UMASS LOWELL RESEARCH REACTOR | Technical Report

Anomalous Bimodal Recovery in 2N2222AUB Transistors Following Pure Gamma Exposure

Implications for the FNI Dosimetry Program

Author: Daniel Gill
Date: May 2026

Sample:	2N2222AUB ($n = 9$), matched prior-irradiation history
Reset:	24-hour 180 °C anneal prior to this run
Exposure:	12,400 rad Co-60, alanine/TLD-verified at grid position
Anneal:	2 hours at 80 °C, per ASTM E1855-15 §8.1.6

1 Executive Summary

Nine 2N2222AUB transistors were exposed to 12,400 rad of pure Co-60 gamma radiation under highly uniform geometry, then subjected to the standard 2-hour 80 °C anneal per the UMLRR transistor dosimetry protocol (based on ASTM E1855-15). The post-anneal HFE measurements reveal a striking and unexpected result: rather than scattering around a single mean residual damage value, the devices split cleanly into **two populations** separated by a factor of 19 in residual $\Delta(1/h)$. Four devices (SN 108, 114, 135, 146) show near-complete recovery (mean $\Delta(1/h) = 0.0026$), while five devices (SN 102, 105, 137, 138, 147) show poor recovery (mean $\Delta(1/h) = 0.0488$).

This bimodal recovery, combined with the magnitude of residual damage in the poor-recovery cluster, has three significant implications for the dosimetry program:

- The standard 2-hour 80 °C anneal does not uniformly remove gamma-induced defects across nominally identical devices. At a gamma dose corresponding to a 1×10^{14} n/cm^2 FNI run, the residual signal misinterpreted as neutron damage averages 22.8% of the true neutron fluence — far higher than the program assumes.
- Pre-irradiation HFE values do not predict which devices will recover well (Pearson $r = -0.06$, $p = 0.88$); the matched-HFE$_0$ screening currently used cannot detect this variability in advance.
- The bimodal recovery pattern is the most plausible mechanism for the wild scatter (−50% under to +300% over) observed in low-power (∼1E11–1E12 n/cm^2) FNI dosimetry runs, where the residual gamma contribution becomes comparable to or larger than the actual displacement-damage signal.

This report documents the experimental setup, calibration basis for the 12,400 rad target dose, full statistical analysis, candidate physical mechanisms, and recommended follow-up actions.

2 Background and Motivation

Routine neutron dosimetry at the UMLRR uses the gain-degradation method per ASTM E1855-15. Transistors are exposed in the Fast Neutron Irradiator (FNI) box, the change in reciprocal HFE is measured before and after irradiation, and the 1 MeV(Si)-equivalent neutron fluence is computed from the Messenger-Spratt damage constant K_n. A 2-hour 80 °C anneal is applied between irradiation and post-irradiation HFE measurement to remove the unstable component of the damage signal (ASTM E1855 §8.1.6, citing Fleming et al. 2010) and isolate the stable displacement-damage contribution attributable to neutrons.

The program produces reproducibility within ±8% at neutron fluences of 1×10^{13} to 3×10^{14} n/cm^2. At lower fluences (1×10^{11} to 1×10^{12} n/cm^2), reproducibility degrades dramatically: individual transistor readings range from approximately 50% under to 300% over the expected fluence based on the reactor power integral and the FNI calibration constant. The scatter is roughly symmetric about the expected value, with batch behavior mixed between tight-but-offset runs and broadly scattered runs.

An initial working hypothesis (the "over-anneal" hypothesis) proposed that the 2-hour 80 °C bake removes a fluence-dependent fraction of the displacement-damage signal at low fluence, where the unstable defect component is a larger share of the total. To establish a baseline for the gamma contribution that survives the standard anneal — independently of any displacement damage — a Co-60 exposure was conducted at the gamma dose corresponding to a 1×10^{14} n/cm^2 FNI run.

3 Experimental Setup

3.1 Device Inventory

Nine 2N2222AUB transistors (TO-18 case, JANTX-grade military spec) were drawn from the UMLRR dosimetry stock. All nine had matched prior-irradiation histories: previously exposed to a single equivalent fluence, then reset using the standard 24-hour 180 °C protocol. Post-reset HFE$_0$ values were measured within an acceptable tolerance band of the original pre-irradiation specification. HFE$_0$ baselines were re-measured the day before the Co-60 exposure (mean = 149.46, $\sigma = 8.07$, range 131.6 to 157.2).

3.2 Co-60 Cave Geometry

The UMLRR Co-60 irradiation facility houses a planar source approximately $12'' \times 12''$. Transistors were arranged in a 3×3 grid spanning a few inches across, centered on the source. The transistor grid spans approximately ±10% of the source linear dimension on either axis, ensuring negligible inverse-square or solid-angle variation across positions. Dose uniformity across the 9 positions is estimated at better than ±2% based on the geometric configuration.

3.3 Dose Determination: 12,400 rad

The target dose of 12,400 rad was selected to match the in-pile gamma dose accompanying a 1×10^{14} n/cm^2 neutron run at the FNI position. The mapping between neutron fluence and accompanying gamma dose in the FNI is determined from a linear fit through four MCNP-derived calibration points (see Figure 1):

$$D_\gamma = 1.24 \times 10^{-10} \cdot \Phi_n \quad \text{[rad per (n/cm}^2\text{)]} \tag{1}$$

The Navy-Cal anchor at $\Phi_n = 1 \times 10^{14}$ n/cm^2 gives $D_\gamma = 12{,}433.77$ rad, rounded to 12,400 rad for the cobalt exposure target. The actual delivered dose was verified using alanine/TLD dosimetry at the exact 3×3 grid position used for the transistors, on the day of the run. The 1 Mrad/hr nominal cave dose rate corresponds to a planned exposure of approximately 45 seconds. Dose-rate differences between in-pile gamma ($\sim 10^2$–10^3 rad/s) and Co-60 cave delivery ($\sim$280 rad/s) are noted; dose-rate sensitivity is expected to be small for the 2N2222AUB (not a known ELDRS-prone device).

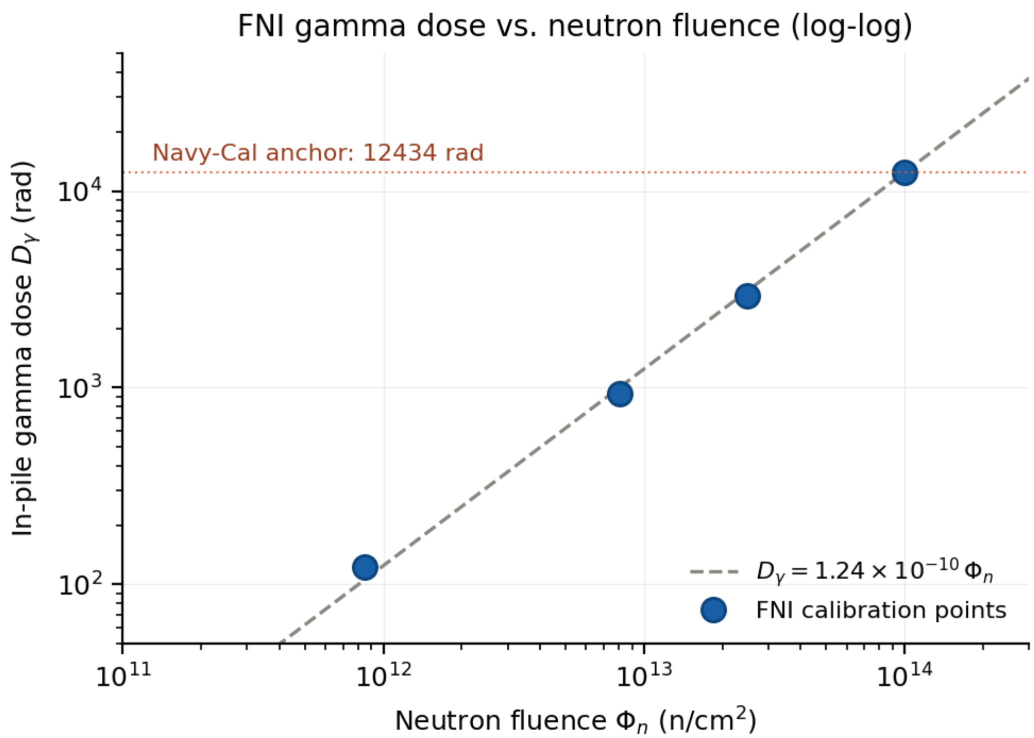

Figure 1. FNI gamma dose vs. neutron fluence. Linear fit through four MCNP calibration points; the Navy-Cal anchor at $\Phi = 1\text{E}14$ gives 12,433.77 rad, which is the target dose for the Co-60 baseline experiment. Power-law exponent fits at 0.967, consistent with linearity.

3.4 Procedure

It is important to note that the DAN calculator's reported "fluence" for a pure gamma exposure is a diagnostic quantity, not a physical neutron fluence. It represents the equivalent neutron fluence that would produce the observed HFE shift if the damage were interpreted as displacement damage. For this experiment, no neutrons are involved; the reported "apparent fluence" is used to directly quantify the magnitude of gamma residual contamination in the standard dosimetry protocol.

Step	Description
1	HFE_0 measurement on all 9 transistors (day before exposure), Keithley 2400/220 at $I_C =$ 2.4 mA, $V_{CE} = 10$ V via GPIB-USB / NI-488.2 v17.0, single reading per device.
2	Devices placed in 3×3 grid, source-centered, in Co-60 cave.
3	Exposure to 12,400 rad (alanine/TLD-verified at the exposure position).
4	Transferred immediately to pre-heated oven; anneal at 80 °C for 2 hours per ASTM E1855-15 §8.1.6.
5	Single post-anneal HFE measurement per device using identical Keithley protocol. No pre-anneal HFE measurement was performed (current-injection contamination of unstable defect populations).
6	$\Delta(1/h)$ computed per device. Apparent neutron-equivalent fluence reported by the DAN fluence calculator using $K_n = 1.24 \times 10^{-15}$ cm^2.

4 Results

4.1 Per-Device Measurements

SN	HFE_0	HFE post	$\Delta(1/h)$	Apparent Φ (n/cm^2)	Cluster
114	131.64	111.61	0.0014	1.10E+12	Good
135	155.88	107.77	0.0029	2.31E+12	Good
146	157.19	107.84	0.0029	2.35E+12	Good
108	156.56	104.12	0.0032	2.59E+12	Good
105	152.29	58.07	0.0107	8.59E+12	Poor
147	145.97	30.30	0.0261	2.11E+13	Poor
137	144.34	16.62	0.0532	4.29E+13	Poor
138	152.01	12.26	0.0750	6.05E+13	Poor
102	149.27	11.68	0.0789	6.36E+13	Poor

Table 1. Per-device measurements sorted by post-anneal $\Delta(1/h)$.

4.2 Descriptive Statistics

4.3 Distribution Shape and Bimodality

Three lines of evidence support a bimodal interpretation rather than continuous scatter:

- **Visual gap in $\Delta(1/h)$:** the values fall in two well-separated ranges: 0.0014–0.0032 (good cluster) and 0.0107–0.0789 (poor cluster), with no values between 0.0033 and 0.0107 (Figures 2 and 4).
- **Shapiro-Wilk normality test:** rejects normality on linear $\Delta(1/h)$ ($W = 0.792$, $p = 0.0165$). Normality is also rejected on apparent Φ ($W = 0.792$, $p = 0.0165$). The log-transformed apparent fluence is closer to normal ($W = 0.881$, $p = 0.161$) but still consistent with a multiplicative two-population model.

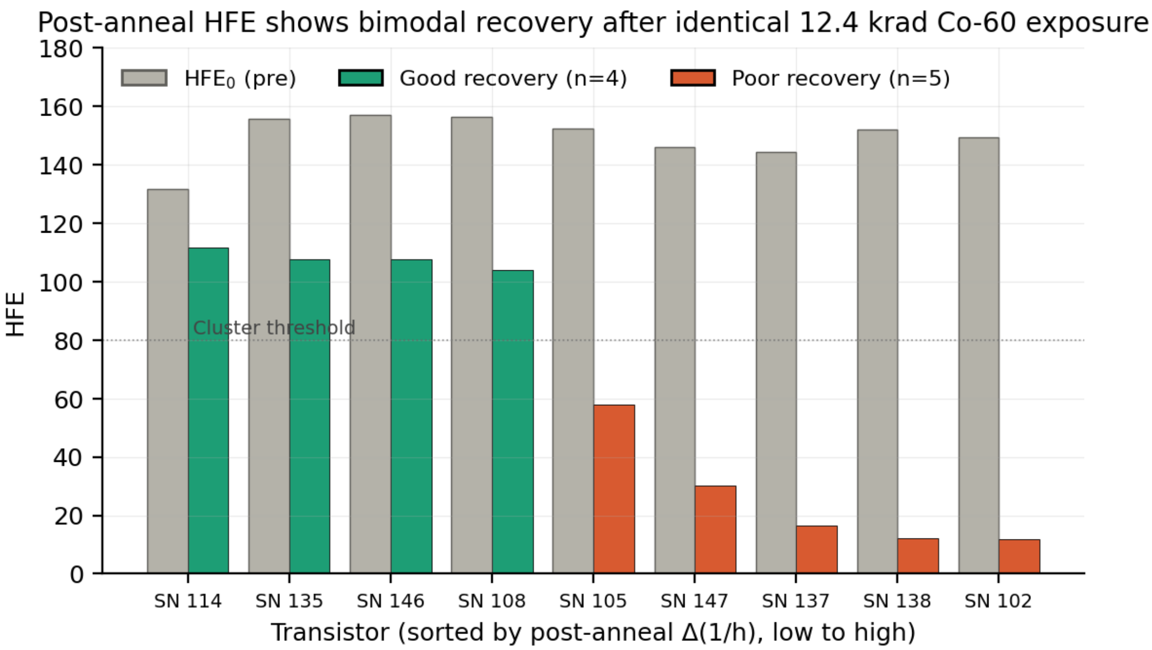

Figure 2. Pre- and post-anneal HFE for all 9 transistors, sorted by post-anneal $\Delta(1/h)$. Pre-irradiation values (gray) cluster tightly; post-anneal values split cleanly into two populations separated by an ~80–100 HFE-unit gap.

- **Bimodality coefficient:** BC = 0.62. Values above ~0.555 suggest a bimodal distribution (Pfister et al. 2013). Sample size limits statistical power, but the cluster ratio (19×) is far larger than any plausible unimodal variance interpretation.

4.4 Two-Sample Tests Between Clusters

All three tests confirm the two clusters are statistically distinct populations rather than tails of a single distribution. Levene's significance also indicates the variance itself differs between clusters (poor cluster $\sigma = 0.0299$ vs. good cluster $\sigma = 0.0008$, a 36× ratio).

4.5 HFE_0 as a Predictor (or Not)

An obvious quality-screening question is whether pre-irradiation HFE_0 predicts which cluster a transistor will fall into. If so, the program could pre-screen devices and reject those destined for the poor-recovery cluster. The data argue strongly against this:

5 Candidate Physical Mechanisms

Several mechanisms could plausibly produce bimodal post-anneal recovery in identically-irradiated devices. The following hypotheses are ordered by current best estimate of likelihood, given what we know about the device population and the experimental controls.

Metric	Value	Notes
n	9	
Mean apparent Φ	2.279E+13 n/cm^2	Across all 9 devices
Std dev (σ)	2.602E+13	95% CI: 1.76E+13 to 4.98E+13
σ/μ	114.2%	Coefficient of variation
Median	8.592E+12	
IQR	4.059E+13	$Q_1 = 2.35\text{E}+12$, $Q_3 = 4.29\text{E}+13$
Min / Max	1.10E+12 / 6.36E+13	Ratio = 57.9×
% of 1E14 reference	22.79%	Gamma residual as fraction of full FNI signal

Table 2. Descriptive statistics for the gamma-only batch.

Test	Statistic	p-value	Result
Welch's t-test (means)	$t = -3.91$	0.0173	Cluster means differ
Mann-Whitney U	$U = 0$	0.0159	Distributions differ
Levene's test (variances)	$W = 7.61$	0.0282	Variances differ

Table 3. Two-sample tests comparing good-recovery and poor-recovery clusters.

5.1 H1: Residual oxide trap inventory from prior irradiation history

Mechanism: The 24-hour 180 °C reset anneal recovers approximately 70% of displacement damage from prior irradiation (temperature-accessible neutron-induced defects only). It does *not* fully restore oxide trapped charge or interface state density. Devices that accumulated different oxide damage during their prior life will now respond differently to new gamma exposure: those with higher pre-existing oxide trap density saturate more gamma defects in deeper traps not accessible to the 2-hour 80 °C bake.

Predictions and consistency with data:

- Predicts bimodal recovery if the prior-history distribution is bimodal (e.g., some devices were used in the program longer than others before this batch). *Consistent with observation.*
- Predicts no correlation with HFE_0 (oxide damage and displacement damage are governed by different physics, and the reset anneal restores HFE_0 via different mechanisms). *Consistent with observation* ($r \approx 0$).
- Predicts no correlation with position in the 3×3 grid (history is a device property, not a position property). *Cannot be tested without position-to-SN logging from the run.*
- **Test:** Repeat the Co-60 exposure with verified-fresh (never-irradiated) 2N2222AUB devices. If fresh devices show tight, single-cluster recovery, H1 is supported. If they bimodally cluster, H1 is contradicted.

5.2 H2: Intrinsic oxide quality variation within the manufacturing lot

Mechanism: Even within a single manufacturing lot, oxide thickness and interface state density vary device to device. Devices with intrinsically higher interface state density retain more gamma-induced

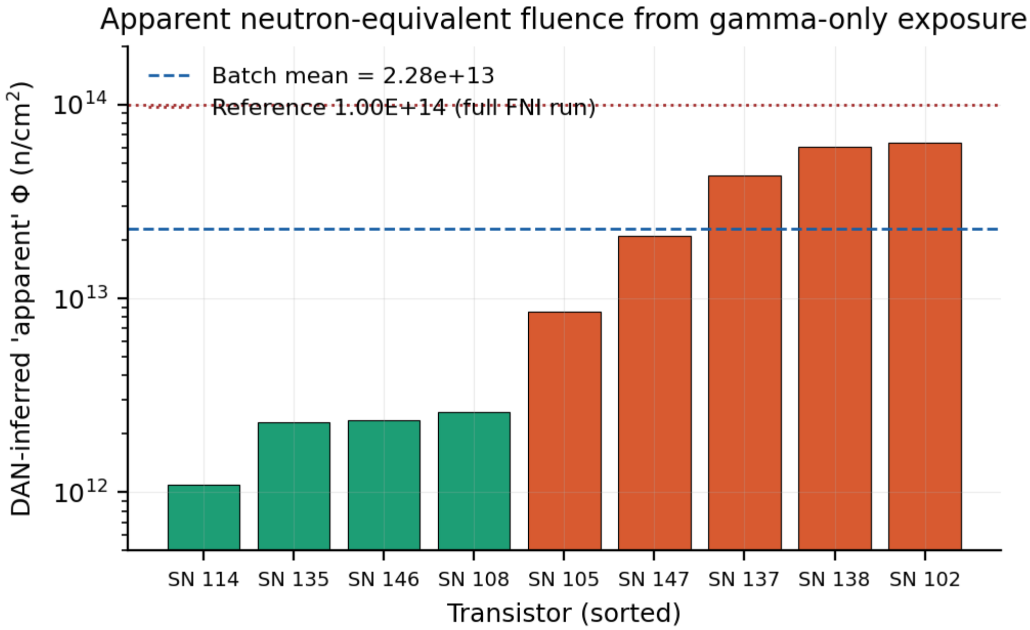

Figure 3. Apparent neutron-equivalent fluence per device (log scale). Even the best-recovered devices (green) show $\sim 10^{12}$ n/cm^2 of apparent-neutron signal from pure gamma; the worst (orange) approach 10^{14}, indistinguishable from a full 1E14 FNI run.

Correlation	**Coefficient**	***p*-value**
Pearson $r(\mathrm{HFE}_0, \Delta(1/h))$	-0.057	0.885
Pearson $r(\mathrm{HFE}_0, \log\Phi)$	$+0.028$	0.943
Spearman $\rho(\mathrm{HFE}_0, \Delta(1/h))$	-0.117	0.765

Table 4. Correlation tests between HFE_0 and post-anneal residual damage.

defects post-anneal regardless of prior history.

Predictions and consistency with data:

- Predicts bimodal recovery if the lot has a bimodal oxide quality distribution. *Plausible but harder to justify a priori* — lot-internal distributions are typically unimodal with Gaussian-like tails. Would require specific manufacturing process split or screening rejection that left a residual bimodal stock.
- Predicts no correlation with HFE_0. *Consistent.*
- **Test:** Same as H1 (fresh-device experiment). If fresh devices show bimodal clustering, H2 is supported over H1.

5.3 H3: 2-hour anneal duration is insufficient at this temperature

Mechanism: The 2-hour 80 °C protocol may be on the kinetics elbow for gamma-defect recovery in some devices. Small device-to-device variations in defect activation energy (e.g. due to slight

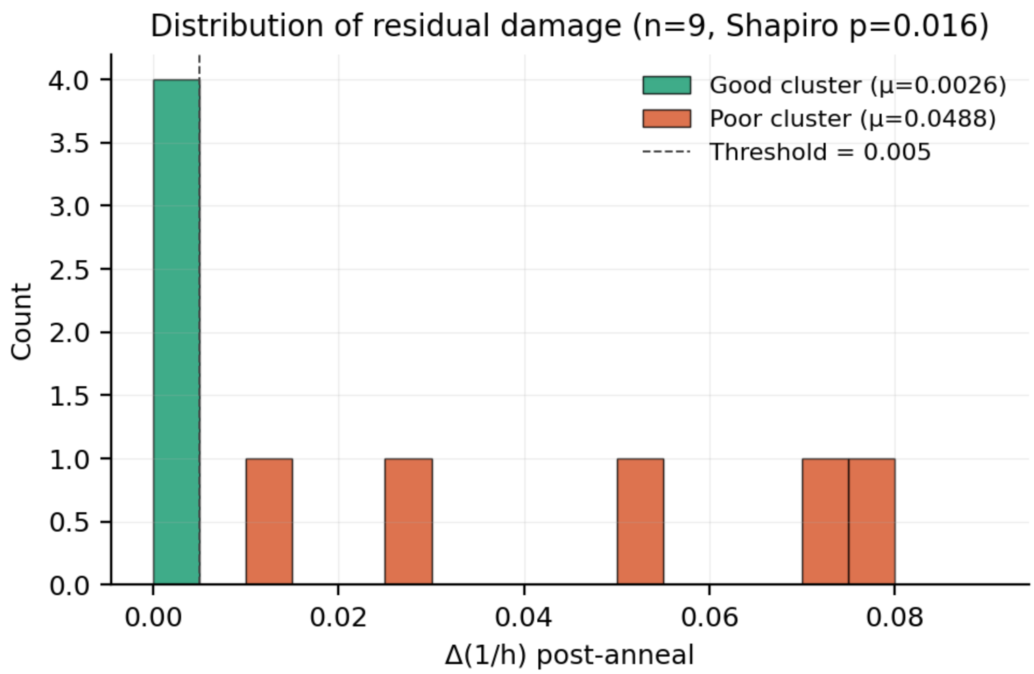

Figure 4. Histogram of post-anneal $\Delta(1/h)$ with cluster overlay. The vertical dashed line marks the empirical threshold (0.005) separating the two populations. No values fall in the gap between 0.0033 and 0.0107.

differences in interface state pinning energies) could push some devices to nearly complete recovery and others to only partial recovery within the same protocol time.

Predictions:

- Predicts that extending the anneal duration would collapse the poor cluster toward the good cluster.
- **Test:** Repeat the experiment with 4-hour or 6-hour anneal at the same temperature. If the poor cluster collapses, H3 is supported. If both clusters persist, H3 is contradicted.

5.4 H4 (rejected): Dose non-uniformity in the cobalt cave

Mechanism: If the 3×3 grid spans a region of significant dose gradient, devices at different positions receive different actual doses, producing scatter in residual $\Delta(1/h)$.

Status: REJECTED. The 3×3 grid spans approximately $\pm 10\%$ of the source linear dimension on either axis, far from the cave edges. Inverse-square and solid-angle effects are negligible at this scale (estimated dose uniformity $>98\%$). Furthermore, a $58\times$ range in apparent fluence cannot be produced by a geometric gradient that would have to be near ± 3 orders of magnitude across a few inches — physically impossible for the cave configuration. Even a $2\times$ dose variation across positions could not produce the observed cluster separation ($19\times$).

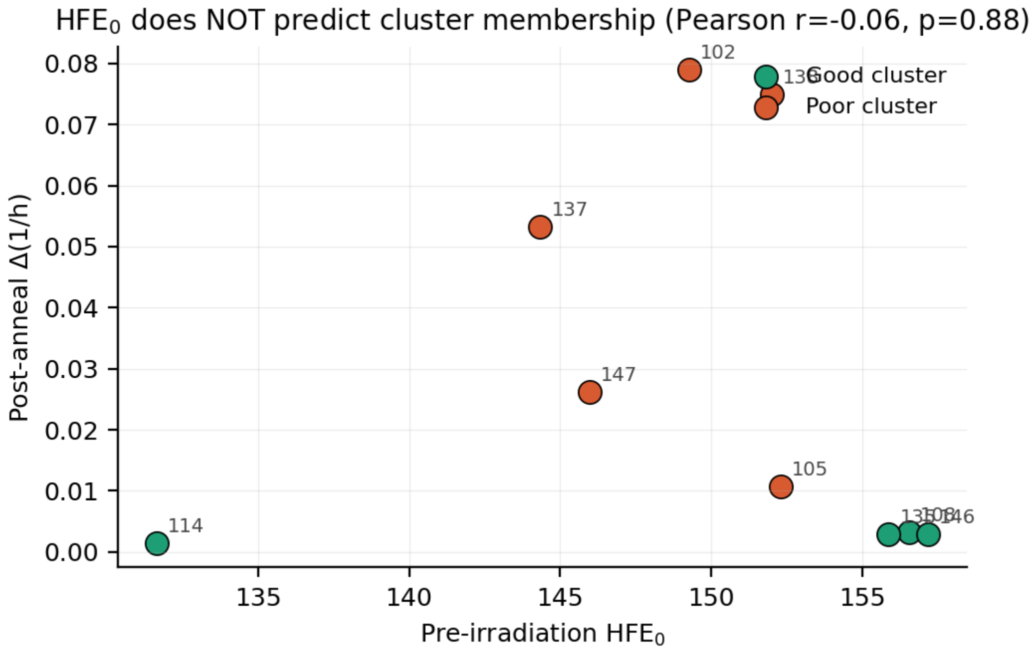

Figure 5. Pre-irradiation HFE_0 versus post-anneal $\Delta(1/h)$. The two clusters overlap in HFE_0 space. Pearson r is essentially zero and not significant, indicating that HFE_0 screening cannot identify devices likely to recover poorly.

5.5 H5 (rejected): Measurement instrumentation artifact

Status: REJECTED. The same Keithley 2400/220 setup, same GPIB-USB adapter, same NI-488.2 v17.0 driver, same `uml_dos_rev3.py` configuration, and same operator measured all 9 devices in a single sitting. Instrumentation drift cannot produce a 19× cluster ratio across an hours-long measurement session, nor produce a discrete bimodal split rather than continuous drift.

6 Implications for the FNI Dosimetry Program

6.1 Quantifying the gamma contamination at high fluence

The mean apparent fluence from pure gamma exposure at 12,400 rad is 2.28×10^{13} n/cm^2, which is 22.8% of the 1×10^{14} n/cm^2 reference. This means that at a full FNI 1E14 run, approximately one-quarter of the inferred neutron fluence may actually be residual gamma damage that the standard anneal failed to remove. This is far larger than the program has historically assumed (typically <5% gamma residual).

However, this estimate carries important caveats. The mean is heavily influenced by the poor-recovery cluster ($n = 5$); the good-recovery cluster shows residuals around 2×10^{12} n/cm^2 (∼2% of 1E14), which is consistent with the program's historical assumption. If the device population can be screened to consist entirely of good-recovery devices, the standard protocol may remain valid for high-fluence work. If not, all reported fluences carry an additional ∼5–25% uncertainty from gamma residual.

6.2 Mechanism for the low-fluence scatter problem

At low neutron fluence (1E11–1E12 n/cm^2), the in-pile gamma dose is correspondingly low (12–124 rad) but it does not scale away the bimodal recovery problem. The *fraction* of gamma residual that survives the anneal is approximately the same. What changes is the ratio of gamma residual to true neutron signal:

FNI Φ_n	**In-pile D_γ**	**Good-cluster gamma residual (apparent Φ)**	**Ratio to true Φ**
1E14 n/cm^2	12,400 rad	∼2E12 n/cm^2	∼2%
1E13 n/cm^2	1,240 rad	∼2E11 n/cm^2	∼2%
1E12 n/cm^2	124 rad	∼2E10 n/cm^2	∼2%
		Poor-cluster gamma residual (apparent Φ)	**Ratio to true Φ**
1E14 n/cm^2	12,400 rad	∼6E13 n/cm^2	∼60%
1E13 n/cm^2	1,240 rad	∼6E12 n/cm^2	∼60%
1E12 n/cm^2	124 rad	∼6E11 n/cm^2	∼60%

Table 5. Projected gamma residual contamination by fluence and cluster. Scaling assumes the fractional recovery is gamma-dose-independent at the protocol's validated range. This assumption is untested below ∼1 krad and is itself a potential confound.

If both clusters exist in the device population in roughly equal proportions (as in this $n = 9$ sample: 4 vs 5), then at low fluence:

- Good-cluster devices read approximately their true neutron fluence with a small (∼2%) gamma offset.
- Poor-cluster devices read their true neutron fluence *plus* a gamma residual that may be 60% of the in-pile gamma dose, which at low fluence is comparable to or larger than the true neutron signal.

This produces exactly the pattern observed in field data: approximately half the devices read close to truth, half read wildly high. The mix-and-match between batches reflects which cluster happens to dominate any given small-N batch. The symmetric $-50\%/+300\%$ scatter pattern is consistent with (true + zero) and (true + large-positive) admixtures.

6.3 Why the problem is invisible at high fluence

At 1×10^{14} n/cm^2, the true displacement-damage signal produces ΔHFE on the order of 100–150 units in a 2N2222AUB with $HFE_0 \approx 150$. The poor-cluster gamma residual contributes an additional ∼60% offset, producing apparent fluences of 1.6E14 instead of 1.0E14. This is a 60% systematic high bias, not noise — but because batches are typically averaged across devices and normalized to a calibration constant K_n derived from these *same* mixed-cluster devices, the systematic offset has been folded into K_n itself. The program is self-consistent at high fluence even though the underlying assumption (no gamma residual) is wrong.

7 Recommended Follow-Up Actions

7.1 Immediate next experiment: fresh-device Co-60 baseline

Procure at minimum 9 (preferably 12–15) 2N2222AUB transistors with verified zero prior irradiation history, ideally sequential serial numbers from a single procurement lot. Repeat the 12,400 rad Co-60 exposure and 2-hour 80 °C anneal protocol used in this experiment. Analyze the resulting $\Delta(1/h)$ distribution.

- If fresh devices show tight, single-cluster recovery near the current good-cluster mean, H1 is confirmed: device history is the root cause, and the program should consider fresh-device rotation or a more aggressive reset protocol.
- If fresh devices also show bimodal clustering, H2 is supported and the issue is intrinsic to the part. The program needs to develop a pre-screening test or accept the bimodal variance as a fundamental limit.

7.2 Anneal-duration sensitivity test

Take 6 additional devices from current stock, expose to 12,400 rad, and anneal at 80 °C for 4 hours instead of 2. Compare the $\Delta(1/h)$ distribution to the current data. If the poor cluster collapses toward the good cluster, the protocol can be fixed by extending anneal time. If both clusters persist, the poor-cluster devices have defects inaccessible at 80 °C and either a higher-temperature anneal or device screening is needed.

7.3 Historical-data re-analysis

Pull historical 1E13–3E14 dosimetry data from the program archive (50+ data points should suffice) and examine the distribution of inferred fluences relative to nominal. If the bimodal-recovery hypothesis is correct, the historical distribution should show either: (a) a fat low tail (good-cluster devices undercutting the mean) and a thin high tail (poor-cluster devices overshooting), or (b) two-peak structure in the residuals from the calibration line. Either would corroborate the bimodal mechanism with substantially larger N than the current 9-device baseline.

7.4 Cancel or restructure the planned 24-device Option B experiment

The previously planned 12×2 reactor experiment (Run 1 at 2 hours, Run 2 at shorter anneal time) was designed to test the "over-anneal" hypothesis at 1E12 n/cm^2. Given the bimodal-recovery finding, that experiment is **unlikely to be interpretable**: the within-batch variance will be dominated by the cluster split, not by anneal-time effects. The cluster ratio ($\sim 19\times$) will swamp any anneal-time signal at $n = 12$. Restructure as follows:

- **Option B-revised:** First, pre-characterize each of the 24 devices via a low-dose Co-60 exposure (e.g. 1,240 rad) and 2-hour anneal to identify each device's cluster membership. Then assign devices to Run 1 and Run 2 such that each batch has an equal mix of good and poor cluster devices. The anneal-time test then becomes a within-cluster comparison with the cluster as a documented covariate.

- **Option C:** Procure fresh devices first, then run the FNI experiment with verified-history-clean stock. The data will be much cleaner and the over-anneal hypothesis (or its absence) will be cleanly testable.

7.5 Documentation update

Once the fresh-device experiment is complete and a mechanism is identified, the UMLRR transistor dosimetry procedure document should be updated to reflect (a) the recommended device screening or pre-characterization protocol, (b) the validated anneal protocol, and (c) the gamma-residual uncertainty budget to be folded into reported fluence uncertainties. The capstone report and any published documentation referencing the FNI calibration should be reviewed for affected results.

8 Limitations of This Analysis

Several limitations should be acknowledged when interpreting these results:

- **Sample size.** $n = 9$ is small for definitive distribution identification. The bimodality coefficient (BC = 0.62) exceeds the unimodal threshold but the Shapiro-Wilk p-value on linear $\Delta(1/h)$ (0.017) is significant at $\alpha = 0.05$ yet not robust against multiple-testing concerns. A larger sample ($n \geq 20$) would be needed to establish bimodality at high confidence. However, the observed cluster ratio is too large to be unimodal variance under any plausible distribution.
- **Device history matching.** "Matched" was defined here as identical prior Φ and identical reset protocol, with HFE_0 in an acceptable band. It is possible that history matching was imperfect (e.g. different reset cycles for different devices in their lifetime), introducing the bimodal pattern via a confounded variable. The fresh-device experiment will resolve this.
- **Position-to-SN mapping.** The positions within the 3×3 grid were not recorded against serial numbers in this run. While dose uniformity is excellent, this prevents a post-hoc check that cluster membership is uncorrelated with position. Future Co-60 runs should log the assignment.
- **DAN apparent-fluence interpretation.** The K_n used by DAN to compute apparent fluence from pure-gamma $\Delta(1/h)$ is the neutron damage constant. Treating gamma damage as if it were neutron damage produces a diagnostic quantity, not a physically meaningful fluence. The numerical values reported as apparent Φ should be interpreted as "the displacement-damage-equivalent signal" rather than a literal neutron fluence.
- **Spectrum and dose-rate differences.** Co-60 gamma energy (mean 1.25 MeV) is harder than the FNI in-pile gamma spectrum (mean $\sim$0.57 MeV from MCNP). Dose rate in the cave ($\sim$280 rad/s) differs from in-pile ($\sim 10^2$–10^3 rad/s). For 2N2222AUB these effects are expected to be small (<10%), but they have not been independently characterized at UMLRR.

9 Conclusions

Nine matched-history 2N2222AUB transistors exposed to 12,400 rad of Co-60 gamma radiation and processed through the standard UMLRR 2-hour 80 °C dosimetry anneal show statistically significant bimodal recovery. Four devices recover to a residual $\Delta(1/h)$ of 0.0026 (good cluster);

five devices recover to 0.0488 (poor cluster), a 19× ratio. The split is not explained by dose non-uniformity, instrumentation, or pre-irradiation HFE_0. The most plausible mechanism is residual oxide trap inventory from prior irradiation history that survives the 24-hour 180 °C reset, producing device-dependent gamma defect response.

This finding is the most plausible explanation for the wild scatter (−50% to +300%) observed in low-power FNI dosimetry runs at 1E11–1E12 n/cm^2: the residual gamma contamination, which is a small relative offset at high neutron fluence, becomes comparable to or larger than the true neutron signal at low fluence, and its bimodal character produces the observed batch-level variance pattern.

The recommended path forward is a fresh-device Co-60 baseline to distinguish H1 (history-dependent) from H2 (intrinsic to part), followed by an anneal-duration sensitivity test to evaluate H3 (kinetics-limited recovery). The planned 24-device FNI Option B experiment should be paused until the device population is characterized; otherwise the bimodal variance will swamp the anneal-time signal it is designed to detect.

References

ASTM E1855-15, *Standard Test Method for Use of 2N2222A Silicon Bipolar Transistors as Neutron Spectrum Sensors and Displacement Damage Monitors*, ASTM International, 2015.

Messenger, G. C. & Spratt, J. P., *The Effects of Neutron Irradiation on Germanium and Silicon*, Proc. IRE 46(6), 1958.

Fleming, R. M., Lang, D. V., Seager, C. H., Bielejec, E., Campbell, J. M., *Effects of high pressure on bistable defects in irradiated silicon*, J. Appl. Phys. 108, 063716 (2010).

MIL-STD-883 TM 1017.3, *Test Method Standard, Microcircuits, Method 1017.3: Neutron Irradiation*, Department of Defense.

MIL-STD-883 TM 1019, *Ionizing Radiation (Total Dose) Test Procedure*.

Pfister, R., Schwarz, K. A., Janczyk, M., Dale, R., Freeman, J. B., *Good things peak in pairs: a note on the bimodality coefficient*, Frontiers in Psychology 4, 700 (2013).

Owusu, F., *MCNP6 Full Core Model of the UMass Lowell Research Reactor*, MS thesis, UMass Lowell, 2019.