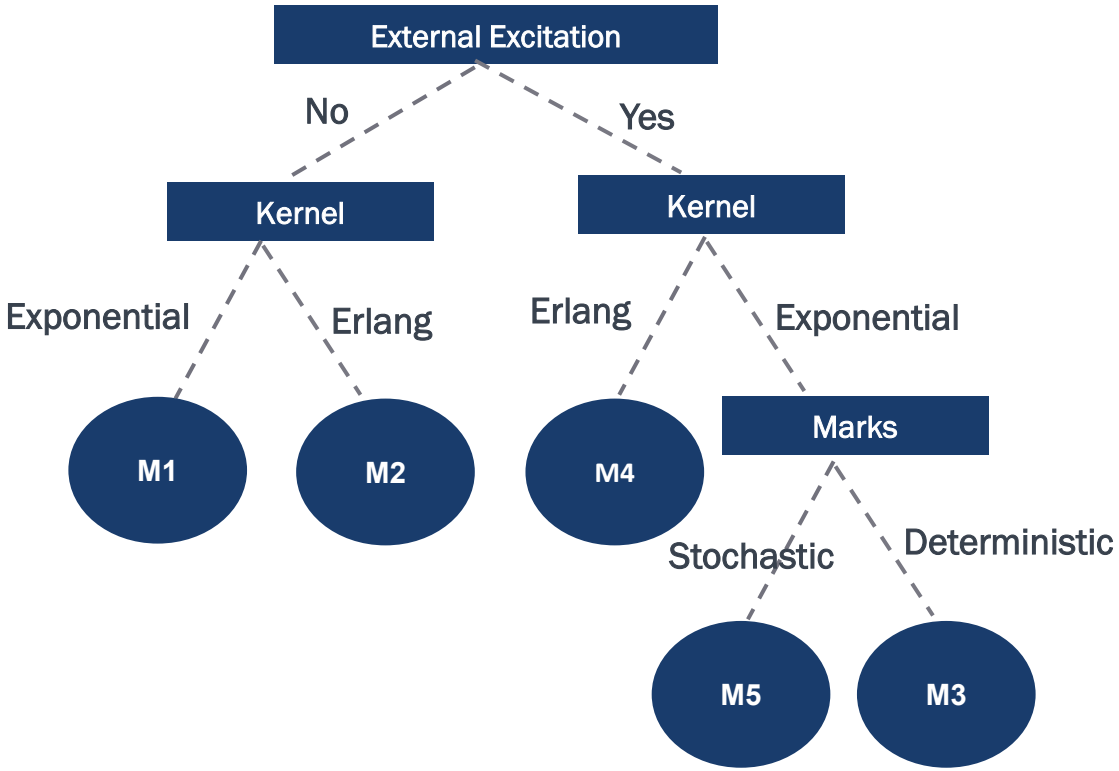

Model	Intensity
M1	$\lambda_t^{(1)} = \lambda_0 + \sum_{t_i<t} m e^{-\delta(t-t_i)}$
M2	$\lambda_t^{(2)} = \lambda_0 + \sum_{t_i<t} m(t-t_i) e^{-\delta(t-t_i)}$
M3	$\lambda_t^{(3)} = \lambda_0 + \sum_{\overline{t_k}<t} \overline{m} e^{-\delta(t-\overline{t_k})} + \sum_{t_i<t} m e^{-\delta(t-t_i)}$
M4	$\lambda_t^{(4)} = \lambda_0 + \sum_{\overline{t_k}<t} \overline{m}(t-\overline{t_k}) e^{-\delta(t-\overline{t_k})} + \sum_{t_i<t} m(t-t_i) e^{-\delta(t-t_i)}$
M5	$\lambda_t^{(5)} = \lambda_0 + \sum_{\overline{t_k}<t} \overline{Y_k} e^{-\delta(t-\overline{t_k})} + \sum_{t_i<t} Y_i e^{-\delta(t-t_i)}$

- For each model, the **parameters** to be **calibrated** are the following : $(\lambda_0, \rho, \overline{m}, m, \delta)$

Calibration of the Hawkes process with external excitation

Calibration results

Kernel	Marks	λ_0	ρ	$\overline{m}$	m	δ	$\|\phi\|$	$\|\bar{\phi}\|$
(1) Exponential	Det.	3.13	-	-	1.54	2.13	72.5%	-
(2) Erlang	Det.	4.07	-	-	2.63	1.33	148.67%	-
(3) Exponential	Det.	2.43	84.79	0.11	1.14	2.48	46.36%	4.43%
(4) Erlang	Det.	1.71	84.79	0.88	41.60	7.53	73.24%	1.55%
(5) Exponential	Stoch.	3.10	84.79	0.12	0.83	2.37	35.21%	5.06%

Kernel	Marks	$\frac{\lambda_0+\rho\cdot\|\bar{\phi}\|}{1-\|\phi\|}$	$\frac{\lambda_0+\rho\cdot\|\bar{\phi}\|}{\|\phi\|}$	$-\ln(\mathcal{L})$	AIC	BIC	p-value KS	
(1) Exponential	Det.	11.38	4.31	6214.01	12434.02	12438.88	0.0165	
(3) Exponential	Det.	11.54	13.35	31243.86	62497.72	62510.39	0.1733	
(4) Erlang	Det.	11.30	4.13	31227.15	62464.30	62476.97	**0.1853**	
(5) Exponential	Stoch.	11.41	20.99	**31208.07**	**62426.14**	**62438.81**	0.1467	

- $\| \phi \|$ is nearly **halved** between the **models with no external excitation** and the models with the **external excitation**
- Among the models with external excitation, the one with an **Erlang kernel** has the **highest p-value in the Kolmogorov–Smirnov** goodness of fit test.
- The model with an **exponential kernel and stochastic marks** achieves the lowest values of the **negative log likelihood, AIC and BIC.**

CREST
CENTER FOR RESEARCH IN ECONOMICS AND STATISTICS

Calibration of the Hawkes Process With External Excitation

Calibration results

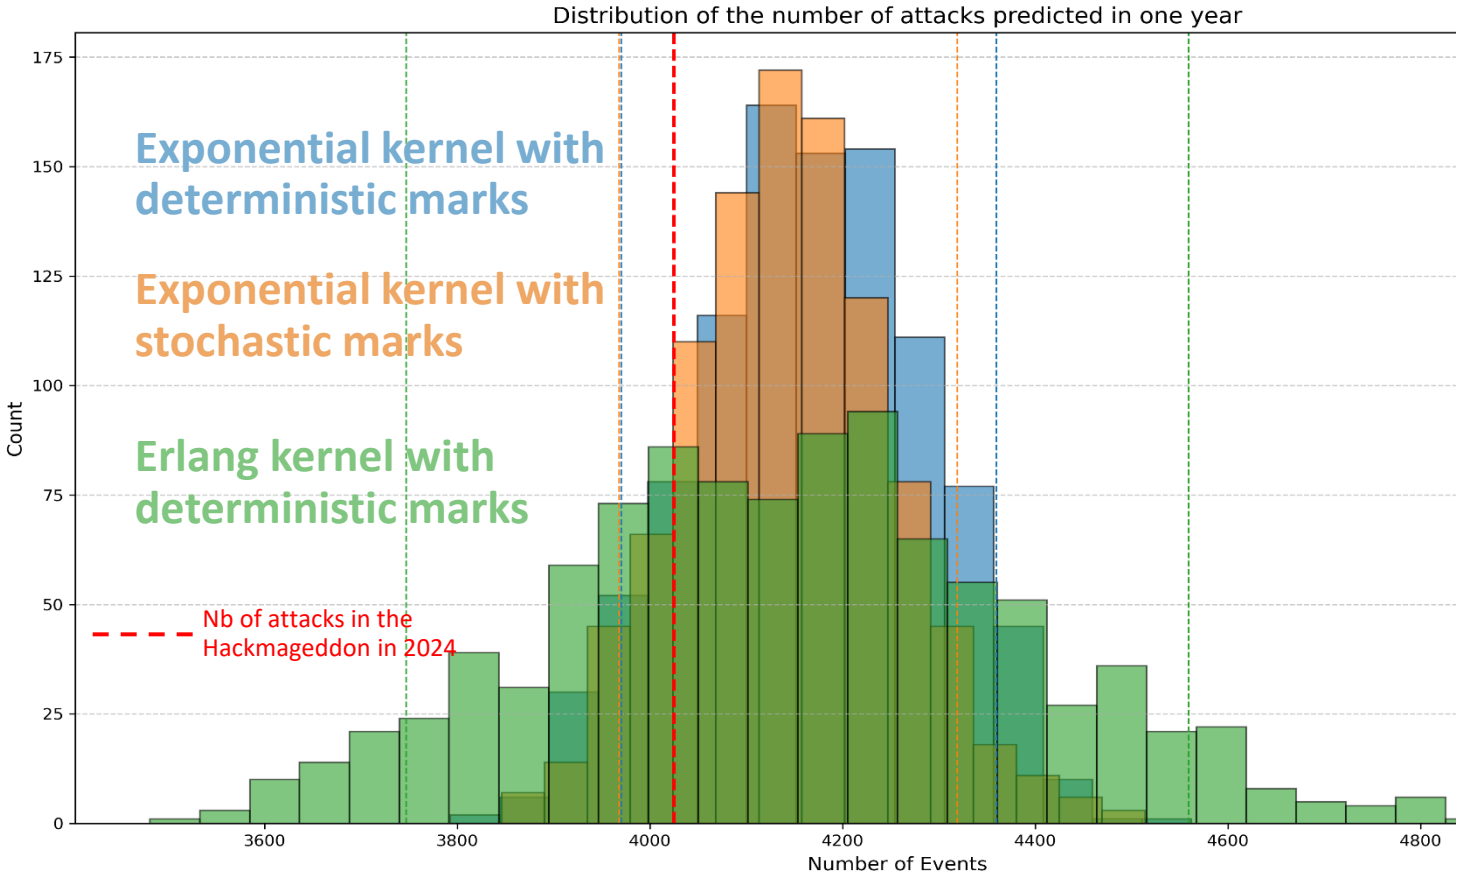

- The simulation is done using **10,000 trajectories** generated using the thinning algorithm
- The distributions seem to **capture the dynamics of cyber attacks in 2024** for the **Hackmageddon** database.
- The distribution based on **the Erlang kernel has a heavier tail**, which is more effective at **capturing extreme scenarios**.

A two-phase Hawkes process

- Continuous time counting process $N_t = \sum_{i \geq 1} 1_{T_i \leq t}$ with the following intensity :

$$\lambda_t = \begin{cases} \underbrace{\lambda_0}_{\text{Baseline intensity}} + \underbrace{\sum_{\overline{T_k}<t} \overline{\phi}(t - \overline{T_k}, \overline{Y_k})}_{\text{External excitation : cyber vulnerabilities}} + \underbrace{\sum_{T_i<t} \phi(t - T_i, Y_i^{bl})}_{\text{Self excitation : cyber attacks}} & \text{if } t \leq \underbrace{\ell}_{\text{Reaction time}} \\ \underbrace{\alpha_0 \lambda_0}_{\text{Reaction parameter}} + \underbrace{\alpha_1 (\lambda_{\ell^-} - \lambda_0)}_{\text{Reaction parameter}} & \text{if } t = \ell \\ \lambda_0 + \alpha_1 \left(\sum_{\overline{T_k}<\ell} \overline{\phi}(t - \overline{T_k}, \overline{Y_k}) + \sum_{T_i<\ell} \phi(t - T_i, Y_i^{bl}) \right) + \sum_{\ell<T_i<t} \phi(t - T_i, Y_i^{al}) & \text{if } t > \ell \end{cases}$$

- **2 phases :**
  - ℓ is the start of the second phase of N_t

A two-phase Hawkes process

Protection measures

- Continuous time counting process $N_t = \sum_{i \geq 1} 1_{T_i \leq t}$ with the following intensity :

$$\lambda_t = \begin{cases} \underbrace{\lambda_0}_{\text{Baseline intensity}} + \underbrace{\sum_{\overline{T_k} < t} \overline{\phi}(t - \overline{T_k}, \overline{Y_k})}_{\text{External excitation : cyber vulnerabilities}} + \underbrace{\sum_{T_i < t} \phi(t - T_i, Y_i^{bl})}_{\text{Self excitation : cyber attacks}} & \text{if } t \leq \underbrace{\ell}_{\text{Reaction time}} \\ \underbrace{\alpha_0 \lambda_0}_{\text{Reaction parameter}} + \underbrace{\alpha_1 (\lambda_{\ell^-} - \lambda_0)}_{\text{Reaction parameter}} & \text{if } t = \ell \\ \lambda_0 + \alpha_1 \left(\sum_{\overline{T_k} < \ell} \overline{\phi}(t - \overline{T_k}, \overline{Y_k}) + \sum_{T_i < \ell} \phi(t - T_i, Y_i^{bl}) \right) + \sum_{\ell < T_i < t} \phi(t - T_i, Y_i^{al}) & \text{if } t > \ell \end{cases}$$

- **Self-excitation :**
  - $\{Y_i^{bl}\}_{i=1,2,\dots}$ is a sequence of iid positive jumps sizes before time ℓ
  - $\{Y_i^{al}\}_{i=1,2,\dots}$ is a sequence of iid positive jumps sizes after time ℓ

A two-phase Hawkes process

Protection measures

- Continuous time counting process $N_t = \sum_{i \geq 1} 1_{T_i \leq t}$ with the following intensity :

$$\lambda_t = \begin{cases} \underbrace{\lambda_0}_{\text{Baseline intensity}} + \underbrace{\sum_{\overline{T_k}<t} \overline{\phi}(t-\overline{T_k}, \overline{Y_k})}_{\text{External excitation : cyber vulnerabilities}} + \underbrace{\sum_{T_i<t} \phi(t-T_i, Y_i^{bl})}_{\text{Self excitation : cyber attacks}} & \text{if } t \leq \underbrace{\ell}_{\text{Reaction time}} \\ \underbrace{\alpha_0}_{\text{Reaction parameter}} \lambda_0 + \underbrace{\alpha_1}_{\text{Reaction parameter}} (\lambda_{\ell^-} - \lambda_0) & \text{if } t = \ell \\ \alpha_0 \lambda_0 + \alpha_1 \left(\sum_{\overline{T_k}<\ell} \overline{\phi}(t-\overline{T_k}, \overline{Y_k}) + \sum_{T_i<\ell} \phi(t-T_i, Y_i^{bl}) \right) + \sum_{\ell<T_i<t} \phi(t-T_i, Y_i^{al}) & \text{if } t > \ell \end{cases}$$

- **Response phase:**
  - Modulating the baseline intensity λ_0 (through $\alpha_0 \in [0,1]$ parameter) and the self-excitation component from past attacks (through $\alpha_1 \in [0,1]$parameter)
  - Reducing the impact of future attacks (after ℓ) through Y_i^{al}

A two-phase Hawkes process

Protection measures

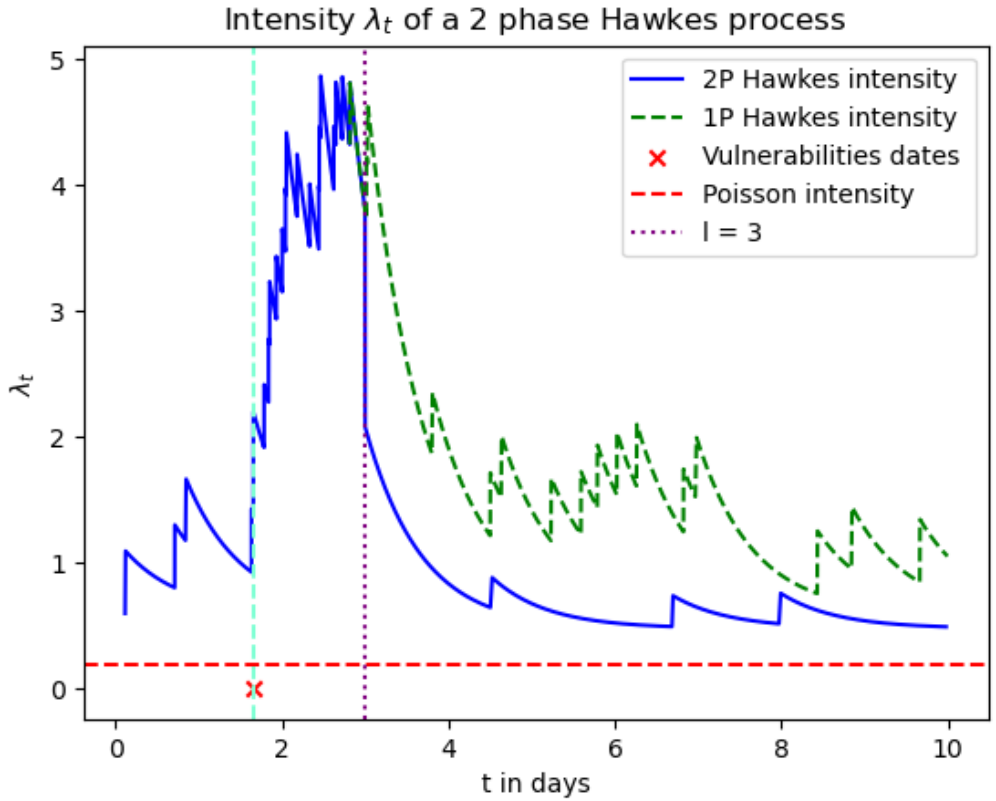

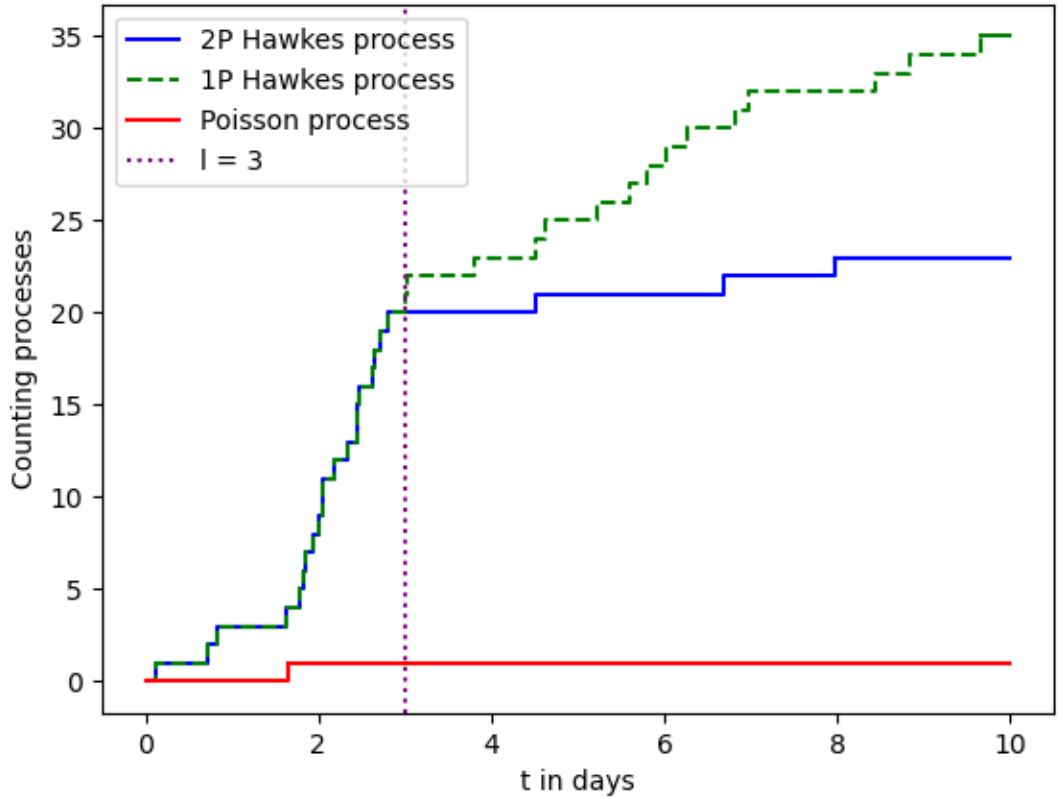

Choice of the assistance capacity

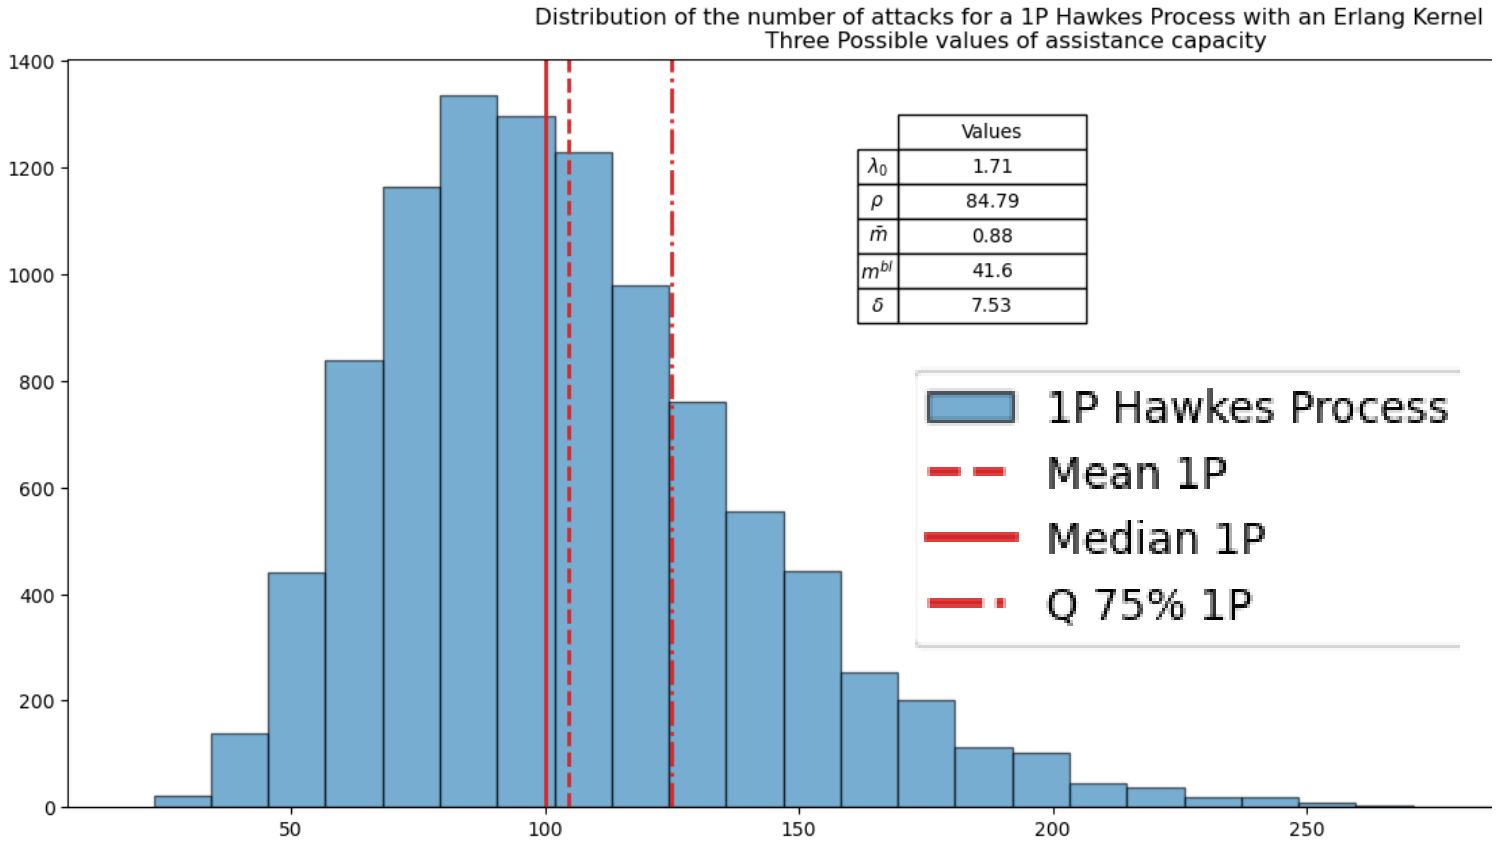

- The goal is to simulate a **cyber pandemic scenario (10 days) with reaction measures** and with **a limited cyber assistance capacity**
- **The insurer's assistance capacity** is chosen using **the distribution of attacks** generated by **the calibrated 1P Hawkes process**, simulated over a **10-day horizon**
- For illustration purposes, we choose the median here. The same analysis was done for the two other quantiles

Choice of the reaction parameters

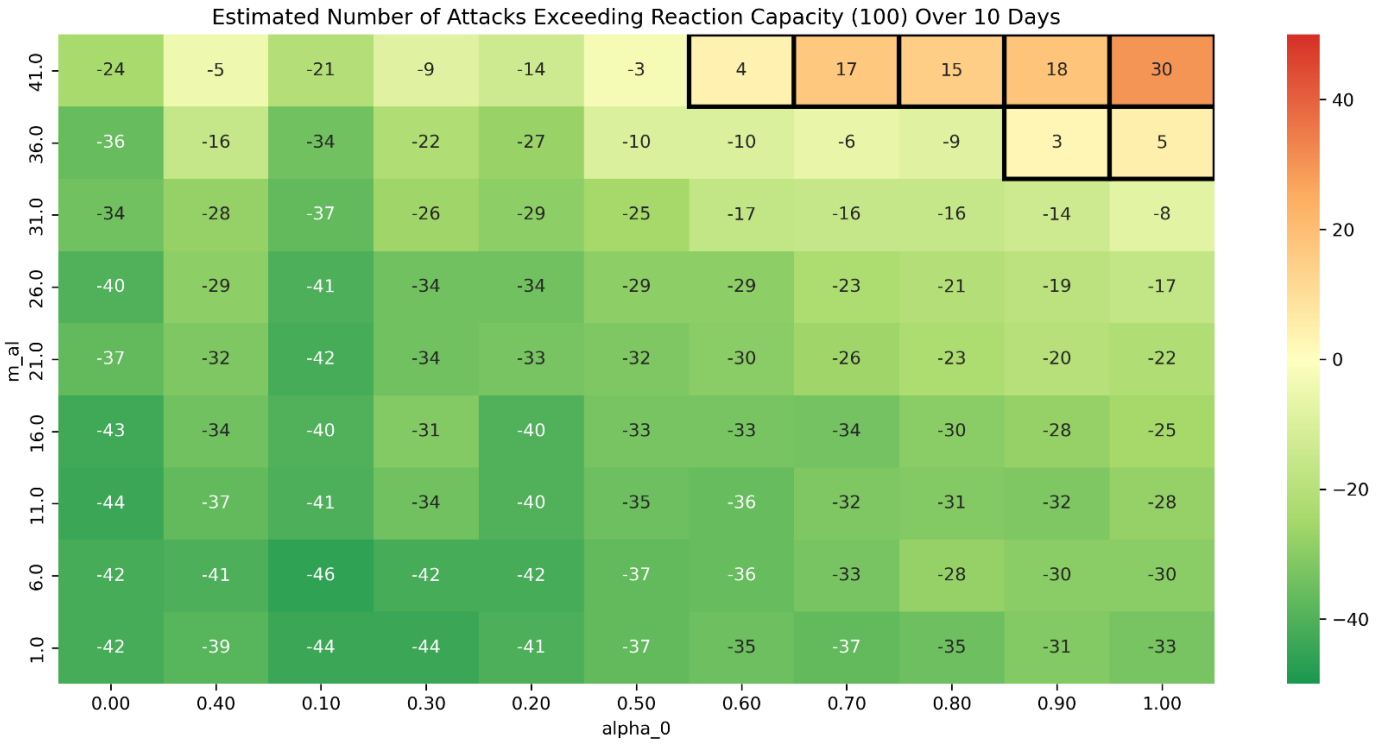

Heatmap for the difference between the 95% quantile of the simulated distribution and the selected assistance capacity (100)

- The goal is to find an **optimal set of reaction parameters** such that **the assistance capacity is not exceeded with 95% probability**
- $\alpha_1 = 1$: a sensitivity analysis shows that its impact is negligible
- A grid search is performed over $(\boldsymbol{\alpha_0}, \boldsymbol{m^{al}})$
- Cells in **green** indicate that the **95% quantile is below the capacity**. Cells in **red with a bold border indicate configurations where the capacity is exceeded.**
- When **the assistance capacity is low**, both $\boldsymbol{\alpha_0}$ and $\boldsymbol{m^{al}}$ (the average contagiousness of events after the introduction of protection measures) **should be reduced simultaneously.**

Future research

- **Future directions**
  - Enrich datasets using **AI/LLMs** (automatic extraction from cybersecurity reports, additional threat-intelligence sources).
  - Model **reaction times as intensity-dependent**, opening the way to **parametric cyber insurance triggers**.
  - Extend the model with **self-exciting vulnerabilities**, other kernels and external excitation for CART algorithms, and **multivariate Hawkes models** after CART group formation.

CREST
CENTER FOR RESEARCH IN ECONOMICS AND STATISTICS