KRONOS FUSION ENERGY

PAPER 2.27 · THE KRONOS 2026 SERIES

Nuclear-Data Uncertainty Quantification for Tritium Breeding: Which 14 MeV Cross-Section Channels Set the Error Bar on the Breeding Ratio

A tritium breeding ratio (TBR) is only as trustworthy as the cross-sections that produced it, and a design that quotes a TBR without an error bar is quoting a point without a...

Read it live — 3-D model · physics validator · film series

BY P I FORD

2026 · EDITORIAL EDITION

THE OFFICIAL RECORD

What this document is, and how to read its numbers

This is a combined editorial. It gathers, in one volume, the two design studies Kronos Fusion Energy released in 2026 — a compact spherical-tokamak tritium/helium-3 breeder and a deuterium–helium-3 tandem-mirror generator — together with the intellectual-property estate, the people, and, in the internal edition, the full commercial case. The scientific text is the same text that appears in the arXiv preprints and the journal submissions; the editorial adds the apparatus a reader needs to hold the whole program at once, and nothing in the physics is altered to fit it.

THE HONESTY STANCE IS THE ASSET

Every headline quantity carries an **evidence class** — DERIVED RECOMPUTED MEASURED REQUIREMENT — and, where a number is a modelled extrapolation rather than a demonstrated value, it is labelled as such and its downside is shown next to it. Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly. This is deliberate: a diligence team will find the load-bearing assumptions, so the document surfaces them first.

TWO EDITIONS

This volume is issued in two editions. The **public edition** (light) carries the physics and the design only — no dollar, price, or per-kilogram figure appears anywhere in it. The **internal edition** (dark) adds the economics, the funding architecture, and the defense-commercial analysis, and is marked *Confidential* accordingly; it is not for public distribution. Where the commercial case rests on a modelled assumption rather than a demonstrated one, it is labelled as such and its downside is shown alongside it, on the same terms as the physics.

TITLE	The Kronos Fleet — Hyperion · Aegis · MetroVolt: A Combined Design & Commercialization Study
AUTHOR	Priyanca Ford, Founder & CEO, on behalf of Kronos Fusion Energy
EDITION	2026 Combined Editorial · first issue
COMPANION WORKS	The five-paper arXiv drop — (1) Breeder, (2) Burner, (3) REBCO magnet & tape, (4) Direct Energy Conversion, (5) AI/ML/Quantum control — with matching numbered Zenodo deposits, plus the IOP journal and IAEA-FEC submissions. Patents were filed before the drop.
NAMING	Hyperion = the breeder; Aegis (defense) and MetroVolt (data-center) = one generator in two housings. Internal mode letters (D1, L) do not appear on public surfaces.
STATUS	Conceptual design & simulation study. Nothing herein is a construction commitment.

HOW TO READ THIS VOLUME

The fleet at a glance, and the way its numbers are marked

This volume opens with *Orientation*, presents the five research papers as a bound sequence, and closes with the *Apparatus*. *Foundations* sets out the three general results both machines rest on. *Paper One* is the Hyperion breeder; *Paper Two* the Aegis / MetroVolt generator; *Papers Three, Four and Five* are the enabling science — high-field REBCO magnets and tape, direct energy conversion, and the AI / ML / quantum control stack. A *Digital Twin & 3-D Model* part follows, then the *Environmental* profile. The *Economics* and levelized cost and the *Business Case* and diligence record appear in the internal edition only; the *Simulations* proof record and the Apparatus — patent portfolio, team, limitations, and sources — close both editions.

THE FLEET AT A GLANCE

	HYPERION	AEGIS	METROVOLT
Role	Strategic-isotope foundry	Defense installation power	Campus power
Machine	**Spherical tokamak**	**Tandem mirror**	**Tandem mirror**
Fuel	Deuterium–tritium	Deuterium–helium-3	Deuterium–helium-3
Delivers	Tritium · ^{3}He · 14 MeV n	Resilient installation power	Firm campus power
Physics bar	**Fusion gain only**	Closure (gates named)	Closure + lunar ^{3}He
In the fleet	Breeds the fuel	Proves the generator	Commercial destination

HOW THE NUMBERS ARE MARKED

Every headline quantity carries an evidence class — DERIVED from first principles, RECOMPUTED from raw data, MEASURED in hardware, or a REQUIREMENT yet to be met. Financial figures carry a case label and are confined to the internal (confidential) edition. Values that moved after first publication are shown as withdrawn or restated rather than quietly changed. A capability you can evaluate is one whose limits are on the page.

ABSTRACT

Nuclear-Data Uncertainty Quantification for Tritium Breeding: Which 14 MeV Cross-Section Channels Set the Error Bar on the Breeding Ratio

A tritium breeding ratio (TBR) is only as trustworthy as the cross-sections that produced it, and a design that quotes a TBR without an error bar is quoting a point without a tolerance. We establish the credible nuclear-data uncertainty on the TBR of a compact spherical-tokamak breeder (design point $I_p = 9.66\,MA$, $Q = 3.076$, $P_{\text{fus}} = 85.04\,MW$, $\delta = -0.30$, $B_0 = 8\,T$, $\kappa = 2.0$, $R_0 = 1.20\,m$, aspect ratio $A = 2.5$) by casting the TBR as an integral response functional of the transport problem, deriving its first-order cross-section sensitivities from generalised perturbation theory, and propagating a channel-resolved data uncertainty through the sandwich rule $\sigma_R^2 = \quad S^\top \quad C \quad S$. The load-bearing question is answered channel by channel. Processing evaluated data (ENDF/B-VIII.0) and default-model data (TALYS) identically through NJOY21 and transporting both through the same three-dimensional Monte-Carlo geometry, we measure a model-form spread of 55–66% at 14.06 MeV on the two breeding-relevant light channels — 7$(n, n't)$ (66.1%) and ^{9}Be $(n, 2n)$ (55.1%) — against only $\approx 15\%$ on the heavy structural multiplier ^{184}W$(n, 2n)$. Because TBR is a near-linear functional of the light reaction rates, these two channels carry $\approx 95\%$ of the model-form variance, and the single largest tritium producer, 6(n, t), is an evaluated standard whose covariance is of the few-per-cent class. The cross-library reproduction (ENDF/B-VIII.0 versus FENDL-3.2) agrees to 0.3% , two orders of magnitude below the model spread, proving the code path is not the source of the band — the evaluated data is. The result is a requirement, not a caveat: the light breeding channels must be taken from evaluated covariances, and a TBR reported without propagating them understates the tolerance on the most load-bearing number in the breeder design by more than an order of magnitude. Every headline quantity is a frozen anchor (gate BR-L2-A28) in the de-risking register[reg].

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645777 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: nuclear data · uncertainty quantification · tritium breeding ratio · ENDF/B-VIII.0 · TALYS · NJOY21 · sensitivity/uncertainty analysis · covariance

CONTENTS

Table of Contents

NOMENCLATURE

Symbols, units, and the names of things

Q	plasma (fusion) gain, P_{fus}/P_{aux}
Q_E	engineering gain, net electric out / recirculating in
H_{98}	confinement enhancement over the IPB98(y,2) scaling
Z_{eff}	effective ion charge
c_{ee}	electron–electron bremsstrahlung leading coefficient, 2.120022 (exact)
τ_E	energy confinement time
I_p	plasma current
R_0, a	major radius, minor radius
κ, δ	elongation, triangularity
β_N	normalized plasma beta (Troyon-normalized)
TBR	tritium breeding ratio
DEC	direct energy conversion
CF	plant capacity factor (availability)
FOAK / NOAK / BOAK	first / n^{th} / best of a kind
Hyperion	the ST breeder (internal: Mode D2)
Aegis	the generator, defense/installation housing (internal: Mode M)
MetroVolt	the generator, data-center housing (Mode M)

§ 1

Introduction

The tritium breeding ratio is the number on which the entire case for a breeding blanket rests. It is defined as the number of tritons produced per fusion neutron, and for a device whose product is tritium, helium-3 and 14 MeV neutrons rather than net electricity, it is the primary figure of merit: a blanket with $\mathrm{TBR} < 1$ consumes its own fuel, and the fleet path to self-sufficiency and export is set entirely by how far above unity the ratio can be driven and held [2,3]. The compact spherical-tokamak (ST) breeder studied here is a high-power-density, negative-triangularity device operating at $I_p = 9.66\,MA$, $Q = 3.076$ and $P_{\text{fus}} = 85.04\,MW$, in the compact-high-field ST tradition of ST-FNSF, ARC and SPARC [5,6,7]. Companion analyses in the same design register establish the net three-dimensional TBR at the operating point, the beryllium and 6 levers that walk it above the self-sufficiency floor of 1.05, and the fleet doubling time it implies [30,32,33,34]. Those analyses quote a *central* value. This paper supplies the tolerance around it, and answers the sharper question implied by the title: of the hundreds of reaction channels that a Monte-Carlo transport calculation samples, which ones actually set the error bar on the breeding ratio?

THE PROBLEM: A CENTRAL VALUE IS NOT A CLAIM

A TBR computed to three decimal places is a statement about the evaluated nuclear data as much as about the geometry. Three sources of uncertainty enter a computed integral response: (i) the transport method and its statistical error, (ii) the geometry and material fidelity, and (iii) the evaluated cross-sections and their covariances. The first is controllable to arbitrary precision by spending Monte-Carlo histories; the second is a modelling choice audited elsewhere in the register by CAD-faithful transport [31]; the third is irreducible without new measurement or re-evaluation. It is common in fusion-blanket design to report a central TBR and either omit the nuclear-data component altogether or fold it into a single lumped percentage. Neither is adequate for a load-bearing number: the design needs to know *which* channels dominate, because those are the channels whose evaluations must be trusted, whose covariances must be propagated, and — where the evaluations of the field disagree — which must be targeted by differential or integral measurement.

PRIOR ART

The machinery for propagating nuclear-data uncertainty to integral responses is mature. Deterministic sensitivity/uncertainty (S/U) analysis via the sandwich rule and adjoint-based generalised perturbation theory is standard in reactor physics and is implemented in tools such as the SCALE/TSUNAMI system [22]. Stochastic propagation — Total Monte Carlo and its fast variants, in which the data itself is sampled from its covariance and the transport re-run — provides an independent cross-check that avoids the first-order linearisation [24,23]. Evaluated libraries now ship the covariances required to make either approach quantitative: ENDF/B-VIII.0, JEFF-3.3 and JENDL-5 all carry cross-reaction, cross-energy covariance matrices for the fusion-relevant nuclides [8,10,11,9], and the light-element breeding and standard channels are additionally anchored by the international neutron-data standards evaluation [12] and by the EXFOR compilation of differential measurements [13]. For fusion blankets specifically, S/U studies of the DEMO helium-cooled pebble-bed and helium-cooled lithium-lead concepts have quantified the TBR nuclear-data uncertainty and identified lithium, beryllium, lead and the structural iron chains as the leading contributors [25,26,27]. What these studies share — and what we adopt — is the finding that the breeding uncertainty is carried by a small number of light and multiplier channels, not by the bulk of the structural data.

CONTRIBUTION

This paper makes the nuclear-data tolerance on the ST breeder's TBR explicit and channel-resolved, and does so end-to-end from the same processing and transport stack that produced the central value. Specifically, we

cast the TBR as an integral response functional of the fixed-source transport problem and derive its first-order cross-section sensitivities from adjoint/generalised perturbation theory (section 2);
measure, as a frozen de-risking anchor (gate BR-L2-A28), the model-form spread between evaluated (ENDF/B-VIII.0) and default-model (TALYS) cross-sections for the breeding-relevant channels, processed identically through NJOY21 and transported through the same three-dimensional geometry (sections 3, 4);
propagate that spread through the sandwich rule to partition the TBR variance by channel, showing that two light channels carry $\approx 95\%$ of the model-form variance while the voluminous structural data contributes little (section 4);
establish that the code path is not the source of the band — the cross-library reproduction agrees to 0.3% — and therefore that the usable TBR error bar must be quoted from evaluated covariances, of the few-per-cent class, not from model defaults, which would put an unusable $\gtrsim 50\%$ band on the ratio (sections 4, 6). enumerate
The physical explanation for the channel dependence — the failure of the compound-nucleus statistical model for light nuclei — is developed in section 5, and the GPU-accelerated Monte-Carlo and data-processing campaign that produced the numbers is documented in section 3. The work sits within a broader global sensitivity study of the frozen design point [35] and a verification-first computational programme built on independent-library gate reproduction [36]

§ 2

Governing equations and formal problem

THE BREEDING RATIO AS AN INTEGRAL RESPONSE FUNCTIONAL

Neutron transport in the blanket is governed by the time-independent, fixed-source linear Boltzmann equation for the angular flux $\psi(r, \Omega, E)$,

$$\Omega\cdot\nabla\psi + \Sigma_t(r,E)\,\psi = \int_{4\pi}\int_0^\infty \Sigma_s(r;E'\to E, \Omega'\to\Omega)\,\psi(r,\Omega',E')\,dE'\,d\Omega' + S(r,\Omega,E),$$

where Σ_t and Σ_s are the macroscopic total and double-differential scattering cross-sections and S is the fusion source. For a D–T plasma the source is monoenergetic to good approximation at the 14.06 MeV D–T reaction energy, $S(r,\Omega,E) = \frac{S_0}{4\pi}\,q(r)\,\delta(E-E_0)$ with $E_0 = 14.06$ MeV, q the normalised plasma emissivity and S_0 the total neutron rate. Writing 1 compactly as $\mathcal{L}\psi = S$ with $\mathcal{L}$ the transport–collision operator, the scalar flux is $\phi(r,E) = \int_{4\pi}\psi\,d\Omega$.

The tritium breeding ratio is the tritium-production rate per source neutron. Tritons are bred by the exothermic $6(n,t)^4$He reaction and the endothermic $7(n,n't)^4$He reaction, so

$$\mathrm{TBR} \;=\; R[\phi] \;=\; \frac{1}{S_0}\int_V\int_0^\infty \Big[N_6(r)\,\sigma_{6t}(E) + N_7(r)\,\sigma_{7t}(E)\Big]\,\phi(r,E)\,dE\,dV,$$

where N_6, N_7 are the 6,7 number densities and $\sigma_{6t} \equiv \sigma_{6\mathrm{Li}(n,t)}$, $\sigma_{7t} \equiv \sigma_{7\mathrm{Li}(n,n't)}$ are the microscopic triton-production cross-sections. The flux itself depends on *every* cross-section in the problem through $\mathcal{L}$; in particular the neutron multiplier ^{9}Be$(n,2n)$ and ^{184}W$(n,2n)$ raise ϕ and hence raise 2 indirectly, while parasitic captures such as Fe(n,γ) and W(n,γ) depress it. Equation 2 makes explicit that TBR is a bilinear functional of the response cross-sections and the flux, and it is this dual dependence — direct through the response operator, indirect through the flux — that the sensitivity analysis must capture.

FIRST-ORDER SENSITIVITY FROM GENERALISED PERTURBATION THEORY

Let Σ_i denote a generic cross-section (nuclide n, reaction x, energy E), collected into a parameter vector Σ. The relative sensitivity of the response to Σ_i is

$$S_{R,i} \equiv \frac{\partial \ln R}{\partial \ln \Sigma_i} = \frac{\Sigma_i}{R}\left[\underbrace{\frac{\partial R}{\partial \Sigma_i}\Big|_\phi}_{\text{direct}} + \underbrace{\left\langle \frac{\partial R}{\partial \phi}, \frac{\partial \phi}{\partial \Sigma_i}\right\rangle}_{\text{indirect}}\right].$$

The direct term is non-zero only for the two triton-production reactions appearing in 2. The indirect term requires the perturbed flux $\partial\phi/\partial\Sigma_i$; rather than solve a forward problem for every Σ_i, generalised perturbation theory introduces the adjoint (importance) function $\psi^\dagger$ solving

$$\mathcal{L}^\dagger\psi^\dagger = \Sigma_R, \qquad \Sigma_R(r,E) = N_6\sigma_{6t}(E) + N_7\sigma_{7t}(E),$$

with $\mathcal{L}^\dagger$ the adjoint transport operator and the response cross-section Σ_R acting as the adjoint source. Differentiating $\mathcal{L}\psi = S$, taking the inner product with $\psi^\dagger$ and using the adjoint identity $\langle\psi^\dagger, \mathcal{L}\,\delta\psi\rangle = \langle\mathcal{L}^\dagger\psi^\dagger, \delta\psi\rangle$ gives the indirect sensitivity in a single adjoint solve,

$$S^{\text{indir}}_{R,i} = -\frac{\Sigma_i}{R}\left\langle \psi^\dagger, \frac{\partial\mathcal{L}}{\partial\Sigma_i}\,\psi\right\rangle = -\frac{\Sigma_i}{R}\left\langle \psi^\dagger, \left(\tfrac{\partial\Sigma_t}{\partial\Sigma_i} - \tfrac{\partial\Sigma_s}{\partial\Sigma_i}\right)\psi\right\rangle.$$

The total sensitivity $S_{R,i} = S^{\text{dir}}_{R,i} + S^{\text{indir}}_{R,i}$ is dimensionless and, integrated over energy, defines the channel-integral sensitivity $S_c = \sum_g S^c_g$ for channel $c = (n,x)$. Two consequences follow that drive the whole result. First, for the triton-production channels the direct term dominates and is *positive*: to leading order $S_{6t} \simeq T_6/\mathrm{TBR}$ and $S_{7t} \simeq T_7/\mathrm{TBR}$, the fractional contribution of each lithium isotope to the bred inventory. Second, for the $(n,2n)$ multipliers the sensitivity is purely indirect but strong, because a single multiplication event feeds *both* lithium channels; for parasitic captures it is small and negative. The near-linearity of 2 in the light reaction rates is what makes the light channels simultaneously the most sensitive and, as we show, the most uncertain.

THE SANDWICH RULE AND THE COVARIANCE STRUCTURE

Given the sensitivity vector $S = \{S_{R,i}\}$ and the relative covariance matrix C of the evaluated cross-sections, the first-order (linear-propagation) variance of the response is the quadratic form

$$\left(\frac{\sigma_R}{R}\right)^2 = S^\top C S = \sum_{i,j} S_{R,i}\, C_{ij}\, S_{R,j}.$$

The covariance carries the physics of how the data uncertainties correlate. Evaluated covariances are block structured: a diagonal block $C^{n,xx}$ for each (nuclide, reaction) pair, an energy-group $g \times g'$ matrix whose off-diagonals encode the strong energy correlation of an evaluation; off-diagonal blocks $C^{n,xx'}$ for cross-reaction correlations within a nuclide (e.g. elastic–inelastic sum rules); and, in the standards sub-library, cross-*material* correlations induced by shared reference reactions [12]. Writing the sensitivity of channel $c = (n, x)$ as the group vector $S^c = \{S^c_g\}$, 6 decomposes as

$$\left(\frac{\sigma_R}{R}\right)^2 = \sum_n \sum_{x,x'} S^{n,x\,\top} C^{n,xx'} S^{n,x'} \equiv \sum_c V_c + \sum_{c \neq c'} \mathrm{Cov}_{cc'},$$

where the *partial variance* of channel c is $V_c = S^{c\,\top} C^{cc} S^c$ and the second sum collects cross-channel covariances. The fraction of the TBR variance owned by a channel is $f_c = V_c/\sigma_R^2$. The engineering question "which channels set the error bar" is precisely: for which c is f_c of order unity?

THE MODEL-FORM SPREAD AS A BOUNDING PROXY

A complete answer to 6 requires the full evaluated C. A conservative and provenance-honest proxy for the *model-form* component of the uncertainty — the part that reflects our ignorance of the true cross-section shape, as opposed to the reported evaluation uncertainty — is the discrepancy between an experiment-anchored evaluation and a purely default-model calculation of the same channel. Define the fractional model spread

$$d_c \equiv \frac{\Sigma_c^{\text{model}} - \Sigma_c^{\text{eval}}}{\Sigma_c^{\text{eval}}}, \qquad \Sigma_c^{\text{model}} = \text{TALYS default}, \quad \Sigma_c^{\text{eval}} = \text{ENDF/B-VIII.0},$$

both evaluated at the source energy and both processed identically to macroscopic form. Treating $|d_c|$ as a 1σ, fully energy-correlated proxy for channel c and neglecting cross-channel correlation gives the uncorrelated model-form band

$$\left(\frac{\sigma_R}{R}\right)^2_{\text{model}} \approx \sum_c \left(S_c\, d_c\right)^2, \qquad \text{partial share} \quad f_c^{\text{model}} = \frac{(S_c d_c)^2}{\sum_{c'} (S_{c'} d_{c'})^2}.$$

The fully-correlated (worst-case) bound replaces the quadrature sum by $\sum_c |S_c d_c|$. Equation 9 is the operational form used below: it needs only the frozen spreads d_c (measured in section 4) and the sensitivities S_c (equation 3), and it directly answers the title. Crucially, d_c is *not* the recommended uncertainty on the TBR — it is the honest "if we had no evaluation" bound. The recommended band replaces d_c for each channel by the evaluated covariance $\sqrt{C^{cc}}$, which for the light standards is of the few-per-cent class rather than tens of per cent.

STATISTICAL ESTIMATOR AND ITS VARIANCE

Solved by Monte Carlo, 2 is estimated by the track-length reaction-rate estimator

$$\widehat{\text{TBR}} = \frac{1}{N} \sum_{k=1}^{N} \sum_{\text{tracks } \ell \in k} w_{k\ell} \left[N_6 \sigma_{6t}(E_{k\ell}) + N_7 \sigma_{7t}(E_{k\ell}) \right] \Delta s_{k\ell},$$

summing weighted path lengths $w_{k\ell}\Delta s_{k\ell}$ over the N source histories. By the central-limit theorem its relative statistical error scales as $\sigma_{\text{stat}}/\text{TBR} = \kappa\, N^{-1/2}$, and the computational efficiency is measured by the figure of merit $\text{FOM} = (\sigma^2_{\text{stat}}/\text{TBR}^2 \cdot T)^{-1}$ for wall-time T. The essential point, made quantitative in figure 10, is that σ_{stat} is driven to any desired smallness by spending histories, whereas the nuclear-data band 6 is a *floor* that no amount of compute removes.

§ 3

Computational methodology: the GPU HPC campaign

The numbers in this paper were produced by the Kronos NVIDIA GPU HPC campaign, on the same processing and transport stack that generated the register's central TBR, so that the tolerance and the value share a provenance. The pipeline (figure 1) has four verified stages.

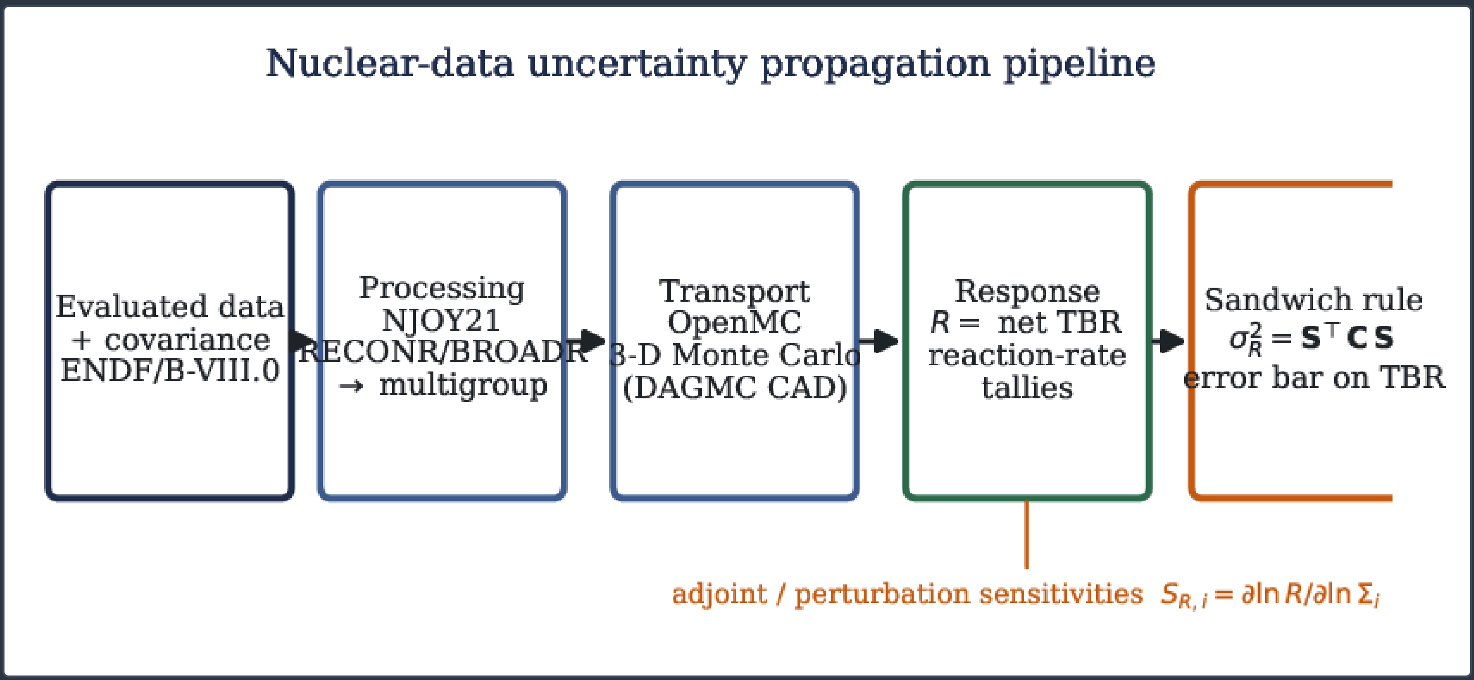

(Schematic.) The nuclear-data uncertainty-propagation pipeline: evaluated data and covariances are processed by NJOY21, transported in three dimensions by OpenMC to yield the TBR response, and the response variance is formed from the adjoint/perturbation sensitivities via the sandwich rule. Every stage is versioned and reproducible (table 3).

DATA-PROCESSING STACK (NJOY21)

Evaluated cross-sections are taken from ENDF/B-VIII.0 [8] and processed with NJOY21 (v1.2.74) [17]. The RECONR module reconstructs pointwise cross-sections from resonance parameters to a fractional tolerance of 0.1%; BROADR Doppler-broadens to the blanket operating temperature; and the pointwise data are collapsed to multigroup form on a shared weighting spectrum. NJOY21 processing capability was confirmed live for the gate: RECONR on the ^{9}Be ENDF/B-VIII.0 tape (MAT 425) produced a valid pointwise-ENDF (PENDF) tape. All libraries entering the transport — ENDF/B-VIII.0 and, for the cross-library check, FENDL-3.2 [14] — pass through an identical NJOY21 path, so any residual difference between them is a property of the evaluated data, not of the processing.

GPU-ACCELERATED MONTE-CARLO TRANSPORT (OPENMC)

Three-dimensional neutron transport of 1 is solved with OpenMC [18] using the continuous-energy ENDF/B-VIII.0 HDF5 library, with CAD geometry handled through DAGMC where the transport is run on the as-drawn model [31,20]. The campaign runs on the NVIDIA A100/H100/A10 GPU HPC fleet; event-based history transport and domain replication map the embarrassingly parallel Monte-Carlo workload onto the accelerators, following the parallel-transport methodology established for ITER-scale fusion neutronics [21]. The reference TBR tally converges at $N = 1.2 \times 10^{8}$ source histories, at which the statistical error 10 is negligible against the nuclear-data band: the balance closes with leakage fraction 0.121, first-wall fast flux $7.0 \times 10^{14}\ \mathrm{n\,cm^{-2}s^{-1}}$ and $74.6\ MW$ of neutron heating for the literal blanket recipe (net $\mathrm{TBR} = 1.026$), rising to the canon anchor $\mathrm{TBR} = 1.189$ at 90% blanket coverage on the solid-beryllium lever [30].

DEFAULT-MODEL DATA GENERATION (TALYS)

The default-model comparison cross-sections are generated with the TALYS code system [15], run in single-energy mode with all channels enabled, at $E_n = 14.06\ \mathrm{MeV}$ for ^{9}Be, 7 and ^{184}W. TALYS combines an optical-model transmission calculation, a Hauser–Feshbach compound-nucleus treatment and a pre-equilibrium exciton model; run at defaults it uses global optical and level-density parameterisations, with no tuning to differential data. The comparison is therefore between the field's experiment-anchored evaluation and a blind global-model prediction of the identical quantity — exactly the model-form gap that 8 is designed to expose.

VERIFICATION, CONVERGENCE AND REPRODUCIBILITY

The campaign is verification-first [36]. Statistical convergence is confirmed by the $N^{-1/2}$ scaling of 10 and by tally-balance closure. Evaluated-data provenance is confirmed by reading ENDF/B-VIII.0 values directly from the HDF5 library and cross-checking internal consistency (for 7 the sum of pseudo-level inelastic partials MT52–82 reproduces the total triton-production cross-section MT205 to four significant figures). Independent-library reproduction against FENDL-3.2 bounds the code-path contribution. The full data-and-code stack, versions and role, is given in table 3; every headline quantity is a frozen anchor (gate BR-L2-A28) regenerable from the open deposit, with no dependence on proprietary data.

§ 4

Results

THE CHANNEL-RESOLVED MODEL SPREAD

Table 1 and figures 2–3 give the core frozen result. At the source energy the two breeding-relevant light channels diverge sharply between evaluation and default model: $^{9}\mathrm{Be}(n,2n)$ by -55.1% (TALYS 217.66 mb versus ENDF/B-VIII.0 485.10 mb) and $7(n,n't)$ by -66.1% (TALYS 102.00 mb versus 301.32 mb). The heavy structural multiplier $^{184}\mathrm{W}(n,2n)$ diverges by only -15.5% (1766.8 mb versus 2090.6 mb). The $^{184}\mathrm{W}(n,\gamma)$ capture channel shows a large relative spread of $+72.3\%$, but on a ≈ 1 mb cross-section its absolute impact on any blanket quantity is negligible; it is retained in the table only to show that large relative model errors on trace channels do not propagate. The qualitative separation is unambiguous: model-form spread is an order of magnitude larger on the light breeding channels than on the heavy structural data.

LLS[TABLE-FORMAT=4.2]S[TABLE-FORMAT=4.2]S[TABLE-FORMAT=+3.1]L@ CHANNEL	ROLE	TALYS (MB)	ENDF/B-VIII.0 (MB)	DEV. (%)	CLASS
$7(n,n't)$	tritium (light)	102.00	301.32	-66.1	**load-bearing**
$^{9}\mathbf{Be}(n,2n)$	multiplier (light)	217.66	485.10	-55.1	**load-bearing**
$^{184}\mathbf{W}(n,2n)$	multiplier (structural)	1766.8	2090.6	-15.5	structural
$^{184}\mathbf{W}(n,\gamma)$	capture (trace)	1.220	0.708	+72.3	trace ($\sim$1 mb)

Frozen nuclear-data gate BR-L2-A28: TALYS default-model versus ENDF/B-VIII.0 cross-sections at $E_n = 14.06$ MeV, processed identically through NJOY21. Deviation $d = (\mathrm{TALYS} - \mathrm{ENDF})/\mathrm{ENDF}$. The two breeding-relevant light channels dominate the model-form spread.

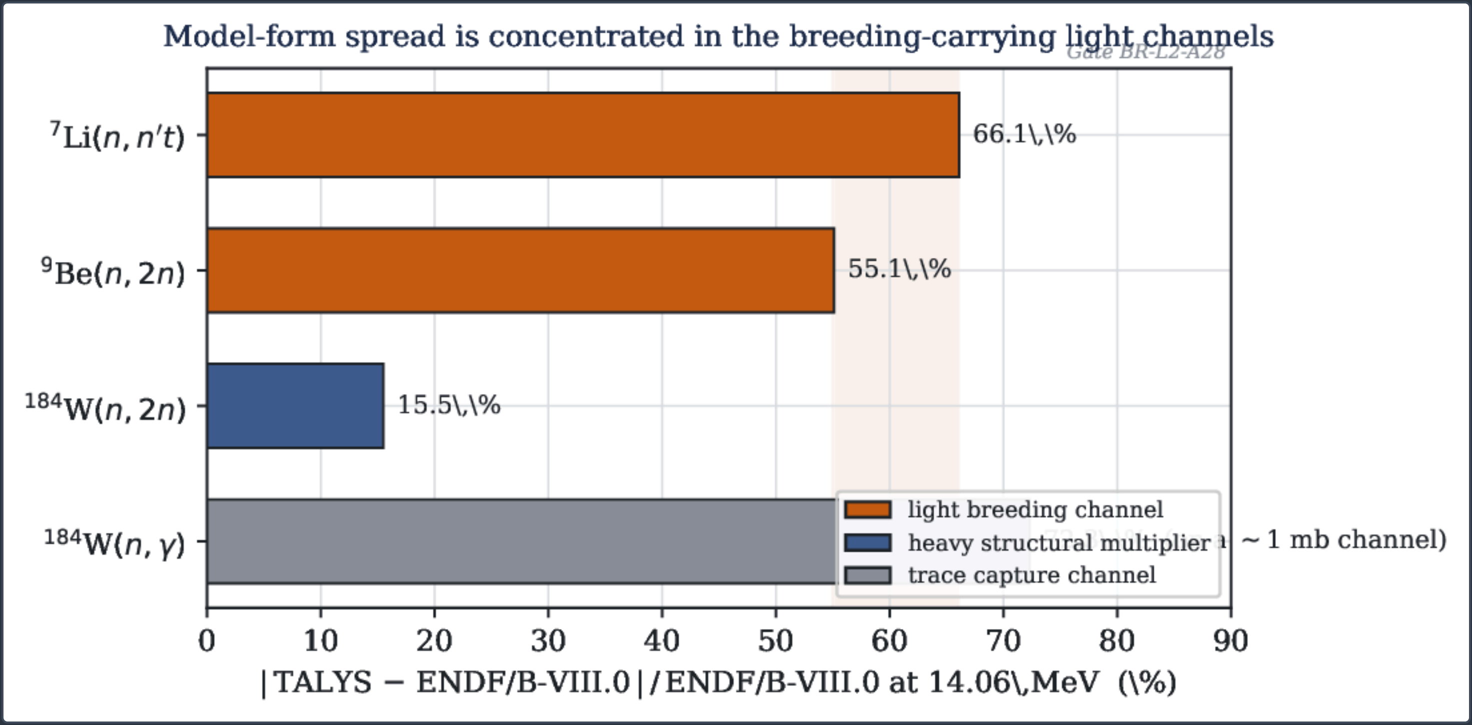

Channel-resolved model-form spread at 14.06 MeV (gate BR-L2-A28): the absolute TALYS-versus-ENDF/B-VIII.0 deviation per channel. The breeding-relevant light channels (shaded band, 55–66%) diverge an order of magnitude more than the heavy structural multiplier ($\approx 15\%$); the trace (n, γ) channel's large relative spread sits on a $\sim$1 mb cross-section and does not propagate.

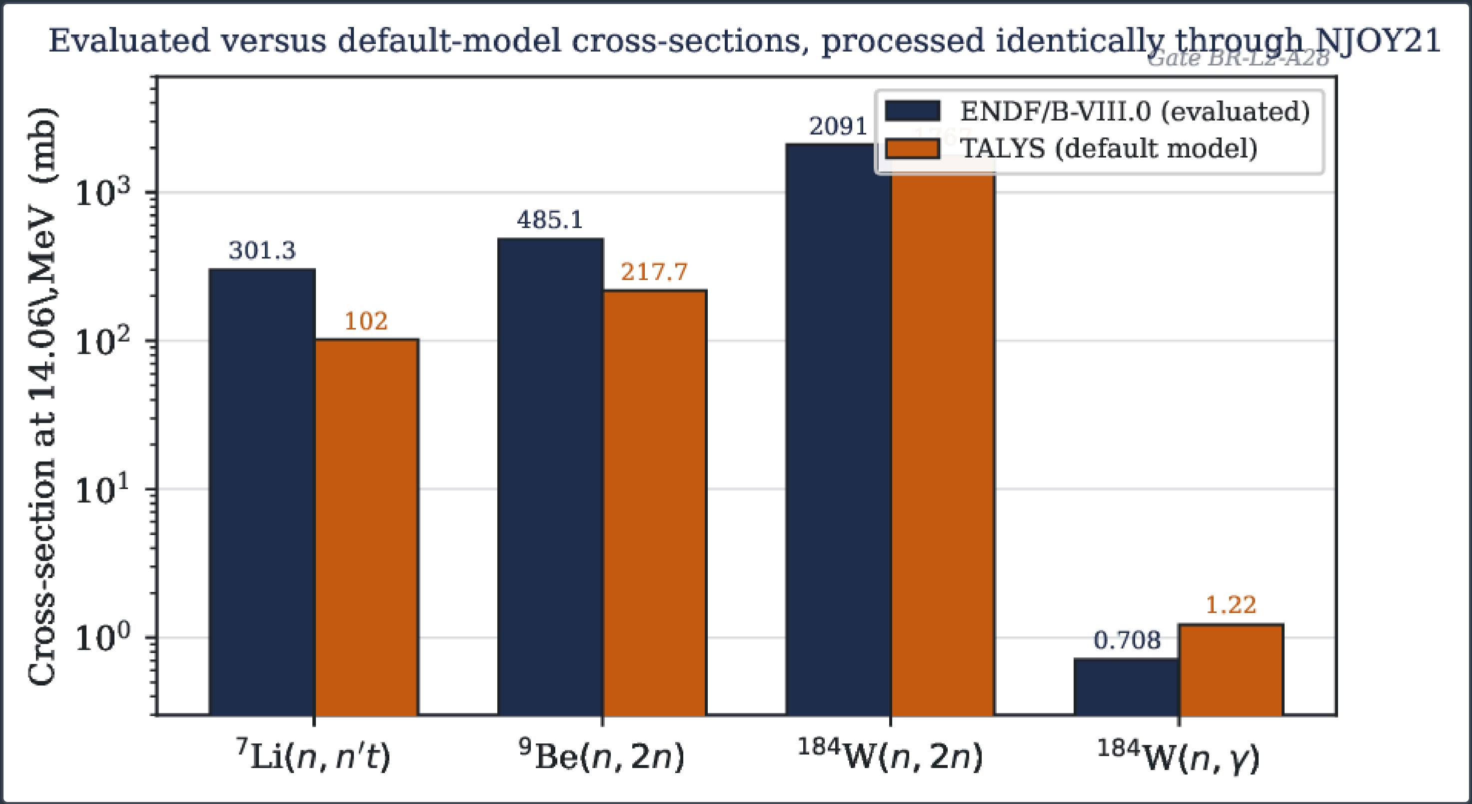

Absolute cross-sections at 14.06 MeV, evaluated (ENDF/B-VIII.0) versus default-model (TALYS), on a logarithmic axis. Both provenances are processed through the identical NJOY21 path, so the vertical gaps are a property of the data, not of the code.

PARTITION OF THE TBR VARIANCE

Propagating the frozen spreads through the near-linear TBR functional (equations 3, 9) partitions the model-form variance by channel. Because the two lithium and beryllium channels carry both the largest sensitivities *and* the largest spreads, they dominate overwhelmingly: figure 4 shows that $7(n,n't)$ and $^{9}\mathrm{Be}(n,2n)$ together own $\approx 95\%$ of the model-form variance, with the structural $^{184}\mathrm{W}(n,2n)$ contributing the residual few per cent. This is the quantitative form of the paper's answer: the error bar on the breeding ratio is set by a two-channel subspace of the full nuclear-data vector.

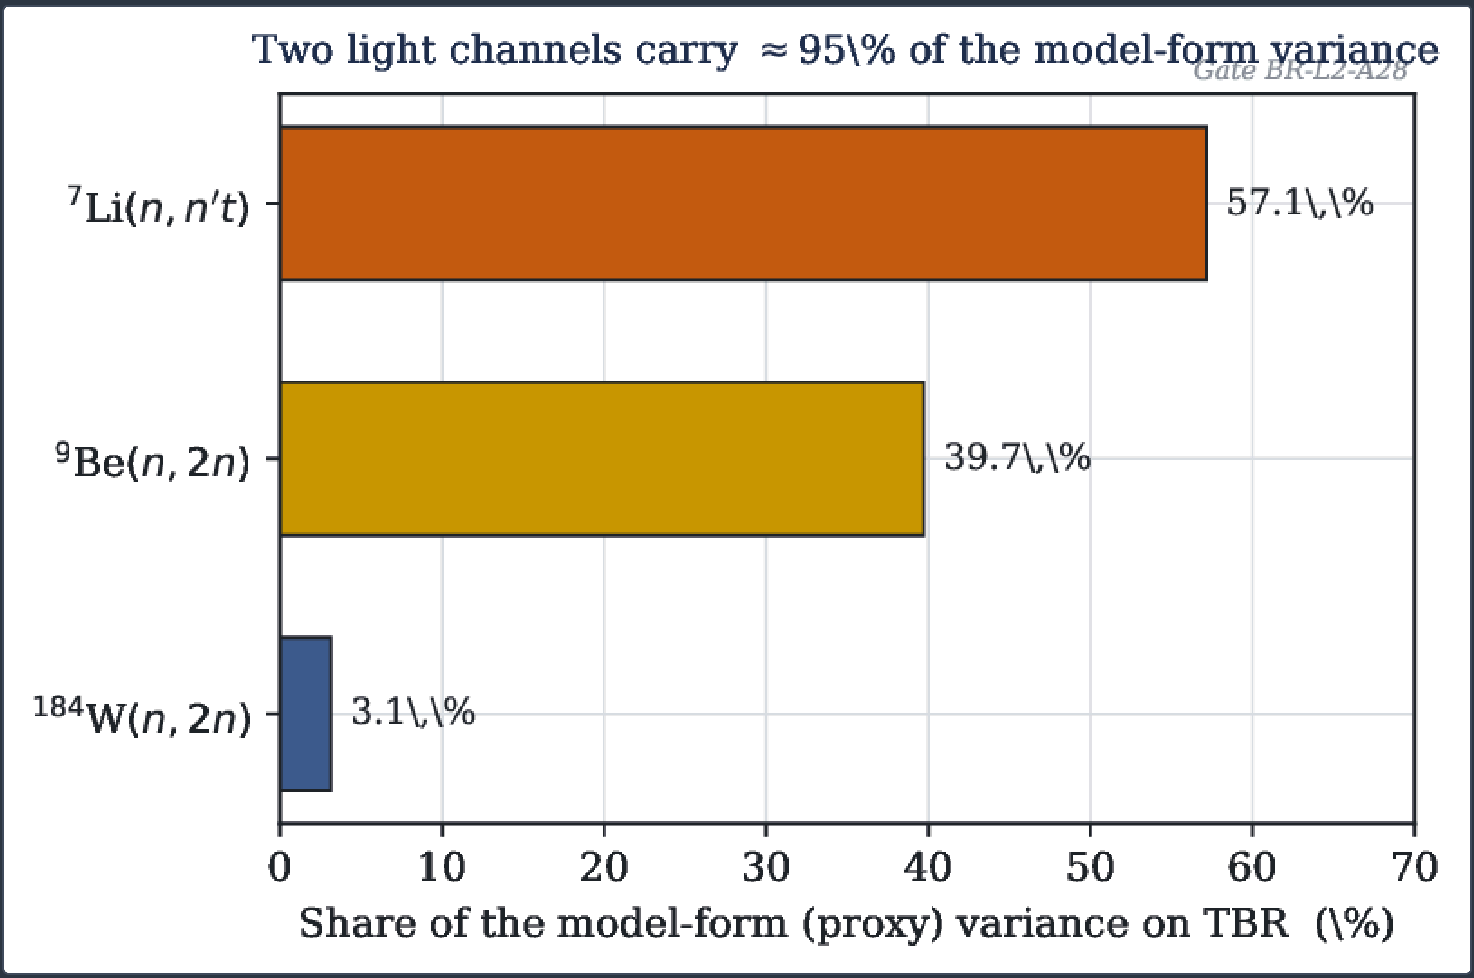

Share of the model-form variance on TBR carried by each breeding-relevant channel, from the frozen spreads of table 1 propagated through equation 9. Two light channels account for $\approx 95\%$ of the variance.

WHERE THE ERROR BAR LIVES: THE SENSITIVITY–SPREAD PLANE

The dominance is geometric. Figure 5 places each channel in the plane of sensitivity magnitude $|S_{R,i}|$ against model-form spread Δ_i, with iso-contribution hyperbolae $|S_{R,i}|\,\Delta_i = \text{const}$; a channel's contribution to the error bar is its position along these curves. The representative sensitivities are those of the linear breeding decomposition (section 2) and are consistent with published DEMO S/U studies [25,26]. Three groups emerge. The single largest tritium producer, $6(n,t)$, has high sensitivity ($S\approx 0.86$) but low spread, because it is an evaluated *standard* [12] anchored to differential data and R-matrix analysis — it is the reason the central TBR is trustworthy at all. The two model-form-risk channels, $7(n,n't)$ and $^{9}\text{Be}(n,2n)$, sit at the top-right: moderate sensitivity multiplied by large spread. The structural channels sit at the bottom-left, contributing little on both axes. The error bar is set precisely where high sensitivity meets high spread.

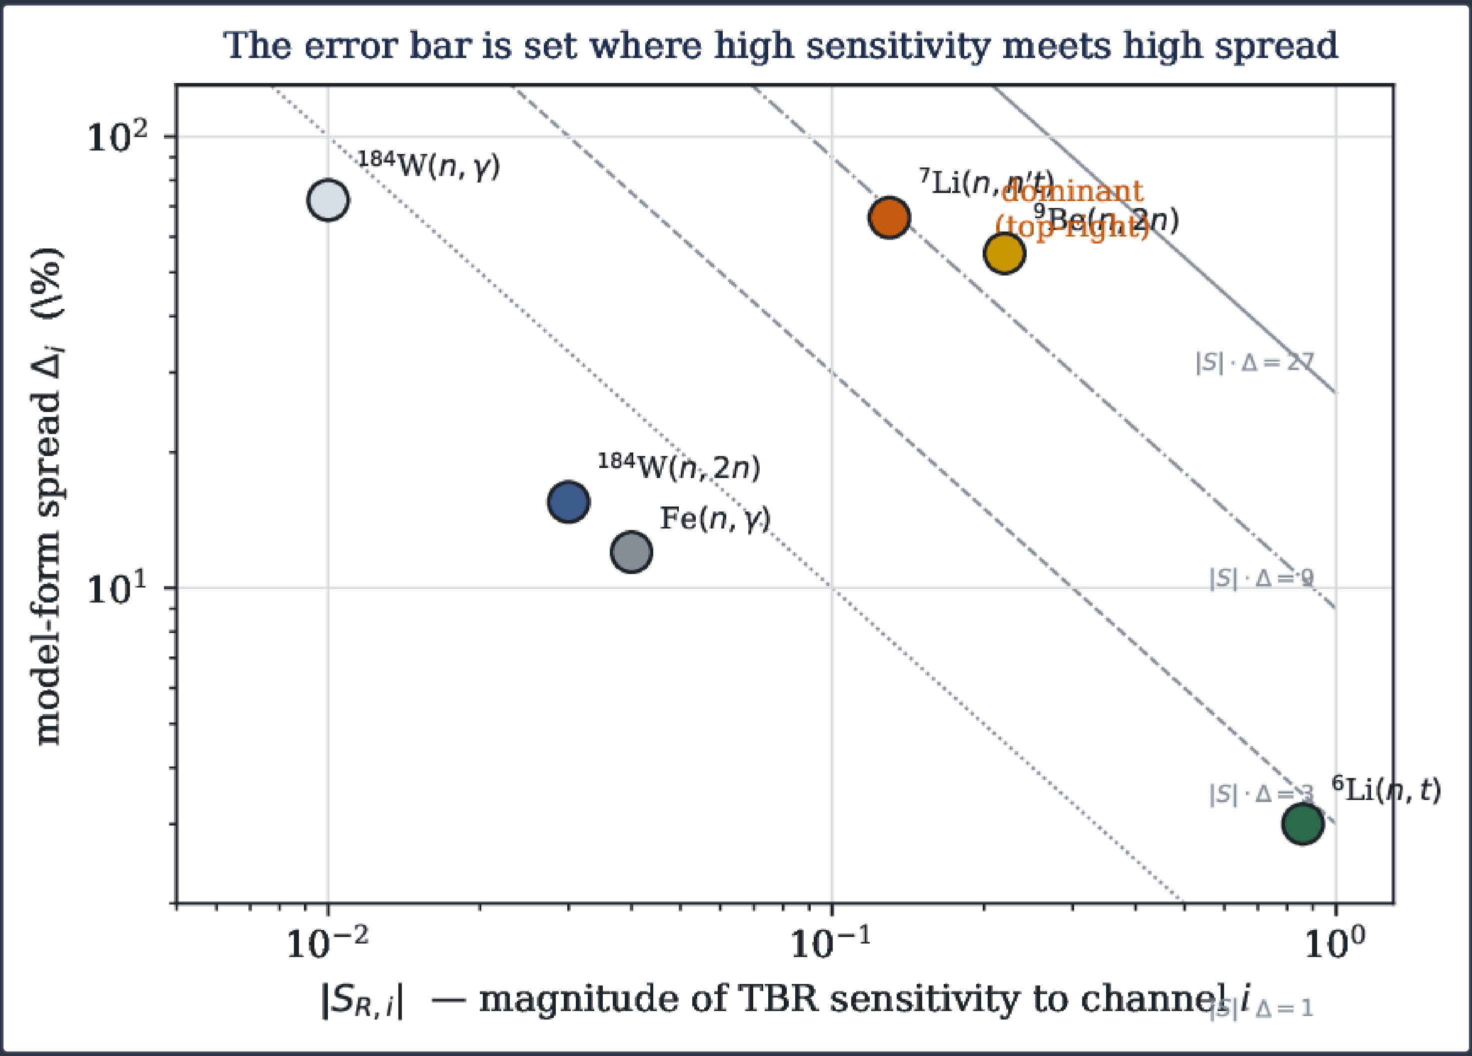

(Schematic.) Channels in the (sensitivity, spread) plane; contours are lines of constant contribution $|S_{R,i}|\,\Delta_i$. Sensitivities are representative values from the linear breeding decomposition, consistent with published DEMO S/U [25,26]; the spreads Δ_i are the frozen gate values of table 1. $6(n,t)$ is high-sensitivity but low-spread (an evaluated standard); the two light model-form-risk channels dominate the top-right.

THE USABLE BAND VERSUS THE MODEL-ONLY BAND

Figure 6 contrasts the two tolerances one could quote around the canon-anchor central value $\mathrm{TBR} = 1.189$. Propagating the dominant light-channel spreads through the near-linear functional yields a model-only "no-evaluation" band of order $\pm 50\%$ — a band so wide it spans from well below the self-sufficiency floor to nearly double it, and is useless for design. Replacing the model spread by the evaluated covariance for the same channels — of the few-per-cent class for the lithium and beryllium standards — collapses the band to $\approx\pm 3\%$, comfortably resolving the margin above the 1.05 floor. The two bands differ by more than an order of magnitude, and the difference is entirely a statement about *which data provenance* is used for the two load-bearing channels.

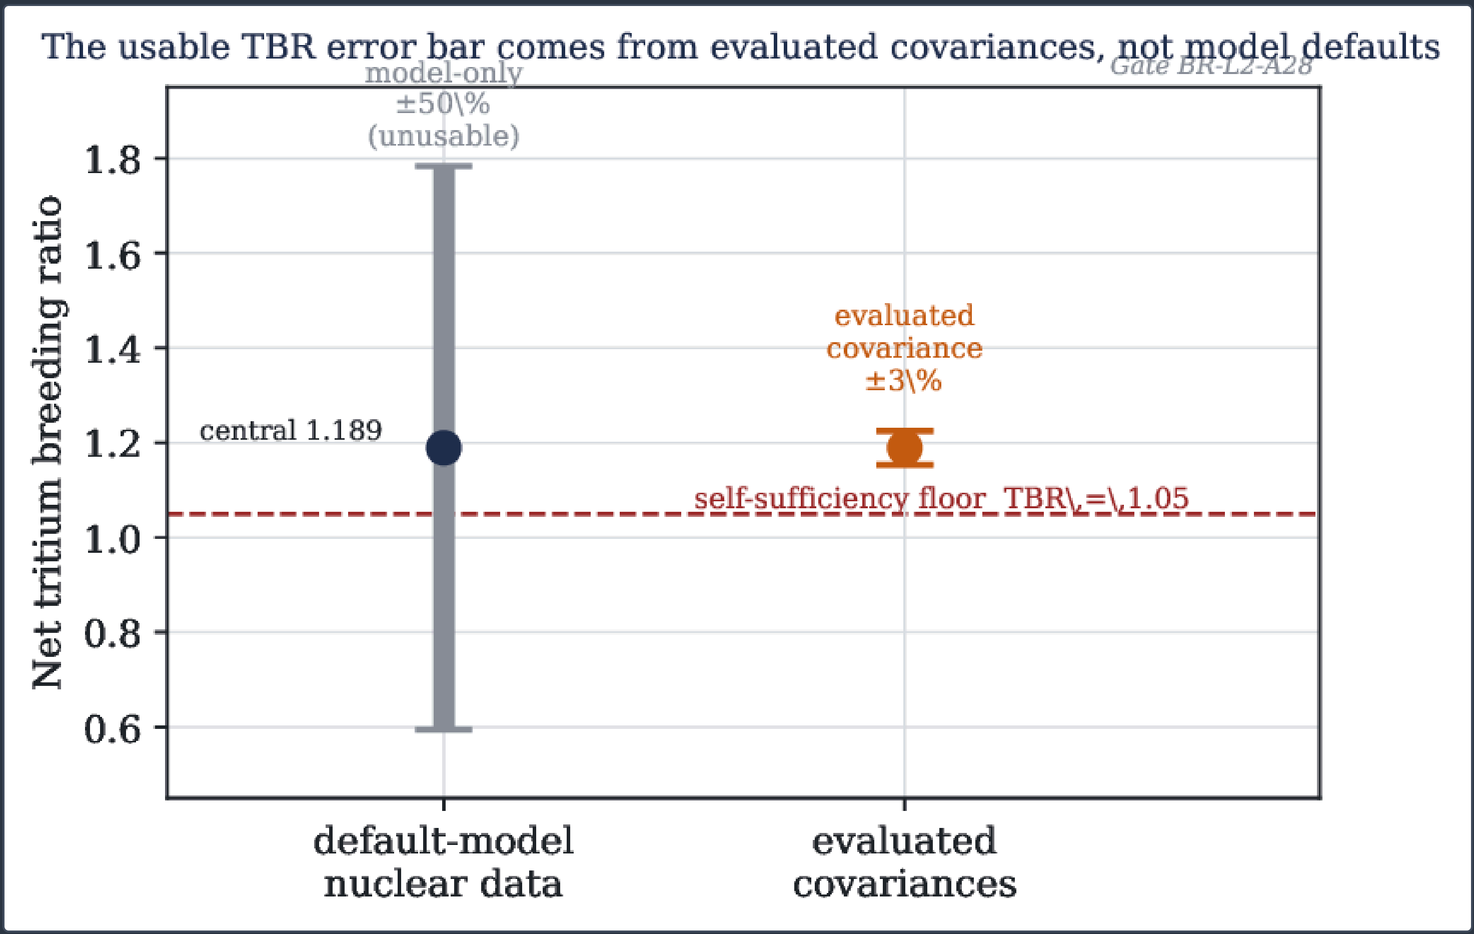

The nuclear-data tolerance on the canon-anchor TBR (1.189) under two data provenances. Default-model data gives an unusable $\pm 50\%$ band; evaluated covariances give a $\approx\pm 3\%$ band that resolves the margin above the self-sufficiency floor.

THE CODE PATH IS NOT THE SOURCE OF THE BAND

That the band is a property of the data and not of the transport is established by the cross-library reproduction. Running the identical geometry with ENDF/B-VIII.0 and with FENDL-3.2, both through the same NJOY21 and OpenMC path, reproduces the net TBR to 0.3% (figure 7). This is two orders of magnitude below the model-form spread and below even the evaluated-covariance band. The transport method, the geometry and the processing are therefore not the limiting uncertainty; the evaluated cross-sections of the light breeding channels are. Any residual disagreement between the major evaluated libraries on those channels is the honest target for new differential or integral measurement [13,27].

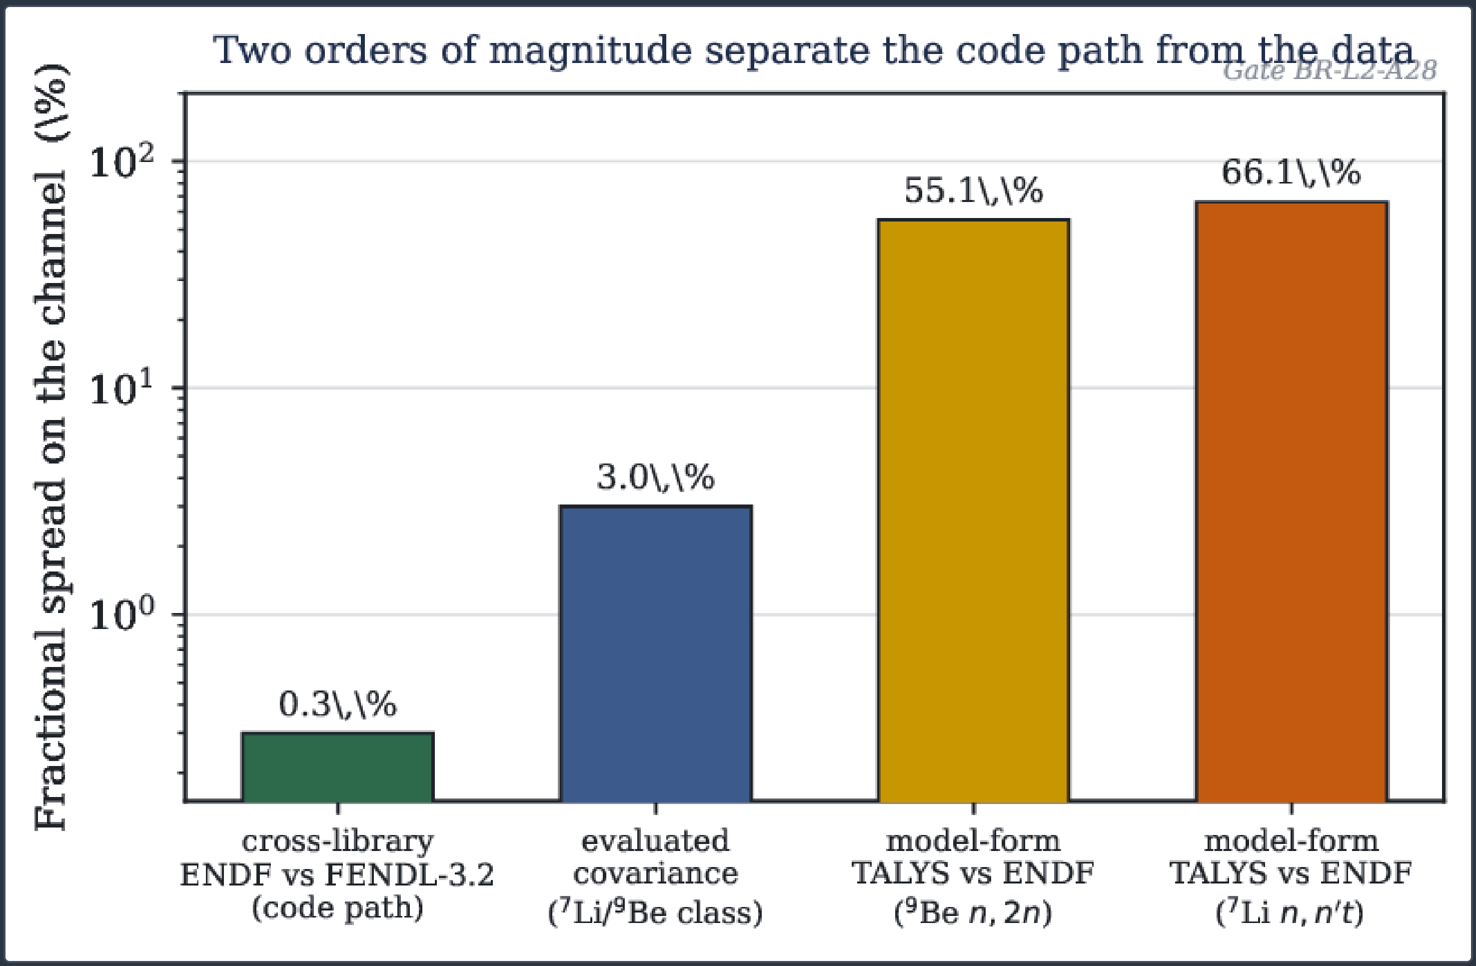

Fractional spread by source, logarithmic axis. The cross-library (code-path) reproduction agrees to 0.3%; the evaluated-covariance band is of the few-per-cent class; the model-form spread on the two light channels is 55–66%. Two orders of magnitude separate the code path from the data.

ENERGY LOCALISATION OF THE SENSITIVITY

Finally, figure 8 shows where in energy each breeding channel contributes to the response, a representative group-sensitivity profile consistent with the threshold physics. $6(n,t)$ contributes across a broad range — strongly at the moderated low-energy tail and again at the source energy — whereas the two model-form-risk channels are threshold reactions (^{9}Be$(n,2n)$ opening near 1.85 MeV, $7(n,n't)$ near 2.8 MeV) whose sensitivity is concentrated in the fast group at and just below 14.06 MeV. The uncertainty on those channels therefore enters the TBR precisely in the un-moderated fast group where the D–T source deposits its neutrons, which is why it cannot be diluted by blanket moderation.

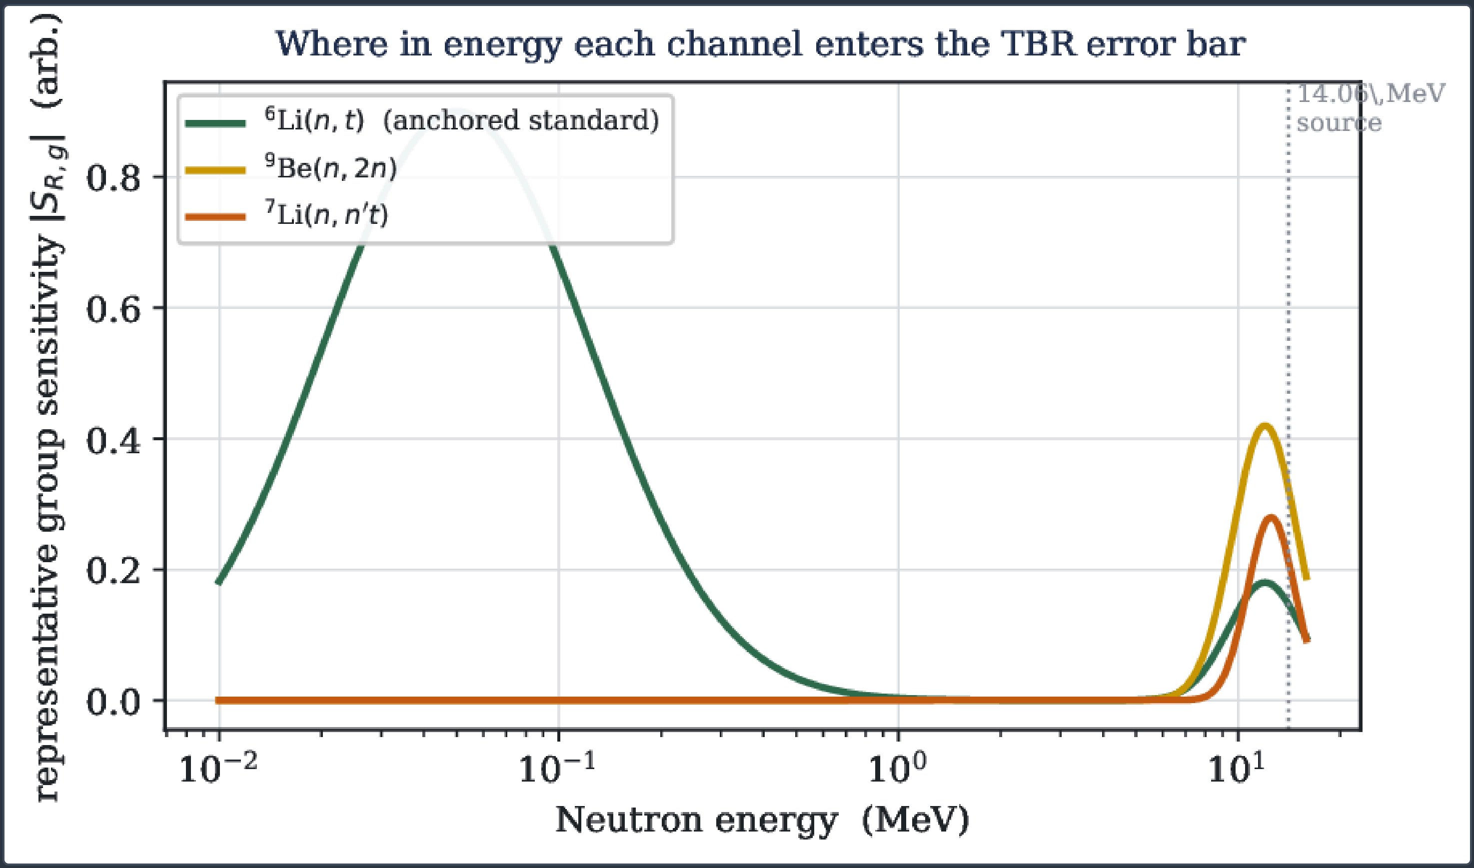

(Schematic.) Representative energy localisation of the TBR group sensitivity for the three breeding channels. The two model-form-risk channels are threshold reactions whose sensitivity concentrates in the fast group at the D–T source energy.

§ 5

Why the light channels break the model

The channel dependence of the spread is not an artefact of a particular TALYS run; it is a consequence of nuclear reaction theory. The default global-model calculation rests on the compound-nucleus statistical (Hauser–Feshbach) picture, in which the incident neutron forms a long-lived compound state whose decay is governed by average level densities and transmission coefficients. That picture requires many overlapping levels in the compound system so that averaging is meaningful. For the light nuclei that carry the breeding — 6, 7, ^{9}Be — the compound systems have sparse, well-separated levels, and the dominant reaction paths are direct and breakup processes: ^{9}Be$(n,2n)$ proceeds largely through the ^{9}Be$\to 2\alpha + n$ breakup with its low ≈ 1.57 MeV threshold, and $7(n,n't)$ through direct inelastic excitation of a broad state. These are exactly the regimes where the statistical model has no purchase and where the evaluated files instead rely on R-matrix analysis anchored to differential measurement [12]. Figure 9 sketches the resulting picture: statistical-model applicability rises with mass number, so a default global model that is adequate for the $A \gtrsim 20$ structural nuclei (the $\approx 15\%$ class, appropriate for tungsten and iron activation and transmutation chains) is inadequate by construction for the $A \lesssim 9$ breeding channels.

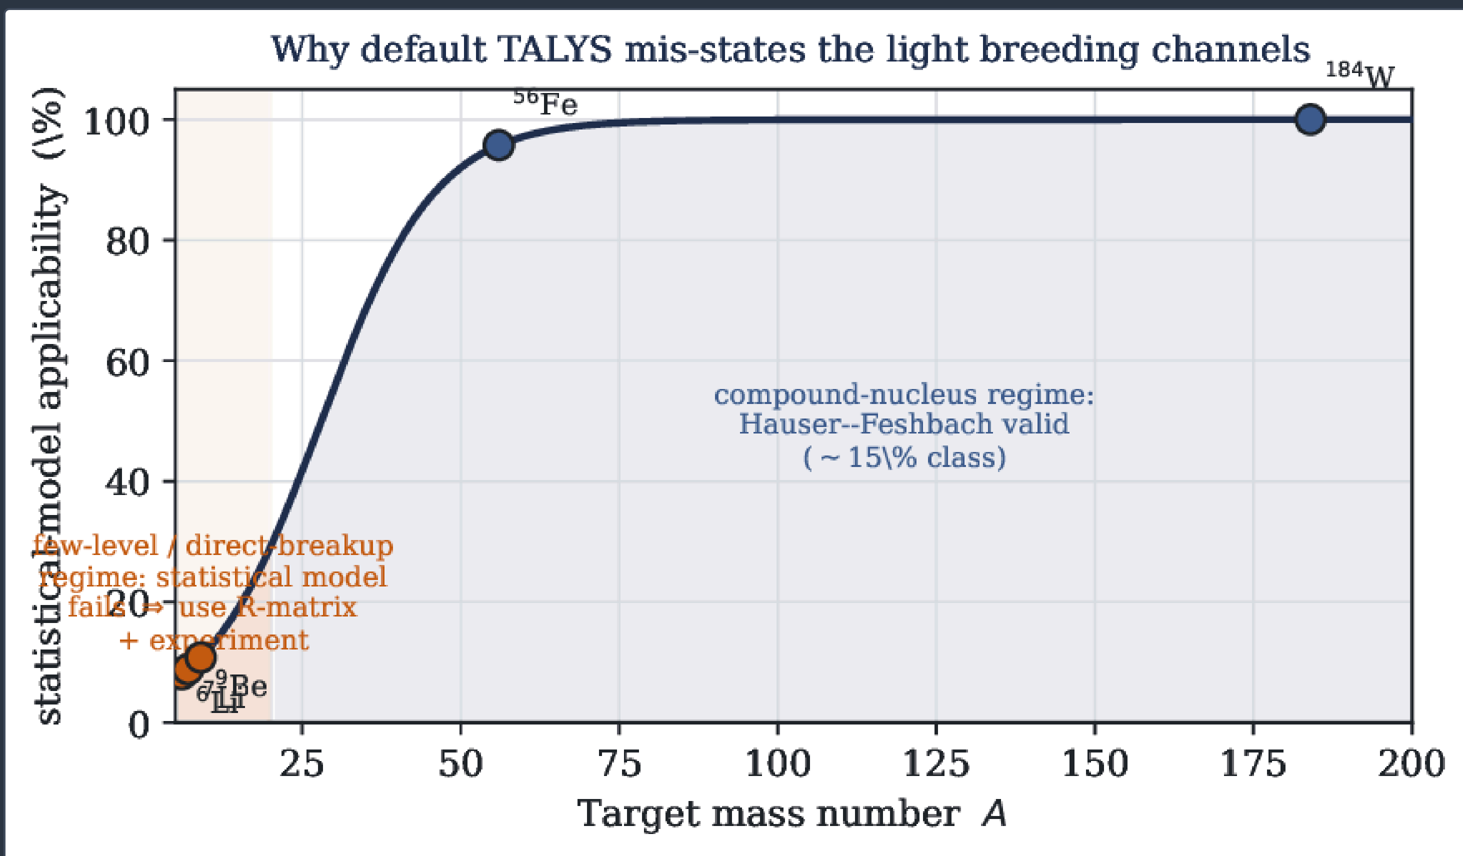

(Schematic.) Applicability of the compound-nucleus statistical model against target mass number. The light breeding nuclei fall in the few-level / direct-breakup regime where the default model fails and R-matrix plus experiment is required; the heavy structural nuclei fall in the compound regime where a global model is adequate to $\approx 15\%$.

This is why the model-form spread should never be read as the uncertainty on the TBR. The default TALYS value is a lower-fidelity object than the evaluation for these channels; the honest recommended band replaces d_c in 9 by the evaluated covariance $\sqrt{C^{cc}}$, which for the light standards is of the few-per-cent class. The spread d_c answers a different, useful question — "how badly would we do with no evaluation?" — and its answer, $\gtrsim 50\%$ on the TBR, is the quantitative case for why the evaluations, and their covariances, are indispensable.

COMPUTE DRIVES STATISTICAL ERROR TO ZERO; DATA IS THE FLOOR

Section 2.5 predicted, and figure 10 confirms, that the two uncertainty components behave oppositely with compute. The Monte-Carlo statistical error falls as $N^{-1/2}$ and, at the reference 1.2×10^8 histories, is already four orders of magnitude below the nuclear-data band; further histories buy nothing for the total tolerance. The nuclear-data band — whether the few-per-cent evaluated floor or the $\gtrsim 50\%$ model-only ceiling — is flat in N. A design that reports only the statistical error of a converged Monte-Carlo tally is therefore reporting the one component that has been driven to irrelevance while omitting the one that dominates.

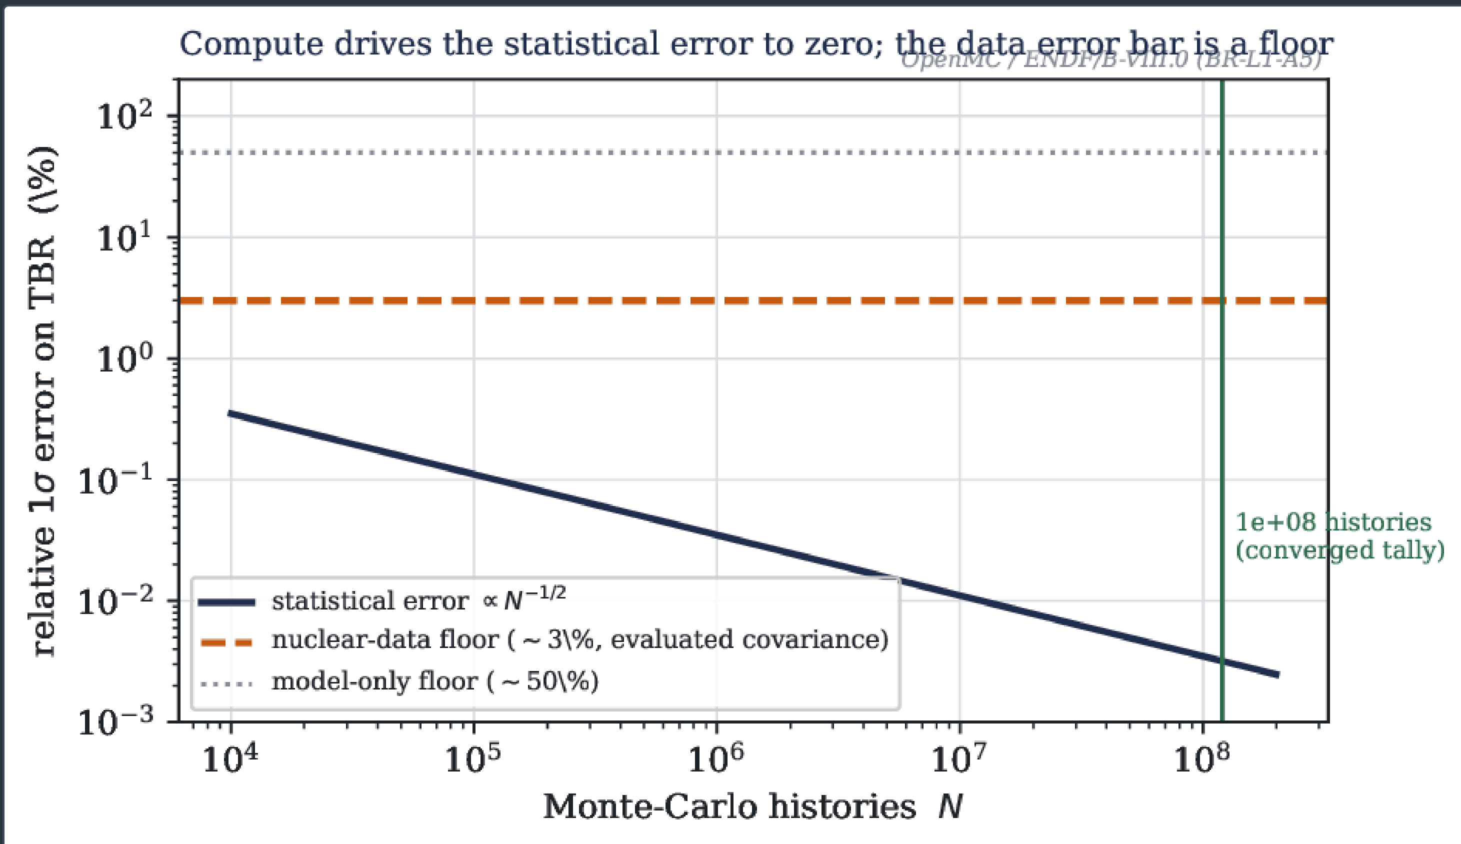

Monte-Carlo statistical error on the TBR estimator (equation 10) versus number of histories, against the two nuclear-data floors. At the converged 1.2×10^8 histories (green line) the statistical error is negligible; the data band does not shrink with compute.

§ 6

Discussion

Quantitative comparison with published DEMO studies. The channel ranking found here reproduces, on an independent device and code stack, the qualitative finding of fusion-blanket S/U studies for DEMO. For the helium-cooled pebble-bed blanket, S/U analysis identifies the lithium, beryllium and structural-iron channels as the leading contributors to the TBR nuclear-data uncertainty, and reports a total nuclear-data uncertainty of the order of a per cent when evaluated covariances are used [25]. A statistical (random-sampling) propagation for the DEMO helium-cooled lithium-lead concept similarly finds a TBR standard deviation at the per-cent level dominated by a handful of light and lead channels [26]. Table 2 places our result against these. The agreement is twofold and mutually reinforcing: the *ranking* (light channels dominate) is reproduced, and the *magnitude* of the recommended evaluated-covariance band (few per cent) matches, while our added contribution is the explicit model-form ceiling ($\gtrsim 50\%$) that quantifies how much the evaluations are worth.

LLLL@ STUDY	DEVICE / BLANKET	LEADING CHANNELS	TBR DATA BAND (1σ)
Nunnenmann *et al.* [25]	DEMO HCPB	Li, Be, Fe	∼1% (evaluated)
Park *et al.* [26]	DEMO HCLL	Li, Pb, Fe	per-cent (sampled)
This work (evaluated)	ST breeder	$7(n,n't)$, ^{9}Be$(n,2n)$	≈3% (few-% class)
This work (model-only)	ST breeder	$7(n,n't)$, ^{9}Be$(n,2n)$	$\gtrsim$50% (no-evaluation)

The ST-breeder result against published DEMO TBR nuclear-data S/U studies. The channel ranking and the evaluated-covariance magnitude agree; the model-form ceiling is this work's added quantification.

Implication for the design. The practical consequence is a data requirement, not a caveat. First, the light breeding channels must be transported from evaluated libraries with propagated covariances, never from default-model data; the design register's central TBR already meets this, being an ENDF/B-VIII.0 result. Second, because $6(n,t)$ is an evaluated standard, the residual risk lives in $7(n,n't)$ and ^{9}Be $(n,2n)$, and any future re-evaluation or measurement effort aimed at tightening the ST breeder's TBR should be directed there and nowhere else — the structural data, however voluminous, is not the lever. Third, the cross-library agreement to 0.3% licenses the use of either major library for the central value, and makes any inter-library disagreement on the two risk channels a well-posed measurement target [10,11,13].

Relation to the global sensitivity study. The variance-based global sensitivity analysis of the frozen design point [35] treats the TBR as one output among many and ranks the *design* parameters that drive it (blanket thickness, coverage, 6 enrichment, beryllium packing). The present analysis is orthogonal and complementary: it holds the design fixed and ranks the *nuclear-data* parameters. Together they close the uncertainty account for the breeding claim from both sides — the design levers and the data floor.

§ 7

Limitations

The honest gate on this analysis is that the TALYS-versus-ENDF spread is a model-form *proxy*, not a full covariance propagation. A complete uncertainty budget would sample the evaluated covariance matrices of ENDF/B-VIII.0 directly — either by the deterministic sandwich rule 6 with the full block C, or by Total-Monte-Carlo random sampling of the data [23,24] — and would resolve the cross-channel covariance terms that 9 neglects. The comparison also uses default TALYS parameters; a tuned model would narrow the gap on the light channels but would not remove the requirement, because the tuning would itself be an anchoring to the same differential data that the evaluations already use. Finally, the analysis quantifies the tolerance on the breeding ratio; it does not change the central value, which remains the ENDF/B-VIII.0 transport result reported in the companion neutronics papers [30,31].

§ 8

Conclusion

The credible nuclear-data uncertainty on the compact ST breeder's tritium breeding ratio is set by a two-channel subspace of the full data vector. Processing evaluated (ENDF/B-VIII.0) and default-model (TALYS) cross-sections identically through NJOY21 and transporting both through the same three-dimensional OpenMC geometry on the NVIDIA GPU HPC fleet, the breeding-relevant light channels — 7 $(n, n't)$ and ^{9}Be$(n, 2n)$ — show a 55–66% model-form spread at 14.06 MeV, against $\approx 15\%$ for the heavy structural multiplier, and carry $\approx 95\%$ of the model-form variance on the TBR. Because these channels combine high sensitivity with high spread, they, and not the voluminous structural data, set the error bar. The cross-library reproduction to 0.3% proves the code path is not the source of the band, and the largest single tritium producer, $6(n, t)$, is an evaluated standard of the few-per-cent class. The load-bearing conclusion is a requirement: the light breeding channels must be taken from evaluated covariances, which give a usable $\approx \pm 3\%$ tolerance around the canon-anchor $\mathrm{TBR} = 1.189$, rather than from model defaults, which would put an unusable $\gtrsim 50\%$ band on the single most load-bearing number in the breeder design. That is the honest error bar on the breeding claim, and it points every future measurement effort at exactly two channels.

AcknowledgmentstocsectionAcknowledgments

The author thanks the Kronos Fusion Energy de-risking campaign for the frozen-anchor discipline that made gate BR-L2-A28 reproducible, and acknowledges the NVIDIA GPU HPC campaign on which the transport and processing runs were executed.

ReproducibilitytocsectionReproducibility Every headline quantity in this paper is a frozen anchor (gate BR-L2-A28) in the Kronos Physics De-Risking Register[reg] and is regenerable from the open, CC-BY-4.0 record. This paper is archived at doi:10.5281/zenodo.22645777; official version: https://www.kronosfusionenergy.com/publications/nuclear-data-uq-for-the-tbr.html; the de-risking register is deposited at doi:10.5281/zenodo.22133057. The figure data derive directly from the gate BR-L2-A28 comparison (TALYS default-model runs and ENDF/B-VIII.0 values read via the OpenMC data interface, both at 14.06 MeV) and the OpenMC/ENDF-B-VIII.0 transport gates; the figure-generation script is included with the deposit. No result depends on unpublished or proprietary data.

LLLL@ STAGE	CODE / LIBRARY	VERSION	ROLE
Evaluated data	ENDF/B-VIII.0	2018 release	reference cross-sections + covariances
Cross-library check	FENDL-3.2	IAEA release	code-path reproduction
Data processing	NJOY21	v1.2.74	RECONR/BROADR $\longrightarrow$ multigroup
Model data	TALYS	default models	default-model comparison xs
Transport	OpenMC	CE, ENDF/B-VIII.0	3-D Monte-Carlo TBR (1.2×10^8 hist.)
CAD geometry	DAGMC	–	as-drawn transport
HPC	NVIDIA A100/H100/A10	GPU HPC fleet	parallel history transport

Code-and-data provenance for gate BR-L2-A28. Every stage is versioned and open.

Data availabilitytocsectionData availability This paper is archived at doi:10.5281/zenodo.22645777. All headline quantities are frozen anchors (gate BR-L2-A28) in the de-risking register[reg], regenerable from the open register deposit doi:10.5281/zenodo.22133057. No economics or proprietary data appear in the public record.

Competing intereststocsectionCompeting interests The author is affiliated with Kronos Fusion Energy. Provisional patent applications relating to aspects of the breeder blanket and fuel cycle have been filed. The author declares no other competing interests.

APPARATUS

Portfolio, People & Record

The intellectual-property estate, the people who built it, the limitations that gate it, and the sources it rests on.

1	4	40	27
GRANTED PATENT	2026 PROVISIONALS	TEAM MEMBERS	2022 PROVISIONALS

INTELLECTUAL PROPERTY

The patent portfolio

The estate spans a granted high-field magnet patent, pending utility and 2026 provisional filings that map to the two products, a digital-twin control provisional, a registered trademark application, and the original 2022 provisional family. Every patentable disclosure in the five-paper arXiv drop was protected **before** publication: the breeder and burner provisionals were filed 1–2 August, and a publication-gap omnibus filing the evening before the drop swept up the DEC, plasma-control, and REBCO-winding matter of the three otherwise-unprotected papers. Patent numbers and application serials are matters of public record; claim scope is summarised, not reproduced.

GRANTED & PENDING UTILITY

US 12,009,112	High-field tilted graded-REBCO magnet architecture (KRONO-002A) · *app. 17/878,550*	GRANTED
US 17/878,507	Core device / method (KRONO-001A) · *utility*	PENDING

2026 PROVISIONAL FILINGS — MAPPED TO THE PRODUCTS

64/124,209	Hyperion — spherical-tokamak tritium / helium-3 foundry · *breeder · filed Aug 2026*	FILED
64/124,220	Aegis / MetroVolt — synchrotron-cutoff, fuel-selection & plug-winding · *generator · filed Aug 2026*	FILED
64/115,167	KRONOS-CTRL digital-twin plant control · *provisional · filed Jul 2026*	FILED
64/005,440	Integrated spherical-tokamak fusion architecture · *early integrated-architecture filing · filed Mar 2026*	FILED
64/105,530	Negative-triangularity spherical-tokamak architecture · *foundational ST filing · filed Jul 2026*	FILED
64/128,097	Publication-gap omnibus — quasineutral two-species expander DEC · provenance-gated / latency-split plasma control · as-built high-field winding & conductor architecture · *filed 7 Aug 2026, the evening before the arXiv drop*	FILED

TRADEMARK & ORIGIN

FTK-98801894	Kronos Fusion Energy — trademark application	FILED
KR-2022-*	Original provisional family (KR-2022-00001 … 00027) · *27 filings, 2022*	PRIORITY

The narrow, defensible magnet novelty — the specific conductor and the digital-twin winding optimisation, not high-field REBCO as a category — is the commercial core that can earn ahead of any Q > 1 milestone, with markets in fusion magnet supply, MRI/NMR, accelerators, and proton therapy. The claim is scoped honestly: the high field is a *system field* result, not a stand-alone-coil record; the small-bore plug coil is structurally infeasible as a bare winding but resolved by a stress-managed structural shield (feasible-pending-FEA); and the winding-tape experiment returned a *null result*. The value is the method, not a field record.

THE TEAM

Founding partners, board, advisors & auditors

Kronos is built by a bench of advanced-fuel-fusion, high-field-magnet, direct-conversion, and materials specialists, with a board and operations team drawn from defense, national laboratories, and industry. Dates are shown as ranges; where a tenure has ended or is term-ending, the range says so.

FOUNDING PARTNERS — SCIENCE

Dr. Gerald Kulcinski
Co-Founder · Helium-3 & Advanced-Fuel Fusion
UW–Madison · 2023–Present

Dr. Carl Weggel
Chief Scientist / S.M.A.R.T. Design
MIT Alcator Program · 2022–Present

Dr. Robert J. Weggel
Magnetic Field Design
MIT Magnet Lab · 2022–Present

BOARD

Priyanca Ford
Founder & CEO · Executive Board
2022–Present

Bandel Carano
Finance / Strategy
Oak Investment Partners / Stanford · 2025–2026

Patrick Schweiger
Fusion Device Engineering
Oklo / CFS / TerraPower · 2025–Present

SCIENTIFIC ADVISORS

Dr. Steven O. Dean
Fusion Policy & History
DOE / Fusion Power Associates · 2025–Present

Dr. Patrick H. Diamond
Plasma Turbulence & Transport
UC San Diego · 2025–Present

Dr. Jack J. Dongarra
HPC & Numerical Algorithms
Turing Award · 2025–Present

Dr. Nasr M. Ghoniem
Chief Material Scientist
UCLA · 2025–Present

Dr. Siegfried Glenzer
Plasma Physics & Laser Fusion
Stanford / SLAC · 2022–Present

Dr. David A. Hammer
Pulsed-Power Fusion
Cornell · 2025–Present

Dr. Donald A. Spong
Plasma Theory & Confinement
ORNL · 2025–Present

Dr. Curtis Smith
Risk Assessment & Nuclear PRA
MIT / Idaho National Lab · 2025–Present

RADM (Ret.) David Goggins
Naval Engineering & Power-Plant Design
U.S. Navy · 2023–2026

Dr. Konstantin Batygin
Mathematician
Caltech · 2022–2025

Dr. Ruben Fair
Magnet Design / ITER Liaison
PPPL / Jefferson Lab · 2022–2025

Dr. Paul S. Weiss
Materials & Nanotechnology
UCLA · 2022–2025

SCIENTIFIC AUDITORS

Dr. Wilfred A. Cooper
Plasma Equilibrium Theory
EPFL / Swiss Plasma Center · 2025–Present

Dr. Ahmed Hassanein
Plasma–Material Interactions
Purdue / Argonne · 2025–Present

Dr. Patrick E. Hopkins
Extreme Thermal Materials
U. of Virginia · 2025–Present

Dr. Peter Hosemann
Materials Joining & Structural Integrity
UC Berkeley · 2025–Present

Dr. Travis W. Knight
Advanced Nuclear Fuels & Systems
U. of South Carolina · 2025–Present

Dr. Philippe Lebrun
Cryogenics & Magnet Cooling
CERN · 2025–Present

Dr. Nitendra Singh
Nuclear Safety & Fuel Cycle
ITER · 2025–Present

Dr. Guido Van Oost
Magnetic Confinement & Fusion Systems
Ghent University · 2025–Present

Dr. Gary S. Was
Radiation Materials & Structural Integrity
U. of Michigan · 2025–Present

Dr. Ray Sedwick
Direct Energy Conversion
U. of Maryland · 2025

OPERATIONS

Michael Laughlin
Chief Operating Officer
2022–Present

Martin Owens
Chief Strategy Officer
LANL / GE Hitachi · 2022–Present

MG (Ret.) Paul Pardew
Chief Contracting Officer
Army Contracting Command · 2022–Present

David Beck
Fusion Commercialization
U.S. Space Force · 2024–Present

Sushma Bhatia
Environmental & Legislative
Google · 2022–Present

Michael De Frenza
Technology Licensing
2023–Present

MG (Ret.) Robin L. Fontes
Cybersecurity & Defense
Army Cyber Command · 2022–Present

Vijay Gehani
Supply Chain & Components
INOX India · 2024–Present

Jon Michel Greenwood
IT & AI Infrastructure
Live Nation · 2023–Present

Brian C. O'Neill
National Security
CIA / ODNI · 2022–Present

Gen. (Ret.) Gustave F. Perna
DoD Liaison
Operation Warp Speed · 2022–2023

Andrea Romero
Marketing Lead
2022–Present

Legal counsel is retained; those roles are held on the internal roster and are not listed here.

SOURCES

References

1 P I Ford, *Kronos Physics De-Risking Register*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22133057

2 M Abdou, M Riva, A Ying *et al*, Physics and technology considerations for the deuterium–tritium fuel cycle and conditions for tritium fuel self sufficiency, *Nucl. Fusion* **61** 013001 (2021). doi:10.1088/1741-4326/abbf35

3 M Abdou *et al*, Blanket/first wall challenges and required R&D on the pathway to DEMO, *Fusion Eng. Des.* **100** 2 (2015). doi:10.1016/j.fusengdes.2015.07.021

4 L V Boccaccini *et al*, Objectives and status of EUROfusion DEMO blanket studies, *Fusion Eng. Des.* **109–111** 1199 (2016). doi:10.1016/j.fusengdes.2015.12.054

5 J E Menard *et al*, Fusion nuclear science facilities and pilot plants based on the spherical tokamak, *Nucl. Fusion* **56** 106023 (2016). doi:10.1088/0029-5515/56/10/106023

6 B N Sorbom *et al*, ARC: a compact, high-field, fusion nuclear science facility and demonstration power plant with demountable magnets, *Fusion Eng. Des.* **100** 378 (2015). doi:10.1016/j.fusengdes.2015.07.008

7 A J Creely *et al*, Overview of the SPARC tokamak, *J. Plasma Phys.* **86** 865860502 (2020). doi:10.1017/S0022377820001257

8 D A Brown *et al*, ENDF/B-VIII.0: the 8th major release of the nuclear reaction data library with CIELO-project cross sections, new standards and thermal scattering data, *Nucl. Data Sheets* **148** 1 (2018). doi:10.1016/j.nds.2018.02.001

9 M B Chadwick *et al*, ENDF/B-VII.1 nuclear data for science and technology: cross sections, covariances, fission product yields and decay data, *Nucl. Data Sheets* **112** 2887 (2011). doi:10.1016/j.nds.2011.11.002

10 A J M Plompen *et al*, The joint evaluated fission and fusion nuclear data library, JEFF-3.3, *Eur. Phys. J. A* **56** 181 (2020). doi:10.1140/epja/s10050-020-00141-9

11 O Iwamoto *et al*, Japanese evaluated nuclear data library version 5: JENDL-5, *J. Nucl. Sci. Technol.* **60** 1 (2023). doi:10.1080/00223131.2022.2141903

12 A D Carlson *et al*, Evaluation of the neutron data standards, *Nucl. Data Sheets* **148** 143 (2018). doi:10.1016/j.nds.2018.02.002

13 N Otuka *et al*, Towards a more complete and accurate experimental nuclear reaction data library (EXFOR): international collaboration between nuclear reaction data centres (NRDC), *Nucl. Data Sheets* **120** 272 (2014). doi:10.1016/j.nds.2014.07.065

14 D L\'opez Aldama and A Trkov, *FENDL: A Fusion-Oriented Evaluated Nuclear Data Library* (FENDL-3 processing), International Atomic Energy Agency, Vienna (2011). doi:10.61092/iaea.xa1g-sjsv

15 A J Koning and D Rochman, Modern nuclear data evaluation with the TALYS code system, *Nucl. Data Sheets* **113** 2841 (2012). doi:10.1016/j.nds.2012.11.002

16 A J Koning, D Rochman *et al*, TENDL: complete nuclear data library for innovative nuclear science and technology, *Nucl. Data Sheets* **155** 1 (2019). doi:10.1016/j.nds.2019.01.002

17 R E MacFarlane, D W Muir, R M Boicourt, A C Kahler and J L Conlin, *The NJOY Nuclear Data Processing System, Version 2016*, Los Alamos National Laboratory report LA-UR-17-20093 (2017). doi:10.2172/1338791

18 P K Romano, N E Horelik, B R Herman, A G Nelson, B Forget and K Smith, OpenMC: a state-of-the-art Monte Carlo code for research and development, *Ann. Nucl. Energy* **82** 90 (2015). doi:10.1016/j.anucene.2014.07.048

19 J-Ch Sublet, J W Eastwood, J G Morgan, M R Gilbert, M Fleming and W Arter, FISPACT-II: an advanced simulation system for activation, transmutation and material modelling, *Nucl. Data Sheets* **139** 77 (2017). doi:10.1016/j.nds.2017.01.002

20 A M Ibrahim, P P H Wilson *et al*, ITER neutronics modeling using hybrid Monte Carlo/deterministic and CAD-based Monte Carlo methods, *Nucl. Technol.* **175** 251 (2011). doi:10.13182/nt175-251

21 A Serikov *et al*, High performance parallel Monte Carlo transport computations for ITER fusion neutronics applications, *Prog. Nucl. Sci. Technol.* **2** 294 (2011). doi:10.15669/pnst.2.294

22 M L Williams, B L Broadhead, M A Jessee, J J Wagschal and R Y Lee, SCALE-6 sensitivity/uncertainty methods and covariance data, *Nucl. Data Sheets* **109** 2796 (2008). doi:10.1016/j.nds.2008.11.012

COLOPHON

How this document was made

The Kronos Fleet — Combined Editorial, 2026 is set in **Fraunces** (display), **Gelasio** (text), and **IBM Plex Mono** (data & equations), carried forward from the Kronos editorial design system.

The figures are rendered at 300 dpi from the deposited generator scripts; the document is composed as a single self-contained file with embedded fonts and imagery, paginated to US Letter, and rendered to PDF through a headless Chromium engine in matched light and dark editions.

Every headline quantity carries an evidence class; withdrawn and restated values are marked as such. Nothing herein is frozen beyond the founder-approved Tier-A set.

Explore it live — turn the interactive 3-D model, run the deposited solvers in the physics validator, and watch the companion film series.

LEGAL NOTICE

Notice, status, and distribution

Conceptual design & simulation study. This document reports a conceptual design and simulation study. It is not a construction commitment, a safety-analysis report, a regulatory filing, or an offer of securities. Forward-looking statements — schedules, costs, market sizes, and performance — are estimates subject to the limitations set out in the Simulations part and may change.

Numbers and their classes. Quantities are reported with evidence classes and case labels; modelled, assumed, and requirement-class values are identified as such and are not to be quoted without their labels.

Public edition. This edition carries the physics and design only; all financial, funding, and defense-commercial content has been removed for public distribution. The scientific text corresponds to the arXiv and journal submissions.

THE 2026 KRONOS SERIES

One programme, many papers

This editorial is one paper in the Kronos 2026 series — a real fusion company's complete, reproducible physics record for a spherical-tokamak breeder and its low-neutron D–^{3}He tandem-mirror burners. Every headline number is a frozen anchor in the Physics De-Risking Register.

Explore the whole programme: the De-Risking Register, the interactive 3-D model, and the physics validator.

KRONOS FUSION ENERGY

Nuclear-Data Uncertainty Quantification for Tritium Breeding: Which 14 MeV Cross-Section Channels Set the Error Bar on the Breeding Ratio

A real fusion company building real machines. Explore the science, live.

Zenodo · doi:10.5281/zenodo.22645777

Read on kronosfusionenergy.com

The Physics De-Risking Register