TALK II · COMPARATIVE SIMULATION RESULTS

nEXO background reduction for shallower depth

How far can we push the cosmogenic ^{137}Xe background down at Yemilab?

Lin Si

Stanford University

CUP, Daejeon · September 2026

CONTEXT

Where this talk sits

Three talks, one argument: nEXO can live at Yemilab if the cosmogenic ^{137}Xe is brought under control.

TALK 0 · GIORGIO GRATTA

nEXO

The experiment and the physics case. At SNOLAB, cosmogenics are a sub-dominant background.

TALK I · EUNJU JEON

Proposed mitigation

The menu of options: add a neutron poison to the xenon, or put boron around the vessel.

TALK II · LIN SI — THIS TALK

Simulation results

Which options actually work, by how much, and what they cost us elsewhere.

The question I am answering: Eunju and Yeongduk proposed several things we could do. I ran them. Here is what each one buys, which combination we prefer, and what we have not yet paid for.

THE PROBLEM

Why we have to do anything at all

Yemilab is 2700 m.w.e. Per muon the physics is the same as at SNOLAB — there are just far more muons.

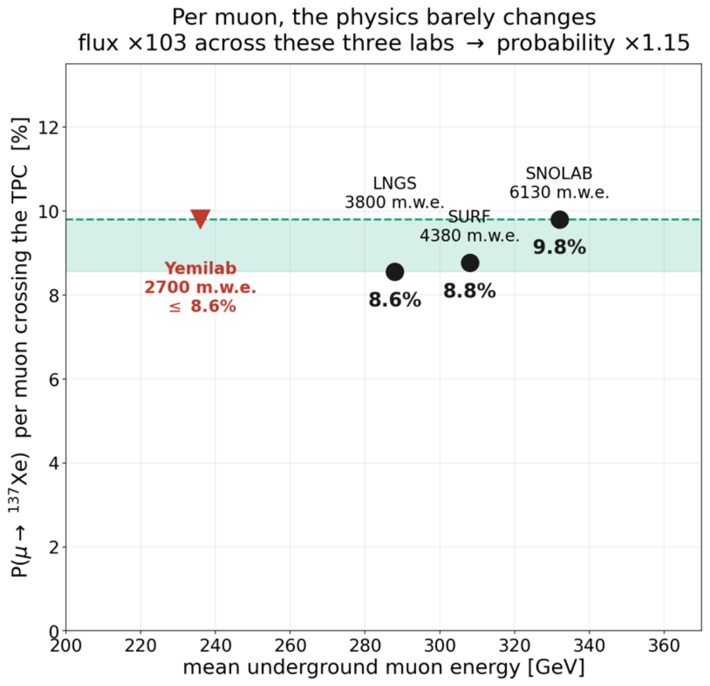

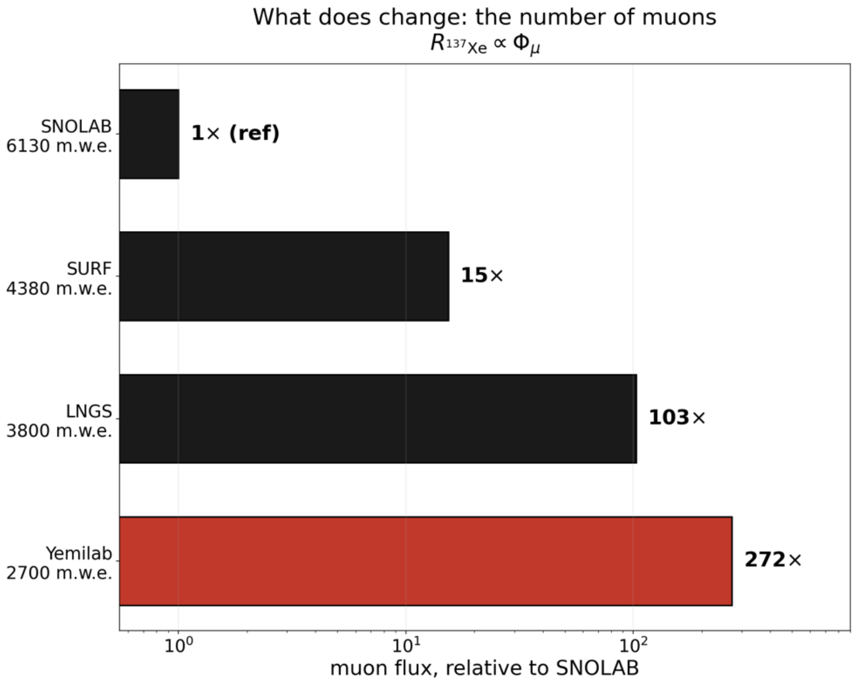

272×

the SNOLAB muon flux

- Per muon, the chance of making a ^{137}Xe is 8.6–9.8 % at SNOLAB, SURF and LNGS — the flux across those three labs changes by 103×, the probability by 1.1×.
- It drifts down with the muon energy: 332 GeV at SNOLAB, 288 at LNGS. Yemilab is 236 GeV, so it should sit at or below 8.6 %.
- So the ^{137}Xe rate is set by one thing: how many muons there are. Yemilab has 272× the SNOLAB flux.

What 272× costs in sensitivity

Half-life reach against S_X, the ^{137}Xe rate relative to the nEXO baseline. Curve reconstructed from Soud Fig. 4.14 — solid where the background paper covers it, dashed beyond.

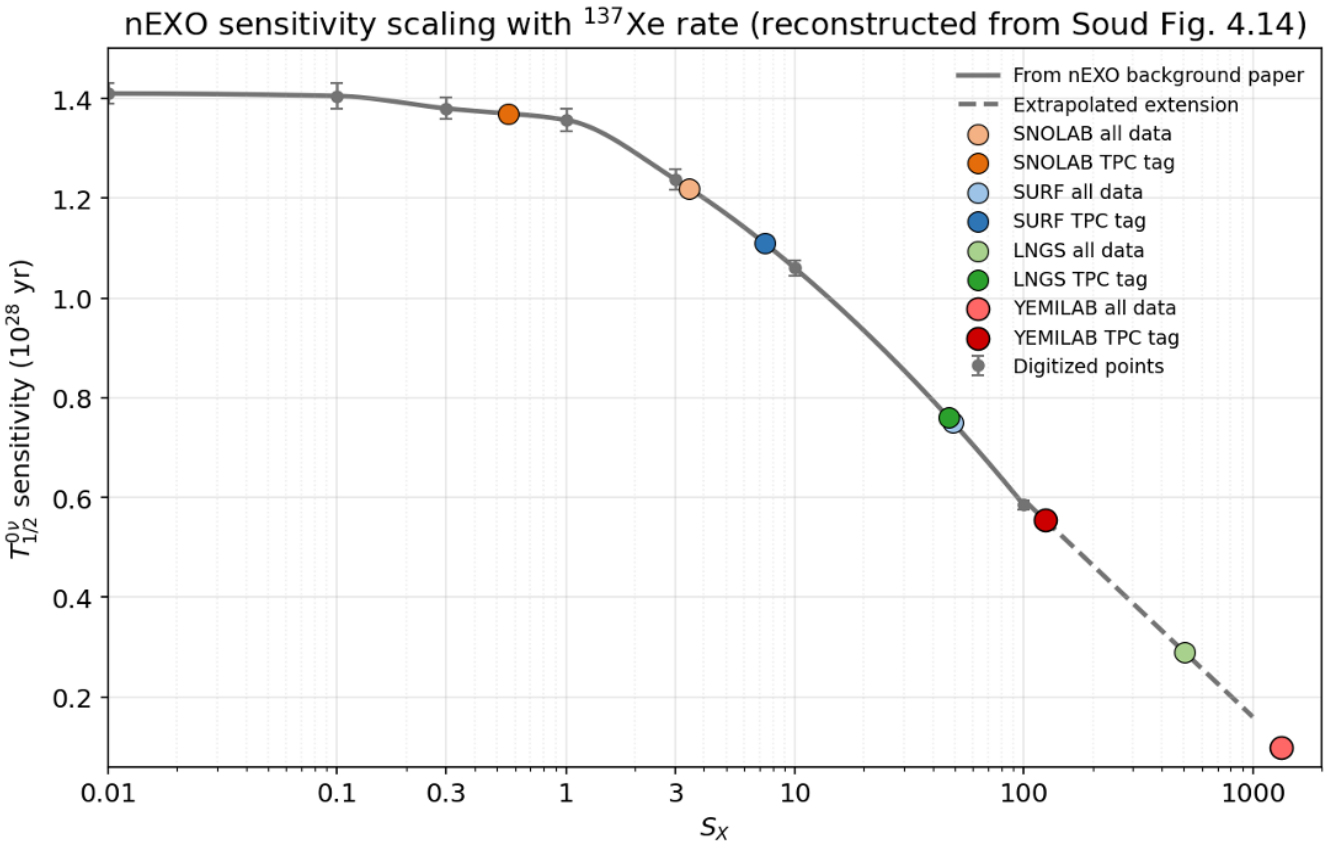

- **Four sites, one curve.** Light markers are all data; dark markers are after the TPC tag, the veto on muon-induced neutron captures (backup slide).
- **SNOLAB, tagged, is the baseline.** T½ ≈ 1.37 $\times 10^{28}$ yr — the number nEXO quotes.
- **Yemilab falls a long way.** Untagged it reaches ≈0.09 $\times 10^{28}$ yr. Even with the tag it only recovers to ≈0.55.
- **So a factor ≈2.5 is still missing.** That is the gap the rest of this talk works on — with the xenon and with the shield.

Note: none of the mitigation in this talk is in this plot. These are the untreated rates, with and without the tag.

Simulation setup, and the check that it is right

Mei–Hime underground spectrum + Gaisser surface flux → nexo-offline Geant4 → count ^{137}Xe atoms.

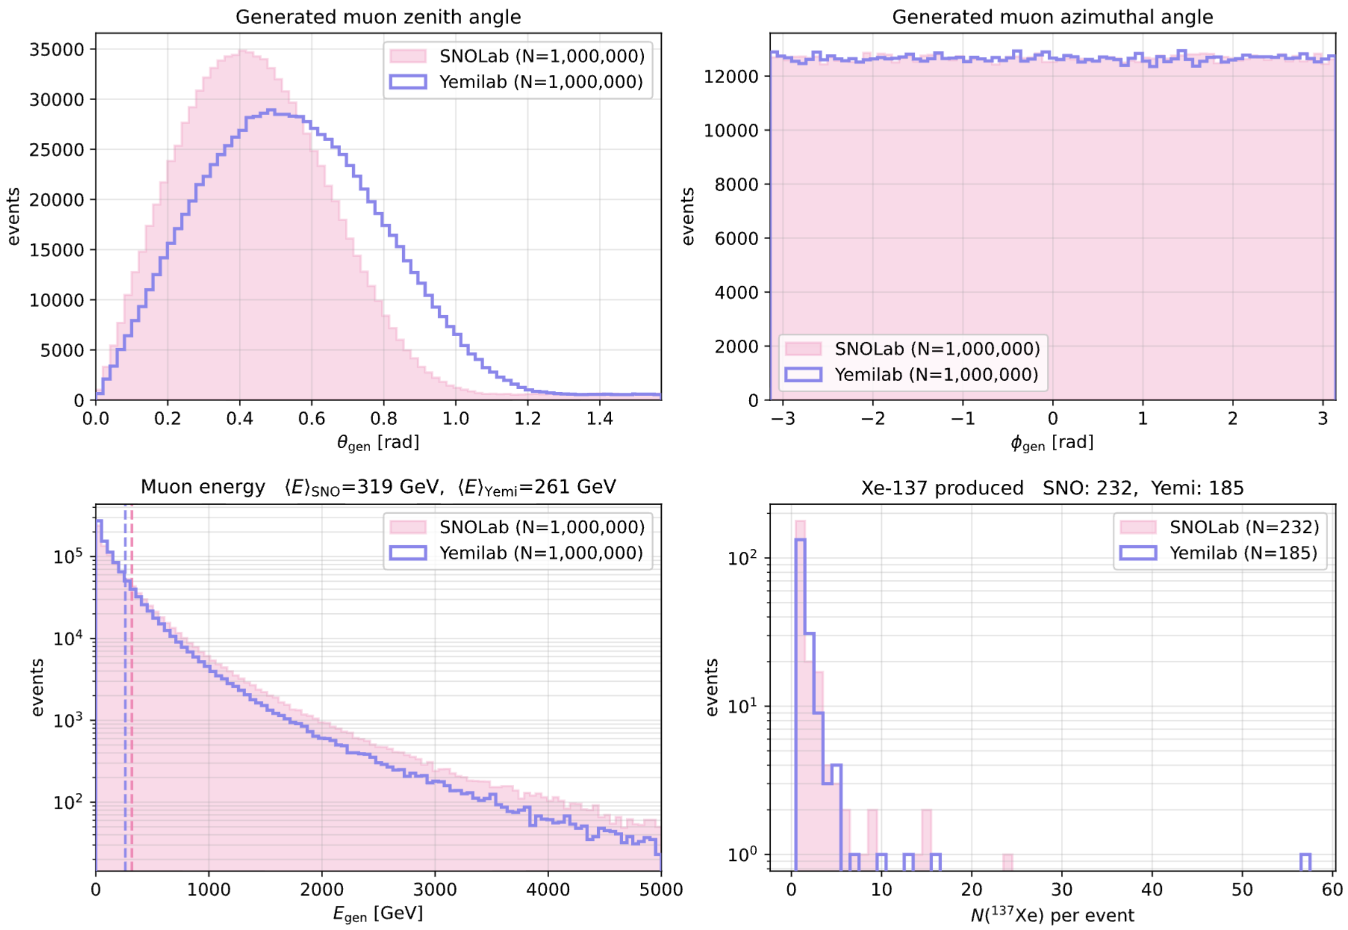

- **Same generator, two depths.** SNOLAB and Yemilab muon samples built with identical machinery, only the overburden changes.
- **Per-muon yield agrees.** 413 ± 47 at SNOLAB vs 357 ± 68 at Yemilab, both from 10^6 muons — the difference is flux, not production.
- **Statistics have grown since then.** The baseline is now 8×10^6 muons, 406 ± 25 per 10^6 μ; the shield and combined runs are 4×10^6 each.
- **Errors are file-to-file SEMs.** ^{137}Xe comes in showers, so √N would be optimistic by 1.2–3.5×. Metric throughout: atoms per 10^6 muons.

406 ± 25

baseline Yemilab yield [/10^6 μ]

Half the parent neutrons are born inside the vessel

This is what fixes where the shield goes, and why the shield on its own cannot be enough.

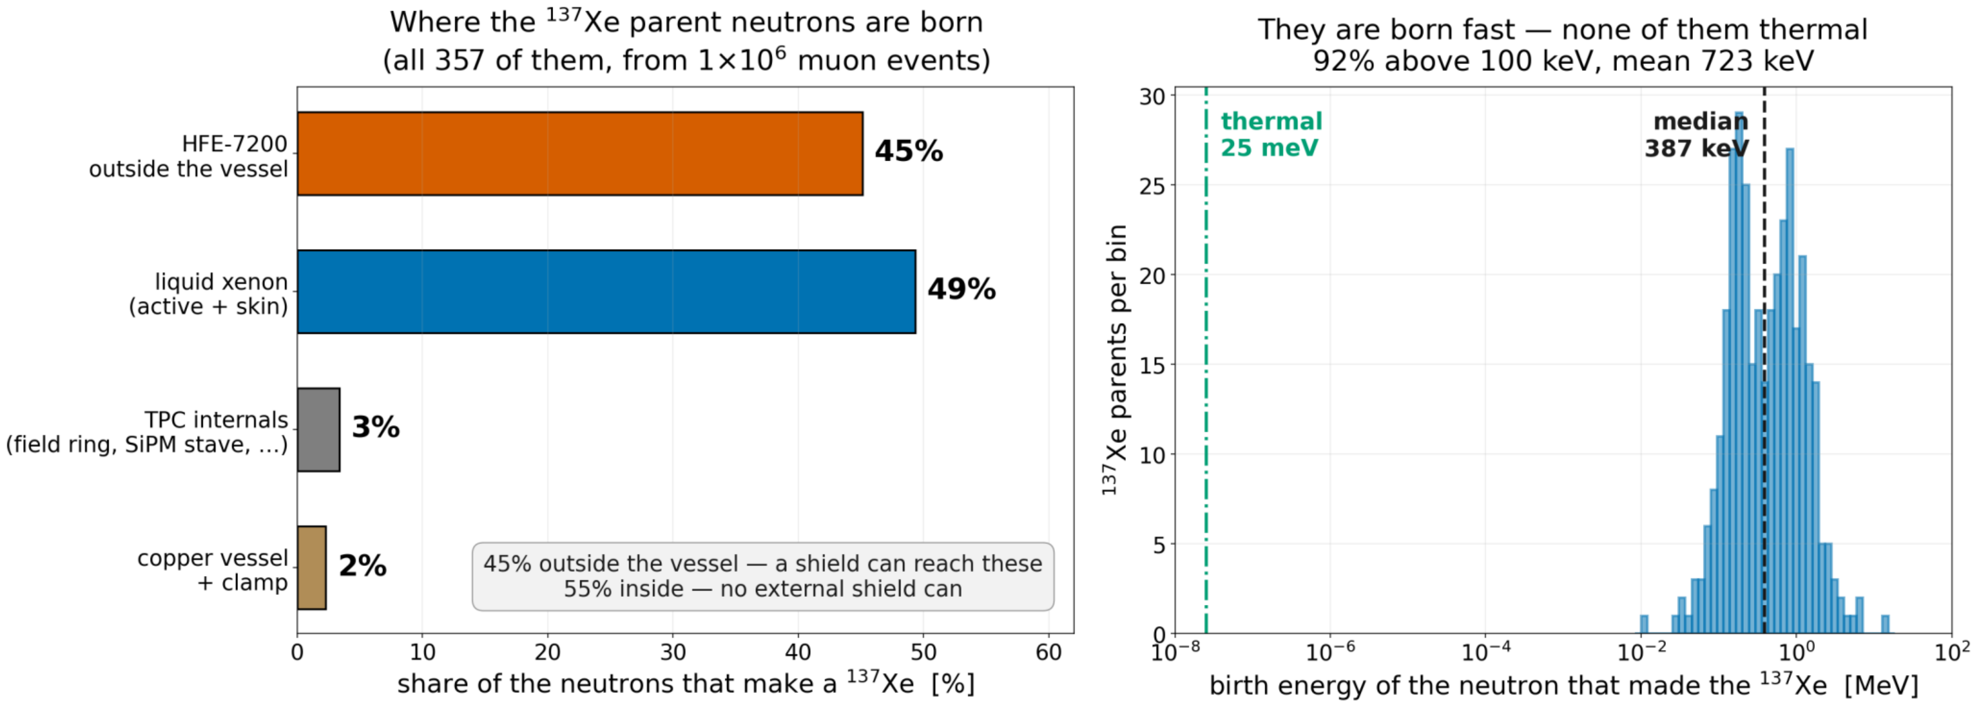

Two consequences. 45 % of the parent neutrons are born in the HFE just outside the copper, and an absorber on the vessel sees all of those. The other 55 % are born inside the vessel, nearly all in the xenon itself, and no external shield can reach them — that is why a second lever is needed. They are born fast and only slow down in the HFE, so the absorber has to catch thermal neutrons, not moderate them.

The result, up front

Two independent levers. One is worth 3.6×, the other 4.0×; together, 10.4×.

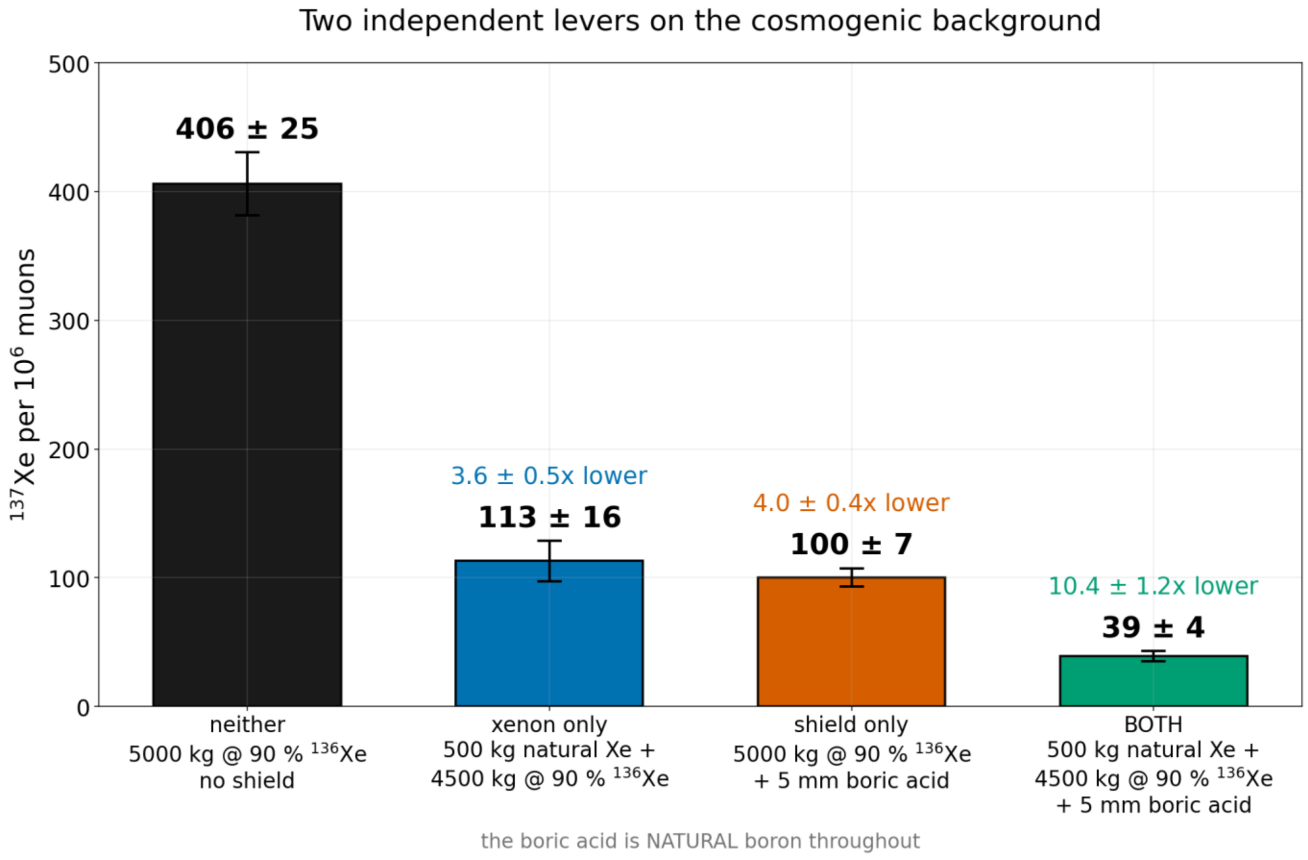

Xenon side — 500 kg of natural xenon inside the 5 t.

406 → 113 ± 16, a factor 3.6 ± 0.5.

Shield side — 5 mm of natural boric acid on the vessel.

406 → 100 ± 7, a factor 4.0 ± 0.4.

Combined — both levers at once.

406 → 39 ± 4 per 10^6 μ, a factor 10.4 ± 1.2.

LEVER 1

The xenon itself

Dope the target with a neutron poison so that a thermalised neutron is captured by something other than ^{136}Xe.

Why ^{131}Xe — and what thermal arithmetic predicts

Proposed by Eunju Jeon and Yeongduk Kim, following Pocar's internal-shielding note (2021).

Isotope	Thermal capture σ [b]	Role in the mixture
^{136}Xe	0.26	the 0νββ target — and the thing making ^{137}Xe
^{134}Xe	0.26	residual from the centrifuge, inert
^{131}Xe	85	the poison: ≈330× the ^{136}Xe cross-section
^{124}Xe	165	the original idea (Lenardo) — but far more expensive

Thermal capture budget. In the baseline mixture 90 % of thermal captures land on ^{136}Xe. Add a few mass-% of ^{131}Xe and that drops to 4–8 % — pure thermal arithmetic then predicts a factor 12–22 less ^{137}Xe.

BUT — the caveat that survives

The parent neutrons are born at a median 390 keV, and half of them are born inside the xenon itself. Those are captured before they ever fully thermalise, so a thermal branching ratio does not describe them.

Only a full transport simulation can say what the poison is really worth.

The four recipes I simulated

All are 5 t total. Every kg of poison is a kg of ^{136}Xe you do not have.

	Baseline	C	E	F	G
Source blend	5000 kg @ 90 % ^{136}Xe	realistic enriched spec, residual ^{131}Xe from the centrifuge	300 kg @ 70 % ^{131}Xe + 4700 kg @ 90 % ^{136}Xe	300 kg @ 80 % ^{131}Xe + 4700 kg @ 80 % ^{136}Xe	500 kg natural Xe + 4500 kg @ 90 % ^{136}Xe
^{131}Xe (poison)	0 %	5.4 %	4.2 %	4.8 %	2.1 %
^{136}Xe (signal)	90.0 %	84.6 %	84.6 %	75.2 %	81.9 %
Predicted from thermal σ	1× (ref)	≈22×	≈17×	≈21×	≈12×

Read the last two rows together. F buys the most poison but gives up the most signal — ^{136}Xe falls from 90 % to 75 %. G is the cheap option: 500 kg of ordinary natural xenon brings about 105 kg of ^{131}Xe with it, and costs the signal 8 percentage points.

^{131}Xe is the poison

21 % of natural xenon, σ = 85 b — 330× the ^{136}Xe thermal cross-section.

^{136}Xe is the price

Every recipe trades signal atoms for capture cross-section. F trades the most.

G is the outlier

It reaches only 2.1 % ^{131}Xe — but the feedstock is a commodity, not an enrichment campaign.

Measured: every recipe gives 3.6–5.4×, not 12–22×

Four xenon recipes, no shield, 10^6 muons each, against the 406 ± 25 baseline.

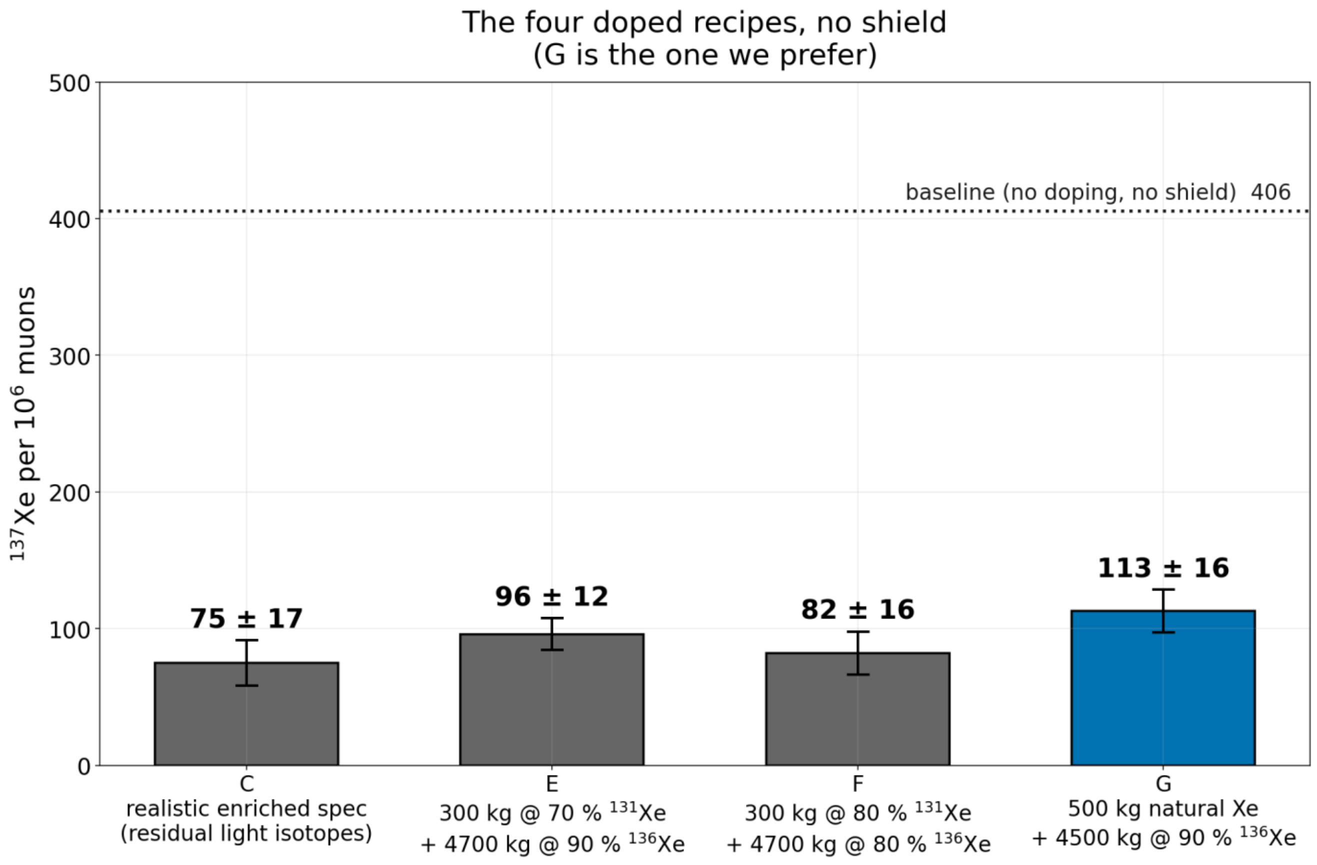

- **Measured reduction.** C 5.4×, E 4.2×, F 5.0×, G 3.6×. Thermal arithmetic had said 22 / 17 / 21 / 12×.
- **They barely separate.** C to G span 75–113 per 10^6 μ, about 1σ apart. Physics does not dominate the choice here, so cost weighs more.
- **Why the shortfall.** Half the parent neutrons are born inside the xenon and are captured before they thermalise. No amount of ^{131}Xe helps there.
- **What we carry forward.** Recipe G: 406 → 113 ± 16 per 10^6 μ, a factor 3.6 ± 0.5.

The xenon alone is not enough: no recipe reaches the order of magnitude the thermal estimate promised. We need the shield too.

So we pick the cheap one: recipe G

The four recipes land close together, so cost and signal loss carry more of the decision.

	C / E / F — enriched ^{131}Xe	G — natural xenon
Poison feedstock	300 kg of 70–90 % ^{131}Xe	500 kg of ordinary natural Xe
Unit price (indicative)	≈ 71 $/g	≈ 0.5 $/g
Cost of the poison	≈ 21 M$	≈ 0.25 M$
^{136}Xe retained	75 – 85 %	81.9 %
Measured reduction	4.2 – 5.4×	3.6×

≈ 85×

cheaper poison, for about 70 % of the reduction. Prices are 2026 quotes — an estimate, not a bid.

Two further arguments for G

- Natural xenon is a commodity. No enrichment campaign, no new supplier.
- ^{136}Xe loss is 8 points, against 15 for recipe F.

LEVER 2

A boron layer outside the vessel

Catch the neutrons after the HFE has thermalised them, before they reach the xenon.
We simulated borated polyethylene and boric acid, and compared the two.

It is the boron that works, not the plastic

Scanned 0–100 mm: borated polyethylene, natural boric acid, 95 % ^{10}B boric acid, and a boron-free null test.

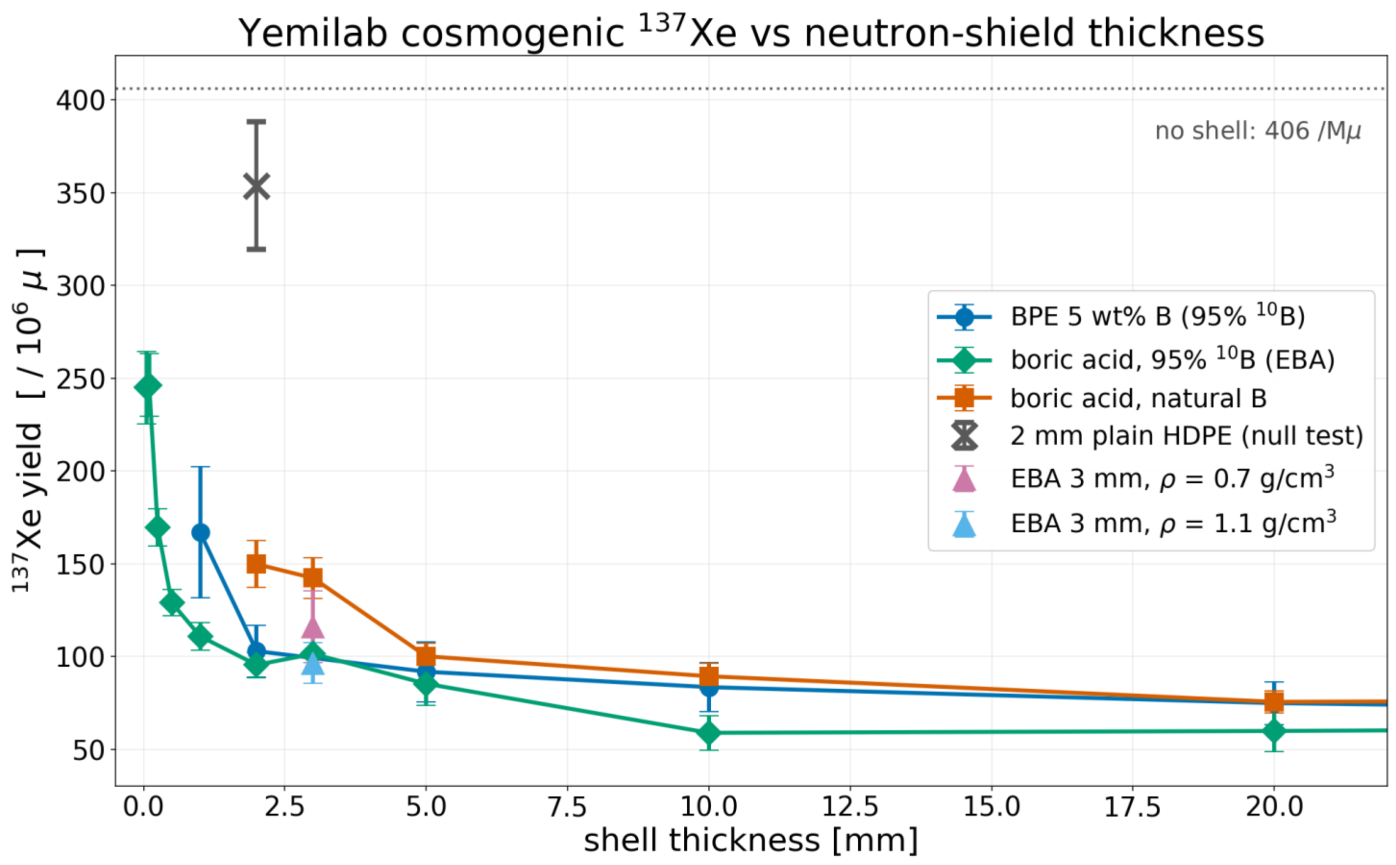

- **The null test is the proof.** 2 mm of plain polyethylene gives 354 ± 34 — the 406 baseline, untouched. The plastic on its own does nothing.
- **So we do not need the plastic.** Boric acid does the same job, without a bulk polymer inside the cryostat that would itself have to be assayed. Natural boric acid is our baseline.
- **What matters is the ^{10}B areal density.** BPE, natural B and 95 % ^{10}B track one curve in $\tau = \Sigma_a t$. Densities ρ = 0.7 / 0.9 / 1.1 g/cm³ at 3 mm differ little.
- **Steep drop, then a plateau.** Nothing below ≈60 per 10^6 μ is reachable with any shell — the floor from the neutrons born inside.

Why so thin is enough: the HFE has already slowed the neutrons down. The shell only has to absorb, not moderate.

The working point, and the geometry that was simulated

5 mm of natural boric acid, ρ = 0.9 g/cm³, standing 30 mm off the copper vessel.

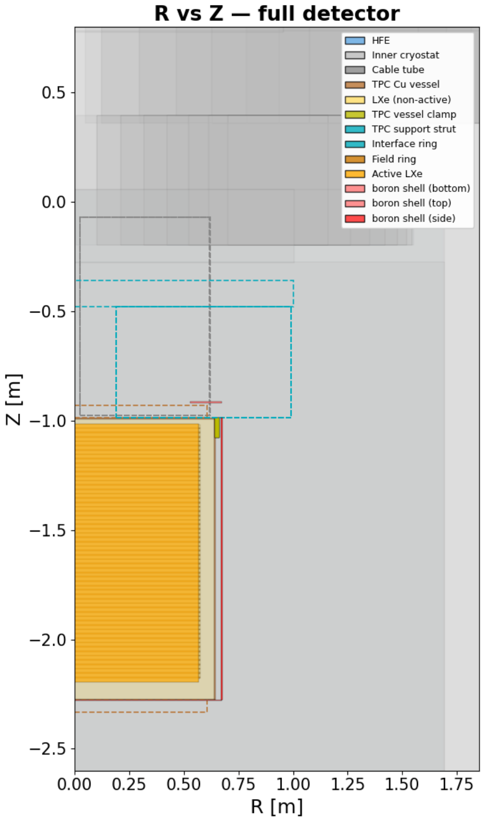

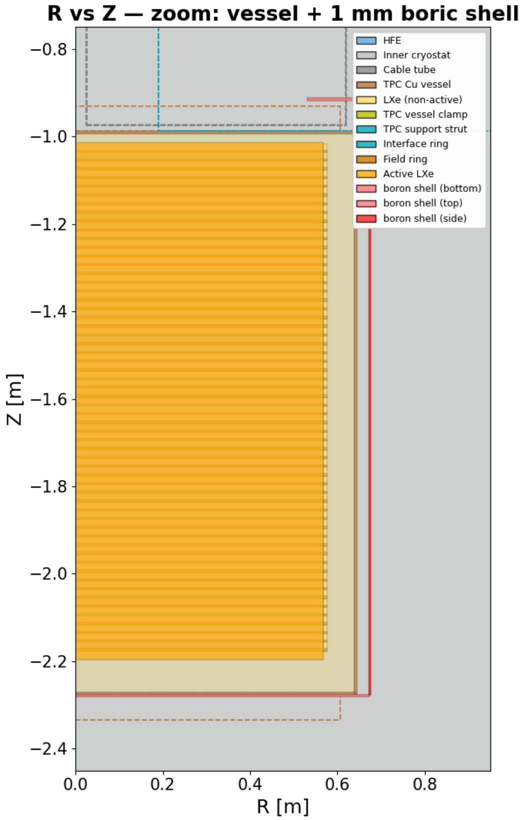

The shell drawn here is 1 mm, thin enough to see the three pieces. What the figure fixes is the geometry and the position: 30 mm off the copper vessel, the same in every run.

100 ± 7

^{137}Xe per 10^6 μ

(from 406)

4.0×

reduction, xenon unchanged

- **Three pieces.** R_in 671.5 mm, H_half 645.1 mm, R_top,in 530 mm — side 4.90 + bottom 1.28 + top 0.48 kg per mm.
- **At 5 mm that is 33.6 kg.** Small compared with anything else in the cryostat.
- **Deliberately simplified.** The real structure outside the vessel is complex; this is a 3-piece shell with a 3 cm radial gap, drawn straight from the Geant4 geometry.
- **Good enough for a reduction factor,** not yet an engineering drawing.

Natural or enriched? Same shielding, five times the mass

1 mm of 95 % ^{10}B boric acid and 5 mm of natural boric acid are the same shield to a neutron.

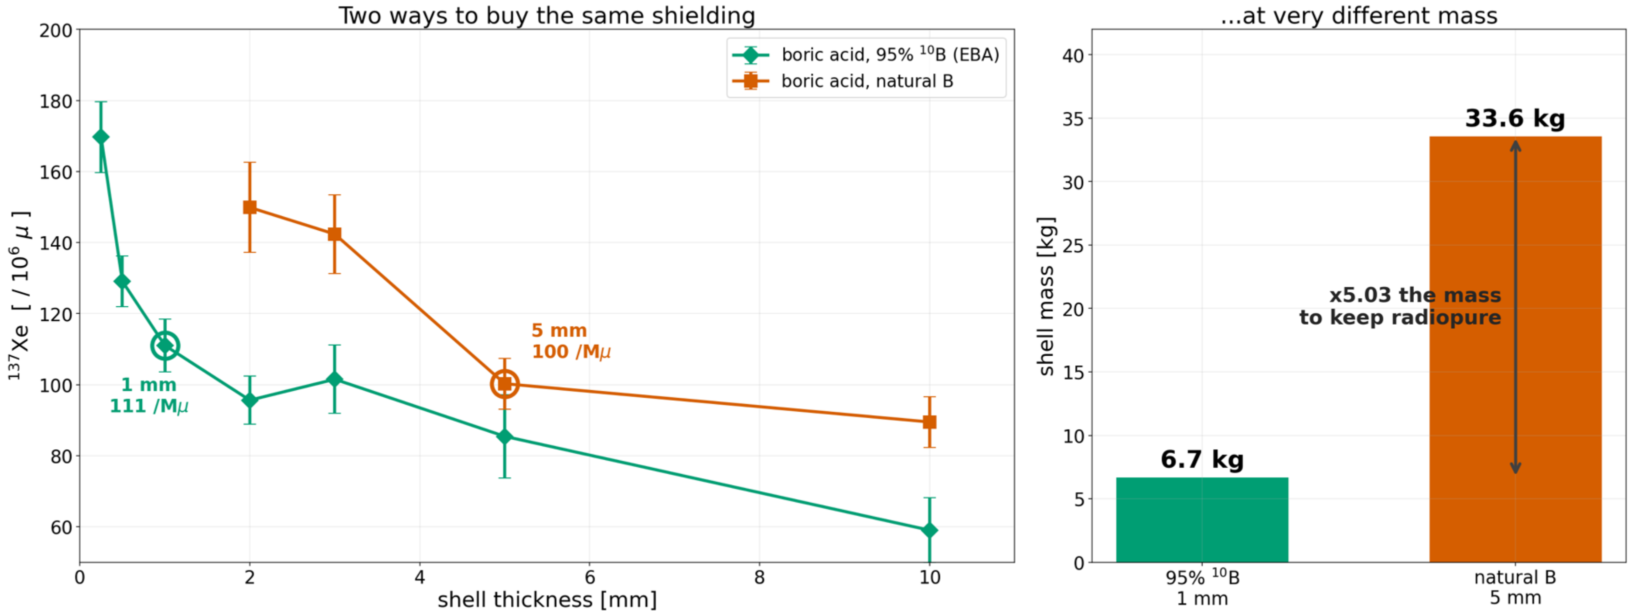

Same ^{10}B optical depth

$\tau = \Sigma_a t = 3.43$ (1 mm enriched) vs 3.35 (5 mm natural). That is the only thing the neutron responds to — and the measured yields agree: 111 ± 8 vs 100 ± 7.

What differs is mass

6.7 kg vs 33.6 kg. The radiopurity requirement is per kilogram, so the enriched shell is 5.03× easier to keep clean — and 5× less to mount against the vessel.

We simulate the natural shell as the conservative baseline; enrichment stays open.

The shield has to hug the vessel

Giorgio asked whether standing it further out would relax its radiopurity requirement. It costs too much ^{137}Xe.

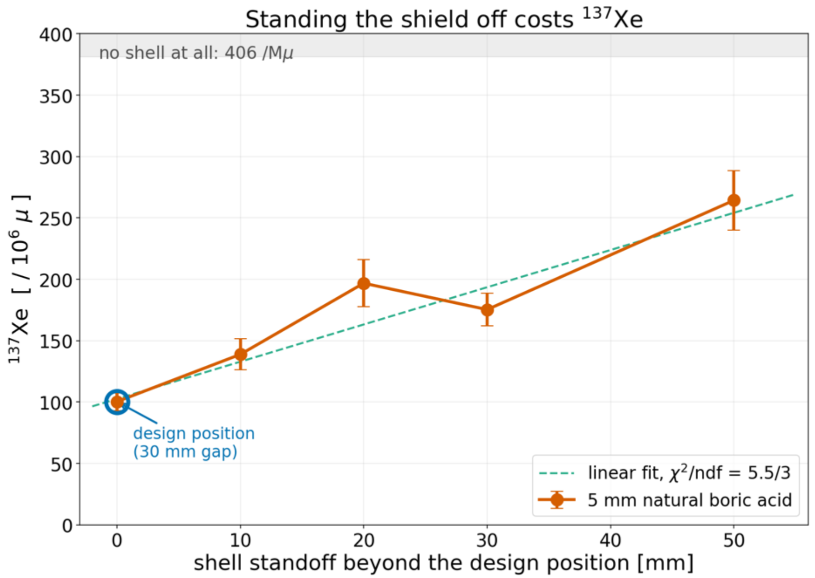

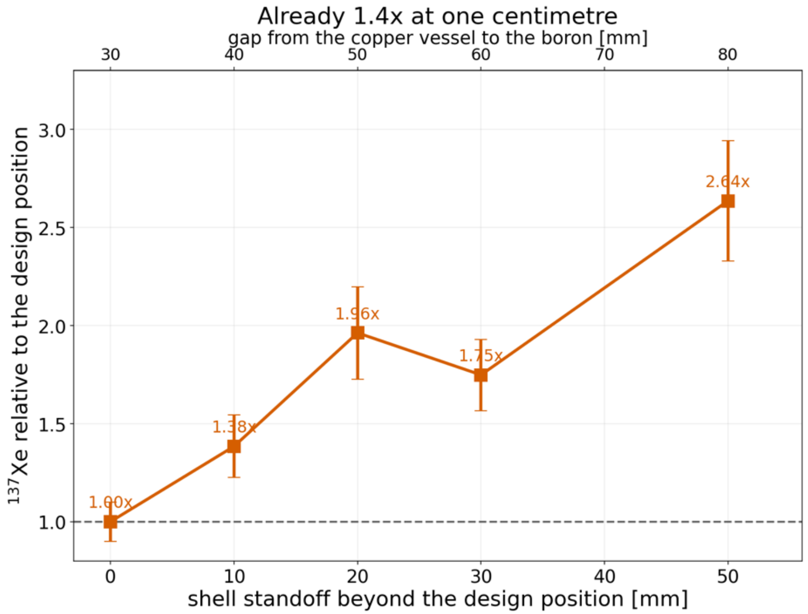

1.38×

more ^{137}Xe per extra centimetre of gap

2.64×

at an 8 cm gap — and no thickness recovers it

Five separate cosmogenic productions, identical except for the shell position. This also sets the installation tolerance.

Setting the budget from nEXO's sensitivity curve

Fig. 15 of arXiv:2106.16243 — half-life vs background scaling. Panel (c) scales every γ component: ^{222}Rn ^{214}Bi plus the detector materials.

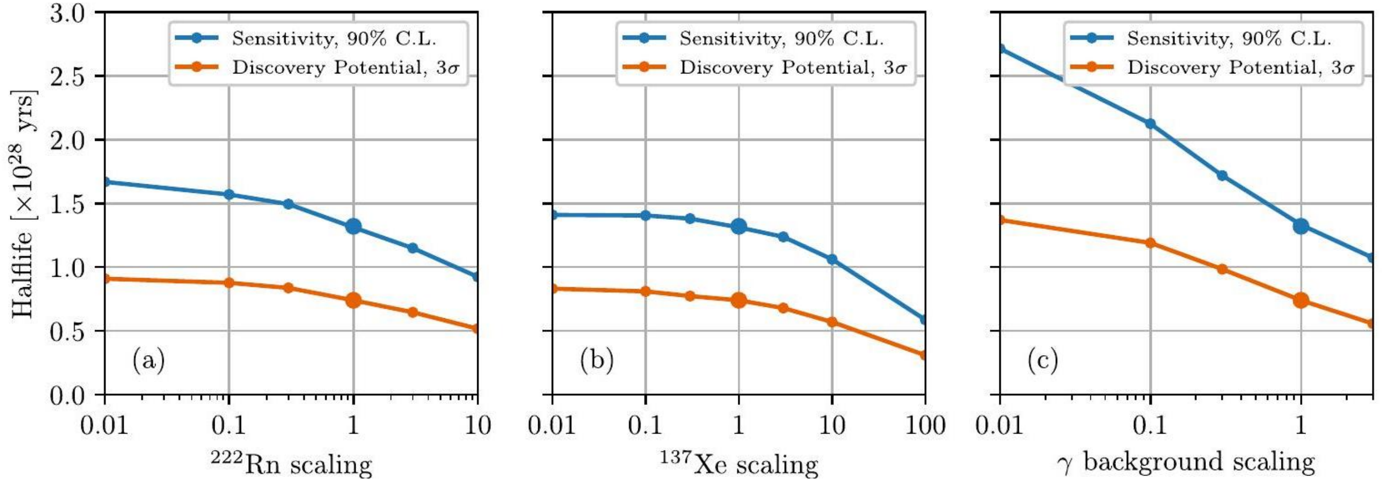

The numbers

- Total ROI background: 7×10^{-5} cts/(FWHM·kg·yr) → 0.14 counts/year in the inner 2000 kg (Sec 5.3).
- γ is ≈95 % of that total, so the shell's ROI counts add straight onto B — panel (c) is the axis the shell moves along.
- The paper's own fit (Sec 5.2) is $T½ \propto B^{-0.22}$, so $\Delta T/T = (1 + \Delta N/B)^{-0.22} - 1$.

The budget: 0.01 counts/(FWHM·yr) for the shell in the inner 2 t. That is +7 % of B, and costs 1.5 % of the half-life reach. The 1.5 % is the threshold I chose for a negligible component — my choice, not an agreed nEXO number. For scale: +1 % of B costs 0.2 %, doubling the γ background costs 14 %.

A_max = budget / (ROI rate per µBq/kg of shell).

The requirement: ≈1 µBq/kg on ^{214}Bi, ≈8 on ^{228}Th

Per chain, if it were the sole contaminant, at the design position, in the inner 2 t. Bar labels are µBq/kg.

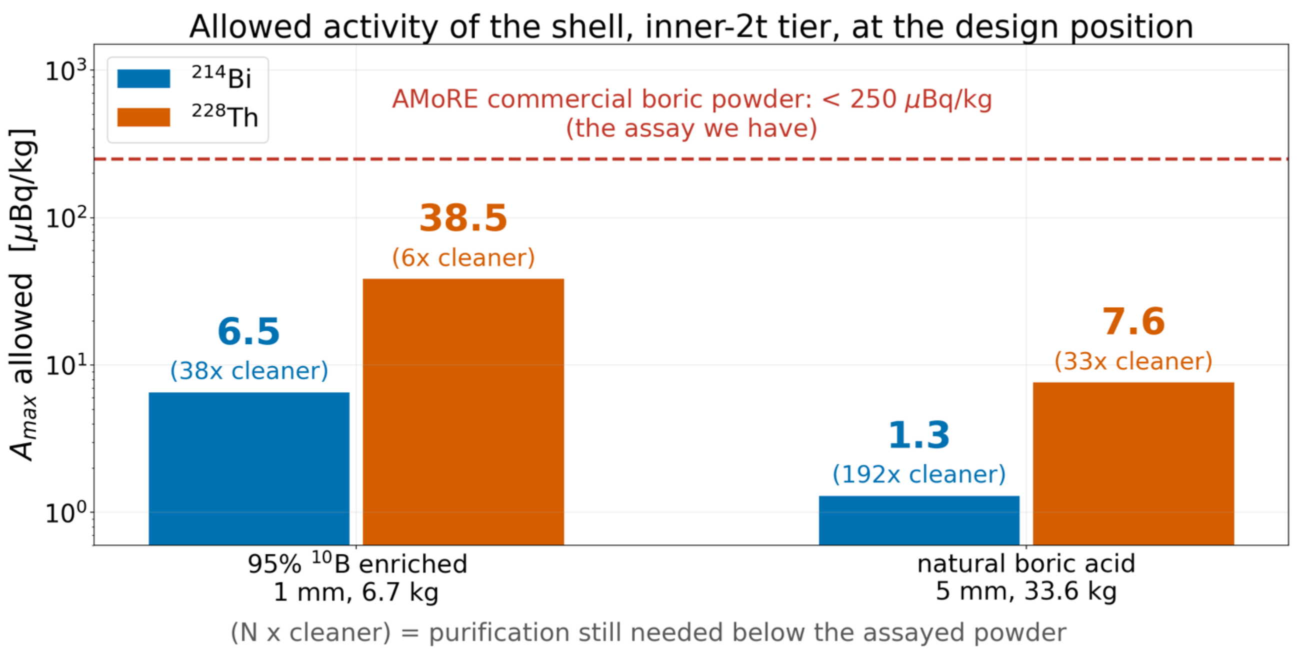

Reading the plot

- ^{214}Bi drives it — 2448 keV is already inside the ROI, so nothing has to happen to it.
- ^{228}Th is ≈6× looser: the 2615 keV line mostly lands outside the ROI.
- The natural shell needs it 5.03× tighter than the enriched one, purely because it weighs 5.03× more.

Is it reachable? The assay we have is AMoRE's commercial boric powder, < 250 µBq/kg, so we need ≈30–200× better than that. Following Olga's assessment that level is achievable for a purified powder — but it has to be measured, not assumed.

A_max is proportional to the 0.01 counts/year budget and inversely proportional to the shell mass.

PUTTING IT TOGETHER

10.4×, and what it does not include

The two levers are independent, so they multiply. The sensitivity number does not follow yet.

RESULT

Where that leaves Yemilab

Recipe G + a 5 mm natural boric-acid shell: 406 → 39 ^{137}Xe per 10^6 muons.

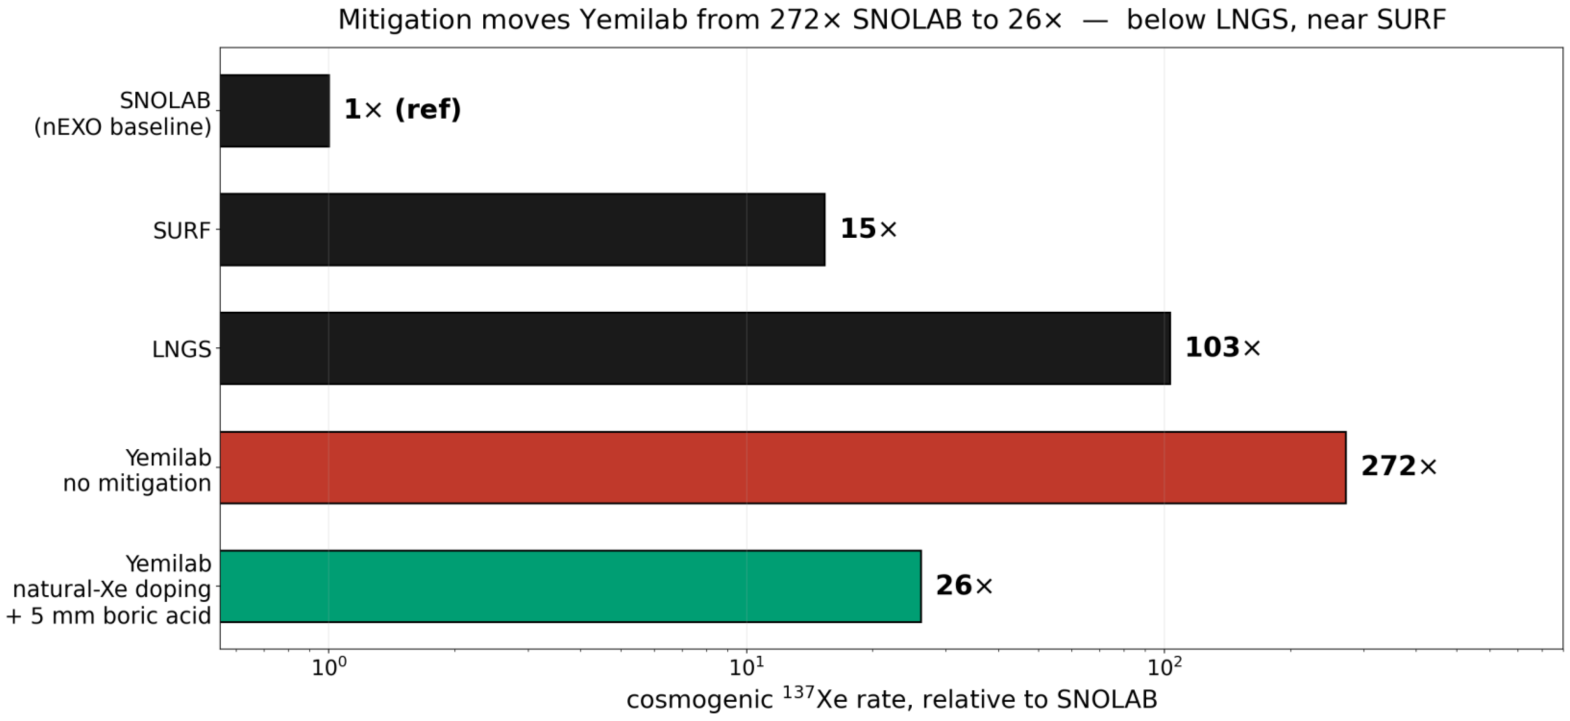

10.4×

combined reduction

(3.6× × 4.0× = 14 predicted)

The levers really are independent

One removes neutrons before they reach the xenon; the other changes what the xenon does with the ones that get through. A clean product would give 28 ± 5 per 10^6 μ; we measure 39 ± 4, which is 1.4 ± 0.3 times that, 1.4σ away. Some overlap is the direction physics predicts, but it is not established.

Mitigated Yemilab lands at ≈ 26× SNOLAB — under twice the SURF cosmogenic level, and well below LNGS.

Two ways to push the shield further

Both working points sit at the knee of the curve, not at the floor. Neither route is simulated yet — these are the next runs.

1 Go thicker — if the powder can be made cleaner

≈ 6.8×

10 mm of 95 % ^{10}B boric acid gave 60 ± 10 per 10^6 μ, against 4.0× at the present working point

We chose 1 mm enriched / 5 mm natural because that is where the curve turns over, not because it is a floor. Thickness keeps buying — but mass is what buys it. 10 mm of enriched powder is ≈67 kg, so A_max tightens tenfold, to ≈0.65 μBq/kg on ^{214}Bi.

So this is a radiopurity decision, not a neutronics one. If the purification reach is better than we assumed, the thickness is there to take.

2 Go closer — if the cryostat allows it

1.38× / cm

the measured penalty per centimetre of extra standoff (slide 17) — the same slope should work inward

The shell sits 30 mm off the copper vessel today. Moving it out costs 1.38× per centimetre with no plateau, and the mechanism behind that — the volume of HFE the shell encloses — points the same way going in.

The limit here is cryogenic and mechanical clearance, not neutron physics. It is worth asking how much of the 30 mm is really needed, and whether the top and bottom pieces could come in further than the side.

Both are cheap to run. A thickness scan at 5 / 8 / 10 / 15 mm enriched, and a gap scan inward of 30 mm.

What this number does not yet include

Everything below is in progress. None of it is a caveat on the 10.4× itself — they are the next pieces of work.

No T½ sensitivity number yet.

Recipe G costs 8 points of ^{136}Xe (90 → 81.9 %) and the shell adds a γ line item. Signal loss, ^{137}Xe gain and shell γ have to enter one background-and-sensitivity fit before a half-life can be quoted.

The ^{137}Xe tagging efficiency is not folded in.

The nEXO background model already vetoes cosmogenic ^{137}Xe with the muon tag. How well that works at 272× the muon rate has not been re-evaluated.

The cost of the muon veto is not evaluated.

A time-and-space veto around each tagged neutron capture removes ^{137}Xe, but it also removes livetime. At 272× the SNOLAB muon rate nobody has estimated that cost.

No thermal or mechanical study of the shell.

Heat load, containment of a hygroscopic powder, outgassing into the HFE, and how it is mounted with a 30 mm tolerance are all untouched.

The shell radiopurity is a requirement, not a result.

≈1.3 µBq/kg on ^{214}Bi needs ≈200× purification below the powder we have assayed, and then an assay to prove it.

SS selection uses an RMS proxy, not the nEXO DNN.

Worth ≈1.5× on the radiopurity requirement; the cosmogenic reduction factors do not depend on it.

Summary

Two levers, worth 3.6× and 4.0×.

A neutron poison in the xenon — 500 kg of natural xenon inside the 5 t; a thin boron layer outside the vessel — 5 mm of natural boric acid.

Together: 10.4 ± 1.2×.

406 → 39 ^{137}Xe per 10^6 muons, measured in one combined run — Yemilab at ≈26× SNOLAB, under 2× the SURF level.

All four xenon recipes land within ≈1σ.

75–113 ^{137}Xe per 10^6 μ. Physics does not pick the winner, so cost does: 500 kg of ordinary natural xenon.

The shield must hug the vessel, and be very clean.

1.38× more ^{137}Xe per cm of standoff; ^{214}Bi below ≈1.3 μBq/kg — reachable, per Olga, but it must be assayed.

Not yet a sensitivity number.

^{136}Xe drops 90 → 81.9 % and the shell adds γ. One combined background fit, plus shell engineering, is the ongoing work.

Backup

Supporting material.

Why: inside the shield, the nearer sink wins

Neutrons fired from inside the cavity, so they never have to cross boron at all.

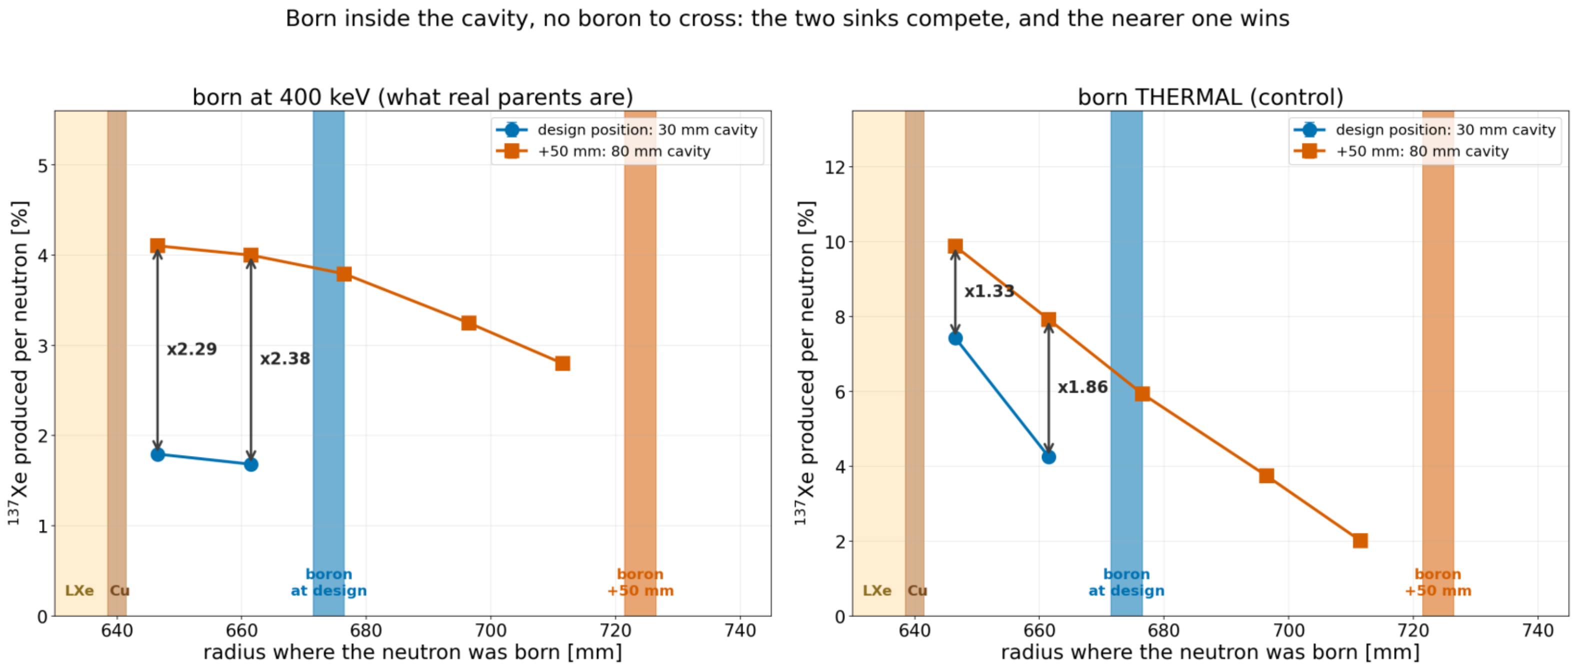

Not an attenuation effect

One crossing of 5 mm absorbs 96.5 % of a thermal neutron but only 0.1 % of a 400 keV one. Standing off does add moderating path — but the cliff travels with the shell, so that is not the reason.

It is a competition of volumes

A neutron that thermalises inside the shell has no boron to cross. It ends on whichever sink is nearer — the boron, or ^{136}Xe. Moving the shell out grows the enclosed HFE from 262 to 723 L, a factor 2.76.

2.76× the volume, 2.64× the yield.

The cut chain, matched to the published nEXO analysis

^{214}Bi (2.20 MeV) and ^{208}Tl (2.61 MeV) sit right at Qββ = 2458 keV, and the shell is inside the cryostat.

Cut	This work	nEXO sensitivity paper (arXiv:2106.16243)
Energy / ROI	σ/E = 0.8 % smear, keep Qββ ± FWHM/2	Sec 3.4 / 4.4 — identical window
SS / MS	energy-weighted 3-D RMS of deposits < 5 mm	DNN discriminator cut at 0.85
Fiducial volume	centroid standoff > 20 mm → 3284 kg	Sec 3.2 — 3281 kg
Reporting volume	standoff > 104 mm → 2007 kg	Sec 4.4 — inner 2000 kg

Cross-check. Fiducial masses reproduce the published values to better than 0.1 %. The SS proxy is an RMS, not the nEXO DNN — a ±2 mm change moves the answer by ≈1.5×, which is my quoted systematic.

How the number is set

- 3×10^8 decays per isotope, thrown through the shell sub-volumes in proportion to their mass.
- Then the same three cuts as the paper, and count what lands in the inner 2 t ROI.
- A_max = budget / (rate per µBq/kg), with the budget from the previous slide.

nEXO sensitivity vs background scaling (paper Fig. 15)

J. Phys. G 49, 015104 (2022) · arXiv:2106.16243

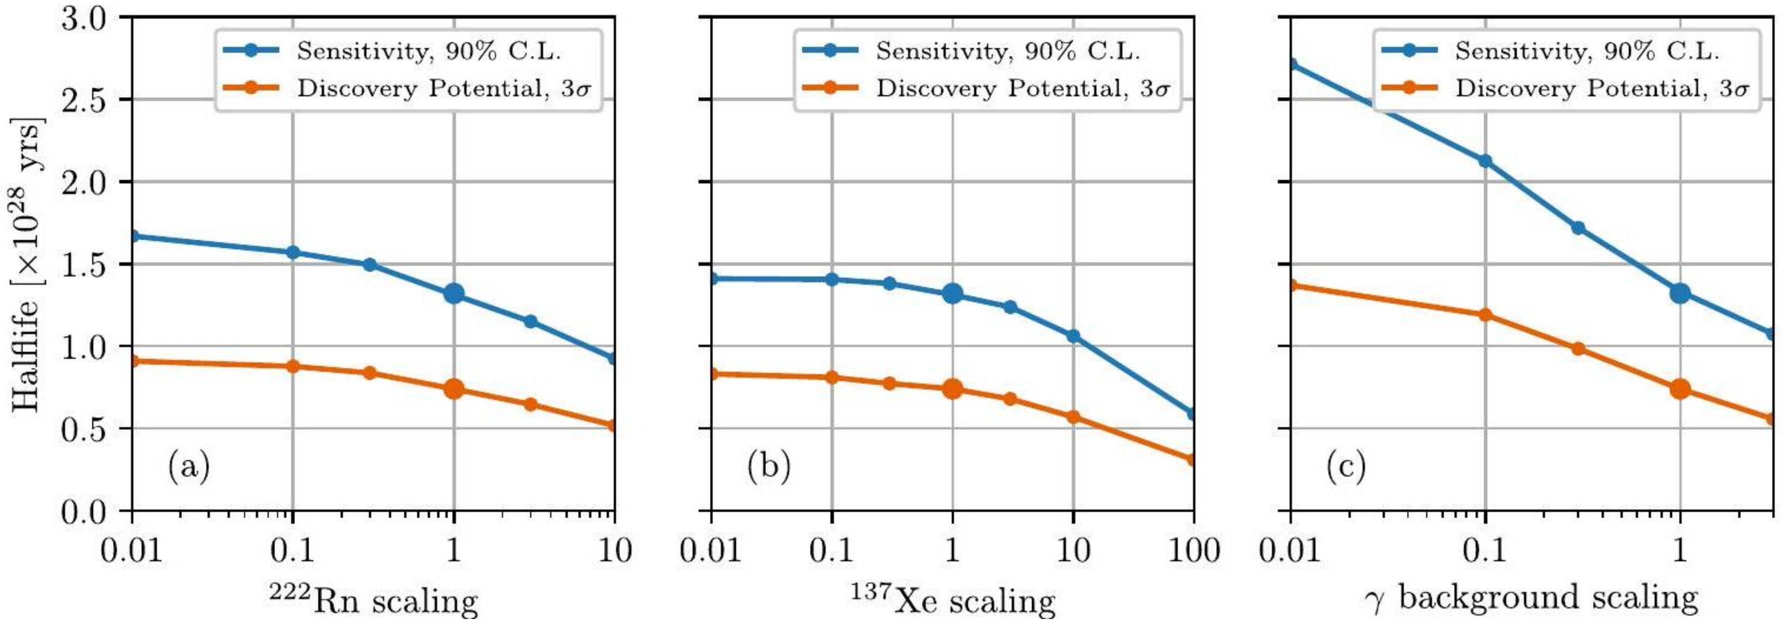

Panel (b) is the ^{137}Xe axis — the lever this talk pulls; scaling ^{137}Xe by 100 takes the 90 % C.L. sensitivity from 1.32 to 0.58 $\times 10^{28}$ yr. Panel (c), the γ axis, is where the shield's own ^{214}Bi / ^{208}Tl enter. The Sec 5.2 fit to these curves, $T_{1/2} \propto B^{-0.22}$, is what turns the shell's ROI counts into a half-life cost, and so fixes the 0.01 cts/(FWHM·yr) budget.

The TPC tag on cosmogenic ^{137}Xe

What the dark markers on the sensitivity plot mean.

- **Where the ^{137}Xe comes from.** Muons make neutrons, and a neutron capture on ^{136}Xe is what makes the ^{137}Xe.
- **The capture tags itself.** It emits a prompt γ, which gives a time and a position.
- **Then veto around it.** Events close to a tagged capture, in time and in space, are cut — that removes part of the ^{137}Xe background.
- **Not used in this talk.** Every number in the main slides is raw ^{137}Xe production, before any veto.

The veto is not free: it also removes livetime, and at 272× the SNOLAB muon rate that cost has not been estimated.

Radioassay cut 1 — energy resolution and the ROI

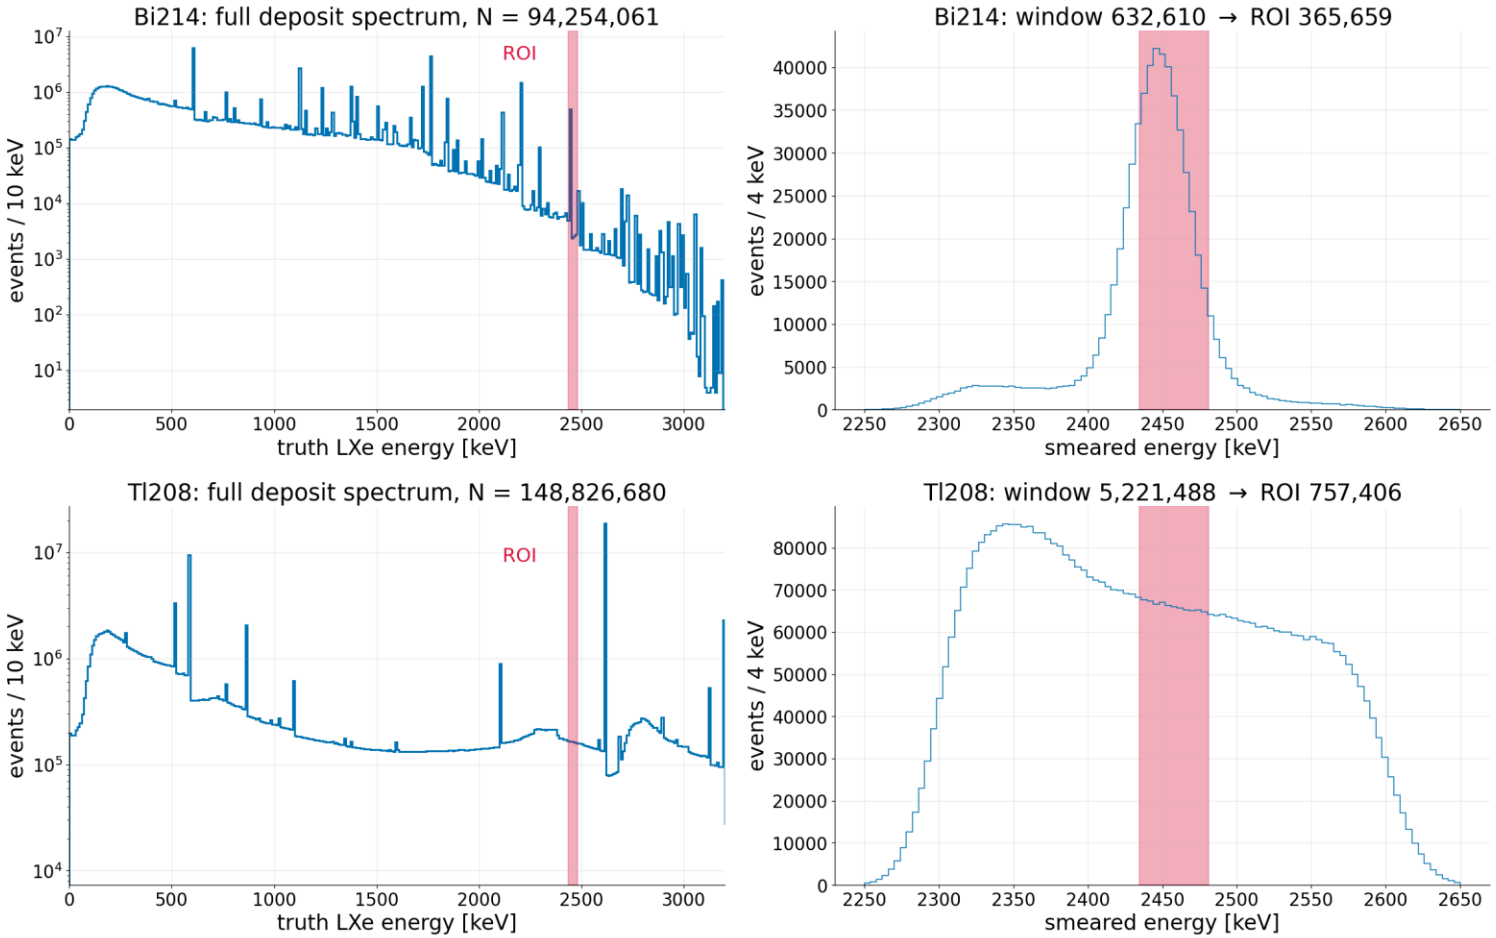

Truth deposits smeared at σ/E = 0.8 %, keeping $Q_{\beta\beta}$ ± FWHM/2 = 2434.6–2481.0 keV. ^{214}Bi: 633 k → 366 k. ^{208}Tl: 5.2 M → 757 k.

Radioassay cut 2 — single-site topology

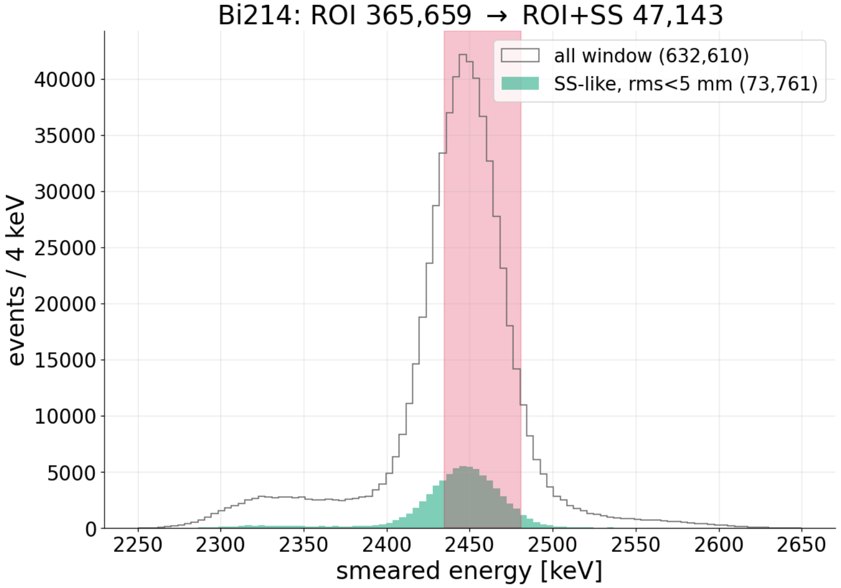

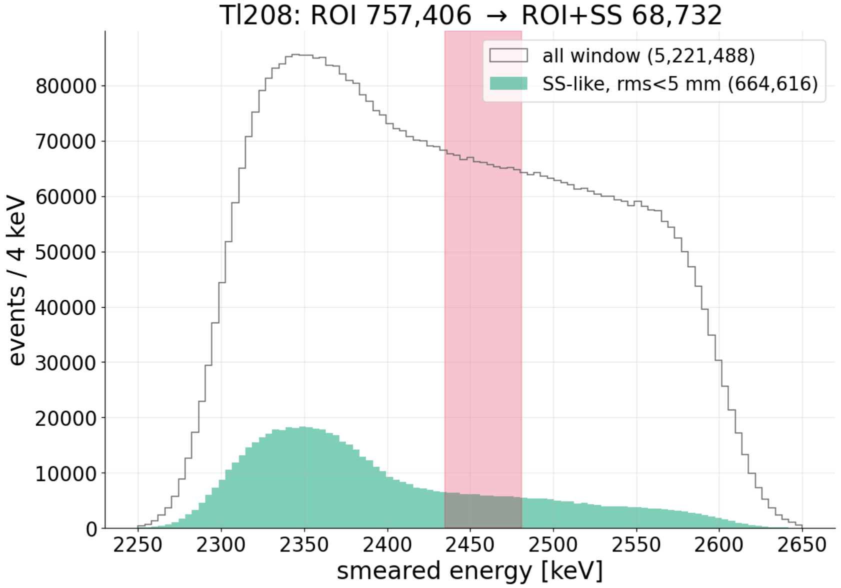

Energy-weighted 3-D RMS of the Geant4 deposits < 5 mm. This is the big reducer: most 2.4–2.6 MeV γ Compton-scatter before photo-absorption. ^{214}Bi 366 k → 47 k, ^{208}Tl 757 k → 69 k.

Radioassay cut 3 — fiducial volume

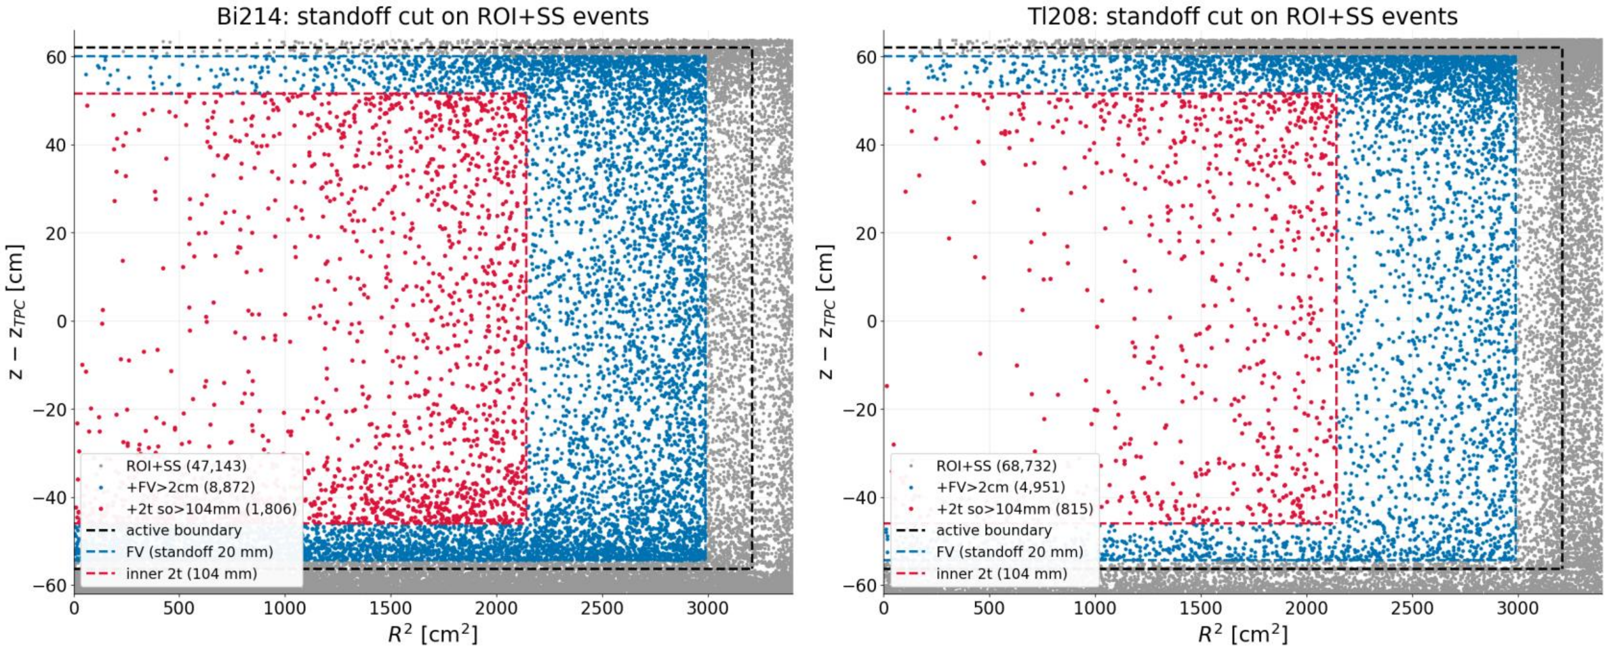

Standoff > 20 mm → 3284 kg (paper: 3281 kg); the inner-2 t reporting metric is standoff > 104 mm → 2007 kg. Surface-peaked background is exactly what the fiducial cut is meant to remove.

Radioassay — survivors and the SS systematic

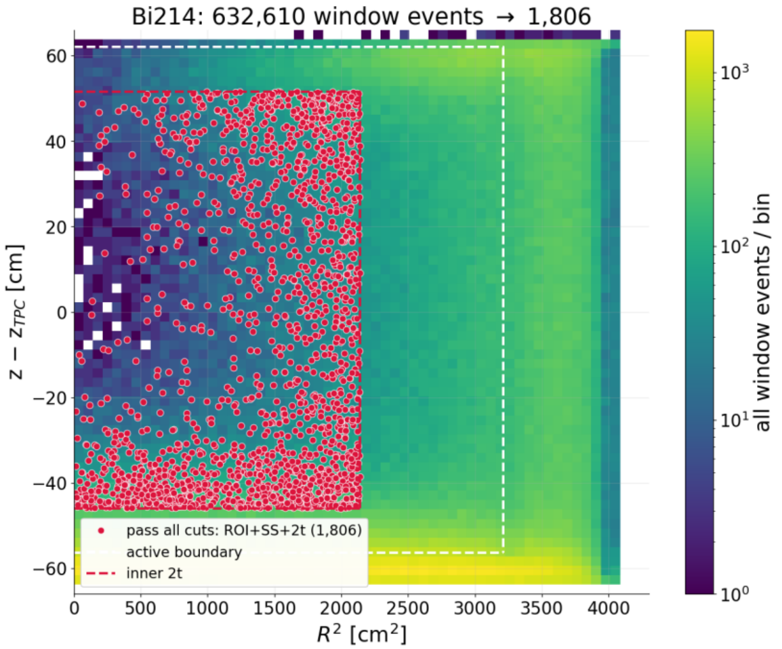

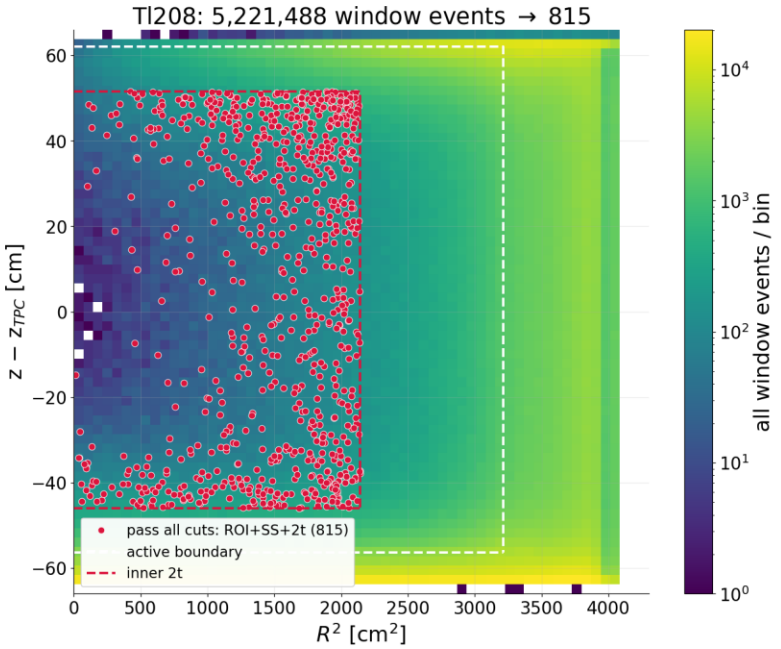

After ROI + SS + inner-2 t: ^{214}Bi 1806, ^{208}Tl 815 out of $3{\times}10^8$ decays each. Counts scale roughly linearly with the RMS threshold — ≈1.5× systematic for a ±2 mm choice.

Fast neutrons walk straight through the shield

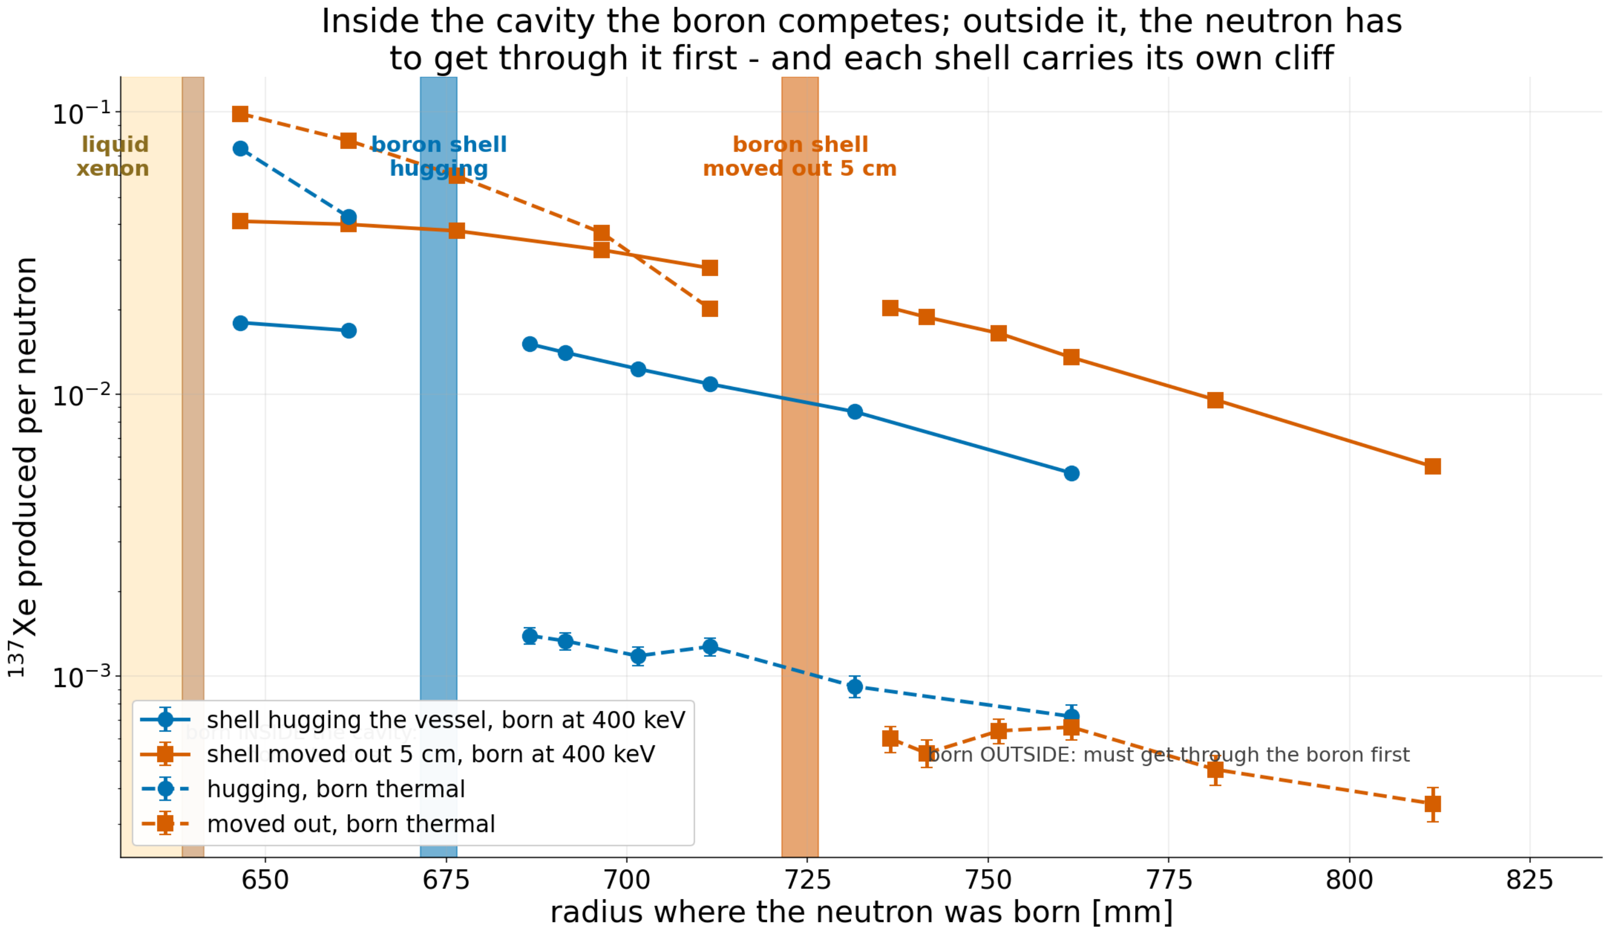

Neutron slabs parked at chosen radii, 150 000 per point; solid = born at 400 keV, dashed = born thermal. One crossing of 5 mm absorbs 96.5 % of a thermal neutron and 0.1 % of a 400 keV one — boron is 1/v.

^{137}Xe yield tracks the HFE volume enclosed by the shield

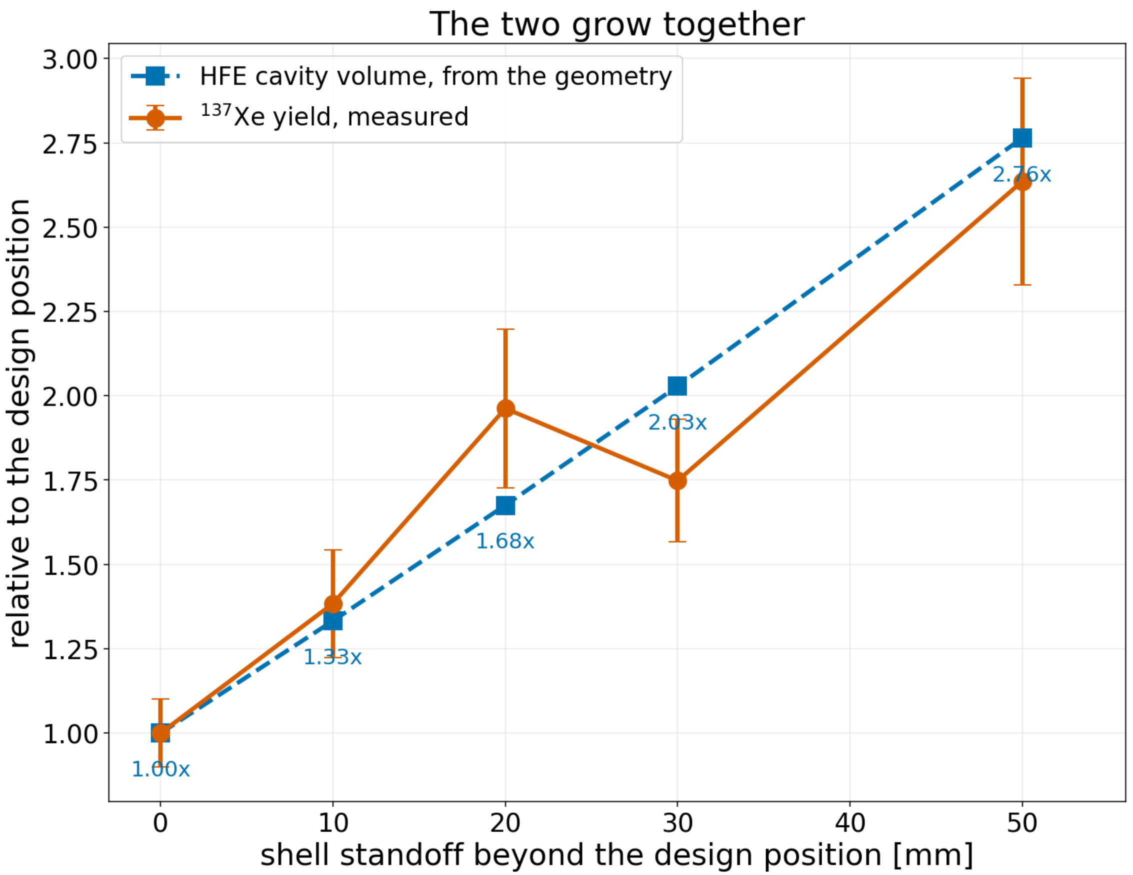

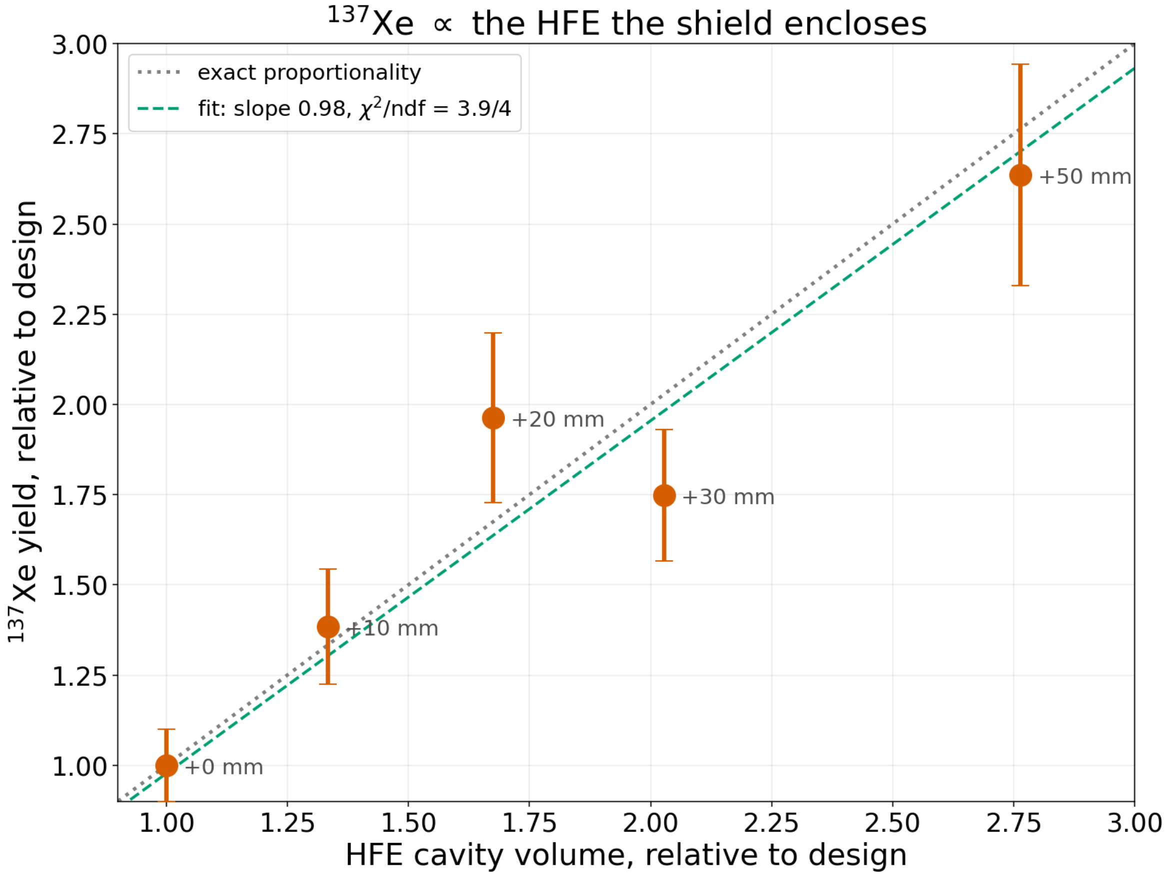

Proportional fit through the origin: slope 0.98, χ^2/ndf = 3.9/4. Five points over a factor 2.8, with volume and gap near-degenerate — consistent with the mechanism, not a derivation of it.

The shield's own γ vs standoff — they disagree, but only near the vessel

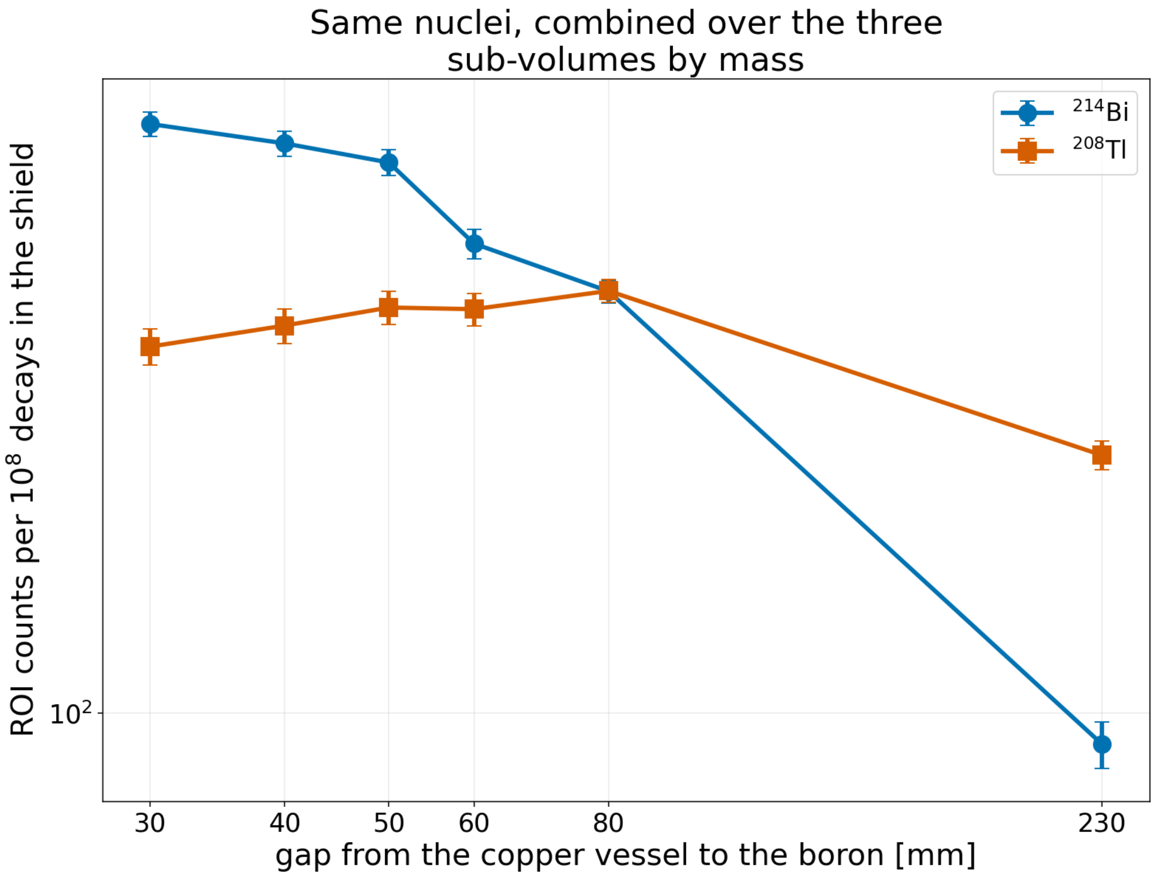

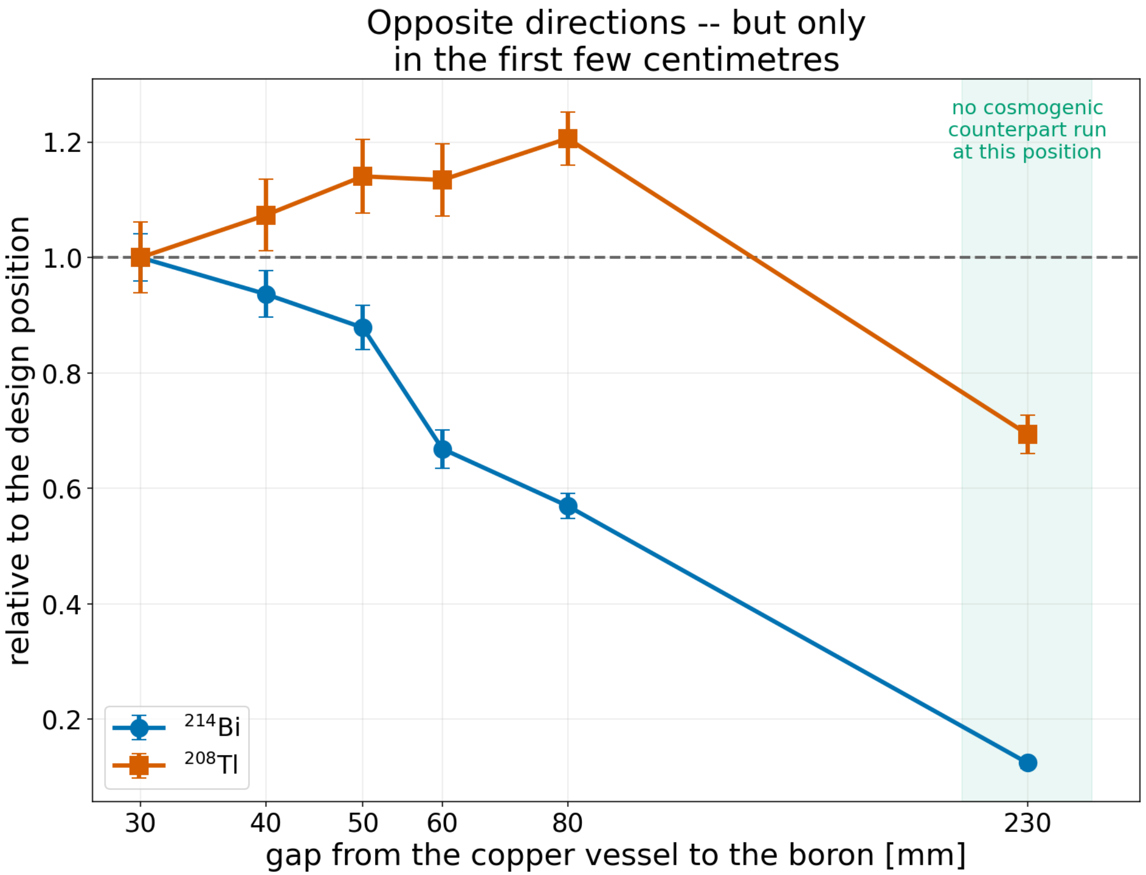

Six positions, 10^8 decays of each isotope distributed over the sub-volumes by mass. Over 30–80 mm ^{214}Bi ROI counts fall to 0.57× while ^{208}Tl rises to 1.21× — but a far point at a 230 mm gap reverses it: ^{214}Bi 0.12×, ^{208}Tl 0.69×. The ^{208}Tl rise is a near-field effect, not a trend.

Where ^{214}Bi and ^{208}Tl part company — cut by cut

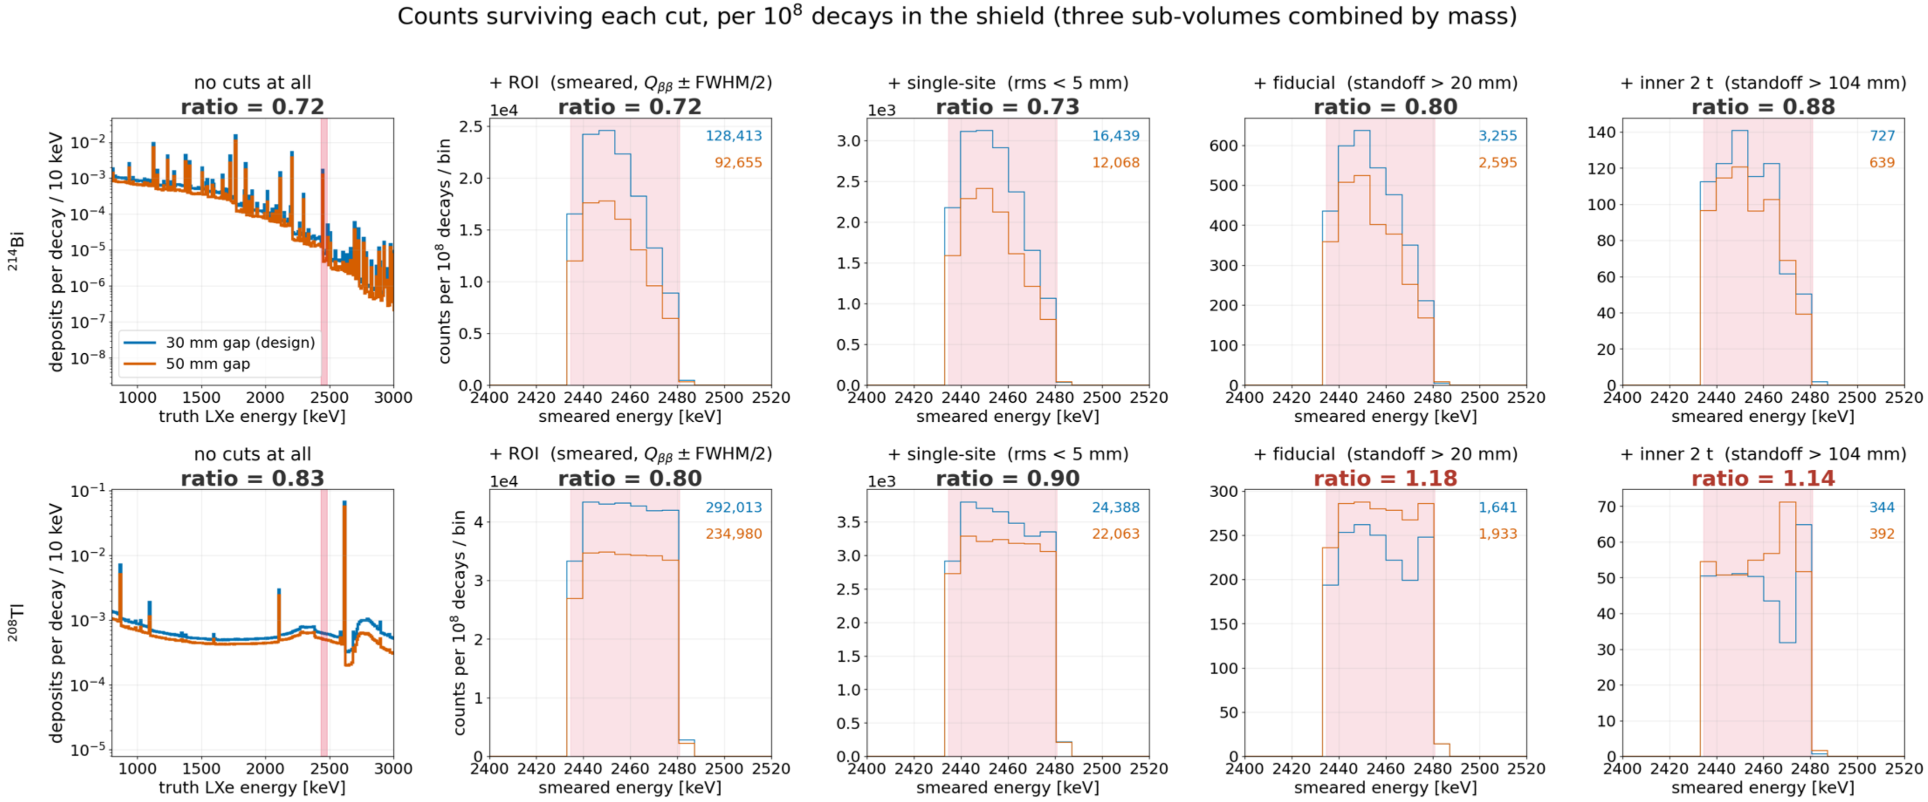

30 mm gap against 50 mm, one cut at a time. ^{214}Bi's 2448 keV is already in the ROI; ^{208}Tl's 2615 keV must shed ≈165 keV — in the xenon (multi-site, cut) or in the HFE/Cu before arriving (single-site, kept). More HFE favours the second.