Bayesian Inference of Delayed Neutron Multigroup Constants: Application to 9 MeV 238U Photofission

Johann Piekar, Aly Elayeb, Adrien Sari, Frédérick Carrel

Active nuclear measurement: from signal to quantification

ACTIVE NUCLEAR ASSAY

Non-Destructive Assay (NDA)

Probe: PHOTONS (γ, X) — LINAC / Accelerator	Probe: NEUTRONS (n) — Generator / Isotopic Source
High-Energy Imaging — Radiography / Tomography	**Neutron Imaging** — Transparency to metals / H
Photofission (API) — Active detection (U, Pu)	**Active Coincidence Counting** — Fissile Quantification (Mass)
Nuclear Resonance Fluorescence (NRF) — Specific Isotopic Signature	**Differential Die-Away (DDA)** — High Sensitivity (Waste)
Photon Activation Analysis (PAA) — Elemental Analysis	**Activation (PGNAA / PFTNA)** — Stoichiometry / Chemical Analysis

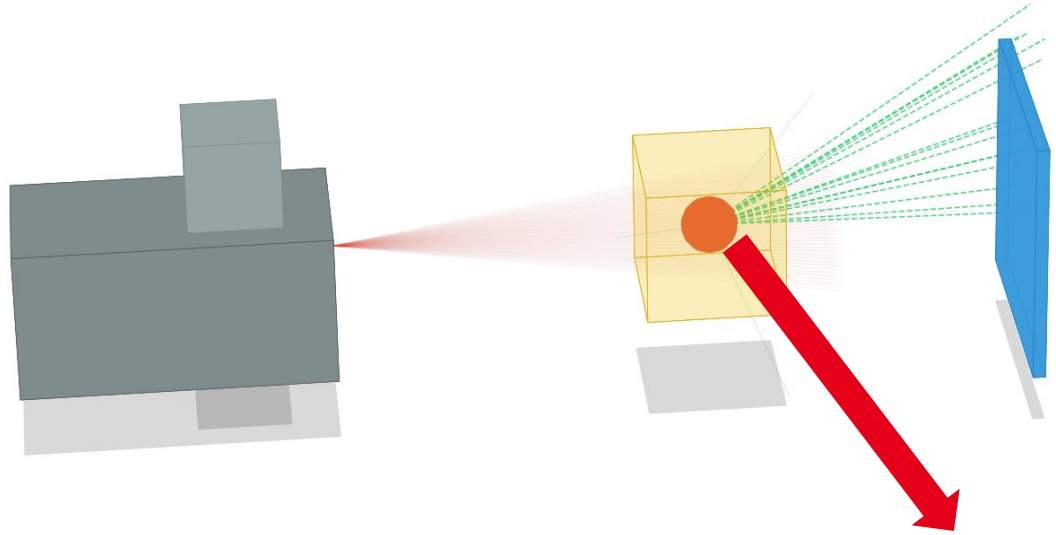

Photon source (Linac, ICS, LPA) → **Object** (waste, U/Pu) → **Detector** (delayed neutrons) → **Signal** (time counts) → **Inverse model** (nuclear parameters) → **Quantity** (mass, ratios, signature)

Nuclear data

$\sigma_{\gamma f}(E)$, $\sigma_{\gamma n}(E)$, a_i, λ_i, covariances

Industrial applications

Nuclear waste characterization

Border control

Active interrogation needs models to transform signals into quantities.

Photofission and delayed neutrons

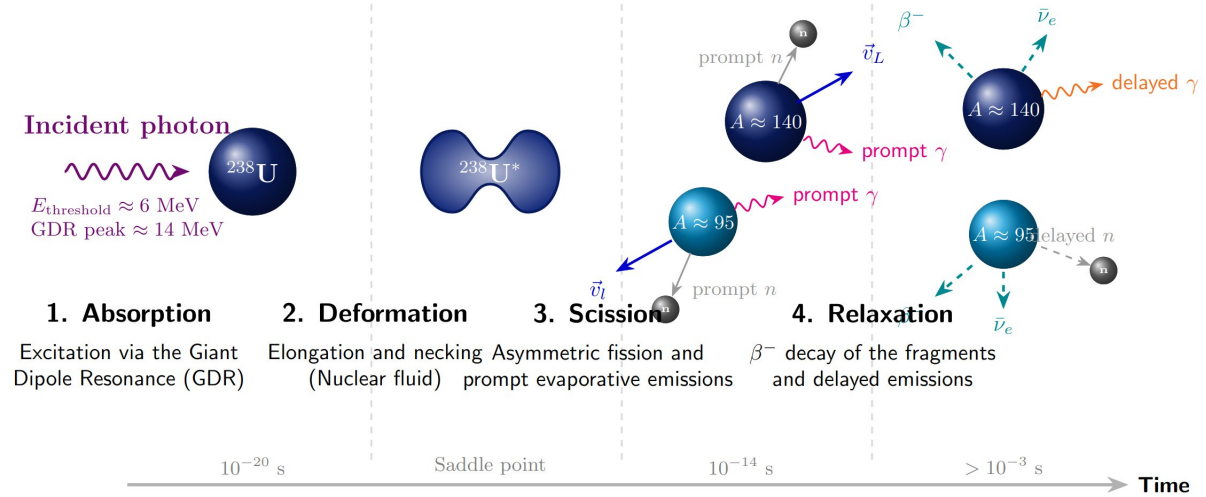

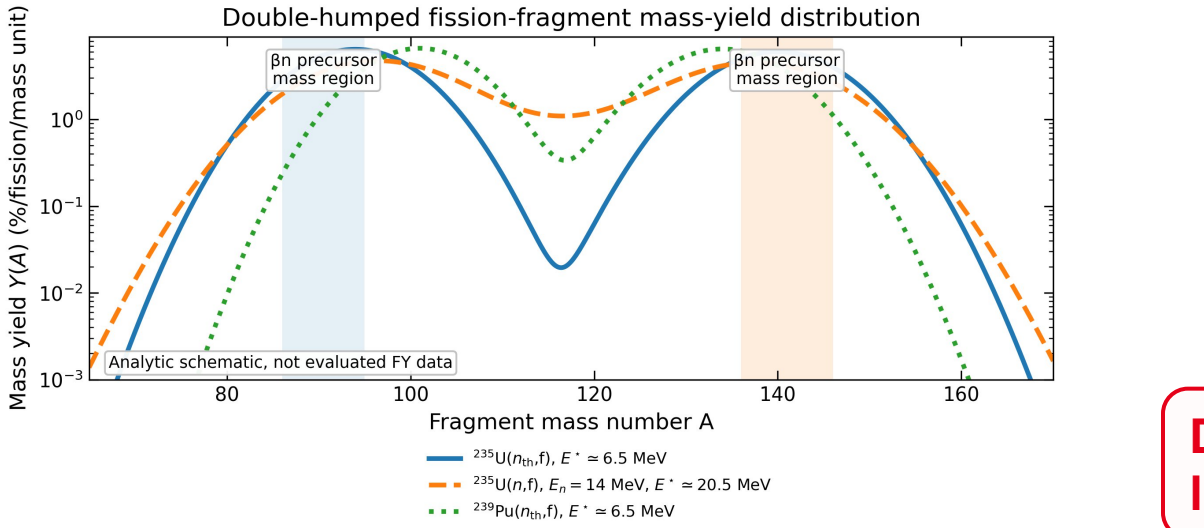

Delayed neutrons as a time signature

Photofission produces neutron-rich fission fragments.

A small fraction of fragments emits neutrons after β decay.

Hundreds of precursors can be represented by a few effective groups.

≈250–300 precursors

≈20–30 dominant contributors

6 or 8 effective groups

Isotopic & Energy Dependence

Variable Yields: Different **actinides** do not produce the same proportions of βn-precursors.
Energy Dependence: Incident photon energy spectrum dictates the final precursor population.
Exploitable Signature: This dual dependence allows for the separation of signatures between different fissile/fertile materials.

Delayed neutrons provide a time signature linked to fission fragment populations.

State of the art

	Reference	Extraction method	Photon source	Actinides	Best use / limitation
GROUP PARAMETERS	**Nikotin & Petrzhak 1965/66**	6-group effective multi-exponential fit	Bremsstrahlung 15 MeV	^{232}Th, ^{235}U, ^{238}U, ^{239}Pu	Historical benchmark; limited methodological detail
	Kull et al. 1970	6-group fit of time spectra	Bremsstrahlung 8 and 10 MeV	^{235}U, ^{238}U	Low-energy benchmark; PNNL average has no uncertainties
	Doré et al. 2006	Sequential fit + ABLA/CINDER summation	ELSA linac <20 MeV	^{238}U, ^{232}Th	CEA PhotoNuc reference below 20 MeV
	Doré et al. ND2007	PhotoNuc synthesis across irradiation/counting windows	ELSA linac 11–19 MeV	^{232}Th, ^{235}U, ^{238}U, ^{237}Np	Best cross-actinide comparison
	Macary et al. 2009	6-group fit of time spectra	ELSA linac 15 and 18 MeV	^{235}U, ^{237}Np	Complements PhotoNuc for fissile/minor actinides
	Sari et al. 2023	6-group fit + 8-group fit with fixed Spriggs half-lives	SAPHIR linac 9 MeV	^{238}U / depleted U 241 g	Closest dataset to 9 MeV active interrogation
NORMALIZATION	**Moscati & Goldemberg 1962**	Total delayed-neutron yield	Betatron 22 MeV	^{238}U, ^{232}Th	Normalization reference; not group constants
	Caldwell & Dowdy 1975	Prompt/delayed multiplicities	Bremsstrahlung 8–13 MeV	^{232}Th, $^{233–236}$U, ^{238}U, ^{237}Np, ^{239}Pu	Normalization reference; not group constants

Current Limitations:

- ***Sparse and legacy literature*** *with highly* ***heterogeneous*** *sources and samples.*
- **Low transparency** in parameter extraction methods.

Main Knowledge Gaps:

- **Complete absence** of **covariance matrices** (modern standard).
- **Beam monitoring:** No beam stability tracking for short-lived groups.
- ***Bremsstrahlung*** **data:** Available data restricted to conventional linac source.
- **8-group framework:** single reference standard (Sari et al. in photofission)

Nikotin
15 MeV
^{232}Th, ^{235}U, ^{238}U, ^{239}Pu
1965/66

Kull
8 and 10 MeV
^{235}U, ^{238}U
1970

Doré et al. Macary et al.
15–18 MeV (ELSA)
^{232}Th, ^{235}U, ^{238}U, ^{237}Np
2006-09

Sari
9 MeV
^{238}U / DU
2023

U-238 relative yields

The dominant group remains near $T_{1/2} \approx 2$ s, but several a_i differences exceed quoted uncertainties.

Observed pattern

- G4 (~2 s) carries ~31–41% across datasets
- The shortest groups G5/G6 vary strongly

21 compatible	1 tension
11 >3σ	Pairwise tests where both uncertainties exist

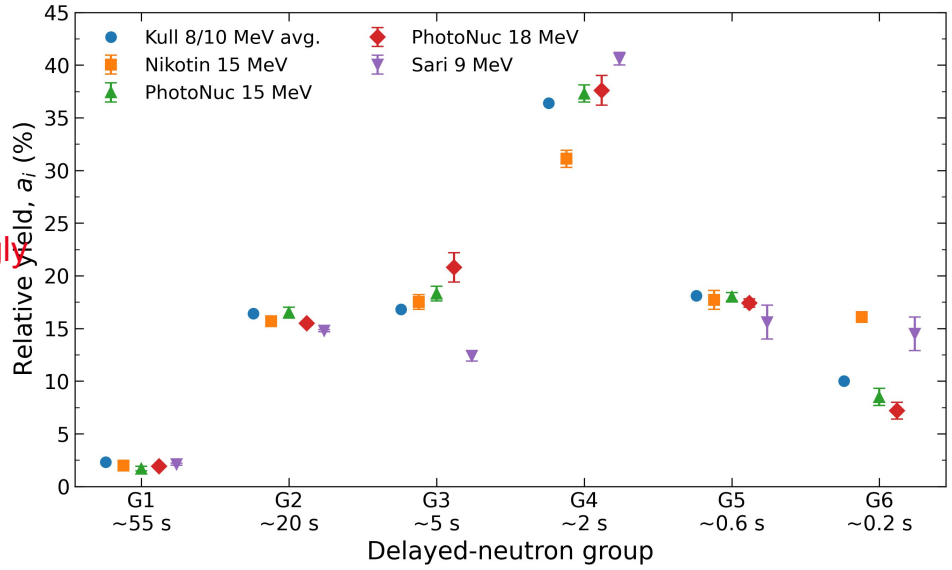

Actinide fingerprints at 15 MeV endpoint

Same-source comparisons are cleaner: changes mainly reflect actinide-dependent fission-fragment yields.

Practical use

Useful for qualitative isotope fingerprints, but absolute quantification needs source-specific calibration.

- **All actinides peak near G4**
- Fertile ^{238}U has a stronger short-lived G5–G6 contribution than ^{235}U and ^{237}Np
- Thorium differs most in the long-lived and short-lived groups

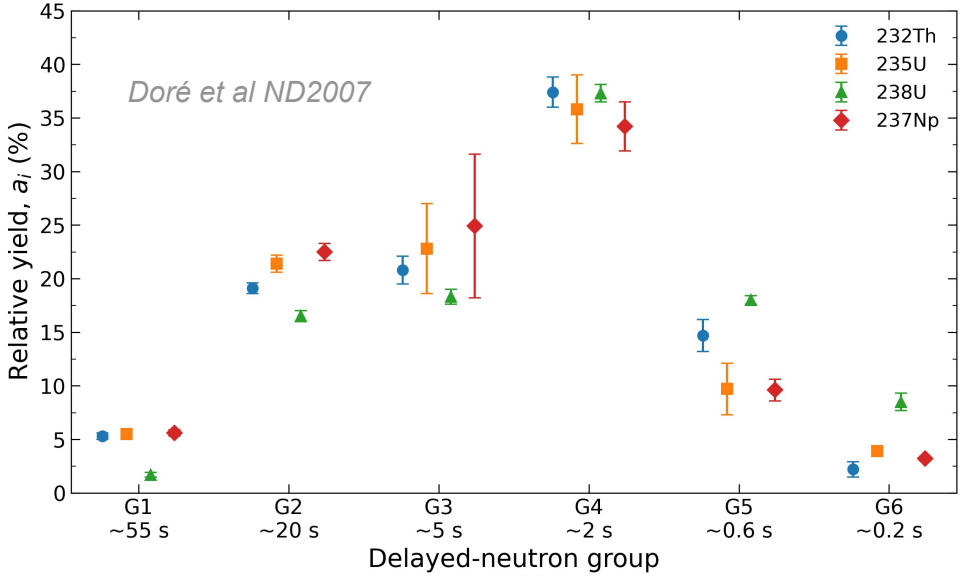

Data: Kull 1970 via PNNL 8/10 MeV average; Nikotin as quoted by Doré et al., ND2007", ; Sari 2023 at 9 MeV.

In this work, we propose...

A pilot 9 MeV 238U photofission experiment with traceable Bayesian inversion.

Scientific motivation

- **Heterogeneous literature**
  Published photofission yields come from different sources, samples and extraction strategies. Full covariance is **not** available.
- **Traceable Bayesian inversion**
  Explicit forward model + joint Bayesian fit, with assumptions and correlations carried through the analysis.
- **Uncertainty-aware output**
  Preliminary group yields reported with posterior intervals and covariance.

Data Pipeline

- **U** — **238U pilot data**
  9 MeV photofission delayed-neutron time spectra
- **t** — **Cyclic windows**
  multi-scale time sampling
- **Σ** — **Forward model**
  Bateman kinetics + pulse-history source normalization
- **p** — **Bayesian inversion**
  MCMC sampling of parameter sets
- **α** — **Preliminary group yields**
  relative α_i, posterior intervals and covariance

Experimental setup

- *Batched disk I/O, PC/FPGA drift monitoring*
- ***Time resolution:** 10 kHz sampling (100 µs bins) grouped into 10 ms network frames.*
- ***Discrimination (-15.7 / -8.0 mV):** hardware amplitude thresholds for pulse detection and event counting across channels.*

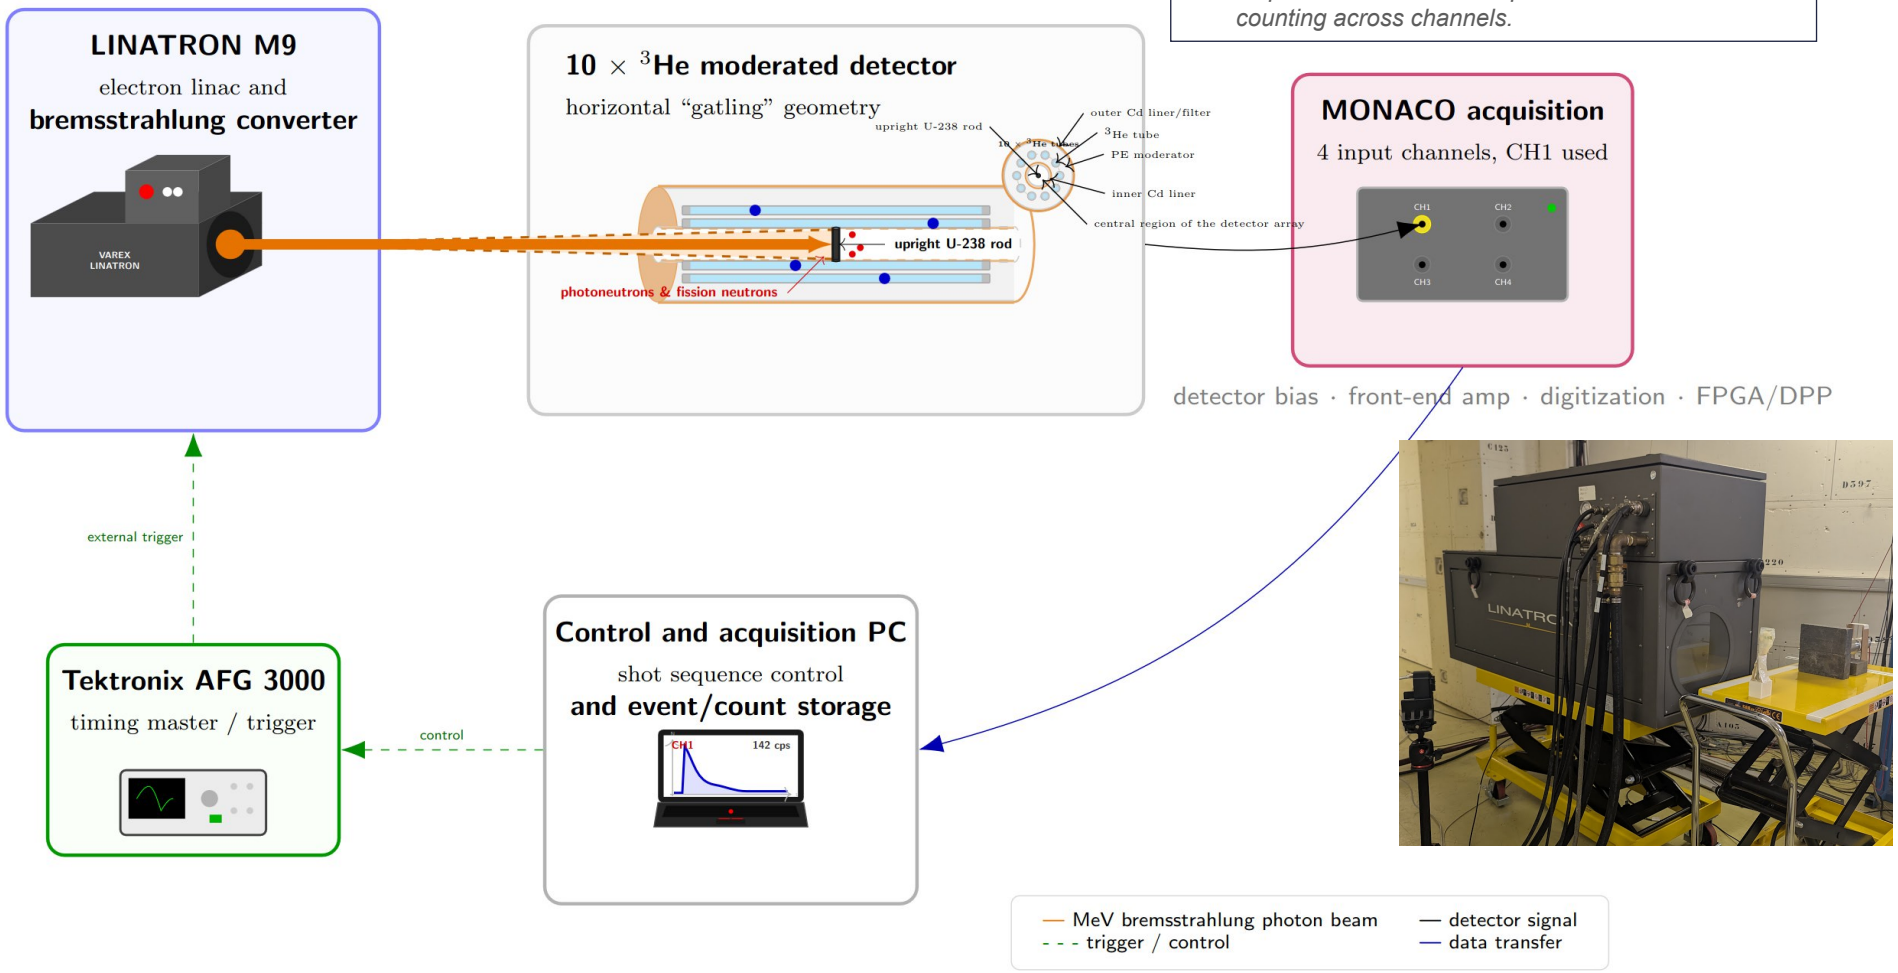

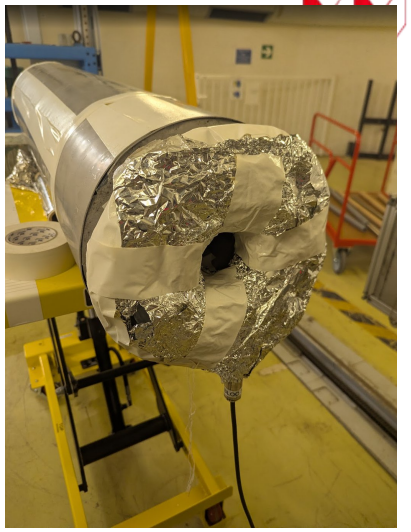

^{3}He moderated setup with state-of-the-art EM shielding.

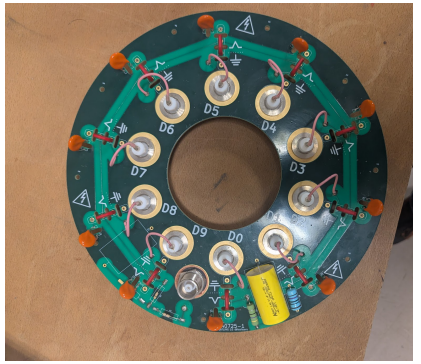

PCB developed to connect multiple ^{3}He detectors

Credit: Louis Garnaud, Wilfrid Husson

Irradiation plan:

Use several cyclic acquisitions to probe different delayed-neutron time scales while preserving statistics.

Target	Cycle		Repeats
Short groups	1 s ON at 200 Hz | 5 s at 1 Hz	~50 min	×500
Medium groups	20 s ON at 200 Hz | 40 s at 1 Hz	~50 min	×50
~~**Long groups**~~	~~5 min ON at 200 Hz | 5 min OFF~~	~~5 h~~	~~×20~~

Long-group mode: no 1 Hz ping during counting, to keep the decay window clean.

1 Hz is out of manufacturer specs → **Beam-monitoring**

Work in progress

Our electronics died

Linac constraint: True beam-OFF (> 1 s) triggers an automatic interlock fault, preventing standard ON/OFF gating.
1 Hz ping solution: Sustained during decay windows to keep the Linac active.

Temporal Structure of the Detected Signal

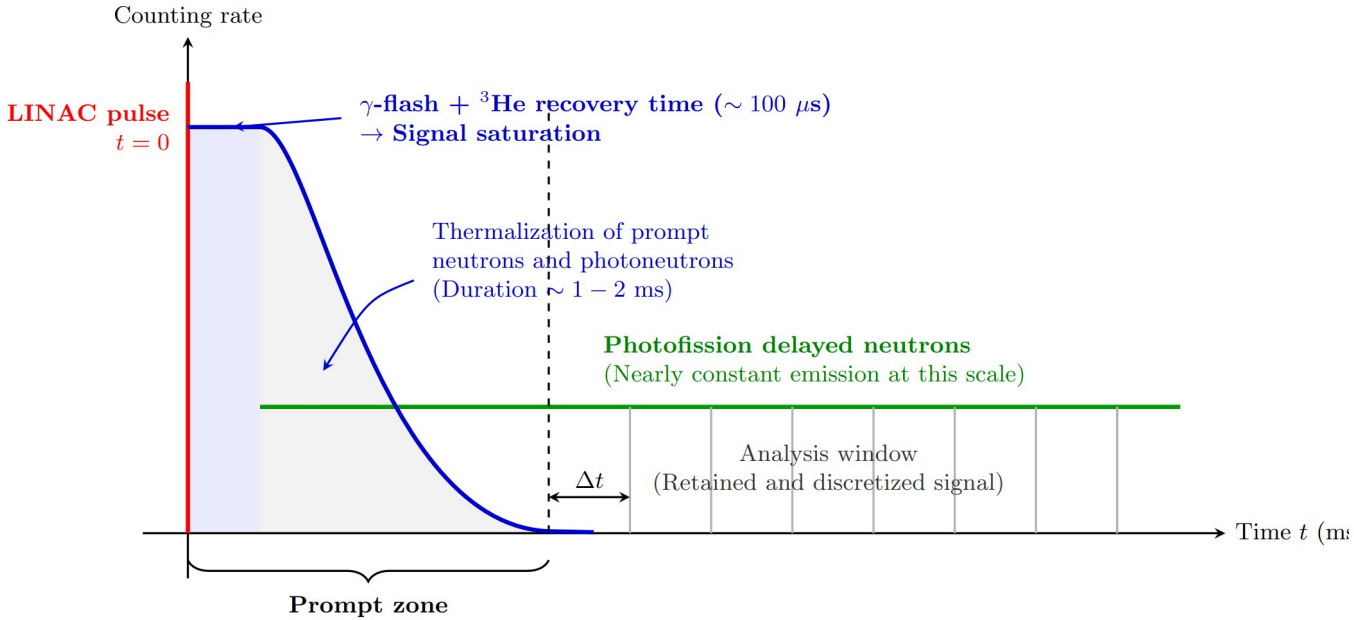

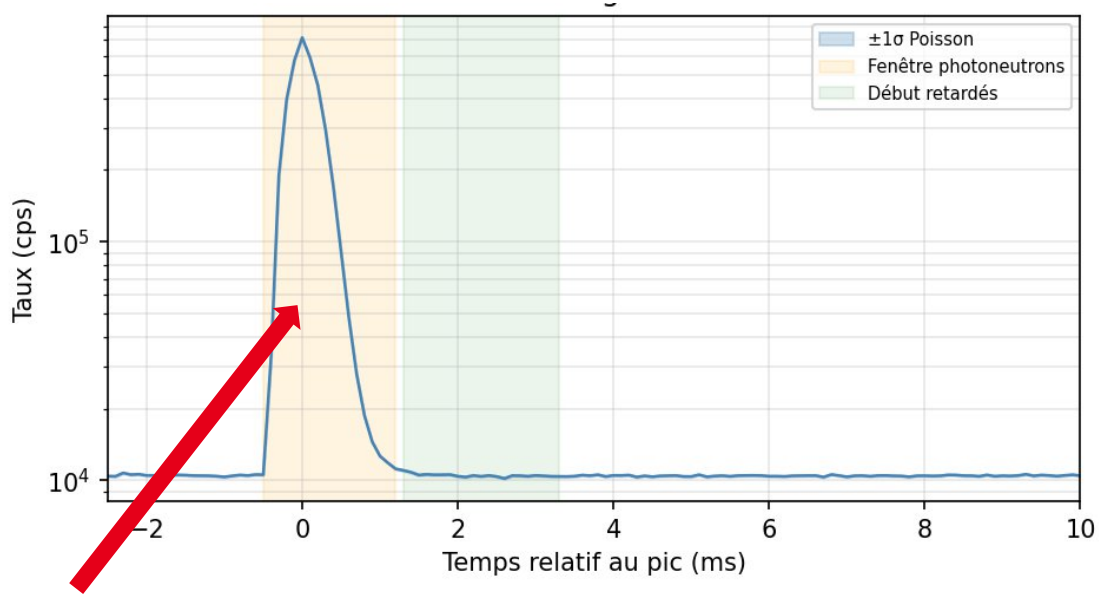

- **Beam normalization**: Prompt signal integral used to monitor LINAC as a proxy for the number of induced photofissions.

Timing chain

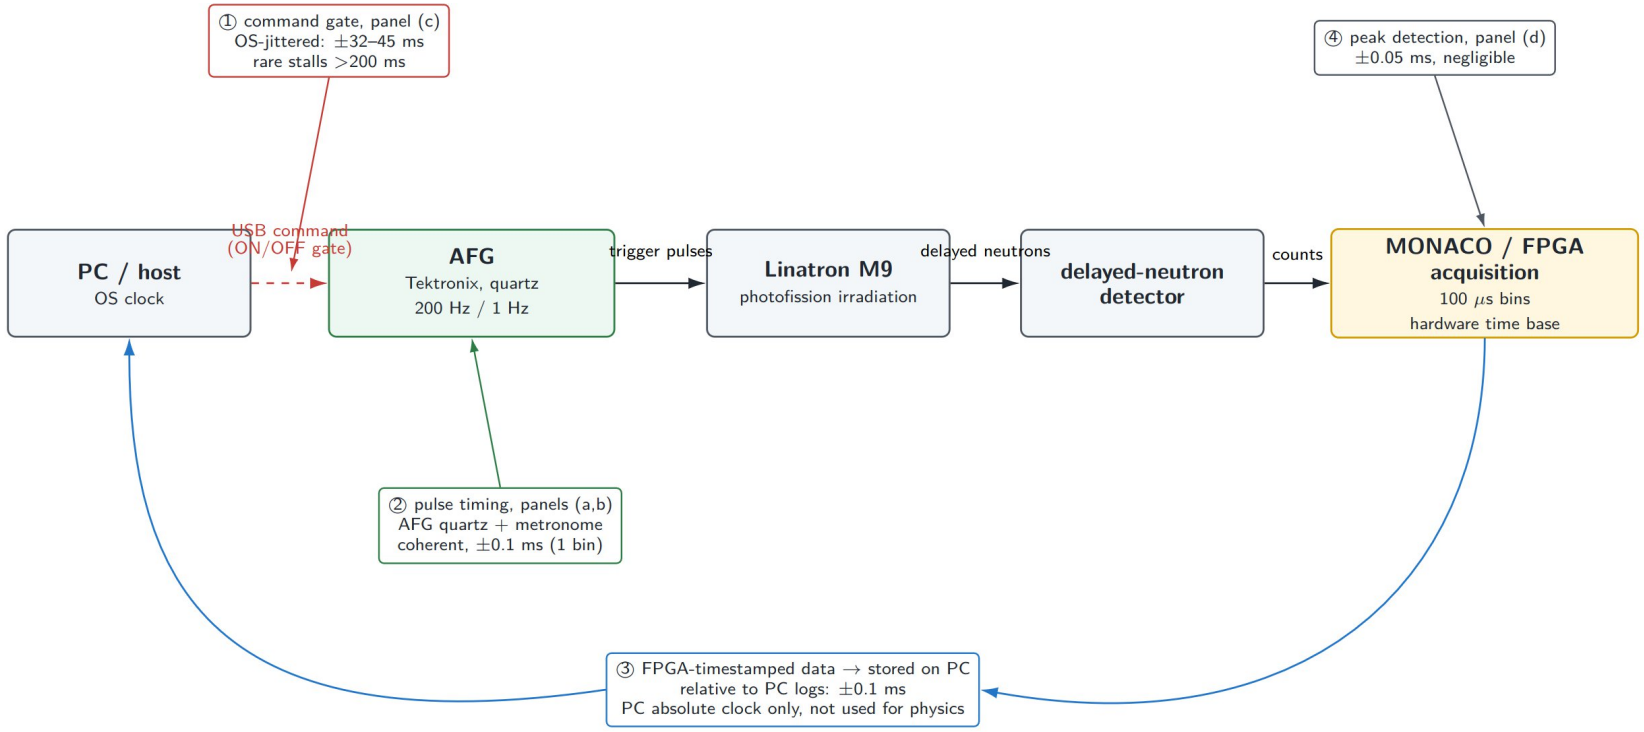

Only the command path is OS-jittered. Pulse generation and counting are hardware-clocked;
physics analysis uses FPGA timestamps and the AFG timing reference.

Timing chain: identifying the dominant latency source

Latency sources

- 1 Hz pulse train is stable once the gate is defined.
- Peak detection jitter is negligible compared with the bin width.
- Dominant uncertainty: software/USB gate latency.
- Heavy-tail delays can reach the 100 ms scale.

Final campaign

Low-latency Python control:

- dedicated core, elevated priority,
  no GC or blocking I/O in critical gates,
  pre-allocated buffers and monotonic timestamps.

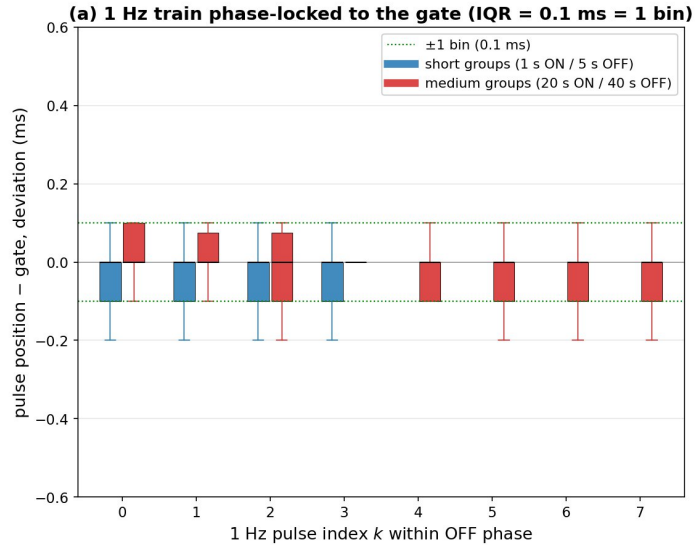

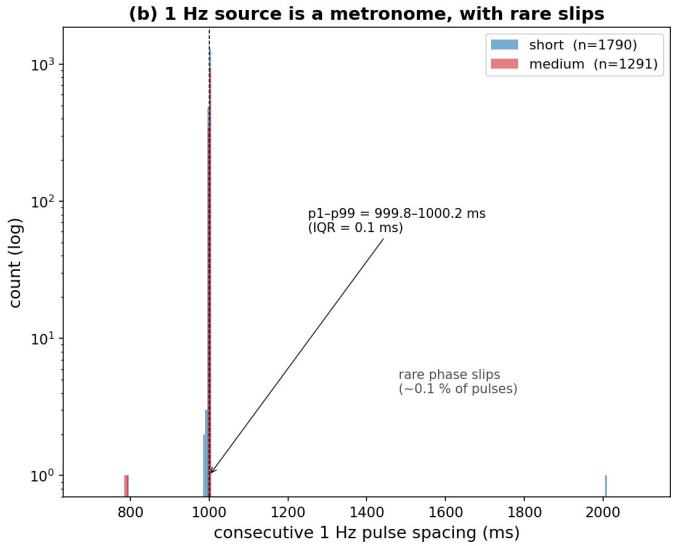

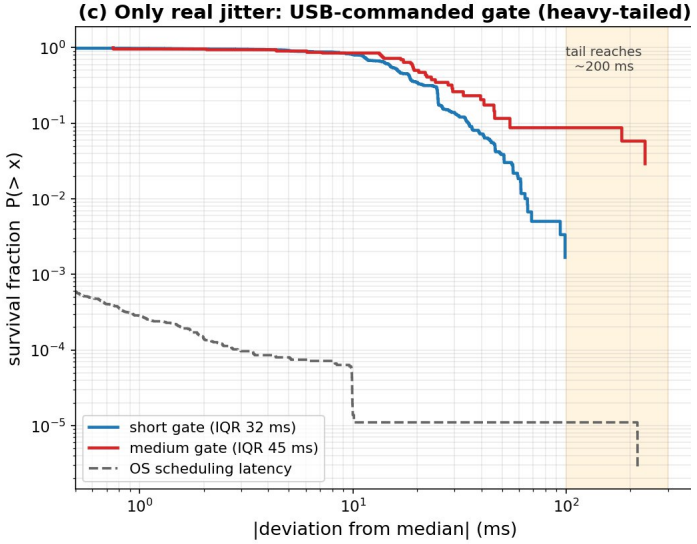

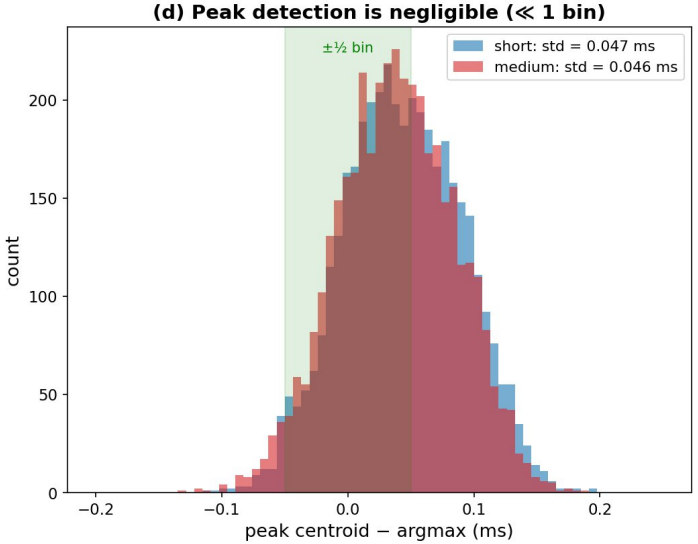

Beam Stability & Monitoring

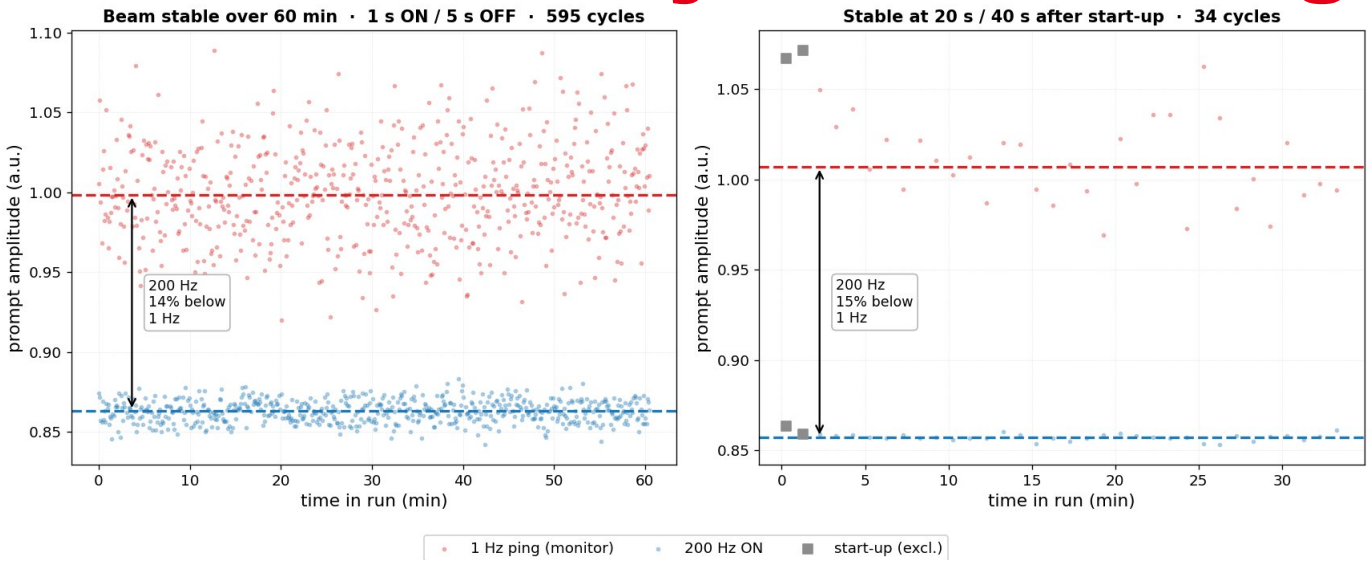

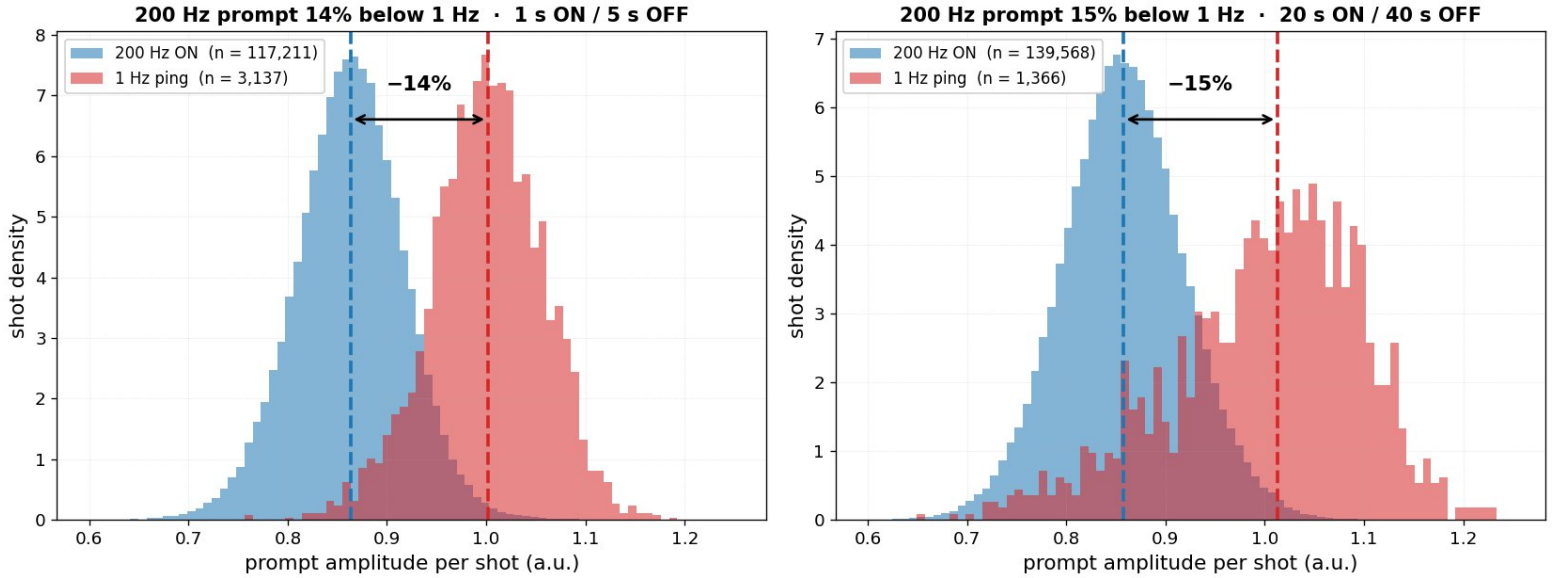

Observations

- **Linac stability:** Constant output during the beam-on phase, limited by Poisson statistics (**negligible drift** of +0.9% over 60 min).
- **Ping phase shift:** Prompt signal is **~15% higher during the 1 Hz** ping phase compared to the 200 Hz irradiation (identical behavior across all runs).

Forward model: full pulse history to expected counts

Full recorded pulse history → Bateman → one joint likelihood over selected bins

1. Full pulse history

(t_k, q_k) Measured pulse times and relative source weights for all pulses.

Not an averaged cycle.

2. Bateman response

$$H_{m,i} = \int_{\text{bin } m} \sum_{k\,:\,t_k < t} q_k \, e^{-\lambda_i (t - t_k)} \, dt$$

Each bin receives contributions from all previous pulses.

No reset at cycle boundaries.

3. Expected counts

$$\mu_{m,r} = b_r \, \text{live}_{m,r} + K_r \sum_i \alpha_i H_{m,i,r}$$

$$y_{m,r} \sim \text{NB}(\mu_{m,r}, \theta_{r,p})$$

All selected bins × all cycles × all runs.

Full-run forward model

(t_k, q_k) → $H_{m,i}$ → $\mu_{m,r}$

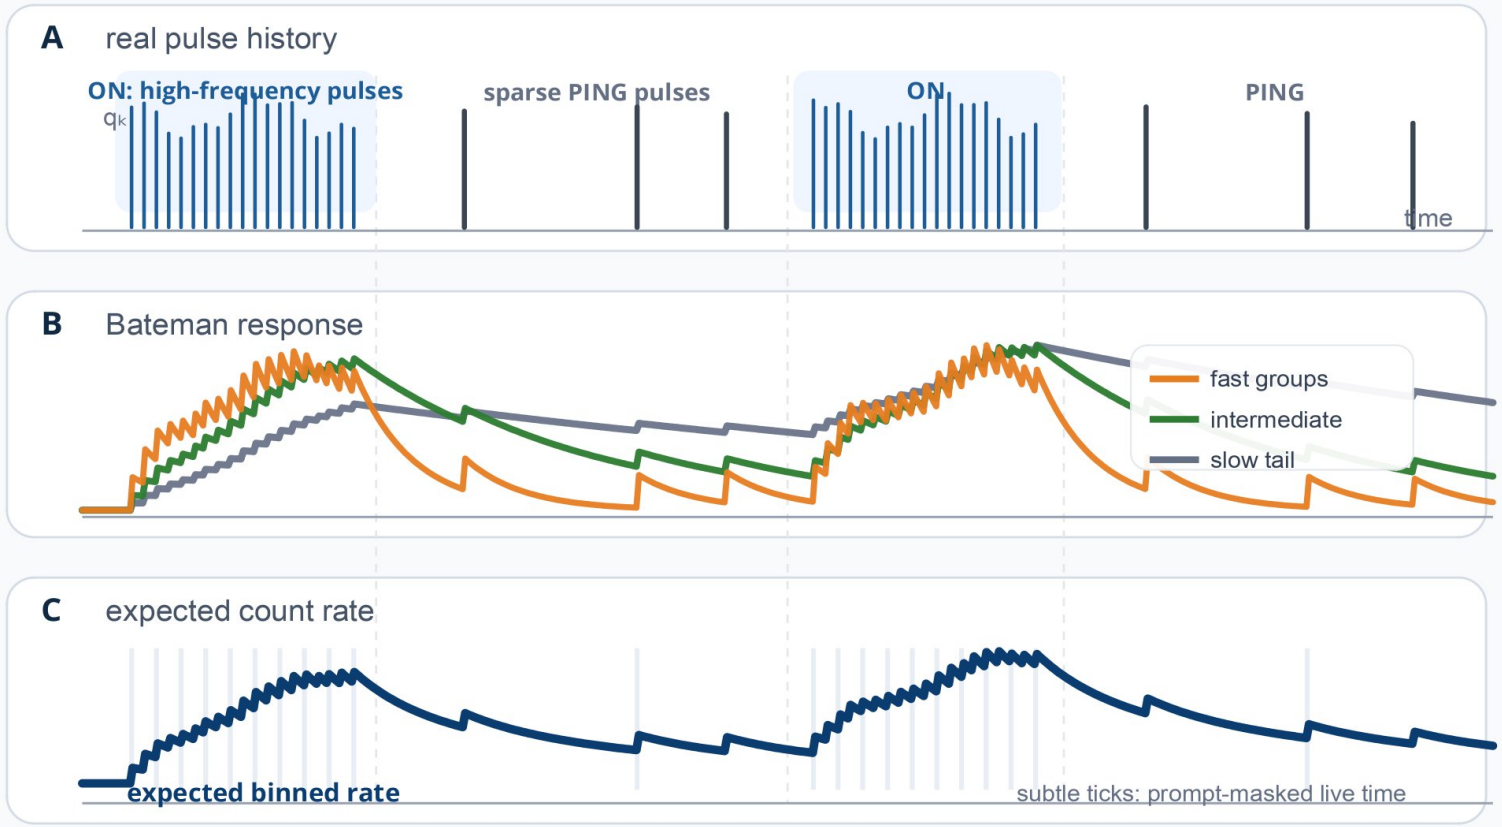

Bayesian inversion: MCMC samples parameter sets $\{\alpha_i, K_r, b_r, \theta\}$ → posterior intervals and correlations

Model reconstruction across time scales

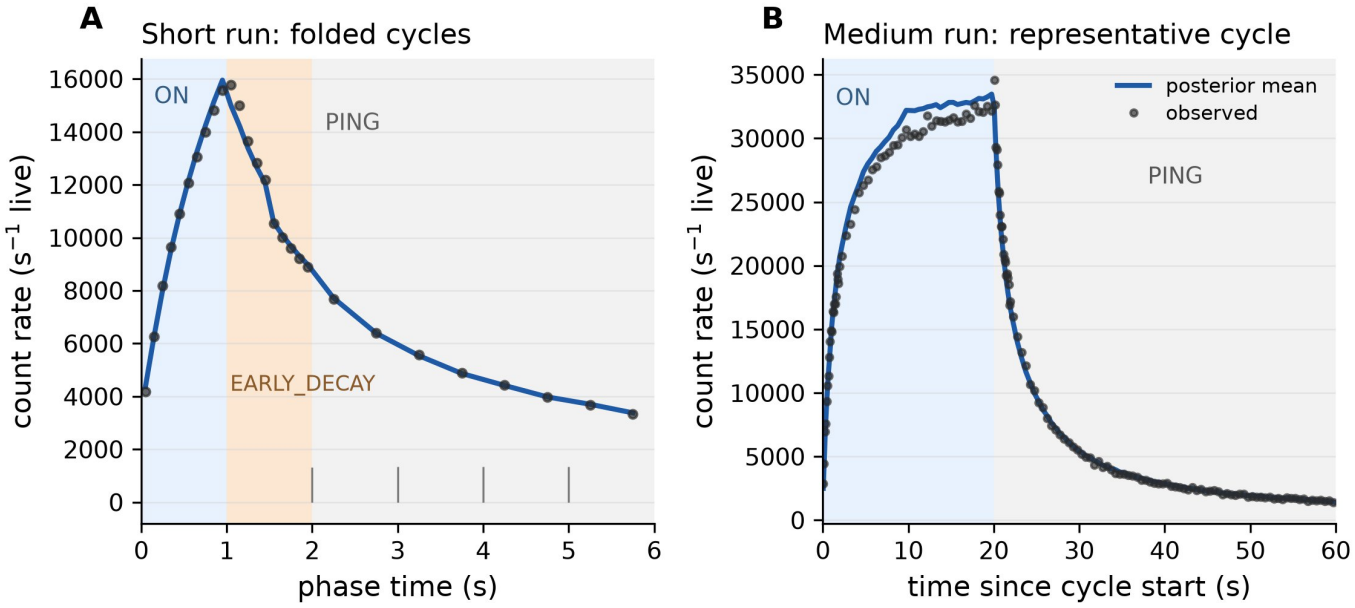

Group contributions and representative build-up / decay

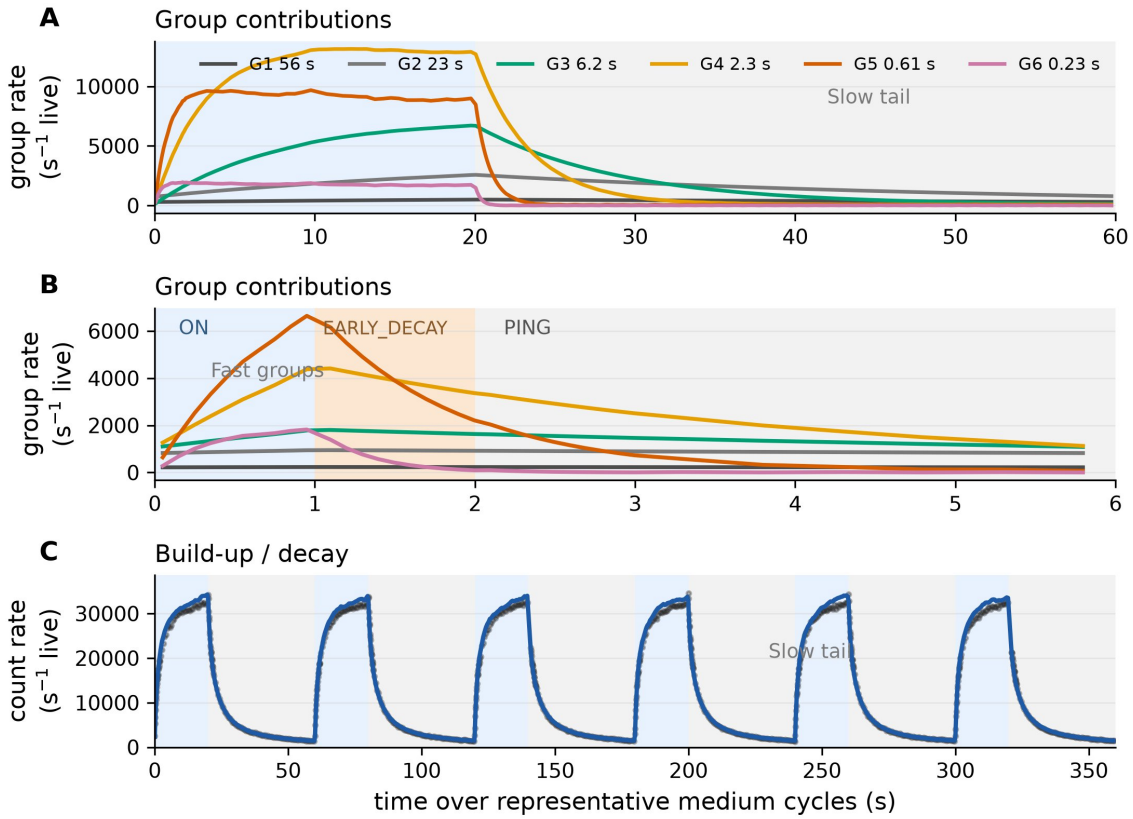

Short window — Early decay carries fast-group information; the G5/G6 split remains sensitive to early-time treatment.

Medium window — Build-up and decay tails constrain intermediate and build-up of slower components.

Results

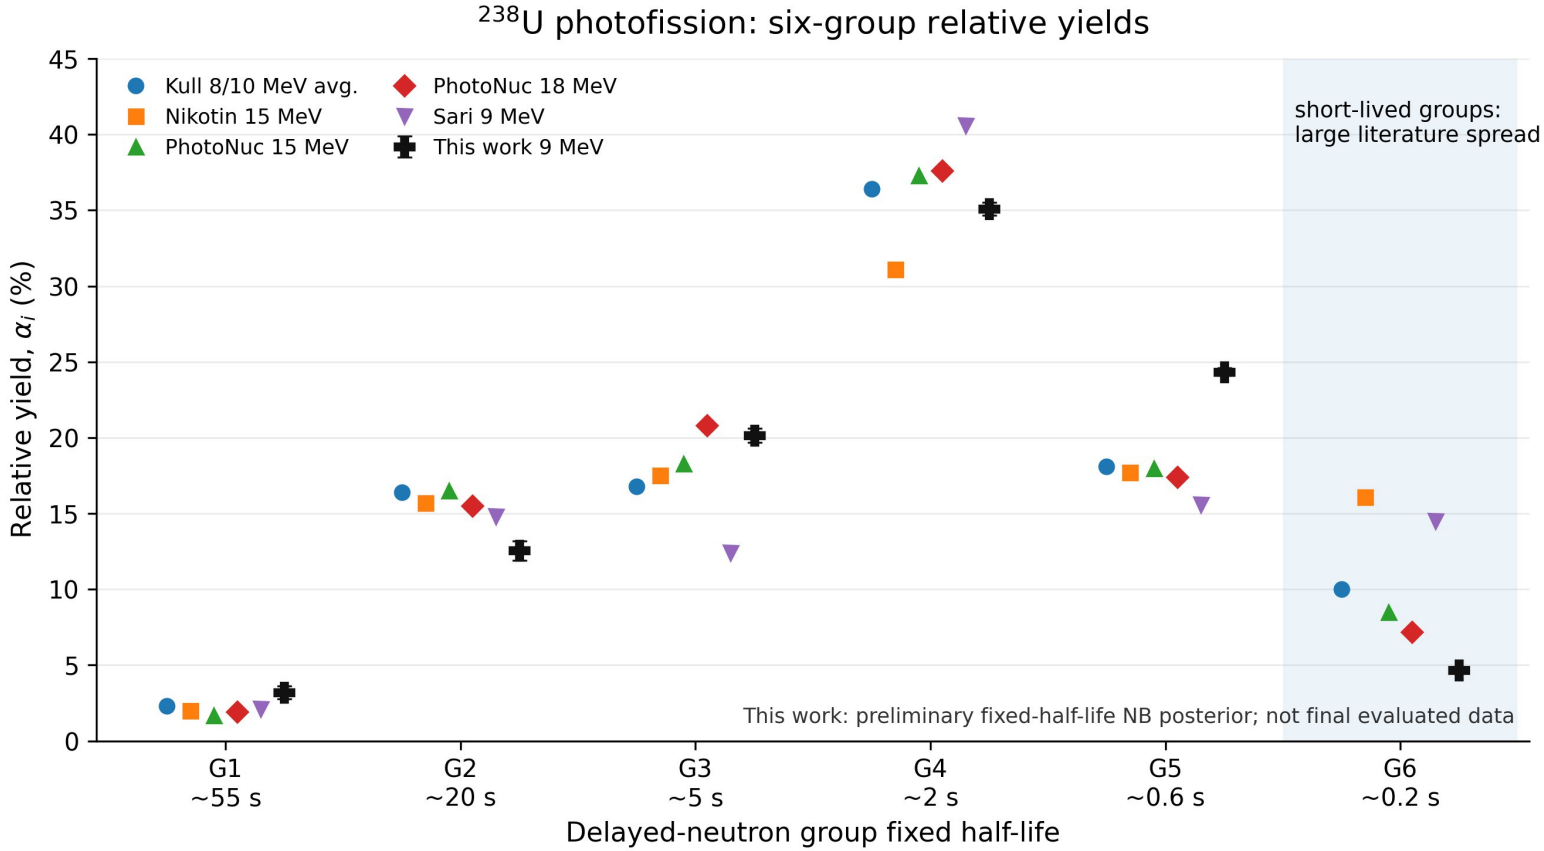

G1+G2≈15.8%
G3+G4≈55.2%
G5+G6≈29.0%

Observations

- Slow-, intermediate-, and fast-group sums are broadly consistent with 238U photofission literature.
- Differences are amplified when splitting neighbouring effective groups.
- Current data constrain the global delayed-neutron time structure.
- **Final group-by-group decorrelation requires additional irradiation windows.**

Discussion: posterior correlations and model checks

1. Parameter identifiability

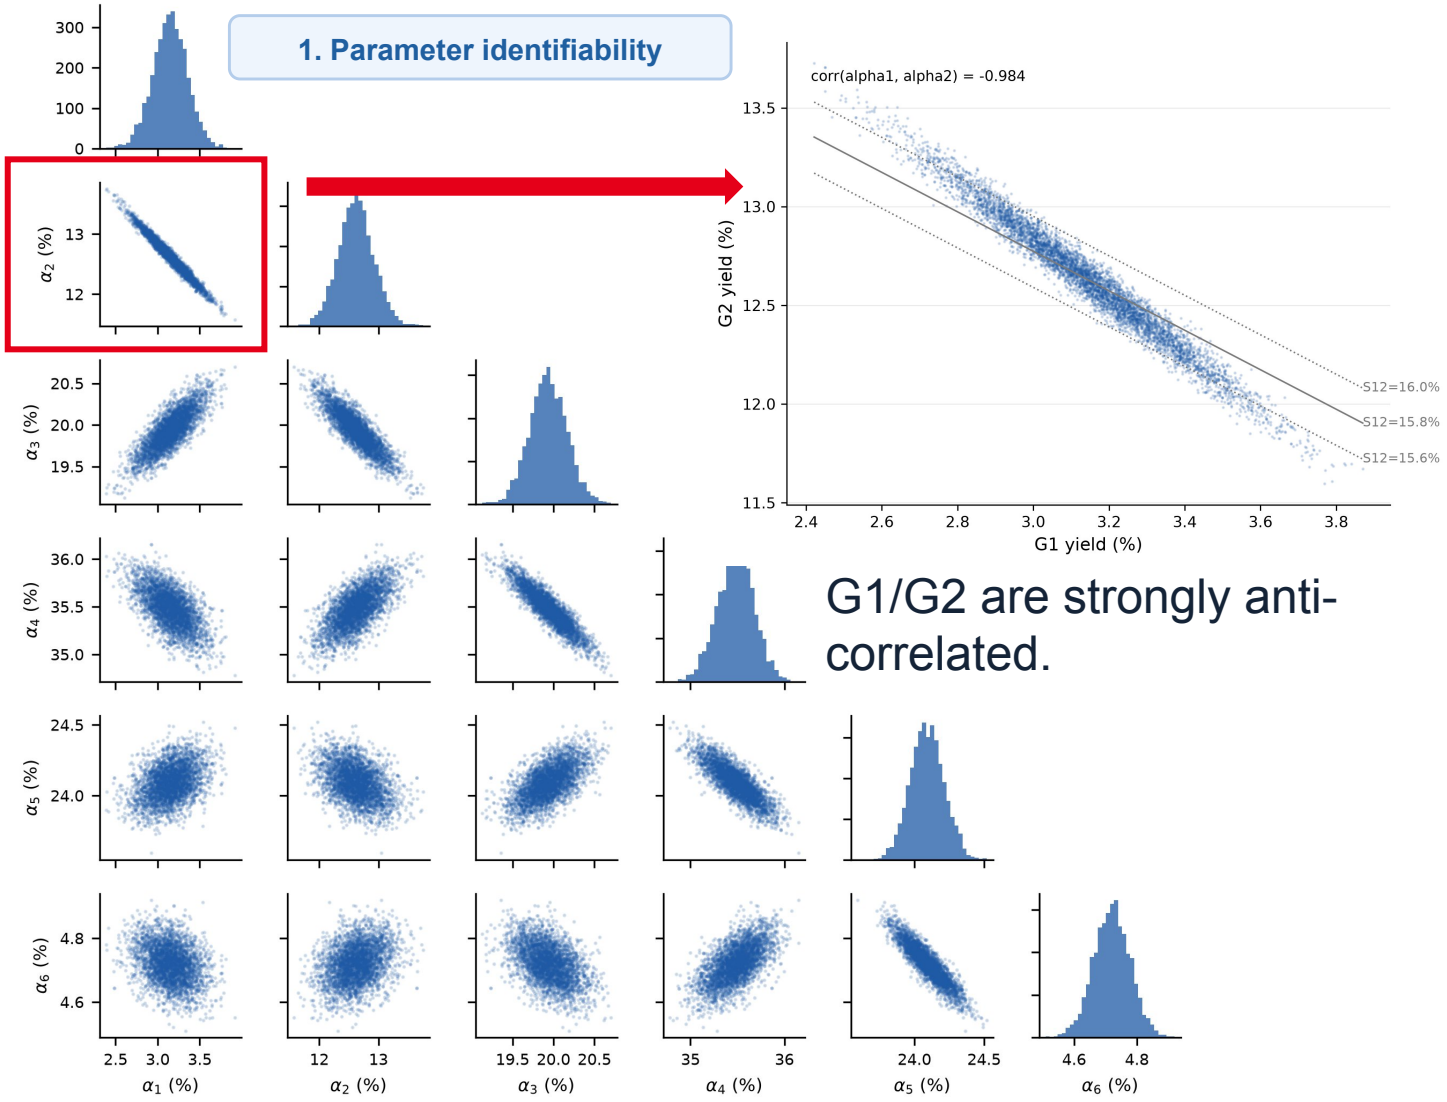

G1/G2 are strongly anti-correlated.

2. Posterior predictive check

Global checks are close to unity.
The early short-cycle window shows the largest dispersion.

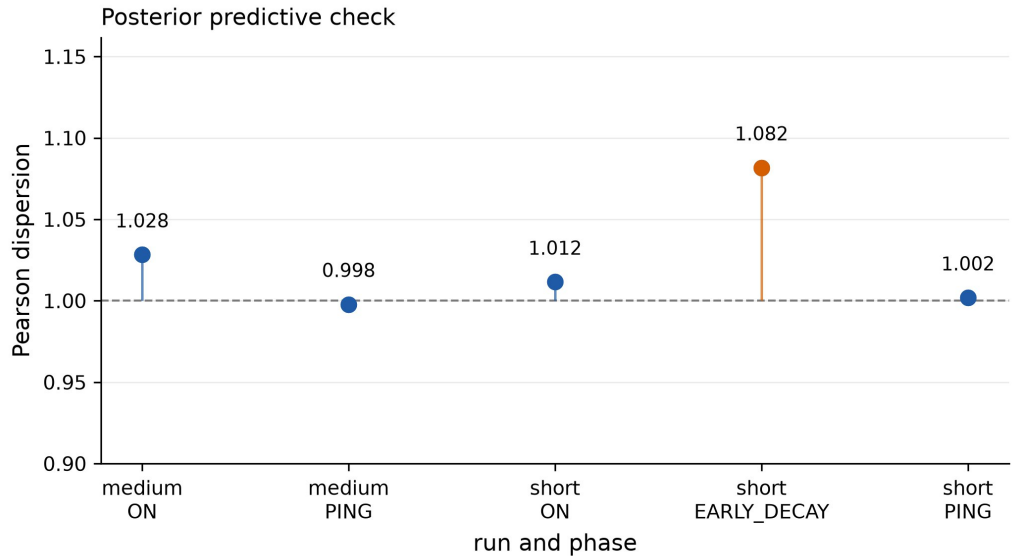

Limits

Validated pilot workflow — not yet a final evaluated nuclear-data measurement.

Inference / identifiability

Slow-group degeneracy

G1/G2 are strongly anti-correlated.
Current data constrain the slow component
$S_{12} = \alpha_1 + \alpha_2$, not a final G1/G2 split.

→ **long irradiation + clean beam-off decay**

Fast timing window

G5/G6 depend on the first seconds after ON.
The early post-ON phase was tested separately;
fast-group yields remain stable.

→ **cleaner short-cycle transition** (1s On/ 1 s Off, AFG internal clock)

Fixed group basis

Fixed 6-group half-lives in this pilot analysis.
Group yields may depend on the effective basis.

→ **compare 6-group vs 8-group inversions**

Experimental / physics systematics

Secondary fast-neutron fission

A few-percent (~3%) contribution may come from secondary fast neutrons in the massive sample.

→ **PHITS estimate + lower-mass references**

Residual ^{235}U contribution

Small isotopic contribution (<1%) from residual ^{235}U content may affect the delayed-neutron mix.

→ **include isotopic composition in forward model**

Detection efficiency model

Current analysis assumes one effective delayed-neutron efficiency for all groups.

→ **test group-dependent spectra / detector response**

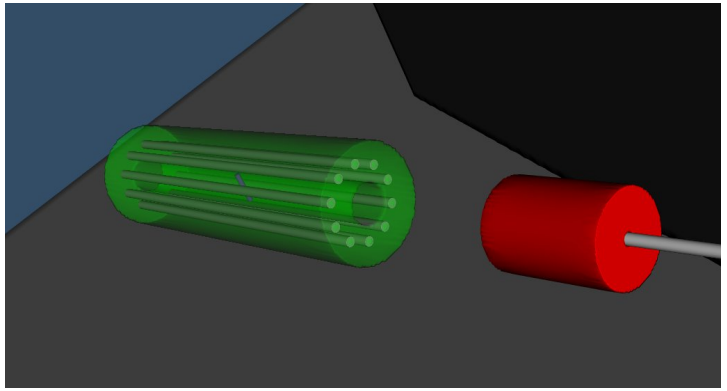

PHITS model used to estimate the fast-neutron contribution

Conclusion & Perspectives

What this pilot demonstrates

- **Traceable inversion workflow**
  for pulsed photofission delayed-neutron data
- **Full pulse-history forward model**
  + Bayesian NB inference
- **Preliminary group relative yields**
  with posterior intervals and correlations
- **Main time-scale components recovered**
  neighbouring-group splits remain preliminary

Next campaign

Long irradiation + clean beam-off decay
→ decorrelate G1/G2

Cleaner short-cycle timing / controlled ON→OFF transition
→ refine G5/G6 early-time constraints

Fixed 8-group comparison
→ assess model-basis effects

Lower-mass / reference samples
→ quantify secondary fast-neutron fission

Additional actinides
→ isotope-dependent delayed-neutron fingerprints

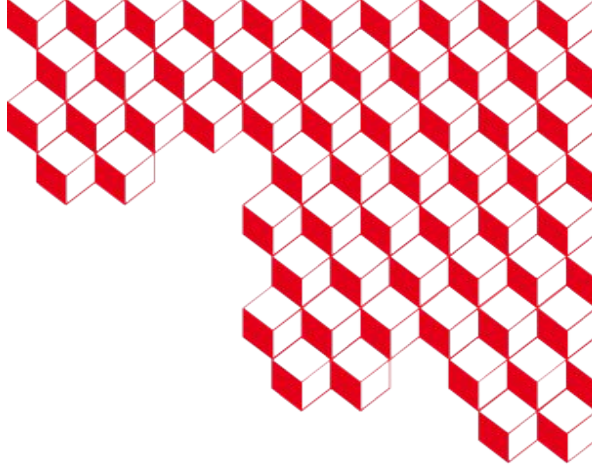

Thanks for your attention !

Acknowledgements: Yoann Moline, Louis Garnaud, Wilfrid Husson

CEA SACLAY
91191 Gif-sur-Yvette Cedex
France
johann.piekar@cea.fr
Standard. + 33 1 69 08 60 00

Récapitulatif

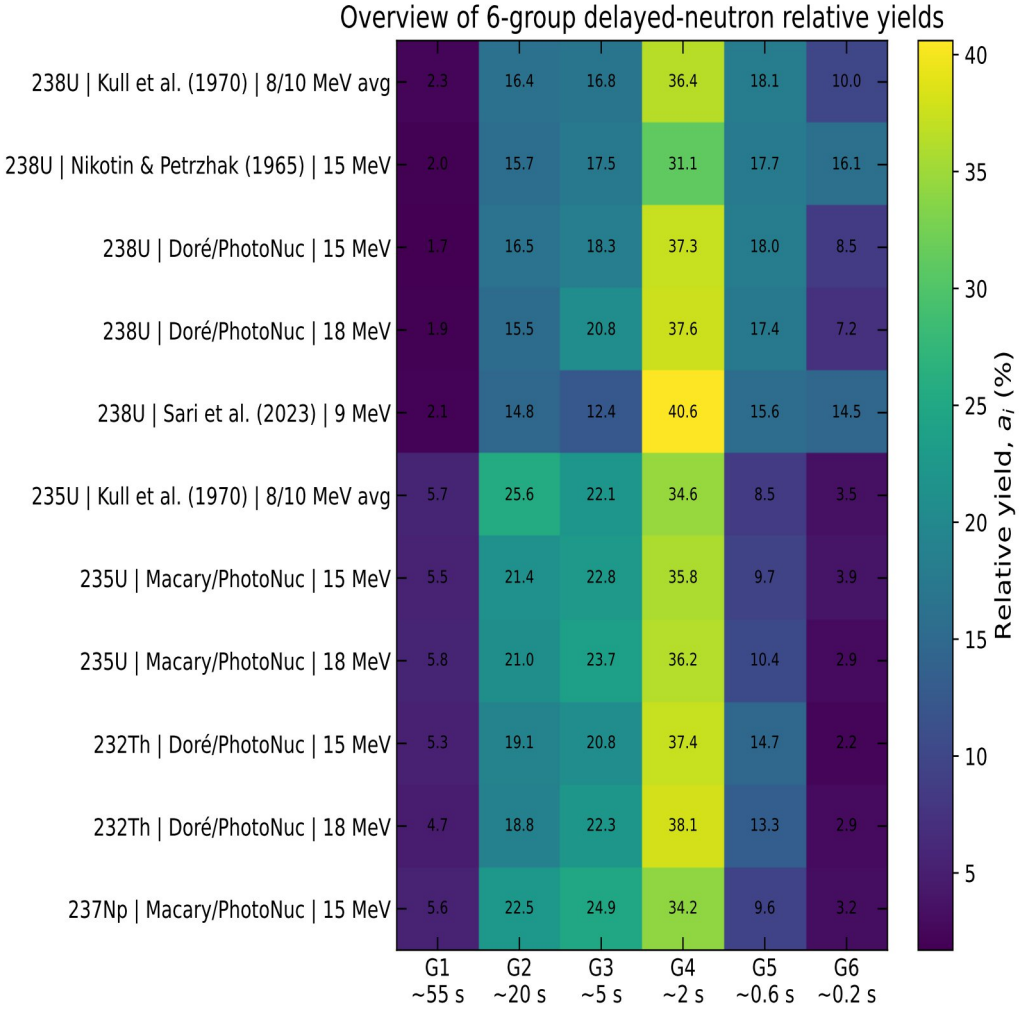

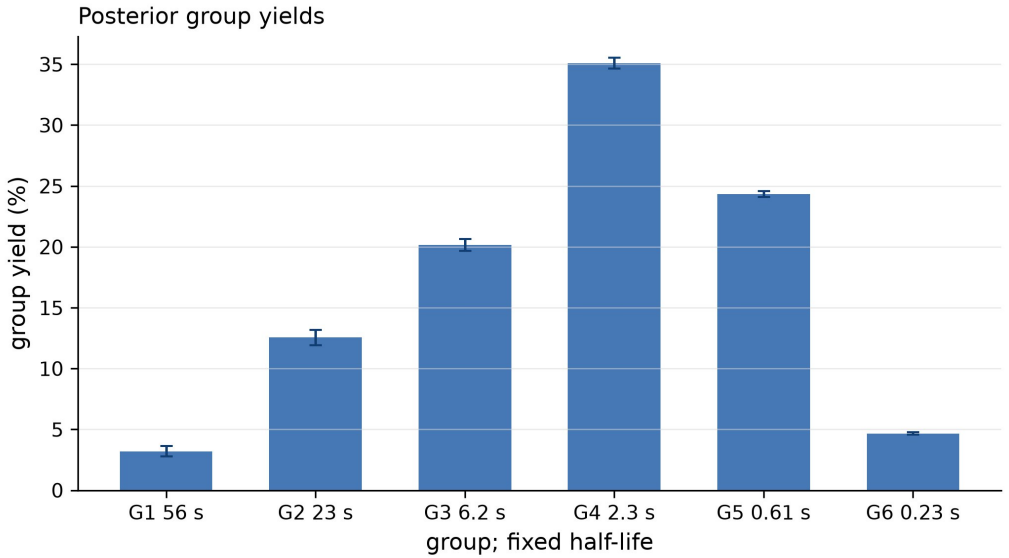

Phase ON ?

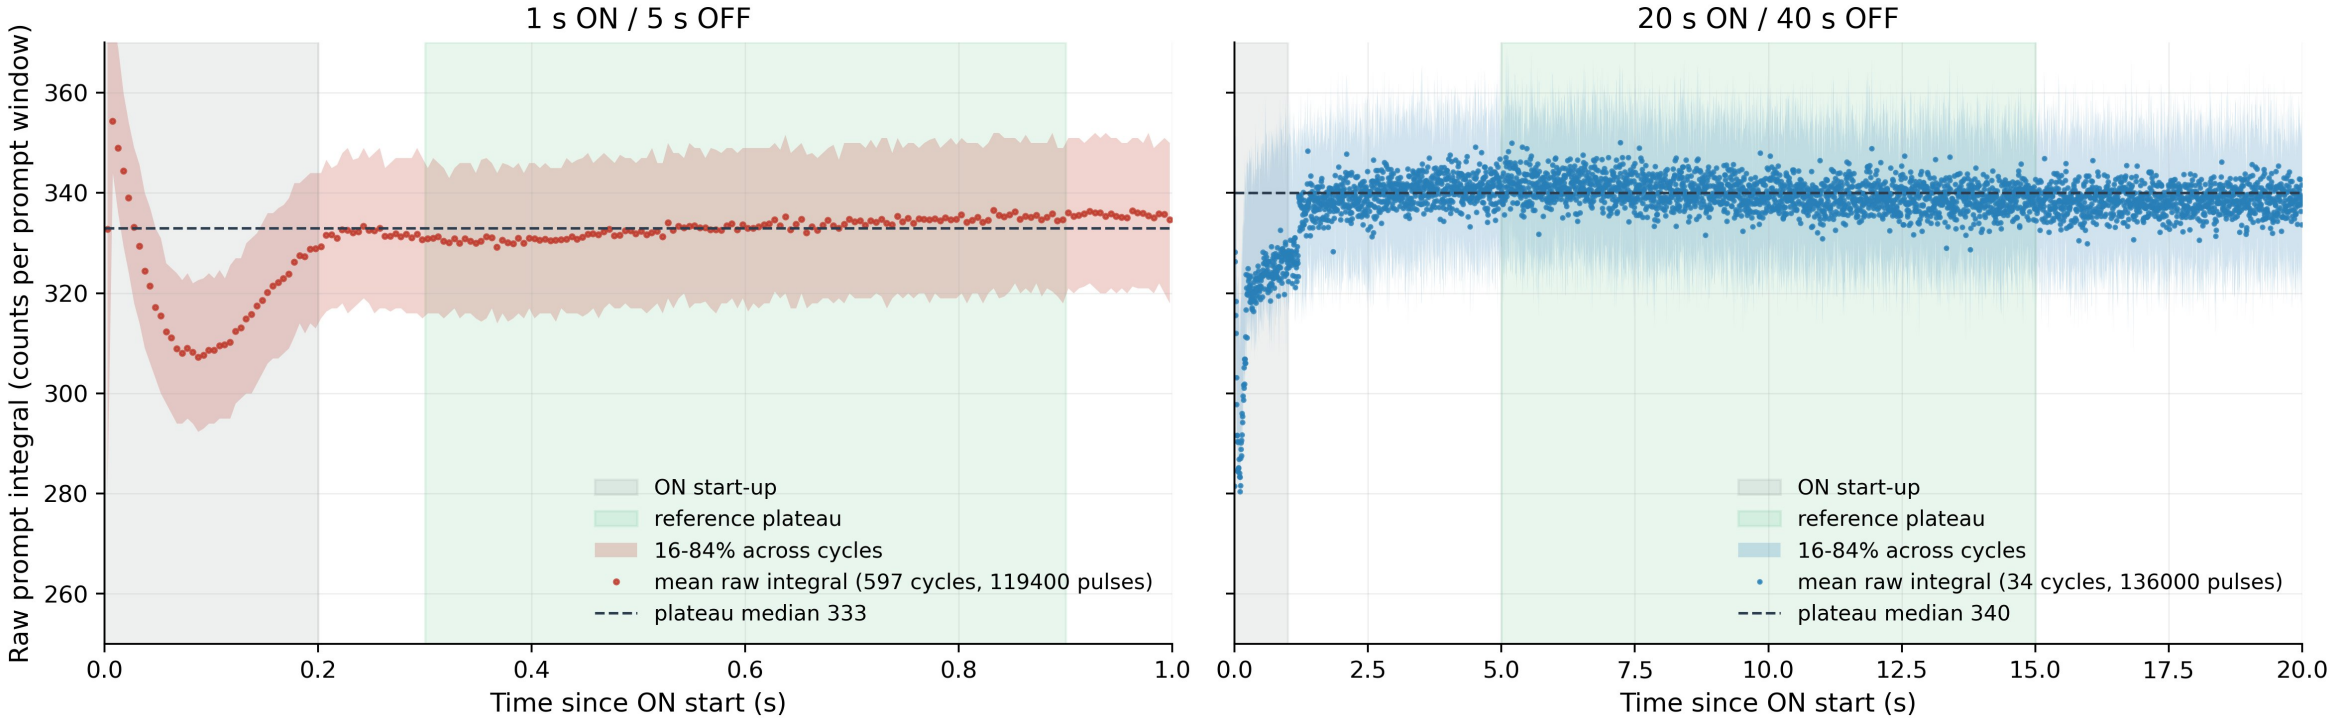
Raw prompt integral during ON: one marker is one 200 Hz pulse
1 s ON / 5 s OFF
20 s ON / 40 s OFF
Raw prompt integral (counts per prompt window)
360
340
320
300
280
260
0.0
0.2
0.4
0.6
0.8
1.0
2.5
5.0
7.5
10.0
12.5
15.0
17.5
20.0
Time since ON start (s)
ON start-up
reference plateau
16-84% across cycles
mean raw integral (597 cycles, 119400 pulses)
plateau median 333
mean raw integral (34 cycles, 136000 pulses)
plateau median 340

MCMC benchmark: uncertainties and covariances

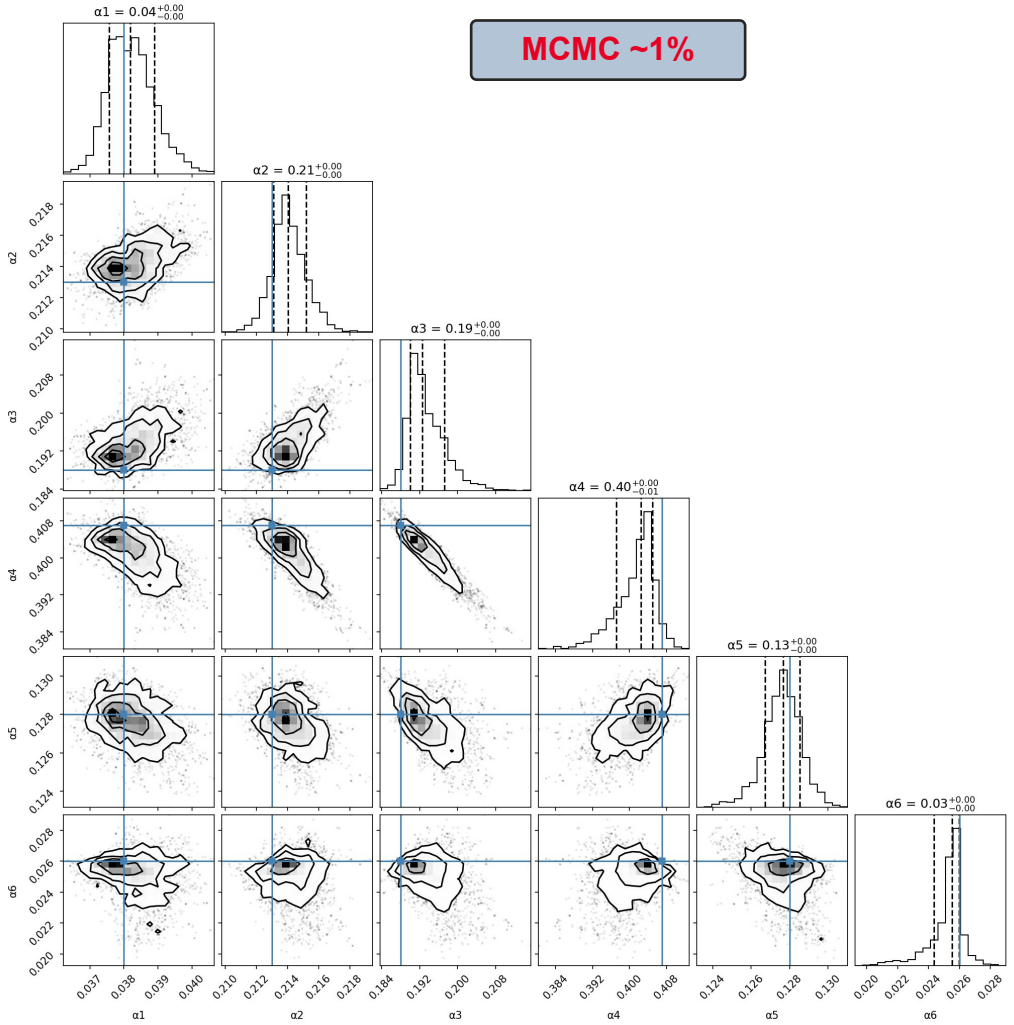

MCMC ~1%

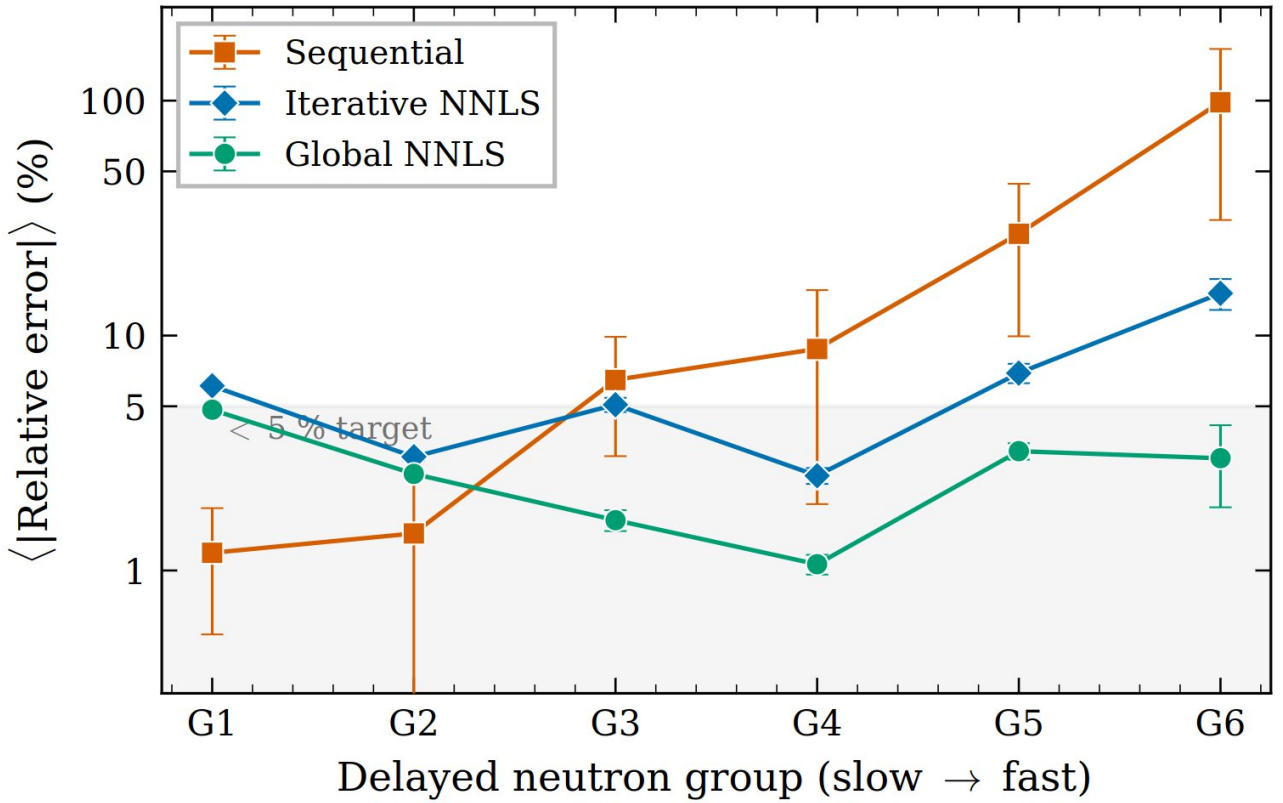

Why this matters for nuclear data

If several combinations of a_i and λ_i explain the data, this is not a failure of the fit: it is valuable information, captured by correlations and covariance matrices.

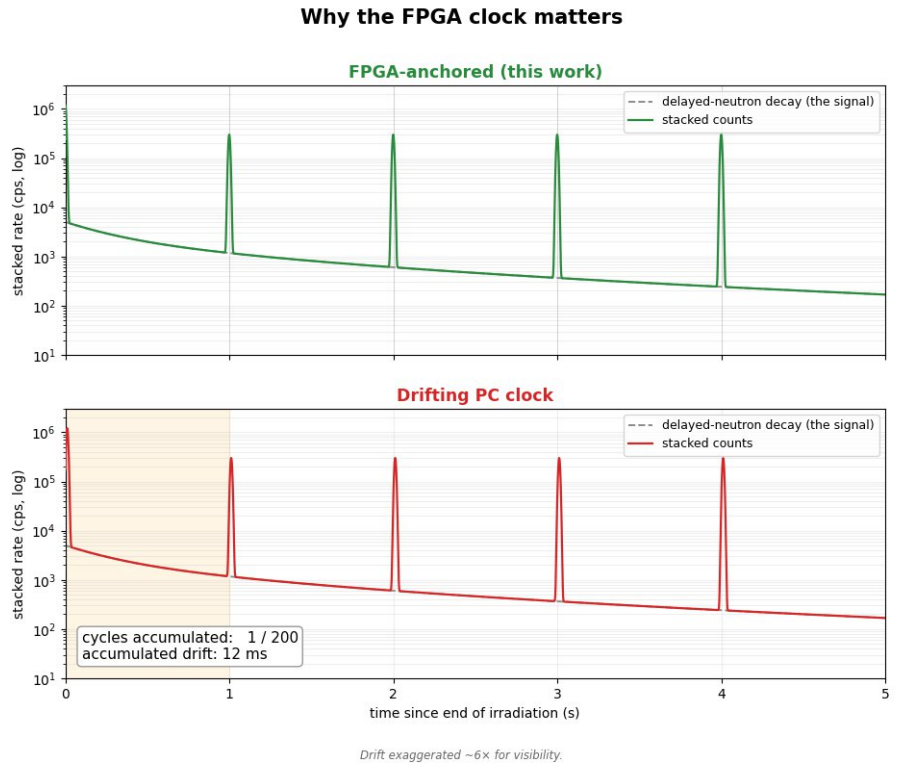
Why the FPGA clock matters
FPGA-anchored (this work)
delayed-neutron decay (the signal)
stacked counts
stacked rate (cps, log)
Drifting PC clock
cycles accumulated: 1 / 200
accumulated drift: 12 ms
time since end of irradiation (s)
Drift exaggerated ~6× for visibility.

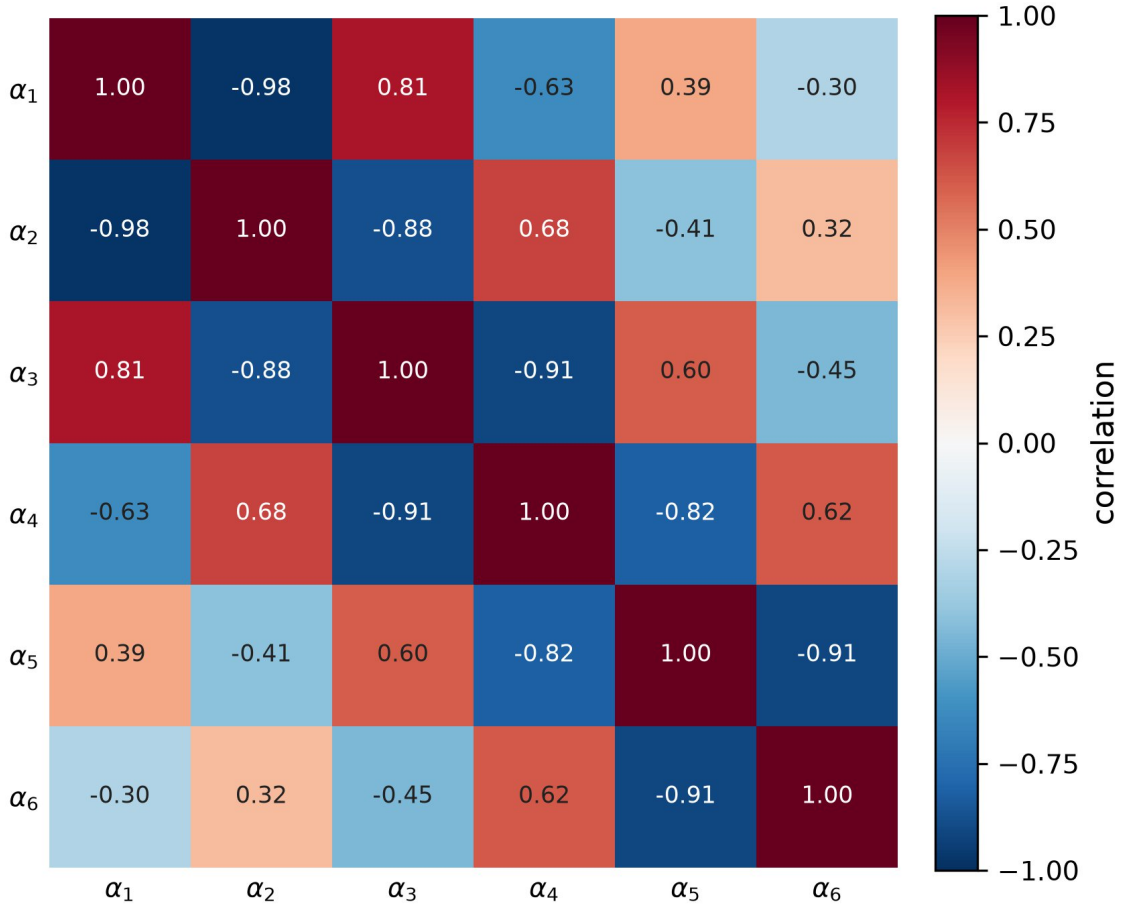
α1 1.00 -0.98 0.81 -0.63 0.39 -0.30
α2 -0.98 1.00 -0.88 0.68 -0.41 0.32
α3 0.81 -0.88 1.00 -0.91 0.60 -0.45
α4 -0.63 0.68 -0.91 1.00 -0.82 0.62
α5 0.39 -0.41 0.60 -0.82 1.00 -0.91
α6 -0.30 0.32 -0.45 0.62 -0.91 1.00
α1 α2 α3 α4 α5 α6
correlation
1.00
0.75
0.50
0.25
0.00
−0.25
−0.50
−0.75
−1.00

Forward model: from source history to expected counts

Signal model used by the forward simulation

$$\mu_k = b_k + \sum_i A_i \int_{\Delta t_k} \left[S(t) \otimes e^{-\lambda_i t} \otimes h_{\mathrm{th}}(t) \right] dt$$

where h_th(t) is the thermalization kernel.

$$n_k \sim \mathrm{Poisson}\left(\int_{\Delta t_k} \left[\epsilon \, (h_{\mathrm{th}} * \textstyle\sum_i a_i \lambda_i N_i(t)) + b \right] dt \right)$$

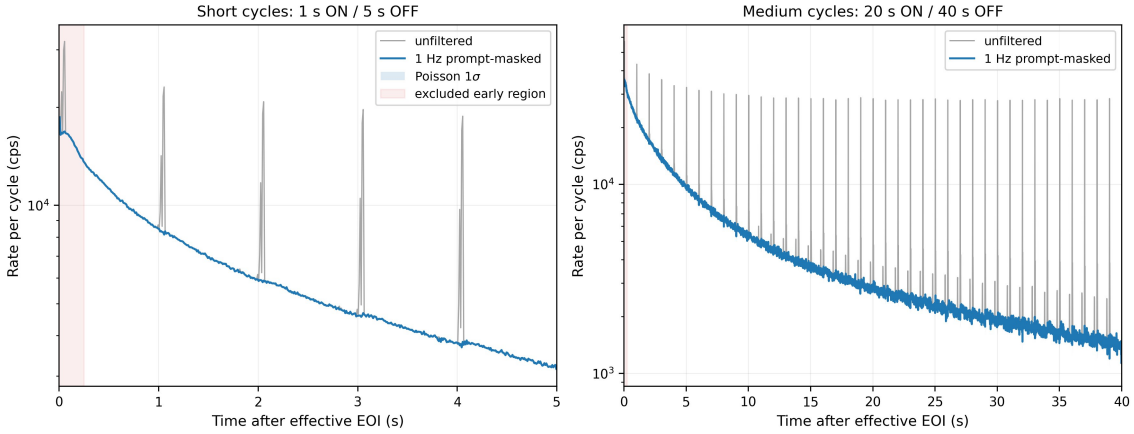

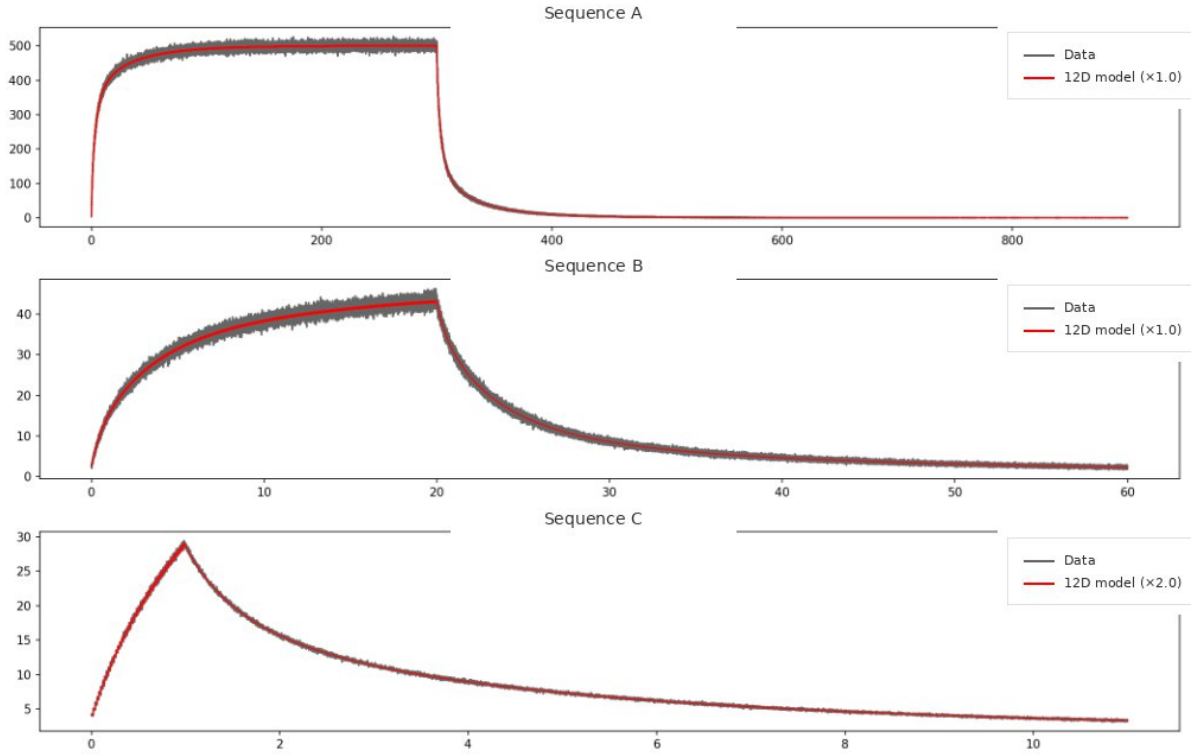

Results

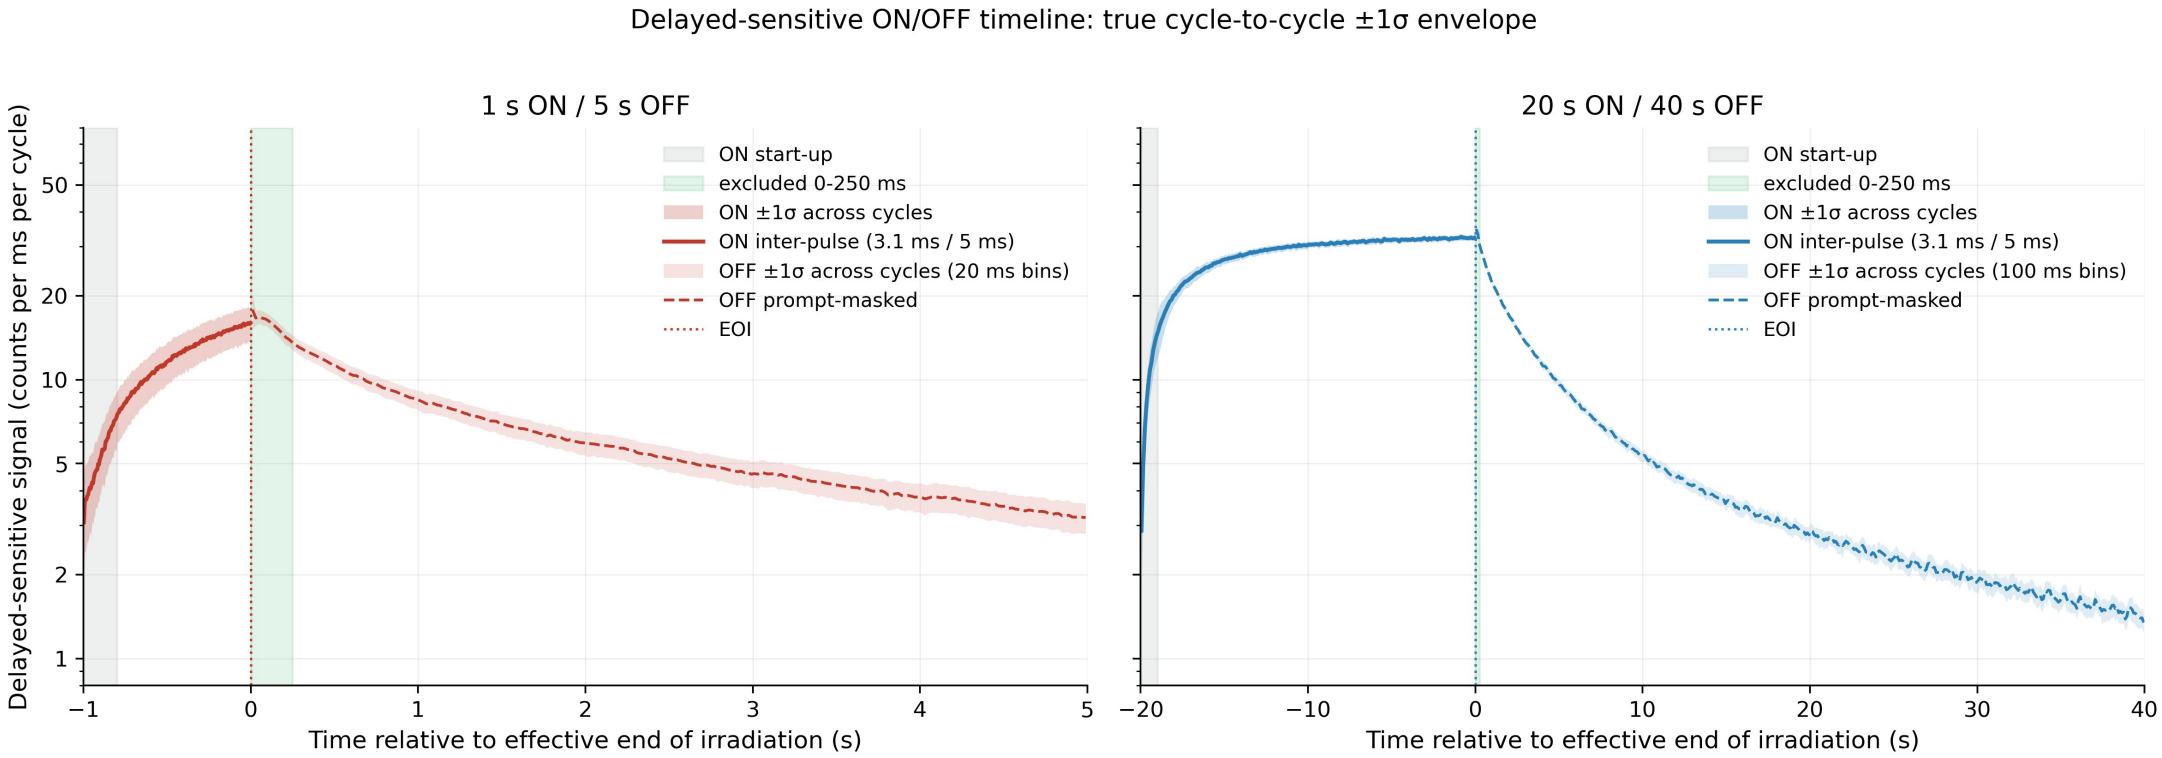

Fit / Inference:

- **ON → OFF transition confirms a viable delayed signal.**
- **Short cycles primarily constrain the fast components.**
- **Long irradiation is required to resolve slow groups and obtain interpretable MCMC covariances.**
- **Ping contribution subtracted via interpolation**

Suggested 15-minute storyline

A possible structure that keeps the audience in the funnel from application to method to experiment.

0–2 min	2–5 min	5–7 min	7–10 min	10–13 min	13–15 min
Context *active interrogation requires a model*	**State of the art** *photonuclear data are heterogeneous*	**Inverse problem** *fit delayed-neutron group constants*	**Experiment** *9 MeV ^{238}U photofission setup*	**Results** *posteriors, correlations, literature comparison*	**Limits** *biases, calibration, next steps*

Opening sentence

"Moving from a detected signal to a quantitative estimate requires an inverse model and reliable nuclear data."

State-of-art framing

Do not present it as bibliography.

Present it as the reason why a transparent fit is needed.

Final loop

Return to active interrogation.

Better data and models make novel sources quantitatively useful.

Cyclic summation

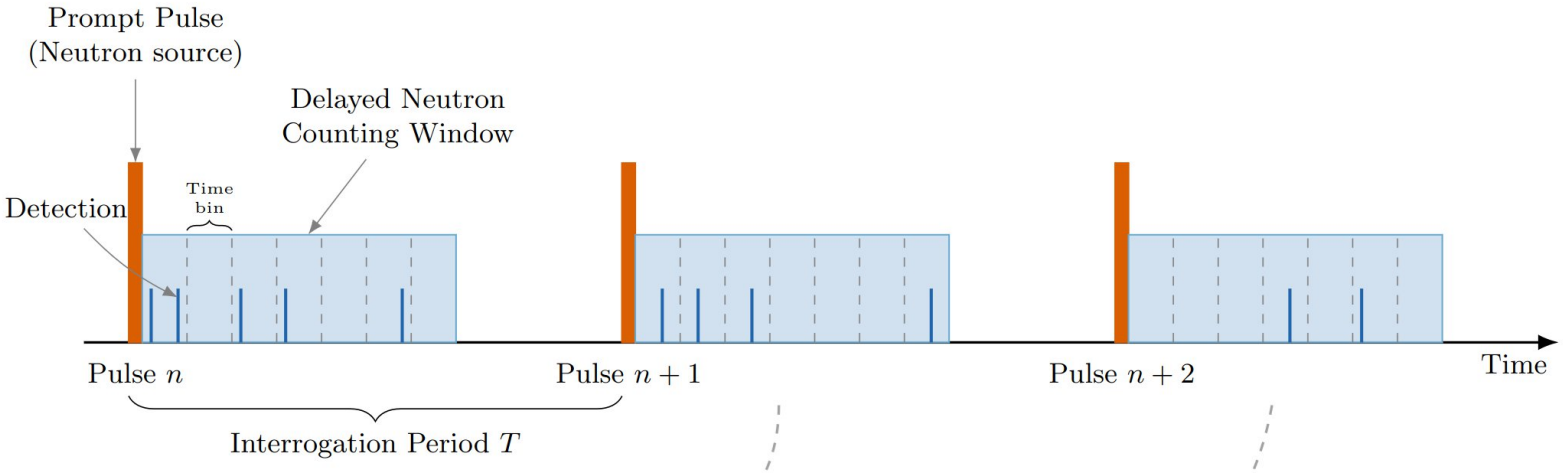

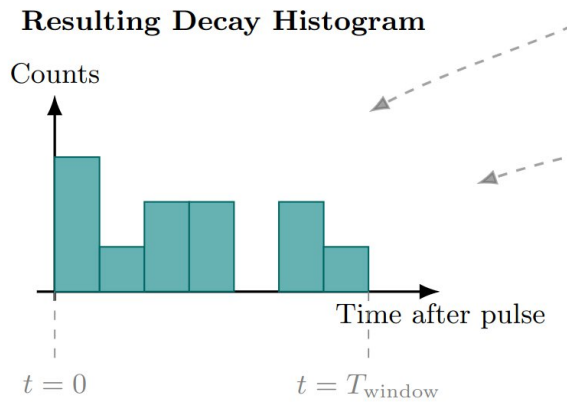

Cyclic Summation:
Events from *all* pulses are aggregated into a single decay curve.

$N_{\text{total}}(t) = \sum_{k=0}^{M} N_{\text{pulse } k}(t)$

Decay constants vary less than yields

TAKEAWAY

λ_i values are more compatible than a_i, but not perfectly invariant because the groups are effective aggregates.

Compatibility result

23 compatible

4 tensions

4 >3σ

- Fitted λ and a are correlated in six-group models
- The 8-group model fixes λ to the Spriggs/JEFF-type basis

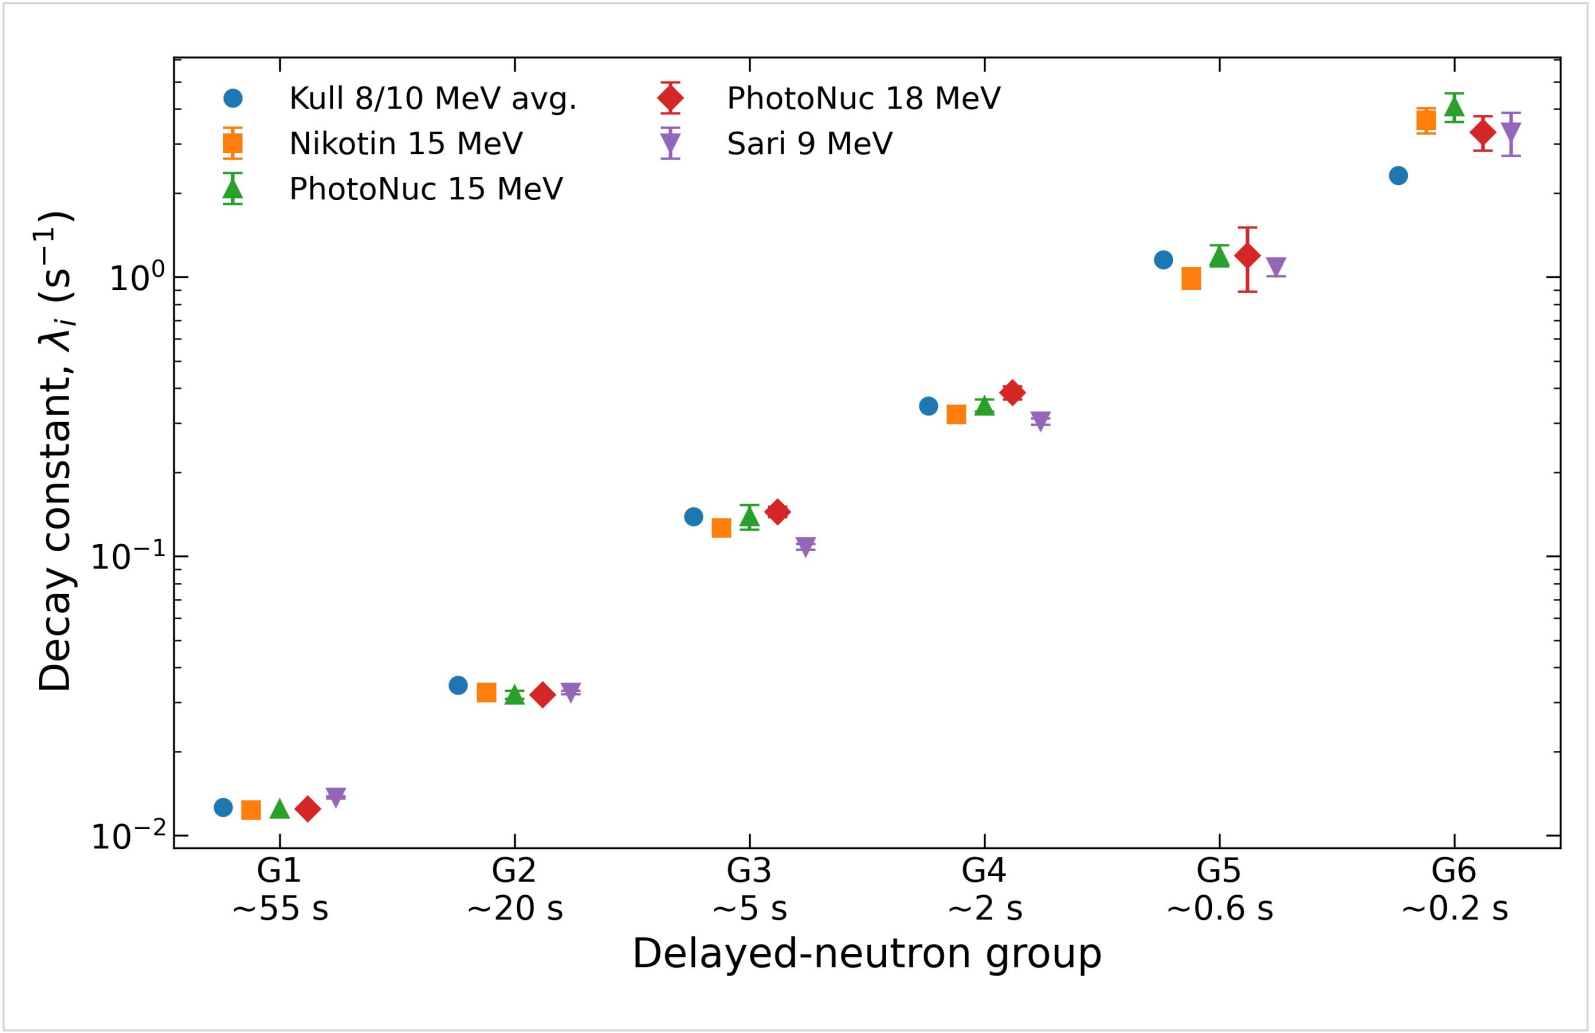

Forward model: full pulse history to expected counts

Real pulse history → Bateman response → joint binned likelihood

1. Full experimental pulse history

Real pulse history: (t_k, q_k) for all pulses

measured timestamps | **relative pulse weights**

Use the complete recorded sequence of pulses, not idealized 200 Hz / 1 Hz cycles.

Timing irregularities and source fluctuations are kept through (t_k, q_k).

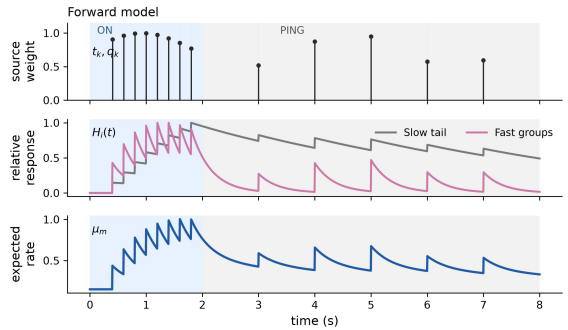

2. Bateman response matrix

$$H_{m,i} = \int_{\text{bin } m} S_i(t)\,dt$$

$$S_i(t) = \sum_{k:\, t_k < t} q_k\, e^{-\lambda_i (t - t_k)}$$

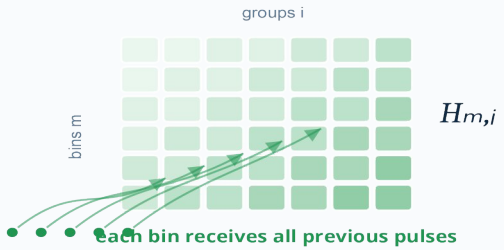

Fixed half-lives → group response $H_{m,i}$.

Cycle-to-cycle build-up is naturally included because the populations are propagated over the full run.

no reset at cycle boundaries

$$H_{m,i} = \int_{\text{bin } m} \sum_{k:\, t_k < t} q_k e^{-\lambda_i (t - t_k)}\, dt$$

3. Expected counts & likelihood

$$\mu_{m,r} = b_r \,\text{live}_{m,r} + K_r \sum_i \alpha_i H_{m,i,r}$$

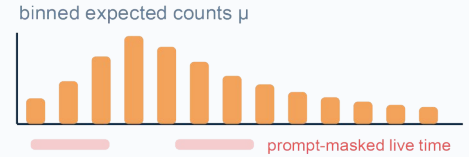

$$y_{m,r} \sim \text{NB}(\mu_{m,r}, \theta_r)$$

or $\theta_{r,\text{phase}}$

α_i: group relative yields K_r: run normalization
b_r: run background $\text{live}_{m,r}$: live time
$H_{m,i,r}$: response from real pulse history

all selected bins × all cycles × all runs

One joint likelihood over the full binned dataset.

Bayesian inversion: MCMC samples $\{\alpha_i, K_r, b_r, \theta\}$ → posterior intervals and correlations

Merci pour votre attention !