Electron Screening in Metal Deuterides

Measurement Systematics, Evolving Target State,
and an Experimental Arbitration Protocol

Jeromie Beasley
Independent researcher · dicipler@gmail.com
SCREEN-PROG-04 · 28 October 2026

AI-assisted: AI tools were used for literature search, computation and drafting; the author directed the work and takes responsibility for its content.

Abstract

Screening parameters inferred from nuclear yields in metal deuterides depend on the response model, the target's preparation and the target's acquisition history. We separate three quantities that are often merged: the instantaneous screening of the target state, a yield-equivalent value within one acquisition interval, and a global constant-model fit to an energy sweep. Fully specified synthetic examples show that the order in which energies are taken, together with a fitted normalization, can shift a constant fit far from the physical range: above the endpoint for a high-to-low sweep and below the starting value for a low-to-high sweep. The underlying drift models are identifiable but severely ill-conditioned in a single sweep, and a sweep that returns over the same energies reduces the start-value error four- to fivefold at fixed fluence. We propose charge-normalized, rapidly interleaved two-energy measurements at controlled accumulated fluence, with independent materials and proton-spectrum diagnostics, and a consistent uncertainty budget of about 7.2 eV per dose step. A third interleaved energy adds a conditional test of yield curvature, with an expected separation of about 4 standard deviations from a homogeneous target for a 150 eV spread, but it cannot by itself separate a screening spread from an additive thermal contribution. The paper reports no new nuclear data; it is a measurement-design analysis with reproducible calculations.

1 Contribution

Inferred screening depends on material state; that is established by prior experiments, not by this paper. Studies of prepared and loaded metallic targets report preparation-sensitive yields and screening fits [2, 3, 4, 7], and comparisons across different reactions and methods are not automatically measurements of the same quantity.

The contribution is downstream of that record: (i) a formalism for acquisition history, the fact that a beam advances the target state while it measures it; (ii) reproducible examples of fitting bias that depends on acquisition order and normalization; (iii) an assessment of practical identifiability, with schedules that improve it; and (iv) a controlled-fluence, interleaved protocol that can test whether the measured response evolves, extended by a third energy as a curvature diagnostic.

2 Scope and commensurability

The subject is low-energy nuclear-yield inference in metal deuterides, including the reported low-energy yield plateau and its thermal-spike interpretation. Company-reported heat, commercial readiness and claims of usable power are outside the evidential basis.

All energies in the calculations and protocol are centre-of-mass energies. For an atomic deuteron incident on a stationary deuteron the nonrelativistic laboratory energy is twice the

centre-of-mass energy; molecular beams need their species and energy per deuteron stated separately.

An inferred U_e is tagged by reaction, detection channel, preparation, depth model, fit window and beam source. A spread across methods is not a calibrated disagreement on identically prepared specimens. The contrast between accelerator fits of about 340 eV on implanted ZrD_2 and a reported bound below 30 eV from a different preparation, source and channel is a reason for arbitration, not by itself a matched contradiction; the ratio of the two is at least about elevenfold, since 30 eV is an upper bound.

3 The electrolysis record as a systematics precedent

Under conventional near-equal DD branching, a mean energy release of 3.65 MeV per reaction implies

$$\dot{N}_{DD} = \frac{1\ \mathrm{J\,s^{-1}}}{3.65 \times 10^{6} \times 1.602 \times 10^{-19}\ \mathrm{J}} \simeq 1.71 \times 10^{12}\ \mathrm{s^{-1}}, \qquad \dot{N}_n \simeq 8.55 \times 10^{11}\ \mathrm{s^{-1}} \tag{1}$$

per watt attributed to those branches. A proposed alternative channel must be assessed with its own branching fractions and radiation transport.

The methodological lesson is coupled systematics: surface chemistry can affect loading, recombination and the measurement response through the same evolving state, and a coherent-looking time dependence need not identify its cause. The largest recent programme makes the point concretely. Project Callisto (2015–2024; more than 500 samples and more than 200,000 sample-hours of calorimetry; reports released 8 September 2026) [12] lists the sources of apparent excess heat it found: recombination, changing calorimeter response, faulty sensors, errors in electrical-power accounting, and the different reaction entropies of hydriding and hydrogen evolution. The last is a chemical term that moves with the material state during a run, the same structure as the beam-target problem below. Full loading was reached in gas loading; the October 2026 review of the reports notes that its electrochemical replications stayed below a loading of 0.9 [13]. Its nuclear detection used high-purity germanium and scintillators, and the 511 keV line in its spectra is attributed to a ^{22}Na source [12]. The programme did no ion-beam or screening work, and the review faults it for never measuring helium and for releasing no raw data [13]. Independent nuclear and materials diagnostics are therefore part of the design proposed here.

4 Measurement record and response model

4.1 Screening dependence

For illustration we use the penetration-factor model

$$R(E, U_e) = \exp\left[-\sqrt{\frac{E_G}{E + U_e}}\right], \qquad E_G = 986\ \mathrm{keV}, \tag{2}$$

the DD Gamow energy, with U_e the usual screening parameter [1]. This is a toy response, not a complete thick-target cross section: S-factors, stopping, depth distributions, resonance terms, exposure and detector response lie outside it. The same model is used throughout the synthetic calculations.

4.2 Selected reported parameters

Host	Reported parameter	Conditions and interpretation
ZrD_2	about 100 eV	Sputtered, low-defect reference in the Szczecin work.
ZrD_2	340 ± 30 eV	Implanted target, proton-yield fit [5].
ZrD_2	up to 500 eV	Oxygen-implanted preparations; not a host constant [4, 6].
Pd / PdD	0.8 ± 0.2 / 1.7 ± 0.2 keV	Beam-loaded versus electrochemically loaded foil; thick-target fits, mainly 2.5–6.5 keV [7].
Ti / TiD	0.3 ± 0.2 / 0.9 ± 0.2 keV	Different loading conditions in the same work [7].
$PdD_{0.05}$	about 400 eV	Implanted target, direct-yield component [6].
$TiD_{1.93}$	about 25 eV	Direct-component screening; the plateau interpretation needs further assumptions [6].

The palladium and titanium values of Karahadian *et al.* come from a rig that loads deuterium electrochemically through the back of a foil while the beam strikes the front; switching loading on raises the fitted screening, and a constant-screening description fails below about 2.5 keV [7]. These values must not be merged with a different host, loading or fit window as though they were repeated measurements of one material.

4.3 Systematics and independent characterisation

Target-density dynamics, contamination, stopping and depth inhomogeneity have long been identified as sources of apparent screening changes [2, 4]. The Szczecin group now states publicly that screening depends strongly on lattice defects and oxidation [10]. A beam that loads its own target is also seen in applied work with 30–60 keV deuterium beams on metals and graphite [11]. No monotonic evolution law follows from preparation differences alone: creation, annealing, loading, loss and surface change compete, and the sign of a dose dependence is an experimental question.

5 Quantitative structure and thermal interpretation

5.1 Thermal penetration-factor leverage

At thermal collision energies, taking $E \to 0$ in (2) gives the penetration-factor ratios below. They are ratios of the factor, not thermally averaged rates.

U_e (eV)	$\sqrt{E_G/U_e}$	$R(0,U_e)/R(0,100)$
100	99.30	1
200	70.21	$10^{12.63}$
340	53.85	$10^{19.74}$
400	49.65	$10^{21.56}$

At a beam energy of 3 keV the same screening values change the factor only by 1.32, 1.92 and 2.24 relative to 100 eV (Figure 1). The thermal leverage must not be read as beam-energy leverage. In the thermal expression of the Szczecin model, $G(U) = \sqrt{E_G/U}\,\exp[-\sqrt{E_G/U}]$ carries a prefactor that falls by 1.844 from 100 to 340 eV and by 2 from 100 to 400 eV.

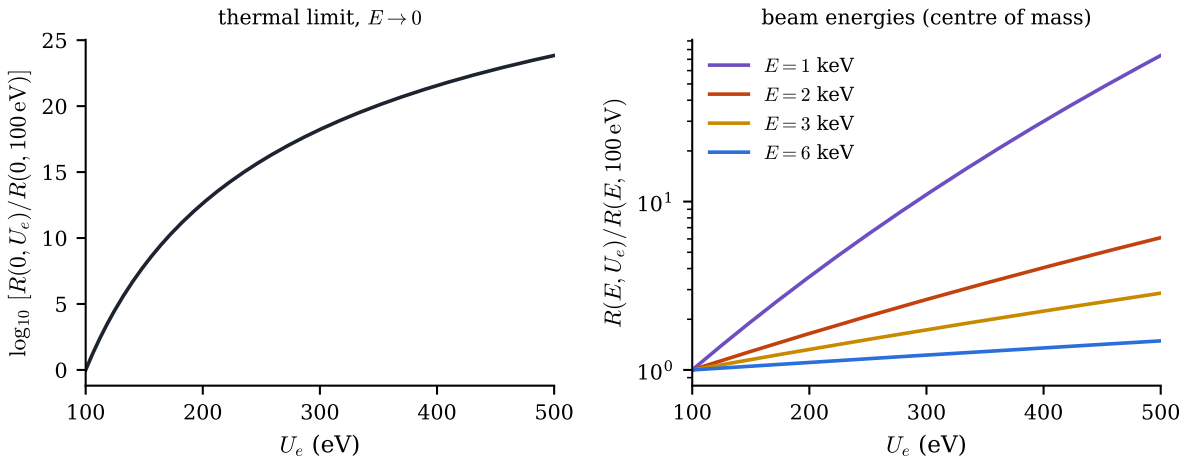


Figure 1: The screening lever. Left: the thermal limit, where a few hundred eV moves the penetration factor by twenty decades. Right: at beam energies the same change is a factor of order one to a few, which is why a beam measurement must control everything else that changes the yield.

5.2 Mobilisation and energy accounting

For a declared activated mobile fraction $f(T) = \exp[-E_A/(k_B T)]$ with $E_A = 0.413$ eV,

$$f(1080\ \mathrm{K})/f(300\ \mathrm{K}) \simeq 1.03 \times 10^5. \tag{3}$$

This is a conditional population ratio, not a yield prediction. Thermal yields depend in addition on the temperature field, track volume and lifetime, densities, cross section and reaction model. A large thermal population factor does not establish coherent concentration of phonon energy or usable power.

5.3 Density powers and the cross-host comparison

Let N_0 be the target deuteron density and $N = fN_0$ the mobile density. The updated thermal treatment uses a mobile–target rate $N_0 N\langle\sigma v\rangle$ [5, 6]: linear in f, quadratic in N_0. A mobile–mobile model instead carries N^2, with a symmetry factor for identical particles. Calling one model "linear" and the other "quadratic" without naming the variable is ambiguous. The source versions differ: arXiv:2409.02112v2 uses $n_m^2/2$ (Eqs. 4 and 17), while v3 uses $n_D n_m$ (Eq. 21), as does arXiv:2605.27438 (Eq. 8). A detected per-ion yield requires $y = \eta \int dt\, dV\, r$, including the track space–time factor, efficiency and branch convention. For uniform tracks with $W = V\tau$ the updated model scales as $\eta W f n_D^2 K$, and any compensating density dependence of W must be shown, not assumed.

The cross-host factors are $f_{Pd}/f_{Zr} = 16.876$ and $G(400)/G(340) = 61.659$, with product 1040.6. With a nominal density ratio $q = n_{Zr}/n_{Pd} = 37$, an effective $f n_D K$ comparison gives 28.12; the updated $f n_D^2 K$ model with equal η and W gives 0.760; the older $f^2 n_D^2 K$ model gives 12.83. These differ in both the density power and the mobile-fraction power.

The measured absolute plateau ratio is $0.0069/0.00026 = 26.54$. The fitted normalization ratio is $14.6 \times 10^6/3.9 \times 10^5 = 37.44$, giving a normalized plateau ratio of 993.5; these are one observation under two normalizations, not two validations. Equating the fitted ratio to a true density ratio gives an equal-W prediction of 0.743, and matching 26.54 would require $W_{Pd}/W_{Zr} \simeq 35.7$. That diagnoses an unspecified factor; it does not rule out thermal models. The fitted normalization also involves stopping and response, since the direct-yield amplitude scales as n_D/C with stopping coefficient C; it is not a measured bulk density. The source's compositions and mass densities give a bulk $n_{Zr}/n_{Pd} \simeq 16.9$, or 20.9 with its rounded loadings,

and neither is the irradiated depth profile. Until stopping, reacting densities and track integrals are reconciled, the near-27 agreement is not an independently normalized validation.

5.4 Titanium and the plateau

The thermal interpretation invokes a larger screening in the hot titanium track than the direct-component value of about 25 eV, with a value near 370 eV used to match the cross-host comparison [6]. That is a model-dependent material-state hypothesis, not an independent measurement. Proton spectral shifts and component amplitudes test it beyond total yields. The same group reports pair emission from a ^{4}He threshold resonance [8], and a recent poster extends palladium measurements with D^+, D_2^+ and D_3^+ beams from 20 down to 3.5 keV, finding a nearly flat yield below 5 keV modelled by thermalized tracks, screening and that resonance [9]. No conclusion below requires the resonance or a particular account of the plateau.

6 Measurement of an evolving state

6.1 The acquisition-history problem

The beam advances fluence while measuring a target whose state may evolve. Let $S(\Phi)$ collect defects, composition, loading and depth profiles, and let a model assign $U_e^{\mathrm{inst}}(\Phi)$ to that state. For a scalar model the exposure-normalized yield during interval i is

$$R_i = \frac{1}{\Delta\Phi_i} \int_{\Phi_i}^{\Phi_i+\Delta\Phi_i} R\big(E_i, U_e^{\mathrm{inst}}(\Phi)\big)\, d\Phi. \tag{4}$$

Exposure, density and detection weights belong inside the integral when they vary. The interval-equivalent value $U_{e,i}^{\mathrm{eq}}$ solves $R(E_i, U_{e,i}^{\mathrm{eq}}) = R_i$. Because R is monotone, it lies within that interval's state range; convexity in this regime places it slightly above the interval mean, never above its maximum.

6.2 Global fitting and normalization

The value returned by a global constant model is different:

$$(\hat{a}, \hat{U}) = \arg\min_{a,U} \sum_i w_i \big[\log R_i - a - \log R(E_i, U)\big]^2, \tag{5}$$

with a a fitted log normalization. Energy, fluence and normalization trade against one another, so $\hat{U}$ can exceed the physical endpoint or fall below the starting value, depending on the order in which energies are taken. With a known, fixed normalization, positive-weight fits of this monotone model cannot place their optimum outside the target-state range.

6.3 A fully specified example

Set $w_i = 1$ and take eight energies in the order $E_i = 6000, 5000, 4000, 3250, 2500, 2000, 1600, 1250$ eV. Assign interval i the fluence bin $[i/8, (i+1)/8]$, integrate (4) with 2000 nodes per interval, and fit (5) with a free normalization:

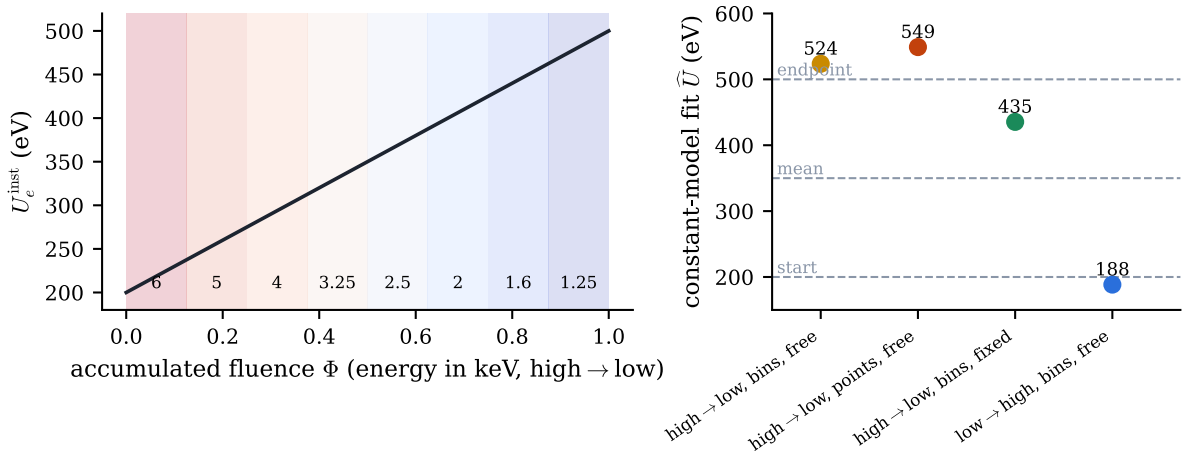


Figure 2: What a constant fit returns depends on how the data were taken. Left: a linear drift with eight energies taken high to low. Right: the constant-model fit under four acquisition models; only the order and the normalization move it outside the start–endpoint range.

Evolution on $0 \le \Phi \le 1$	Mean (eV)	Endpoint (eV)	Constant fit (eV)
$300 + 100\Phi$	350.00	400.00	406.44
$250 + 200\Phi$	350.00	450.00	464.28
$200 + 300\Phi$	350.00	500.00	523.61
$200 + 300(1 - e^{-\Phi/0.2})$	440.40	497.98	528.03
$200 + 300(1 - e^{-\Phi})$	310.36	389.64	407.32
Constant 340	340.00	340.00	340.00

Against the campaign mean the bias in these examples runs from about 16 to 50 percent; it is not a universal number. For the 200 → 500 eV case (Figure 2), fixing the normalization gives 435.43 eV, and reversing the order with a free normalization gives 188.47 eV, below the starting value. Reading the instantaneous state at $\Phi = i/7$ instead of integrating each bin gives 548.98 eV. These are different acquisition models: point sampling stands for a real exposure only when the exposure is short on the scale of target evolution, and otherwise (4) is required. The fit settles near the state the target had while the low-energy points were taken, where the yield is most sensitive to U_e; averaging inside each interval moves it by less than half an eV. A saturating law approaches 500 eV only asymptotically, so "above the endpoint" and "above the mean" are different statements and must not be interchanged. The saturating fits of 607 and 394 eV quoted for this comparison in earlier work depend on an unstated recipe and are carried here as open benchmarks.

6.4 Practical identifiability and alternative schedules

For a point-sampled drift model $\log R_i = a - \sqrt{E_G/(E_i + U_0 + \Delta U\, f_i)}$ the Jacobian columns are 1, b_i and $b_i f_i$ with $b_i = \sqrt{E_G}/[2(E_i + U_i)^{3/2}]$; they are generically independent, and the interval model also has full column rank. High parameter correlation therefore means ill-conditioning, not structural non-identifiability at every precision.

For the integrated descending linear example with uniform log-yield errors of 2 percent, the local standard errors are 30.83 eV on U_0 and 28.34 eV on ΔU, with correlation -0.9931. At 0.2 percent they fall to 3.08 and 2.83 eV with the same correlation. At fixed total fluence, splitting eight exposures into sixteen multiplies count errors by $\sqrt{2}$; in that comparison a descending–ascending schedule gives $\sigma(U_0) \simeq 6.9$ eV and two descending blocks give 6.0 eV, against 30.8 eV for one descending sweep (Figure 3). With a log-linear error profile from 2 percent at 6

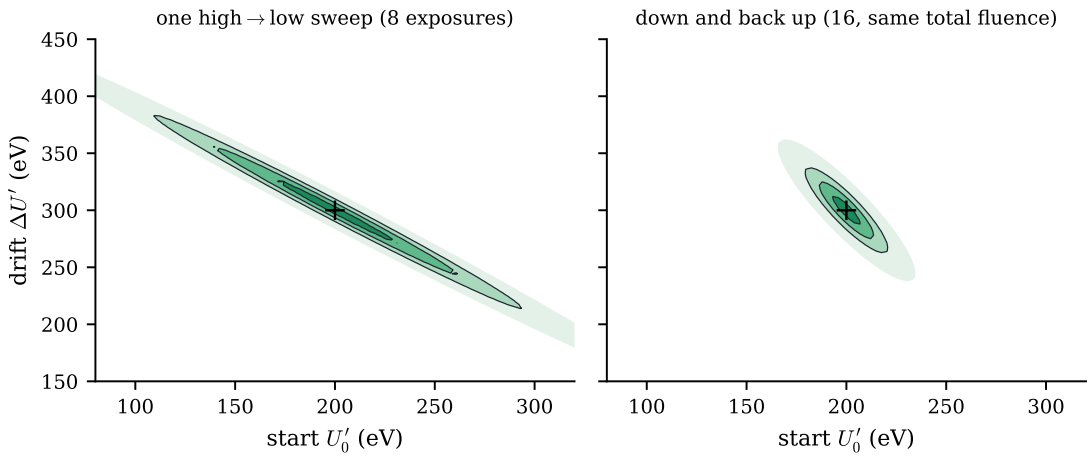


Figure 3: The degeneracy belongs to the sweep, not the physics. $\Delta\chi^2$ contours (1, 4, 9) for start value and drift, uniform 2 percent precision, at the same total fluence. One high-to-low sweep leaves a long valley; returning over the same energies closes it.

keV to 30 percent at 1.25 keV, the corresponding values are about 23.1, 30.3 and 124.6 eV. Under these assumptions the gain is a factor of four to five; switching overhead is not charged. Published curves that lack acquisition order or fluence metadata may remain insufficient for recovery, through missing information and practical conditioning.

6.5 What the examples establish

State differences, within-interval averaging and global-fit confounding are distinct effects. The simulations do not show that they explain the full literature spread, or that real targets evolve monotonically. A residual or runs test on eight points has very low power, so failing to reject is not evidence of a stable target. The calculations are reproducible illustrations of a systematic to test.

7 An interleaved, controlled-fluence protocol

7.1 Separate a common scale from energy-dependent shape

At one energy a changing density scale and a changing screening parameter mimic each other. At 3 keV and $U_e = 340$ eV, (2) gives $\partial \log R/\partial U_e \simeq 2.57 \times 10^{-3}$ eV^{-1}, so a 10 percent scale change corresponds to $\log(1.1)/(2.57 \times 10^{-3}) \simeq 37$ eV in this local model.

7.2 Conditional cancellation with two energies

For charge-normalized yields $y_j(\Phi) = A(\Phi)F_j(U_e)$ with a shared, energy-independent A, the ratio cancels A. For the response (2),

$$\frac{\partial}{\partial U_e} \log \frac{R(2000, U_e)}{R(6000, U_e)} = \frac{\sqrt{E_G}}{2}\Big[(2000 + U_e)^{-3/2} - (6000 + U_e)^{-3/2}\Big] \simeq 3.403 \times 10^{-3}\ \text{eV}^{-1} \quad (6)$$

at $U_e = 340$ eV. The noise cost of the ratio against a single energy depends on the count allocation and the comparator; against a single energy at 3 keV with equal noise it is a factor 1.07. Charge integration is an explicit normalization measurement with its own uncertainty. Density cancels as one scale only if its depth-profile shape stays fixed over the depths sampled

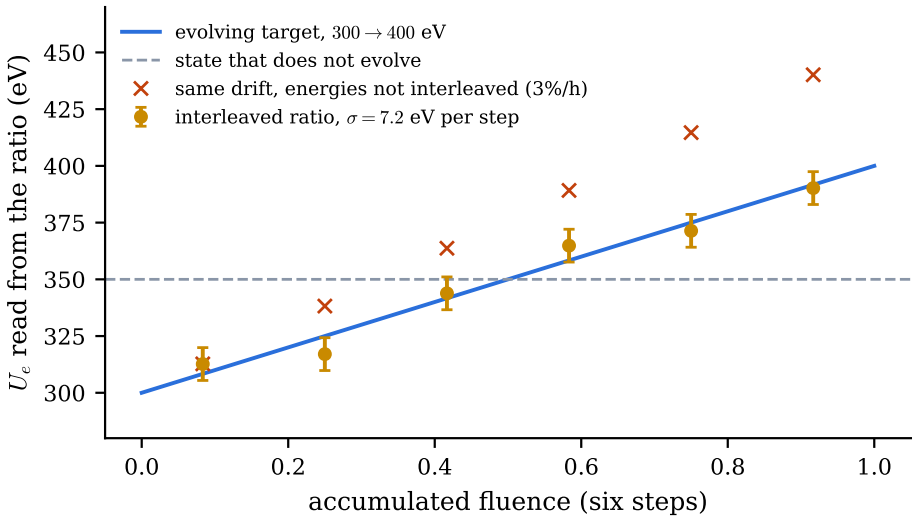


Figure 4: The protocol on one target. Six dose steps, each an interleaved 2/6 keV ratio with the 7.2 eV budget, follow an evolving state; the crosses show what a 3 percent per hour drift forges when the two energies are run one after the other.

by both energies. Energy calibration, stopping, S-factor ratios, channel efficiency, energy spread, background and additive components can all alter the ratio. The S-factor and response ratio enter the calibration of absolute U_e; a fixed response ratio cancels in changes of the log ratio only while the response and weights stay fixed.

7.3 Error budget

Term in one log-ratio estimate	Standard deviation
Counting at 2 keV	0.020
Counting at 6 keV	0.005
Residual detector gain	0.004
Residual dead time	0.002
Spot reproducibility across settings	0.006
Background subtraction	0.005
Stopping-model residual	0.010
Quadrature total for independent terms	0.0246

This gives about 7.23 eV per dose step in the direct-response model. It excludes an absolute S-factor or energy-scale error and is a design assumption, not calibrated accuracy; shared or dose-correlated terms need a covariance model. For six independent, equally spaced steps with the same 7.23 eV error, the mean carries 2.95 eV, the fitted end-to-end linear change 8.65 eV, and a difference of two endpoints 10.23 eV.

7.4 Interleaving and correlated drift

Use a symmetric interleaving sequence such as ABBA, with a cycle short compared with both target evolution and instrumental drift. At fractional log-ratio drift rate r and unmatched time separation τ, the order-of-magnitude residual is $r\tau$. For $r = 3$ percent per hour, $\tau = 5$ min and the sensitivity (6), this is about 0.73 eV; a six-hour unmatched separation gives about 53 eV, a substantial confounder (Figure 4). Stopping and spot terms that evolve with dose share the signal's time signature, and interleaving alone does not prove them absent. Record calibration

histories, beam species, spot, current, detector live time and total diagnostic dose.

7.5 Performance

If a stated precision p is the combined log-ratio error, the local uncertainty is $p/(3.403 \times 10^{-3})$; if instead each of two yields carries error p, it is larger by $\sqrt{2}$:

Precision	Combined ratio error	Two equal independent yields
	local $\sigma(U_e)$ in eV	
1%	2.94	4.16
2%	5.88	8.31
3%	8.82	12.47
5%	14.69	20.78
10%	29.39	41.56

A 240 eV contrast divided by the single-step error of 7.23 eV is about 33; a 100 eV contrast gives about 14. These are local contrast-to-error ratios under an exact reference, not discovery significances for a measured trend; the relevant nonlinear hypotheses should be fitted with the full covariance and reported with confidence intervals and power. A single-session claim also needs facility count rates, switching overhead, background and acceptable dose.

7.6 Fluence accounting

No acquisition interval is exposure-free. Its fluence, and any preparation irradiation, enters the cumulative ledger. Reversing the energy order does not reverse the material history. Independent targets, blocks, reference specimens and repeated preparation are useful controls; a previously visited state cannot be assumed reproducible.

7.7 A third energy as a curvature diagnostic

A third interleaved energy adds a second normalized shape observable. For a fixed family of distributions, such as a lognormal U_e with mean μ and standard deviation s, common sampling weights and no unconstrained additive contribution, the two ratios estimate μ and s. They do not recover an arbitrary distribution, and energy-dependent depth weighting changes the forward map. With two energies alone and μ refitted, the spread is invisible: one ratio is fitted exactly.

The forecast charges the fixed-total-time cost (three settings at $T/3$ each instead of two at $T/2$, so count errors grow by $\sqrt{3/2}$), keeps the covariance from the shared 6 keV reference, and uses the listed residual terms. The statistic is the expected (Asimov) profile likelihood-ratio separation $\sqrt{q_A}$ from the homogeneous null $s = 0$ with μ refitted; a local relative precision s/σ_s at $s = 150$ eV is not a detection statistic and overstates the separation about threefold. For $\mu = 340$ eV and $s = 150$ eV over six steps:

Settings (keV)	Stopping residual independent per step	Stopping shared, one smooth error	Stopping shared, unrelated between ratios
2, 3, 6	4.03	–	2.81
2, 3.5, 6	4.05	4.59	2.68

"One smooth error" is a single stopping-power mis-scaling common to all steps, which moves the 3.5/6 ratio by 0.49 of the 2/6 ratio; "unrelated between ratios" takes equal, uncorrelated stopping offsets on the two ratios and is the pessimistic case. A nonlinear simulation at 2, 3.5 and 6 keV, with common mean and width across six repeats (1000 datasets per case at $s = 150$

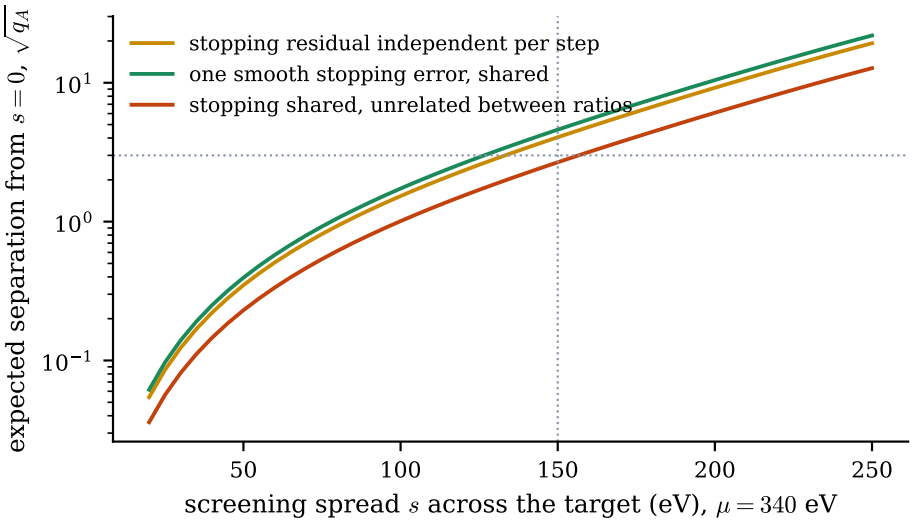


Figure 5: Expected separation of a lognormal screening spread from a homogeneous target, six dose steps at 2, 3.5 and 6 keV, fixed total time, for three treatments of the stopping residual. Dotted: $s = 150$ eV and three standard deviations.

eV, 400 at $s = 0$), gives median $\sqrt{q} = 4.05$ (independent stopping) and 2.70 (shared, unrelated); the fractions exceeding $q > 9$ are 0.86 and 0.37, and the false-rejection rate under $s = 0$ at $q > 2.71$ is 0.07 and 0.05, each with a Monte Carlo error of about 0.01. The forecast is steep in the spread (Figure 5): about 1.5 at $s = 100$ eV and 9 at $s = 200$ eV.

Two ratios cannot distinguish mean, width and a free plateau amplitude. A homogeneous target with an additive constant yield of 10 percent of the direct yield at 2 keV fits a width near 138 eV exactly at the three energies (Figure 6), and opposite 1 percent ratio distortions on a homogeneous target fit a width near 99 eV. A fourth energy supplies a residual check only for the fixed two-parameter shape without the free plateau; more nuisance parameters need more observations or independent constraints, such as proton peak positions. The third energy is therefore a pilot and shape diagnostic. Its width is interpreted only with justified distribution and depth assumptions and a calibrated nuisance structure, and dose dependence is compared in the raw ratios rather than in separately fitted widths.

8 Independent spectral and radiation diagnostics

The thermal literature models a direct screened component plus an additive plateau [5, 6], and its thermal rate contains the screening as well as population, temperature and resonance factors. A thermal model whose screening grows with defects also changes with dose, so neither dose dependence nor broadening by itself excludes thermal tracks, and heating cannot be assumed to cancel in a total-yield ratio. Proton peak positions and component amplitudes, already part of the arbitration proposal, constrain the decomposition.

For a postulated 23.85 MeV channel producing one electron–positron pair per event, one watt corresponds to about 2.62×10^{11} events per second. A predicted 511 keV line also requires the pair branching fraction, positron stopping and annihilation, attenuation, geometry and detector efficiency, and the absence of a line excludes only a quantified detectable flux with stated backgrounds and confidence limits. The calorimetry record above contains no such search: Callisto's germanium spectra show a calibration-source line, and no cell with excess heat was examined for it. This diagnostic is conditional and is not required for the acquisition-history protocol.

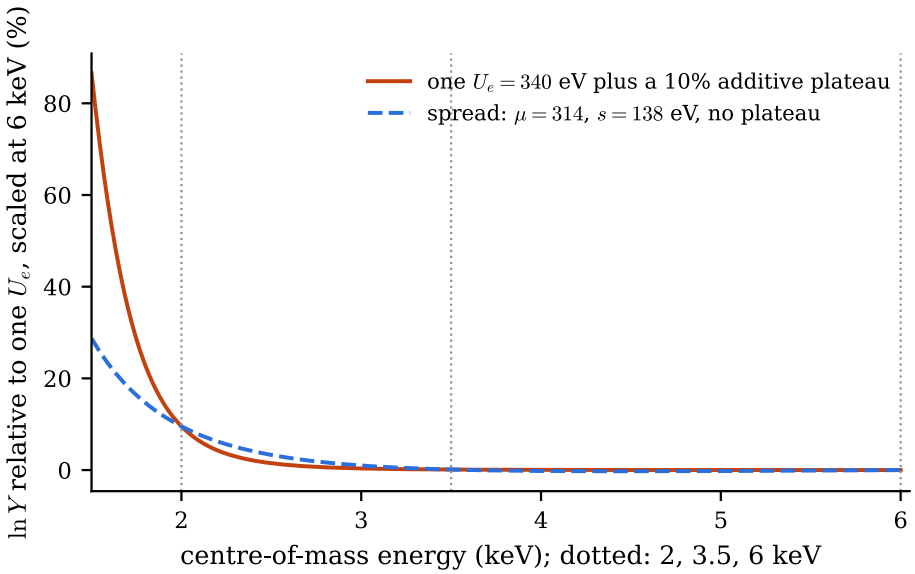


Figure 6: A spread and a plateau give the same three numbers. A homogeneous target at 340 eV with an additive constant yield equal to 10 percent of the direct 2 keV yield matches a spread of $s \simeq 138$ eV exactly at 2, 3.5 and 6 keV; the two part only below 2 keV.

9 From measurement to specification

Treat target preparation as a written, repeatable materials procedure. Useful observables include bulk and surface composition, loading and its depth profile, vacancy signatures, phase, trap populations, thermal history and transport times. Positron annihilation, X-ray diffraction, thermal desorption and surface or depth analysis constrain different parts of that state and are not interchangeable proxies. The analogy with a specification culture in metallurgy is methodological.

10 Arbitration matrix

Arm	Preparation	Bulk decoration	Surface oxidation
1	Sputter-cleaned, inert transfer	reference	minimized
2	Controlled bulk oxygen, cleaned surface	lower dose	minimized
3	Controlled bulk oxygen, cleaned surface	higher dose	minimized
4	Cleaned, then controlled surface oxidation	reference	deliberate
5	Controlled carbon comparison	characterized	characterized
6	Conventional aged preparation	measured	measured

Implantation creates a depth profile, so bulk-alloyed preparation or depth-resolved measurement is needed before the arms differ in one scalar variable. Apply the same charge-normalized, interleaved protocol at recorded fluence, with materials measurements before and after and proton spectra throughout. Where back-side electrochemical loading is available, running the dose-stepped ratio with loading on and off separates defect build-up from deuterium build-up. A distinct dose response in matched arms constrains competing explanations; no single pattern proves bulk decoration without excluding surface and profile effects.

11 Assessment

The paper delivers a formalism for measurement history, reproducible examples of order- and normalization-dependent bias, a diagnosis of ill-conditioning with schedules that relieve it, and a protocol whose limits are stated. Its practical point is to make fluence and target state controlled observables, rather than infer them from an unrecorded energy sweep. The dose-stepped interleaved ratio has not been reported by any group; it is the experiment the record calls for.

12 Open questions

1. Does the calibrated direct response change with controlled fluence, and in which direction?
2. Which preparation, depth and source differences explain the reported spread across methods?
3. What density factors and track space–time integrals enter a cross-host, per-ion plateau comparison?
4. Can thermal desorption and independent phase and defect measurements constrain mobile fractions in decorated regions?
5. Can proton spectra separate direct and thermal contributions at the selected centre-of-mass energies?
6. What beam-spot, stopping, S-factor and energy-calibration floors does a given facility demonstrate?
7. Does the ^{4}He threshold resonance and its radiative signature replicate independently?
8. Which pre-specified energy and fluence schedule gives useful power at fixed total dose and count cost?

Declarations

Funding. No external funding. **Competing interests.** The author declares no competing interests. **Prior publication.** An earlier version of this work, SCREEN-PROG-03, is deposited on Zenodo (doi:10.5281/zenodo.21935286). **AI use.** AI tools were used for literature search, computation, figure preparation and drafting; the author directed the work and takes responsibility for its content. **Data and code.** The scripts and outputs behind every computed number and figure accompany this paper in its Zenodo record (Appendix A).

References

[1] H. J. Assenbaum, K. Langanke and C. Rolfs, Z. Phys. A **327**, 461 (1987).

[2] A. Huke *et al.*, Phys. Rev. C **78**, 015803 (2008), arXiv:0805.4538.

[3] A. Cvetinović *et al.*, "Electron screening in palladium," Phys. Lett. B **838**, 137684 (2023).

[4] A. Kowalska *et al.*, Materials **18**, 1331 (2025), doi:10.3390/ma18061331.

[5] K. Czerski *et al.*, "Observation of thermal deuteron–deuteron fusion in ion tracks," arXiv:2409.02112 (v2 and v3).

[6] K. Czerski *et al.*, "Thermal deuteron–deuteron fusion in metallic targets," arXiv:2605.27438 (2026).

[7] M. E. Karahadian *et al.*, "Enhanced nuclear fusion in the sub-keV energy regime," arXiv:2512.06212 (2025).

[8] K. Czerski *et al.*, Phys. Rev. X (2025), doi:10.1103/chlp-b215; University of Szczecin announcement, 5 November 2025, `https://en.usz.edu.pl/?p=7293`.

[9] M. Thulichery, A. Kowalska, K. Czerski *et al.*, poster, Zakopane Conference on Nuclear Physics 2026, `https://indico.ifj.edu.pl/event/1430/contributions/5959/`.

[10] University of Szczecin, announcement of 5 November 2025, `https://en.usz.edu.pl/?p=7293`.

[11] Chen *et al.* (Nusano), contribution to ECRIS'26, 29 September 2026, `https://indico.jacow.org/event/128/contributions/15943/`.

[12] Project Callisto reports, released 8 September 2026, `https://callisto.report/`; chapters 16 (nuclear detection) and 18 (recommendations).

[13] J.-P. Biberian, E. Storms, F. Celani, P. Werbos, F. Gordon *et al.*, review of the Project Callisto reports, October 2026, `https://callisto.review/review.html`.

A Reproduction

Every computed number and figure comes from these scripts (Python 3, NumPy, SciPy, Matplotlib); each prints the values quoted.

Script	What it produces
`screen04_model.py`	The shared response (2), bin-integrated campaigns, the constant-model fit (5) and the lognormal mixture.
`check_table.py`	The table of Section 6.3, the fixed-normalization, reversed-order and point-sampled values.
`s01`–`s05`	Leverage table, mobilisation factor, per-watt rates and plateau arithmetic.
`s05b_density_powers.py`	The density conventions and normalizations of Section 5.3.
`s08`, `s09`	Sensitivities, error budget, interleaving residuals and the performance table.
`s13_third_energy_forecast.py`	The Asimov separations of Section 7.7 for every stopping treatment and spread.
`s14_identifiability.py`	Rank, Fisher errors and correlations of Section 6.4, and the alternative schedules.
`s16_checks.py`	The uneven-profile schedules, the distortion-induced width and the Monte Carlo of Section 7.7.
`s15_figures.py`	Figures 1–6 and the plateau–spread match.

The core of the forward model is short enough to give here:

```
EG = 986000.0                                  # eV, centre of mass
def lnR(E, U): return -np.sqrt(EG/(E+U))       # equation (2)
def campaign(Uf, Es, point=False, nodes=2000): # equation (4)
    out = []
    for i, E in enumerate(Es):
        if point: out.append(lnR(E, Uf(i/(len(Es)-1)))); continue
        f = (i + (np.arange(nodes)+0.5)/nodes)/len(Es)
        out.append(np.log(np.mean(np.exp(lnR(E, Uf(f))))))
    return np.array(out)
def fit_const(lY, Es, fixed=False):            # equation (5)
    r = lambda U: lY - lnR(Es, U)
    cost = lambda U: np.sum((r(U) - (0 if fixed else r(U).mean()))**2)
    return minimize_scalar(cost, bounds=(0,2000), method='bounded').x
```