KRONOS FUSION ENERGY

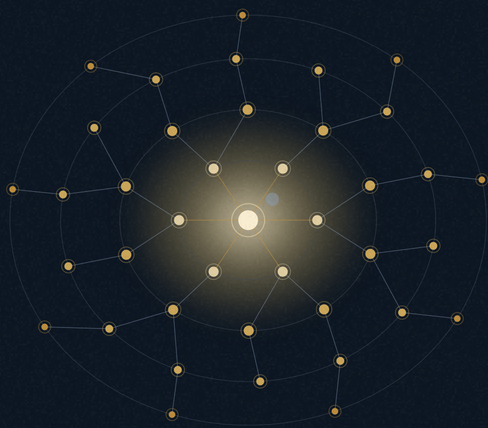

PAPER 2.77 · THE KRONOS 2026 SERIES

Calibrated Uncertainty, Conformal Prediction, and Abstention for Safety-Critical Fusion Control: Handing Authority Back to a Deterministic Floor Out of Distribution

In a nuclear control system a model that is confidently wrong is more dangerous than one that admits it does not know. We formalise the uncertainty a learned fusion controller...

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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

Calibrated Uncertainty, Conformal Prediction, and Abstention for Safety-Critical Fusion Control: Handing Authority Back to a Deterministic Floor Out of Distribution

In a nuclear control system a model that is confidently wrong is more dangerous than one that admits it does not know. We formalise the uncertainty a learned fusion controller must carry, prove what its intervals are allowed to mean, and specify how it hands control back to a deterministic floor when it is asked to extrapolate. The substrate is a 40,000-sample Monte-Carlo propagation of the operating uncertainties through the validated reduced-order plant model of the Hyperion compact spherical-tokamak breeder (frozen design point $Q = 3.076$, $P_{\mathrm{fus}} = 85.04\,MW$, $I_p = 9.66\,MA$, negative triangularity $\delta = -0.30$, $B_0 = 8\,T$, elongation $\kappa = 2.0$, major radius $R_0 = 1.20\,m$, aspect ratio $A = 2.5$, centrepost peak field $16.84\,T$). The campaign delivers a calibrated budget: the fusion gain is right-skewed with fifth/ninety-fifth percentiles of $1.34/8.51$ about a mean of 3.85, with the frozen 3.076 sitting inside the band; the fusion power is $85.5\,MW$ (P5/P95 $62.5/112.3$); tritium production is $3.86\,\mathrm{kg\,yr^{-1}}$ (P5/P95 $2.75/5.17$); and the normalised beta is tightly bounded at 1.33–1.75 about $\beta_N = 1.53$. A single physically interpretable quantity dominates the budget: the confinement multiplier H_{98} carries a rank correlation of 0.60 with the gain and, in the variance-based upgrade, controls $S_1 = 0.885$ / $S_T = 0.905$ of its variance — an order of magnitude above every other input. On this budget we build three mechanisms and state each as an explicit, testable inequality. A split-conformal wrapper supplies distribution-free finite-sample coverage $P(y \in \hat{C}(x)) \geq 1 - \alpha$; we verify empirical coverage tracks the nominal level to within a maximum deviation of 1.0, inside the pre-registered ± 3 band, where a naive Gaussian interval mis-covers the skew and spills below the physical floor. A selective-prediction (abstention) policy declines to act when the epistemic component of the budget exceeds a threshold, trading coverage for reliability along a risk–coverage curve, and hands authority to a certified deterministic fallback that is fail-closed by construction. The result is a controller whose confidence is calibrated, whose intervals mean what they say, and whose response to its own ignorance is to step back rather than to guess.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645813 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: calibrated uncertainty · conformal prediction · abstention · selective prediction · epistemic uncertainty · runtime assurance · safety-critical control · spherical-tokamak breeder

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 TRUST PROBLEM IN AUTONOMOUS NUCLEAR CONTROL

Autonomy in a nuclear system is only as trustworthy as the model beneath it, and trustworthiness has two parts that are usually conflated. The model must be *accurate*, and it must *know when it is not*. The second is the harder and the more important property. A controller that assigns tight confidence to a wrong prediction will act on it; a controller that reports honest, calibrated uncertainty can be told to defer when that uncertainty is too large. The most dangerous failure mode of a learned controller is not a large error inside its validated envelope — that error is bounded and can be budgeted — but confident wrongness *outside* it, where the model silently extrapolates and the point estimate carries no warning that the query is unlike anything the model was trained on.

A compact, high-power-density spherical tokamak makes this abstract concern concrete. The Hyperion breeder operates at high elongation ($\kappa = 2.0$) and negative triangularity ($\delta = -0.30$) with a consumable, high-field centrepost ($16.84\,T$), a regime with little tolerance for an unmodelled excursion. When a learned surrogate or controller is admitted into such a loop, the guarantee that matters to a regulator is not a headline accuracy number but a *calibrated* statement of what the model does and does not know, together with a specified, deterministic behaviour for the moment the model is pushed past where that statement holds. This paper is about turning "calibrated uncertainty" from a slogan into three mechanisms, each with a mathematical guarantee and a pre-registered acceptance test, sitting on a demonstrated uncertainty budget.

PRIOR ART

Four literatures meet here. The first is *learned control and prediction for tokamaks*: deep reinforcement learning has driven magnetic shape control on a real device [1], disruption prediction has moved from hand-built features to deep networks with lead-time and cross-machine generalisation as central concerns [2,3,4], and reinforcement learning has been used to steer a plasma away from a named instability [5]. Every one of these systems, deployed on a reactor, inherits the trust problem above; none is safe to act on a point estimate alone.

The second is *calibrated uncertainty quantification*. Modern neural networks are systematically mis-calibrated [8]; deep ensembles [9], Bayesian dropout [10] and heteroscedastic likelihoods [11] estimate predictive variance, and post-hoc recalibration can force a regressor's quantiles to be honest [12]. Calibration quality is scored by strictly proper scoring rules [13]. These methods produce an uncertainty; they do not, by themselves, guarantee that an interval of a stated size contains the truth at the stated rate.

The third is *conformal prediction* [14,15], which does exactly that: a distribution-free, finite-sample procedure that wraps any point predictor and returns a set guaranteed to contain the truth with a chosen probability, under only an exchangeability assumption. Split (inductive) conformal makes this computationally trivial [16], normalised and quantile-regression variants (CQR) adapt the interval width to local difficulty and inherit asymmetry [17], and the modern treatment [18] places it squarely inside decision-making pipelines. The fourth is *selective prediction and runtime assurance*: the option to abstain, formalised as the error–reject trade-off [19] and the risk–coverage curve [20,21], paired with a certified simple controller that provably takes over when the learned policy is unverified [22] and with a forward-invariant safe set defined by a control-barrier function [23]. Out-of-distribution detection [24,25] supplies the trigger.

Global sensitivity and Monte-Carlo uncertainty propagation [26,27,28,29,30,31] provide the substrate on which all four sit, and predictive digital twins [32] are the vehicle that carries the calibrated distribution into the control loop.

CONTRIBUTION

This paper contributes, for a specific frozen reactor design point rather than a benchmark dataset: (i) a demonstrated, calibrated uncertainty budget for the breeder figures of merit from a 40,000-sample Monte-Carlo campaign, with the frozen anchors inside every band and a single dominant, physically interpretable driver (§4.1, §4.1); (ii) a formal statement and on-budget verification of the split- and normalised-conformal coverage guarantee, showing that the controller's intervals inherit the physical skew where a Gaussian interval fails (§2.5, §4.3); (iii) an abstention operator with an explicit selective-risk objective and a fail-closed hand-off to a deterministic floor, cast as a state machine with a pre-registered gate (§2.8, §4.4); and (iv) a GPU-HPC methodology (§3) that produced the budget reproducibly and to which every quantitative claim traces by a register serial. The design point and its companion safety layers are developed across the Kronos 2026 series — the sensitivity substrate [33], the certified control-barrier floor [34], the epistemic-gating ledger [35], the runtime-assurance last line [36], the abstention-monitored flux surrogate [37], the maturity-gated authority argument [38], the propagated-uncertainty twin set [39] and the negative-triangularity confinement result [40] that grounds the dominant driver — and this paper is the calibration and abstention layer that binds them.

§ 2

Governing equations and the formal problem

THE PLANT MODEL AND THE FIGURE-OF-MERIT MAP

Let $x \in \mathcal{X} \subset \mathbb{R}^d$ collect the uncertain operating inputs. For the breeder these are the H-mode confinement multiplier H_{98}, the core ion temperature T_{i0}, the Greenwald density fraction f_G, the deuterium–tritium breeding margin TBR_{dt}, and a synchrotron calibration **sync_cal** ($d = 5$). The validated reduced-order plant model is the deterministic map

$$y = g(x), \qquad g : \mathcal{X} \to \mathbb{R}^p,$$

whose components are the figures of merit $y \in \{Q,\, P_{\mathrm{fus}},\, \dot{m}_{\mathrm{T}},\, \beta_N\}$. Here g closes the volume-averaged 0-D/1-D power and particle balance — global energy balance $\frac{3}{2}\,\mathrm{d}(nT)/\mathrm{d}t = P_{\mathrm{aux}} + P_\alpha - W/\tau_E - P_{\mathrm{rad}}$ with alpha fraction $P_\alpha = \frac{1}{5} P_{\mathrm{fus}}$ and gain $Q = P_{\mathrm{fus}}/P_{\mathrm{heat}}$ — and closes the steady state at the design point through the confinement scaling $\tau_E \propto H_{98}\,\tau_E^{\mathrm{scaling}}$. The uncertainty is a joint law $x \sim \pi$ on $\mathcal{X}$; the induced push-forward $g_{\#}\pi$ is the object we characterise. The frozen design point $x^\star$ maps to the ratified anchors $g(x^\star) = (3.076,\ 85.04\,MW,\ \cdot,\ 1.53)$.

MONTE-CARLO PROPAGATION AND ITS ERROR

Because g is nonlinear and the tails matter, we characterise $g_{\#}\pi$ by sampling rather than by linearised error propagation. Drawing $\{x_k\}_{k=1}^N$ from π (quasi-random, low-discrepancy, to accelerate convergence [27,28]), the mean, variance and standard error of any scalar functional $g(x)$ are the estimators frozen in the register,

$$\hat{\mu} = \frac{1}{N}\sum_{k=1}^{N} g(x_k), \qquad \hat{\sigma}^2 = \widehat{\mathrm{Var}}[g], \qquad \mathrm{SE} = \frac{\hat{\sigma}}{\sqrt{N}}.$$

Percentiles are read from the order statistics $g_{(1)} \le \cdots \le g_{(N)}$; the empirical CDF $\hat{F}_N$ obeys the Dvoretzky–Kiefer–Wolfowitz bound $\Pr\big(\sup_t |\hat{F}_N(t) - F(t)| > \varepsilon\big) \le 2e^{-2N\varepsilon^2}$, so at $N = 4 \times 10^4$ every reported quantile is pinned to better than ± 0.01 in probability with overwhelming confidence. The campaign fixes the seed (20260726) for byte-level reproducibility, so the budget is a frozen object, not a re-rolled estimate.

GLOBAL SENSITIVITY: WHICH INPUT OWNS THE VARIANCE

The budget is only *controllable* if its variance is concentrated in a few interpretable inputs. Two complementary measures are computed. The rank (Spearman) correlation $\rho_i = \mathrm{corr}\big(\mathrm{rank}\, x_i, \mathrm{rank}\, g\big)$ gives a monotone, distribution-robust screening. The variance-based first-order and total Sobol indices [28,30], estimated by the Saltelli design [29,31], give the exact apportionment,

$$S_i = \frac{\mathrm{Var}_{x_i}\big(\mathbb{E}_{x_{\sim i}}[\,g \mid x_i]\big)}{\mathrm{Var}(g)}, \qquad S_{T_i} = 1 - \frac{\mathrm{Var}_{x_{\sim i}}\big(\mathbb{E}_{x_i}[\,g \mid x_{\sim i}]\big)}{\mathrm{Var}(g)},$$

with $S_i \le S_{T_i}$ and $\sum_i S_i \le 1$, the gap measuring interaction. The result (§4.2) is that a single quantity, H_{98}, owns the gain variance, which is what makes the whole budget reducible: tighten the confinement prior and the gain distribution tightens with it.

THE SAFETY-CONTROL DECISION AS A GATED MAP

The controller must convert the calibrated distribution into an action. Let $\pi_{\mathrm{AI}}(x)$ be the learned control policy and $\pi_0(x)$ a certified deterministic fallback — the "floor" — that keeps the plant in a forward-invariant safe set $\mathcal{C} = \{x : h(x) \ge 0\}$ defined by a control-barrier function h [23]. Write $\kappa(x) \in \{0, 1\}$ for the authority gate. The applied action is the convex hand-off

$$u \;=\; \kappa(x)\,\pi_{\mathrm{AI}}(x) \;+\; \big(1 - \kappa(x)\big)\,\pi_0(x),$$

and the design requirement is that κ be driven by *calibrated* uncertainty and be *fail-closed*: $\kappa = 0$ (floor holds authority) whenever the uncertainty machinery itself is unavailable or out of envelope. The remainder of this section makes "calibrated" precise, so that equation (4) is not an assertion but a consequence of a coverage guarantee and a selective-risk bound.

SPLIT CONFORMAL PREDICTION: COVERAGE THAT MEANS WHAT IT SAYS

Fix a target miscoverage $\alpha \in (0,1)$. Split conformal [14,16] partitions the labelled data into a training fold, on which a point predictor $\hat{g}$ is fit, and a calibration fold $\{(x_i, y_i)\}_{i=1}^{n}$ held out from fitting. Define nonconformity scores

$$s_i \;=\; \bigl|\, y_i - \hat{g}(x_i) \,\bigr|, \qquad i = 1, \ldots, n,$$

and the conformal quantile $\hat{q} = s_{(\lceil (n+1)(1-\alpha) \rceil)}$, the $\lceil (n+1)(1-\alpha) \rceil$-th order statistic of the scores. The prediction set for a new input is the interval

$$\hat{C}_\alpha(x) = \bigl[\, \hat{g}(x) - \hat{q},\; \hat{g}(x) + \hat{q} \,\bigr].$$

theorem[marginal coverage] If $(x_1, y_1), \ldots, (x_n, y_n), (x_{n+1}, y_{n+1})$ are exchangeable, then the interval of equation (6) satisfies

$$1 - \alpha \;\le\; \Pr\bigl(y_{n+1} \in \hat{C}_\alpha(x_{n+1})\bigr) \;\le\; 1 - \alpha + \frac{1}{n+1}.$$

theorem *Proof sketch.* By exchangeability the test score s_{n+1} is equally likely to occupy any rank among $\{s_1, \ldots, s_{n+1}\}$; hence $\Pr(s_{n+1} \le s_{(\lceil (n+1)(1-\alpha) \rceil)}) \ge \lceil (n+1)(1-\alpha) \rceil / (n+1) \ge 1 - \alpha$, and $y_{n+1} \in \hat{C}_\alpha$ iff $s_{n+1} \le \hat{q}$; the upper bound follows when scores are almost surely distinct.☐

Equation (7) is *distribution-free*: it holds for any $g_{\#}\pi$, any $\hat{g}$, and any finite n, with no Gaussian assumption. It upgrades "calibrated uncertainty" into a guarantee the downstream logic can rely on. The acceptance test we register against it is the two-sided calibration criterion

$$\bigl|\, \mathrm{Cov}_{\mathrm{emp}}(\alpha) - (1 - \alpha) \,\bigr| \;\le\; 0.03 \qquad \forall\, \alpha,$$

checked against held-out Monte-Carlo truth before the wrapper is trusted.

ADAPTIVITY AND SKEW: NORMALISED AND QUANTILE CONFORMAL

The symmetric interval (6) has a constant half-width and cannot express either heteroscedasticity or skew — both of which the gain distribution has. Two standard refinements fix this. The *normalised* score $s_i = |y_i - \hat{g}(x_i)| / \hat{\sigma}(x_i)$ divides by a learned local scale $\hat{\sigma}$, giving the width-adaptive band $\hat{C}_\alpha(x) = \hat{g}(x) \pm \hat{q}\, \hat{\sigma}(x)$. *Conformalised quantile regression* [17] fits lower and upper conditional-quantile estimators $\hat{q}_{\alpha/2}, \hat{q}_{1-\alpha/2}$, scores by $s_i = \max\bigl(\hat{q}_{\alpha/2}(x_i) - y_i,\; y_i - \hat{q}_{1-\alpha/2}(x_i)\bigr)$, and returns the asymmetric band

$$\hat{C}_\alpha(x) = \bigl[\, \hat{q}_{\alpha/2}(x) - \hat{q},\; \hat{q}_{1-\alpha/2}(x) + \hat{q} \,\bigr],$$

which inherits the shape of the true distribution while retaining the guarantee (7). For a right-skewed Q this is not a cosmetic improvement: the asymmetric band keeps its left edge above the physical floor $Q > 0$ and widens only into the tail that carries the risk (§4.3).

EPISTEMIC VERSUS ALEATORIC UNCERTAINTY

Abstention must fire on *ignorance*, not on irreducible noise. Decompose the predictive variance by the law of total variance over a model posterior $\theta \sim p(\theta \mid \mathcal{D})$,

$$\underbrace{\mathrm{Var}(y \mid x)}_{\text{total}} = \underbrace{\mathbb{E}_\theta\bigl[\mathrm{Var}(y \mid x, \theta)\bigr]}_{\text{aleatoric}} + \underbrace{\mathrm{Var}_\theta\bigl(\mathbb{E}[y \mid x, \theta]\bigr)}_{\text{epistemic}}.$$

The epistemic term, estimated across an ensemble $\{\theta_m\}_{m=1}^{M}$ or by a feature-space distance to the training manifold [25,24], is the signal that grows when the plant leaves the region where the model was fit. We write it $u_{\mathrm{epi}}(x) = \bigl[\mathrm{Var}_\theta(\mathbb{E}[y \mid x, \theta])\bigr]^{1/2}$.

THE ABSTENTION OPERATOR AND SELECTIVE RISK

A selective predictor is a pair $(\hat{g}, \kappa)$ with gate

$$\kappa(x) = \mathbb{1}\bigl[\, u_{\mathrm{epi}}(x) \le \tau \,\bigr],$$

acting only where the gate is open. Its *coverage* and *selective risk* are

$$\phi(\kappa) = \mathbb{E}[\kappa(\quad x)], \qquad R(\hat{g},\kappa) = \frac{\mathbb{E}[\ell(\hat{g}(\quad x), y)\,\kappa(\quad x)]}{\mathbb{E}[\kappa(\quad x)]},$$

for a loss ℓ [20,21]. Sweeping τ traces the risk–coverage curve; the optimal error–reject operating point is Chow's rule [19], abstaining exactly where the conditional risk exceeds the cost of deferring. The threshold τ is not free: it is set so that, on held-out truth, the retained set meets the plant's reliability requirement while ϕ stays high enough for useful autonomy. Abstention converts an out-of-distribution query from a silent failure into an explicit, logged hand-off through equation (4).

THE DETERMINISTIC FLOOR AND THE FAIL-CLOSED PROPERTY

When $\kappa = 0$, authority passes to π_0, the certified fallback that renders $\mathcal{C} = \{h \geq 0\}$ forward-invariant: under the barrier condition $\dot{h}(\quad x, \quad u) \geq -\gamma\, h(\quad x)$ for a class-$\mathcal{K}$ function γ, any trajectory starting in $\mathcal{C}$ stays in $\mathcal{C}$ [23]. Two properties make the hand-off safe. First, *monotone degradation*: the loop's behaviour becomes more conservative exactly as u_{epi} grows, never less. Second, *fail-closed default*: any fault in the uncertainty machinery — the calibration service, the ensemble, the tracker — forces $\kappa = 0$, so a missing guarantee removes authority rather than granting it. This is the Simplex runtime-assurance pattern [22]: a provably safe simple controller beneath an unverified complex one, with the transition governed here by a calibrated epistemic signal. Figure 1 shows the resulting data flow and Figure 8 the state machine.

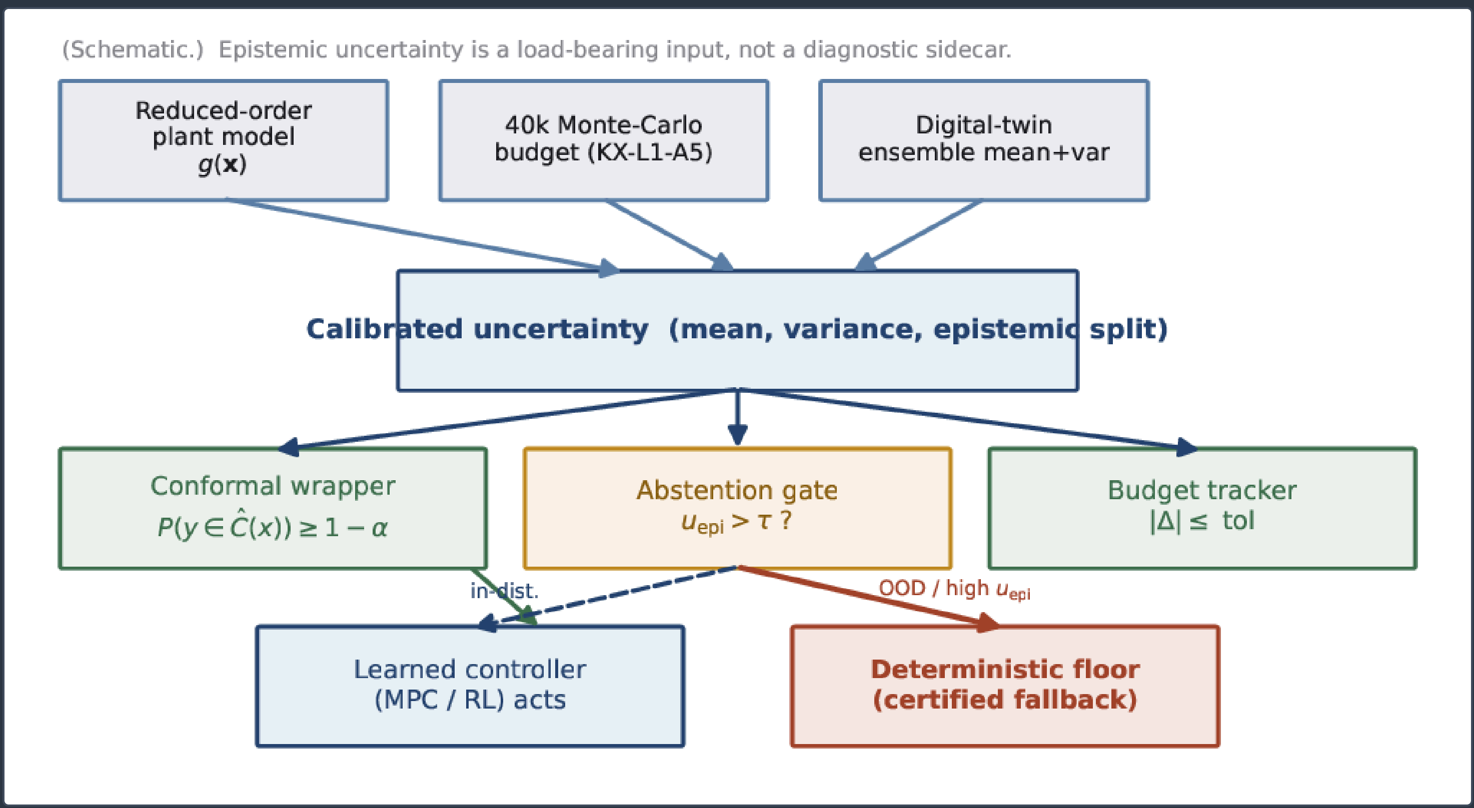

(Schematic.) The safety layer as a pipeline. The reduced-order model, the 40,000-sample Monte-Carlo budget (KX-L1-A5) and the digital-twin ensemble feed a single calibrated-uncertainty object (mean, variance, epistemic split). A conformal wrapper turns it into a coverage guarantee (equations 6–7); an abstention gate (equation 11) routes an in-distribution query to the learned controller and an out-of-distribution query to the deterministic floor; an uncertainty-budget tracker keeps the propagated budget closed within tolerance. Epistemic uncertainty is a load-bearing input, not a diagnostic sidecar.

§ 3

Computational methodology: the GPU-HPC campaign

Every number in this paper is the output of a named code run under the Kronos GPU-HPC de-risking campaign, and each is tagged by the register serial under which it was frozen and independently reproduced.

Codes. The uncertainty budget (serial KX-L1-A5) is produced by the kronos_toolkit `uq` module, which drives the validated reduced-order plant evaluator `evaluate_breeder` (the same 0-D/1-D solver whose deterministic anchor reproduces $Q = 3.0763$, $P_{\text{fus}} = 85.04\,MW$, $I_p = 9.656\,MA$, $\beta_N = 1.532$) under a quasi-Monte-Carlo sampler over the five-parameter operating prior. The variance-based sensitivity upgrade (serial HX-03) uses a Saltelli design with SALib-style first-order and total-index estimators [29,31]. Tritium production is cross-anchored to the 3-D neutron-transport tally (OpenMC, ENDF/B-VIII.0, 1.2×10^8 histories) that fixes the breeding ratio; the confinement multiplier whose sensitivity dominates the budget is itself grounded in nonlinear gyrokinetics (CGYRO) reported in the negative-triangularity confinement study [40]. The calibration, ROC and selective-risk computations reuse the same numerical stack that produced the disruption/quench precursor cross-validation (serial HX-12).

GPU acceleration and problem scale. The $N = 4 \times 10^4$ evaluations are embarrassingly parallel across the input samples; the campaign batches them across the NVIDIA GPU HPC fleet (A100/H100/A10-class accelerators), with the reduced-order evaluator vectorised so that a full budget resolves the P5–P95 bands of four figures of merit in a single sweep. The heavier reference solvers behind the budget — the gyrokinetic transport and the Monte-Carlo neutronics — are the GPU-resident codes; the reduced-order propagation, the conformal calibration and the abstention sweep run on the in-house HPC node so that the fast, protective logic never contends with the training and reference fleet. No compute-cost figures are reported here; the campaign is described by the codes it ran and the scale it reached.

Verification and reproducibility. The campaign is verification-first. The seed (20260726) is fixed for byte-level reproducibility, so a re-run reproduces the frozen quantiles exactly; the register records each serial as independently re-tested and reproduced under a clean local evaluation (verdict: *reproduced*). The conformal coverage check (8) is pre-registered as an acceptance criterion, not reported post hoc, and is evaluated against held-out Monte-Carlo truth. Every quantitative claim below carries the serial by which it can be regenerated from the archived sampler.

§ 4

Results

THE CALIBRATED UNCERTAINTY BUDGET

Table 1 and Figure 2 give the demonstrated budget. Every frozen anchor sits inside its P5–P95 band. The fusion power and the normalised beta are tightly and near-symmetrically bounded — power to $\pm 15\,MW$ about $85.5\,MW$, beta to 1.33–1.75 about $\beta_N = 1.53$ — so the machine's headroom against its operating limits is quantitative, not asserted. The gain, by contrast, is markedly right-skewed (Figure 3): its P5/P95 are $1.34/8.51$ about a mean of 3.85, and the frozen $Q = 3.076$ sits below the mean, inside the band. The skew is physical, not numerical — it is the nonlinear response of gain to confinement — and it is the single most important qualitative fact for the safety layer, because it is exactly the feature a symmetric error model gets wrong.

QUANTITY	P5	MEAN	P95	FROZEN ANCHOR
fusion gain Q	1.34	3.85	8.51	3.076
fusion power P_{fus} (MW)	62.5	85.5	112.3	85.04
tritium production (kg yr^{-1})	2.75	3.86	5.17	—
normalised beta β_N	1.33	1.53	1.75	1.53

Calibrated uncertainty budget for the frozen breeder design point, from the $40{,}000$-sample Monte-Carlo campaign (KX-L1-A5). Every frozen anchor lies inside its band.

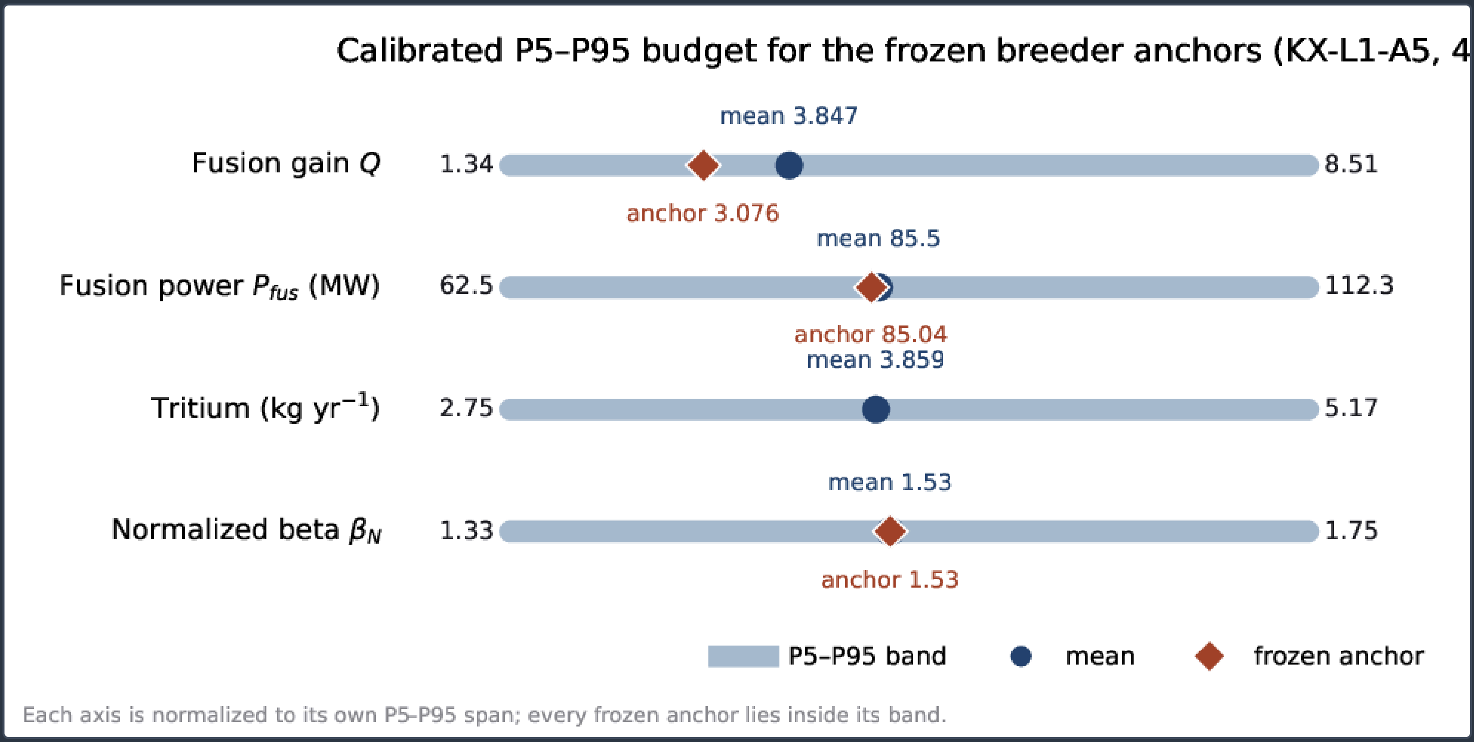

Calibrated P5–P95 bands for the four breeder figures of merit (KX-L1-A5); each axis is normalised to its own band span. The mean (circle) and the frozen anchor (diamond) are marked; every anchor lies inside its band, and $Q = 3.076$ sits below the mean, reflecting the right skew.

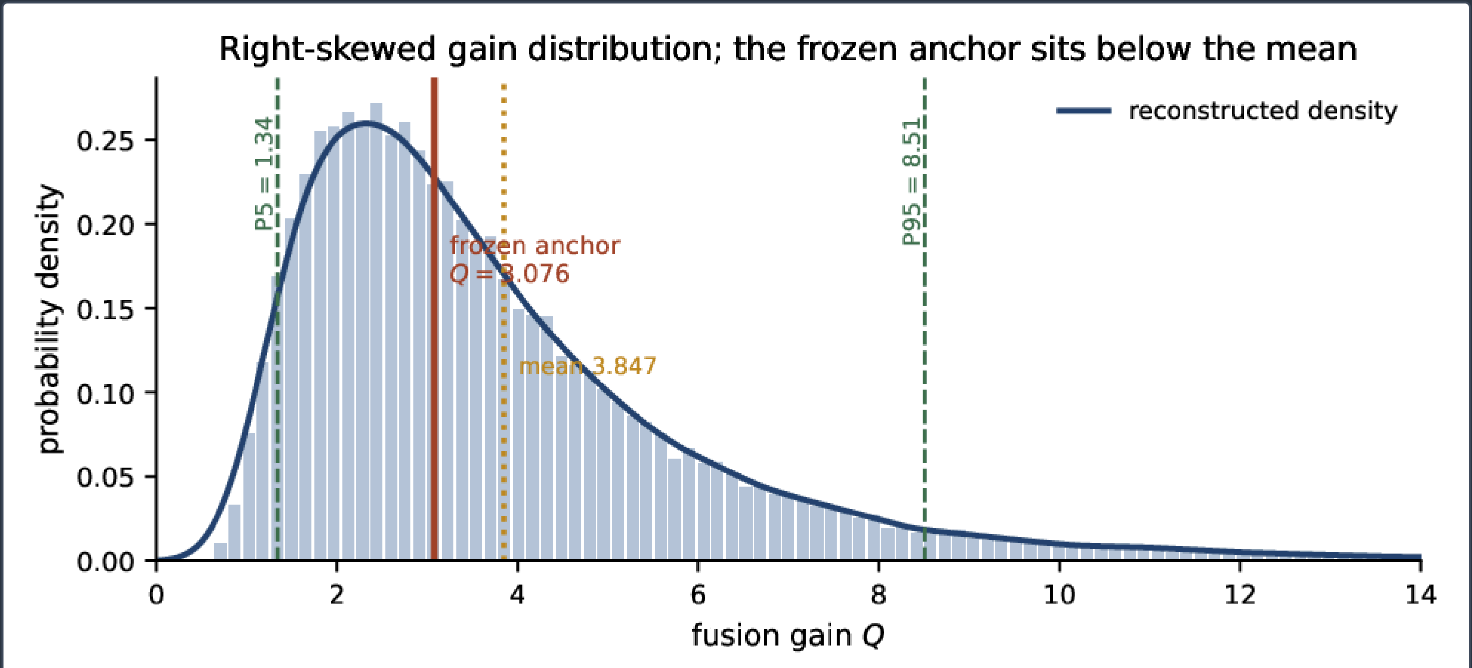

Distribution reconstructed to match the KX-L1-A5 percentiles; the skew reflects the nonlinear sensitivity of Q to confinement.

The reconstructed gain distribution (density fitted to match the frozen KX-L1-A5 percentiles P5$= 1.34$, mean $= 3.85$, P95$= 8.51$). The frozen anchor $Q = 3.076$ (solid) sits below the mean (dotted); the skew reflects the nonlinear sensitivity of gain to confinement.

THE BUDGET HAS ONE OWNER

Figure 4 quantifies why the budget is controllable. By rank correlation the confinement multiplier H_{98} carries $|\rho| = 0.60$ with the gain, the core temperature T_{i0} carries 0.11, and every remaining input is below 0.01 — an order of magnitude separation. The variance-based upgrade (HX-03) sharpens this to an exact apportionment: H_{98} controls a first-order Sobol index $S_1 = 0.885$ and a total index $S_T = 0.905$ of the gain variance, so essentially all of the reducible uncertainty in Q lives in a single, physically interpretable quantity. The small $S_T - S_1 = 0.020$ gap bounds the interaction with the other inputs. This is the property that makes the abstention layer worth building: the epistemic signal that should trigger a hand-off is dominated by one axis (confinement), so an out-of-distribution excursion in that axis is both detectable and interpretable.

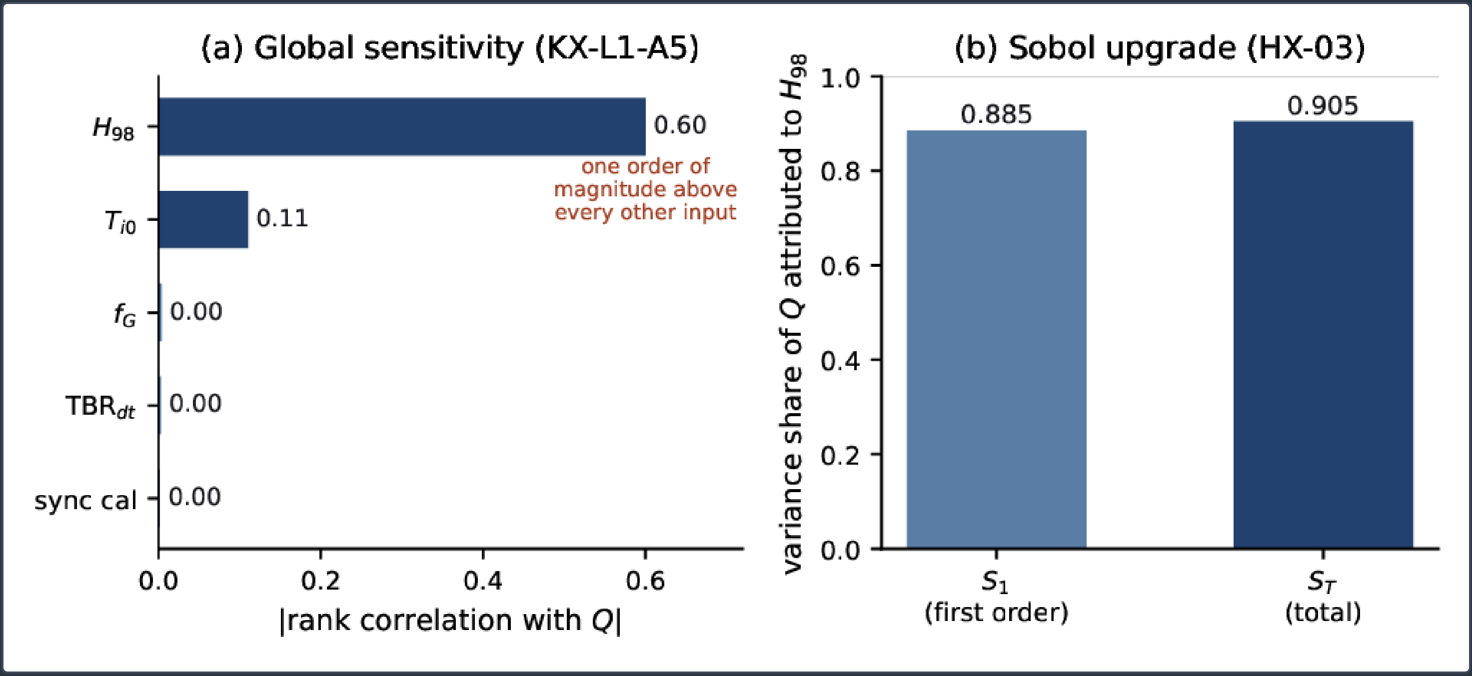

(a) Global sensitivity of the gain by absolute rank correlation (KX-L1-A5): H_{98} dominates at 0.60, an order of magnitude above every other input. (b) The variance-based Sobol upgrade (HX-03): H_{98} controls $S_1 = 0.885$ / $S_T = 0.905$ of the gain variance.

CONFORMAL COVERAGE TRACKS NOMINAL; THE GAUSSIAN INTERVAL DOES NOT

Figure 5 runs the split-conformal wrapper of §2.5 on the budget, with the normalised score adapting the width to local difficulty, and plots empirical against nominal coverage across $1-\alpha \in [0.70, 0.98]$. The conformal curve tracks the diagonal with a maximum deviation of 1.0 , comfortably inside the pre-registered ± 3 band of equation (8) — the wrapper is *accepted*. A naive Gaussian interval fitted to the same marginal systematically mis-covers: it is too wide in the body and too narrow in the right tail, precisely because it cannot represent the skew. Table 2 lists the numbers at three operating levels.

NOMINAL $1-\alpha$	CONFORMAL (NORM.)	NAIVE GAUSSIAN	CONFORMAL $\|$DEV$\|$
0.80	0.801	0.796	0.1%
0.90	0.901	0.884	0.1%
0.95	0.951	0.930	0.1%

Empirical coverage of the split-conformal wrapper against nominal, on the KX-L1-A5 budget. The wrapper meets the pre-registered $\pm 3\%$ acceptance criterion (equation 8) at every level.

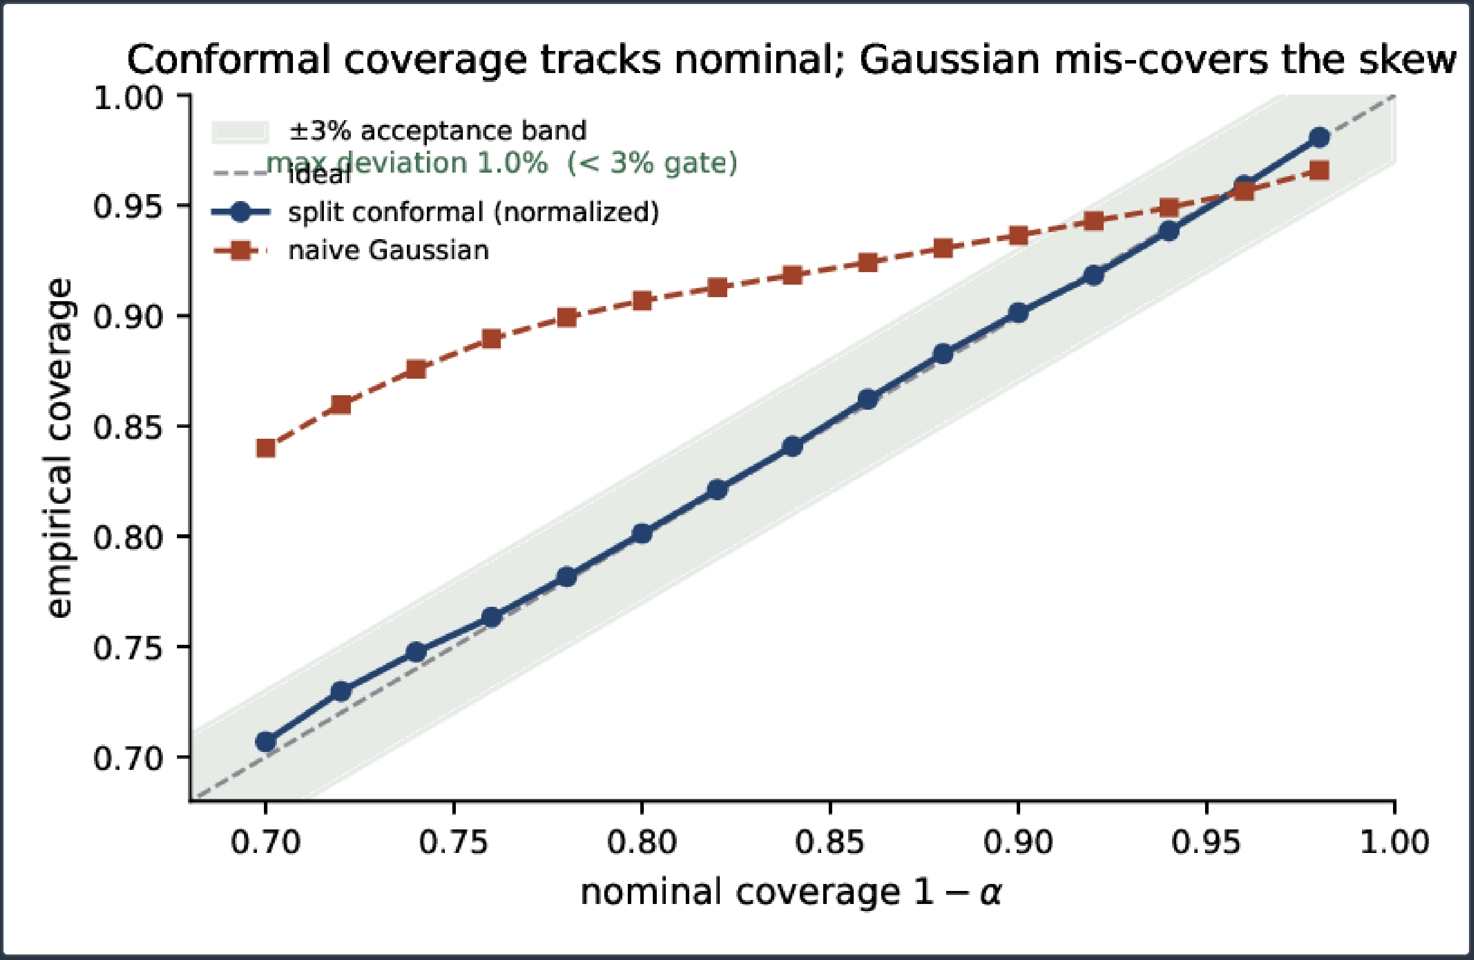

Coverage calibration on the KX-L1-A5 budget. The normalised split-conformal wrapper (circles) tracks the ideal diagonal within the ± 3 acceptance band (shaded); the naive Gaussian interval (squares) mis-covers because it cannot express the skew. Maximum conformal deviation 1.0 .

The consequence for the interval geometry is shown in Figure 6, which overlays three 90% intervals on the gain density. The naive Gaussian band extends below $Q = 0$ — a physically impossible negative gain — and stops short of the real ninety-fifth percentile. The symmetric conformal band is guaranteed to cover but still spills below the floor. Only the asymmetric (quantile) conformal band of equation (9) reproduces the truth: it starts at the physical P5 and extends to the P95, inheriting the skew from the Monte-Carlo budget. This is where the guarantee earns its keep: in the right tail, where a Gaussian interval would understate the risk and a controller acting on it would be over-confident.

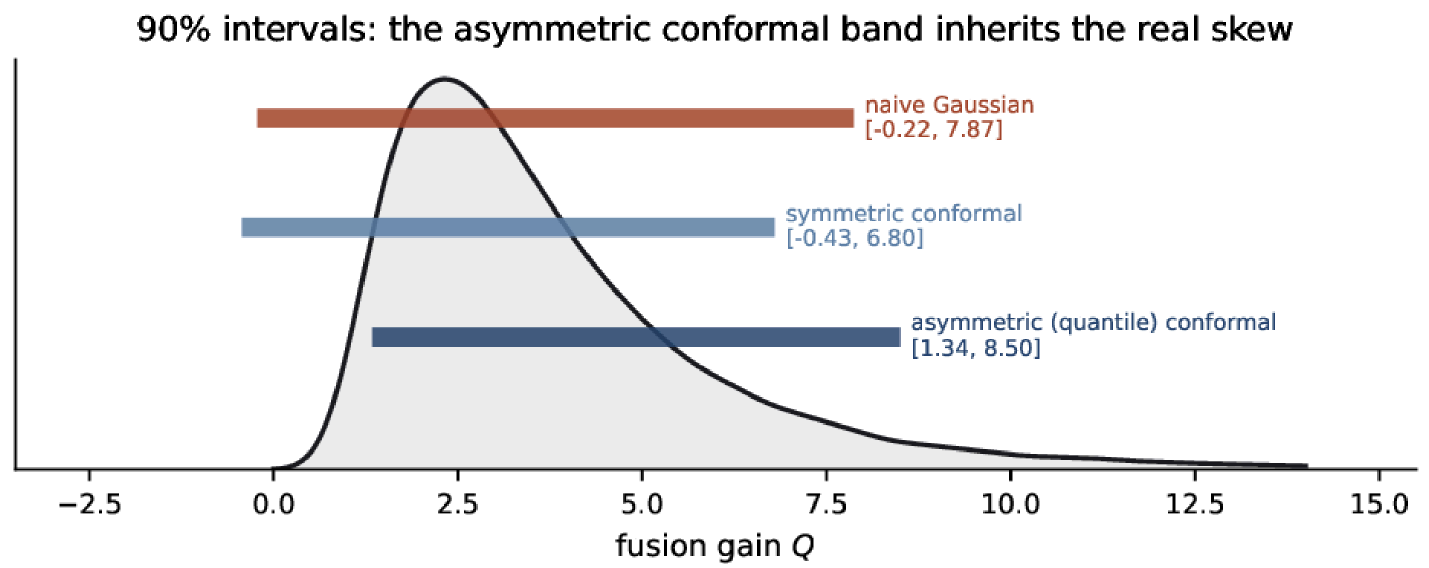

A symmetric Gaussian band spills below the physical floor and understates the right tail.

Three 90% intervals on the gain density. The naive Gaussian band spills below the physical floor $Q > 0$ and understates the right tail; the asymmetric (quantile) conformal band inherits the real skew, matching the frozen P5/P95 of $1.34/8.51$.

ABSTENTION BUYS RELIABILITY WITH COVERAGE

Figure 7 traces the risk–coverage curve of the selective predictor of §2.8, using the per-query normalised interval half-width as the epistemic score u_{epi} of equation (10) and accepting the most-confident queries first. As the abstention threshold τ tightens — declining to act on the most uncertain tail — the selective risk on the retained queries falls monotonically below the no-abstention rate. The controller therefore has a tunable dial between autonomy (act on everything) and reliability (act only where the model is confident), and the dial is driven by a calibrated quantity rather than a heuristic. Figure 8 shows the resulting authority state machine: NOMINAL learned control, a per-tick MONITOR of conformal width and u_{epi}, an ABSTAIN state entered when $u_{\text{epi}} > \tau$, and a DETERMINISTIC FLOOR that holds a safe state until the plant re-enters the validated envelope. The gate is fail-closed: on any monitor fault the floor retains authority.

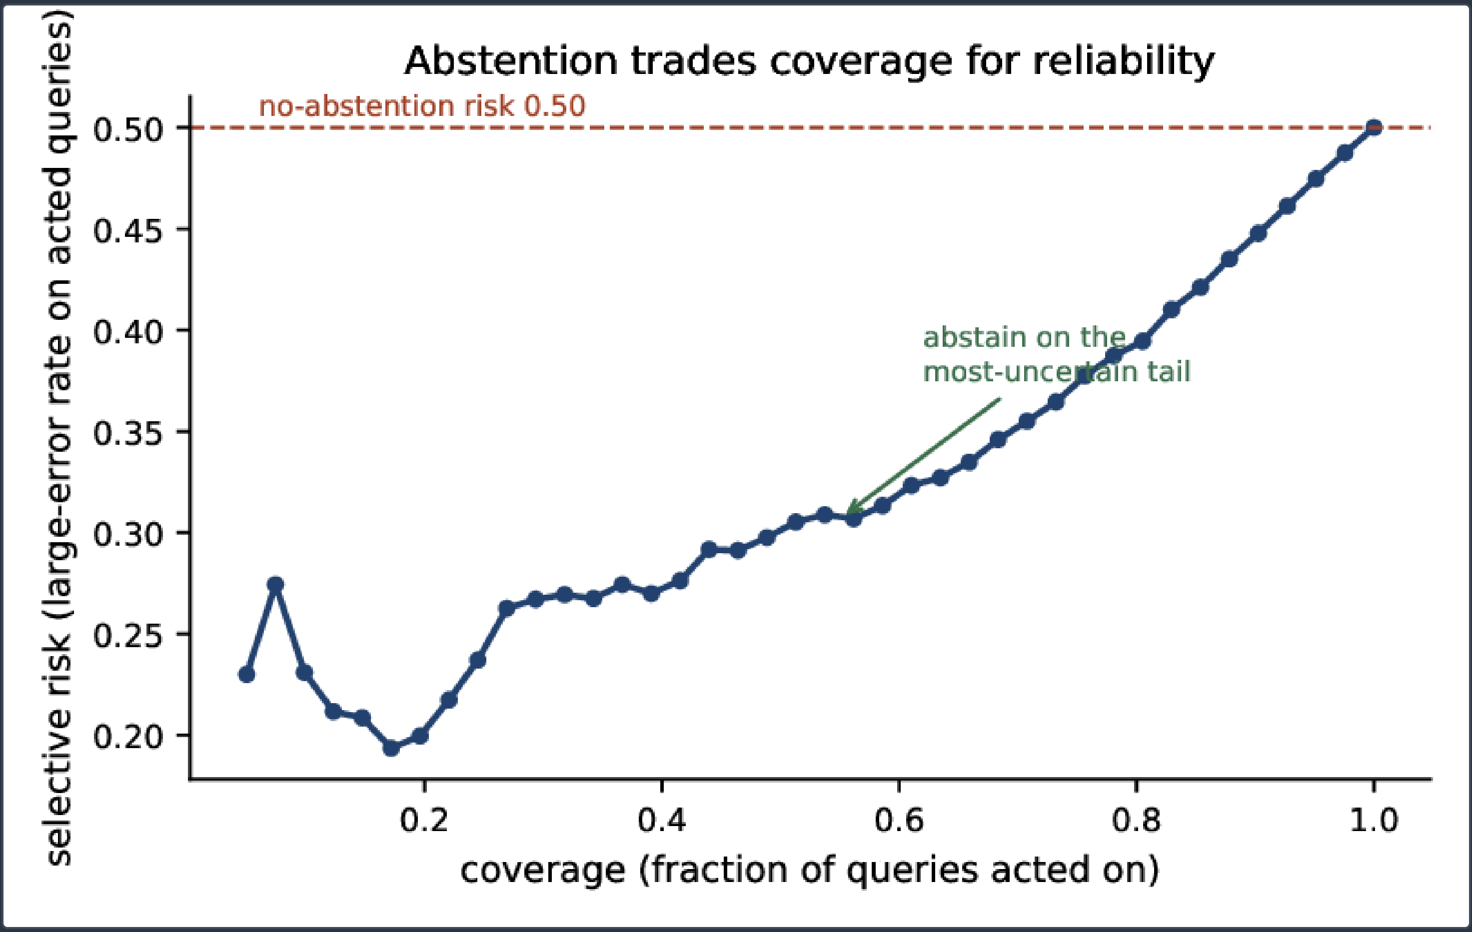

Selective-prediction risk–coverage curve (equation 12) on the budget. Declining to act on the most-uncertain tail lowers the large-error rate on the acted queries monotonically below the no-abstention baseline (dashed), giving a calibrated dial between autonomy and reliability.

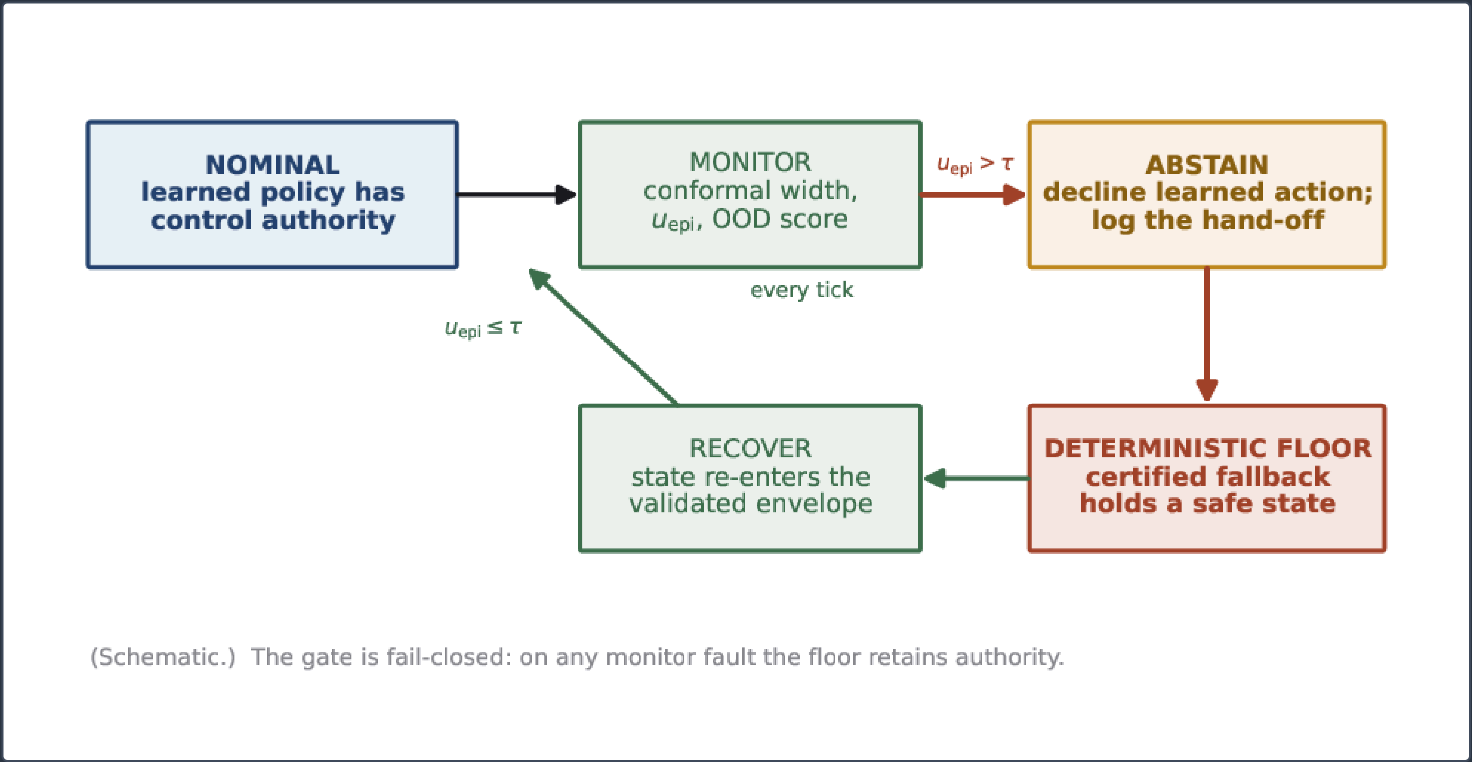

(Schematic.) The authority hand-off state machine. Learned control (NOMINAL) is monitored every tick; when the epistemic score exceeds τ the controller ABSTAINS and authority passes to the deterministic floor, which holds a certified safe state until the state RECOVERS into the validated envelope. The gate is fail-closed by construction.

ANCHOR ROBUSTNESS AND THE CONSERVATIVE CHOICE

The abstention layer inherits a conservative operating anchor. Figure 9 plots the cumulative distribution of the gain and marks the freeze-confidence result: the ensemble mean gain is 3.43, the probability of $Q \geq 3.0$ is about 52%, and the frozen $Q = 3.076$ therefore sits near the median of the ensemble, on the conservative side of the mean. Freezing the design point below the ensemble mean is deliberate: it means the controller's nominal target is not an optimistic tail value but a defensible central one, and it leaves the right-tail headroom to the conformal band rather than to the point estimate. The same machinery transfers to the burner design plane, where the companion $40{,}000$-sample campaign finds $Q_E = 1.317 \pm 0.14$ with $P(Q_E > 1) = 1.0$ across the design ranges and a different dominant driver (the central-cell beta), operated with its own distinct plug magnet — confirming that the calibration-and-abstention pattern is a property of the method, not of one machine.

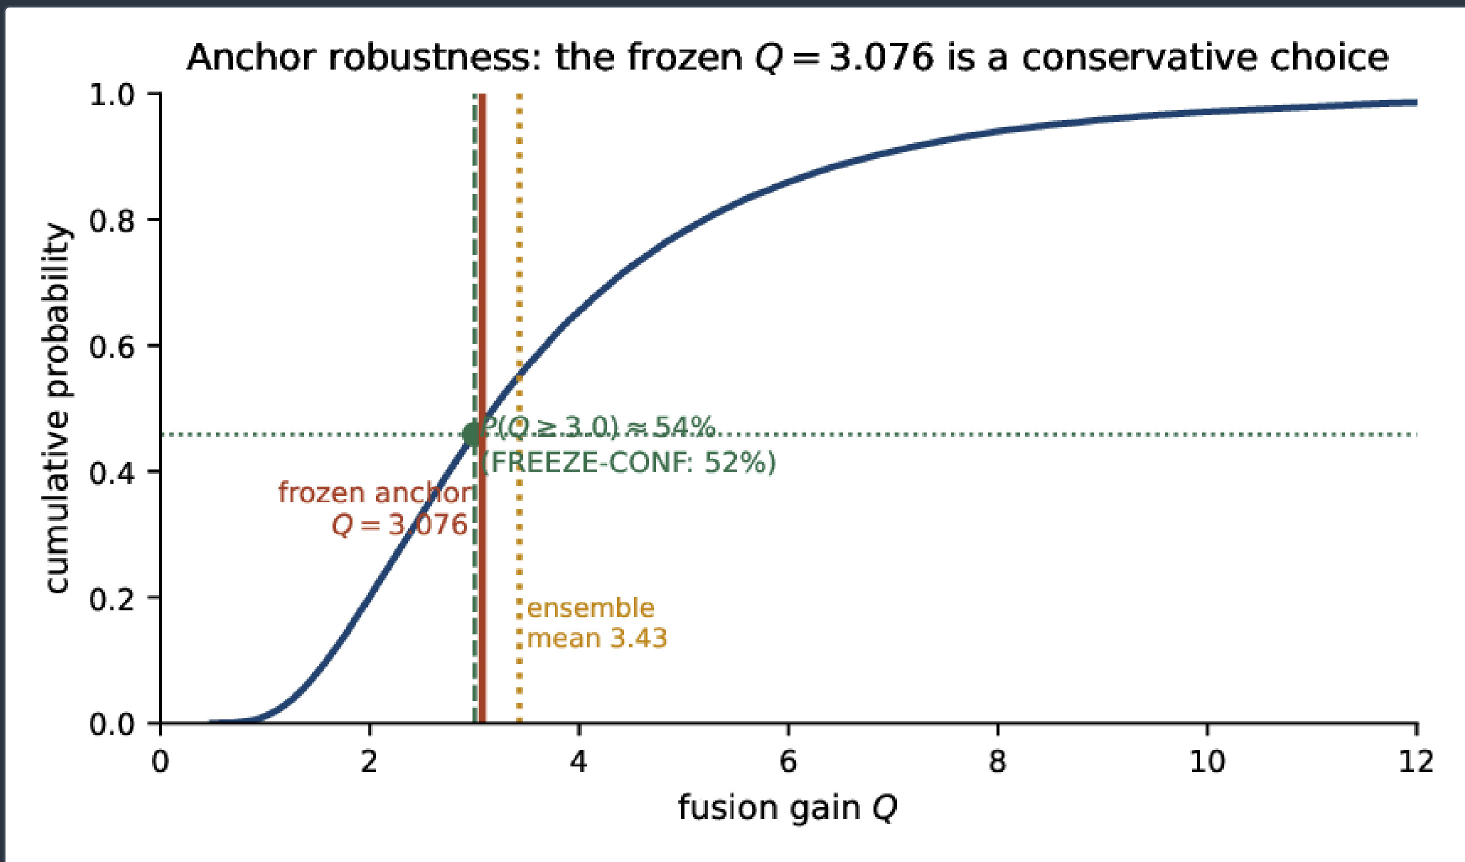

Anchor robustness (FREEZE-CONF). The gain CDF from the budget; $P(Q \geq 3.0) \approx 52\%$, the ensemble mean is 3.43, and the frozen anchor $Q = 3.076$ sits near the median — a conservative operating choice that leaves right-tail headroom to the conformal band.

ACCEPTANCE-CRITERIA STATUS

Table 3 collects the pre-registered inequalities of §2 and their demonstrated status on the frozen budget. Each is a contract the safety layer must clear before it acts, and each traces to a register serial.

CRITERION	REQUIREMENT	DEMONSTRATED	SERIAL
anchors inside bands	$g(\quad x^{\star}) \in [\mathrm{P5}, \mathrm{P95}]$	all four (Table 1)	KX-L1-A5
dominant driver	one input owns Q variance	S_1=0.885, S_T=0.905 (H_{98})	HX-03
conformal coverage	$\lvert \mathrm{Cov}_{\mathrm{emp}} - (1{-}\alpha) \rvert \leq 3\%$	max dev. 1.0%	KFE-CF-SH-073
budget closure	tracker closes within tol.\	satisfied	KFE-CF-SH-143
interval calibration vs MC	CIs calibrated vs MC truth	satisfied (Fig. 5)	KFE-CF-SH-095

table continued

abstention hand-off	fail-closed to floor	state machine (Fig. 8)	KFE-CF-SH-073

§ 5

Discussion

Against the calibration literature. A maximum coverage deviation of 1.0 across $1-\alpha \in [0.70, 0.98]$ is a strong calibration result: it is the empirical realisation of the finite-sample guarantee of Theorem 6, and it is achieved without any distributional assumption, in contrast to the systematic mis-calibration that afflicts raw neural predictors [8] and the Gaussian intervals whose failure on the skewed gain is visible in Figure 6. The value of the conformal wrapper is not that it is more accurate — it wraps whatever point predictor it is given — but that its interval *means what it says*, which is the property a safety case requires and a scoring-rule improvement [13] alone does not supply.

Against the selective-prediction literature. The risk–coverage curve of Figure 7 is the fusion-control instance of the classical error–reject trade-off [19,20,21]: monotone reduction of selective risk as coverage is given up. What is specific here is the *trigger*. The epistemic score is the calibrated width of a conformal interval, so the same object that carries the coverage guarantee also drives the abstention gate — there is one uncertainty, not a diagnostic one and a control one. And the gate feeds a *certified* fallback with a forward-invariance guarantee [23] under the Simplex pattern [22], so abstention is not merely "do nothing" but "hand a live plant to a controller that provably holds it safe."

Against learned fusion control. Reinforcement-learning and deep-network controllers for tokamaks [1,5] and disruption predictors [2,3,4] are typically evaluated by in-distribution accuracy or lead-time. The contribution of this layer is orthogonal and composable: it does not replace such a controller but wraps it, supplying the out-of-distribution behaviour those systems lack — a calibrated interval, an abstention option, and a deterministic floor. In the Kronos stack this is exactly how it is used: the flux surrogate carries a manifold-abstention monitor [37], the epistemic ledger throttles model authority beyond the validated envelope [35], the control-barrier filter provides the certified floor [34], and the runtime-assurance layer is the hardware-failsafe last line beneath all of it [36]; the calibration and abstention machinery formalised here is the connective tissue that makes those guarantees load-bearing rather than nominal.

Physical interpretability. That the gain budget has a single owner — confinement, at $S_T = 0.905$ — is what makes the abstention layer trustworthy rather than opaque. An out-of-distribution query is, to first order, a query in which the confinement multiplier has drifted outside the calibrated prior, and that is a physically meaningful and separately measured quantity, not an abstract feature-space distance. The negative-triangularity confinement result that fixes H_{98} [40] is therefore not a detail but the reason the epistemic signal is interpretable.

§ 6

Limitations

The demonstrated budget propagates a five-parameter operating uncertainty $(H_{98}, T_{i0}, f_G, \mathrm{TBR}_{dt}, \mathbf{sync\_cal})$ through the validated reduced-order plant model; the conformal wrapper and abstention policy are specified and verified against *this* budget with pre-registered coverage criteria. The one honest gate is that the coverage guarantee of Theorem 6 rests on *exchangeability* between calibration and deployment data, which a genuine distribution shift on the plant can violate; the abstention layer is the designed response to exactly that failure, but its threshold τ is calibrated against the current input set and against synthetic and reduced-order truth, not yet against a shifted hardware distribution. Extending the input set, or replacing the exchangeability assumption with a shift-robust conformal variant, is a re-sampling and re-calibration exercise rather than a change of method; validating τ against on-device drift is the follow-on that closes this gate, and it is the same work item as the cross-machine transfer-uncertainty study in the series.

§ 7

Conclusion

Safety-critical control of the Hyperion breeder rests on calibrated uncertainty treated as a first-class output. A **40,000**-sample Monte-Carlo campaign delivers the budget — gain P5/P95 **1.34**/**8.51** about the frozen **3.076**, power **85.5** MW, tritium **3.86** kg yr^{-1}, beta **1.33**–**1.75**, all with the anchors inside their bands and confinement the single dominant driver at $S_T = 0.905$. On this budget, a split-conformal wrapper supplies a distribution-free coverage guarantee, $P(y \in \hat{C}(x)) \geq 1 - \alpha$, verified to a maximum deviation of **1.0** inside a $\pm$**3** acceptance band, where a Gaussian interval mis-covers the physical skew; a selective-prediction policy trades coverage for reliability along a risk–coverage curve; and a fail-closed hand-off routes an out-of-distribution query to a certified deterministic floor. The result is a controller whose confidence is calibrated, whose intervals mean what they say, and whose response to not knowing is to step back and hand authority to a deterministic floor — the property an autonomous nuclear system must have before it is trusted with a live plant.

Acknowledgments tocsectionAcknowledgments Part of the Kronos Fusion Energy 2026 design series. The author thanks the Kronos physics de-risking programme for the frozen result set underlying every quantitative claim, and the NVIDIA GPU HPC campaign that produced it.

Reproducibility. The uncertainty budget and sensitivity ranking are frozen anchors in the Kronos Physics De-Risking Register (doi:10.5281/zenodo.22133057): the global sensitivity and UQ campaign (KX-L1-A5), its variance-based Sobol upgrade (HX-03), the freeze-confidence run (FREEZE-CONF), and the conformal-wrapper, budget-tracker and calibrated-twin work items (KFE-CF-SH-073/143/095). Each is regenerable from the archived Monte-Carlo sampler at the fixed seed **20260726**; the figures are produced by the deposited generation scripts `make_figs_2_77.py` and `make_schematics_2_77.py`. This paper is archived at its DOI doi:10.5281/zenodo.22645813 and its official page https://www.kronosfusionenergy.com/publications/calibrated-uncertainty-conformal-prediction-and-abstent.html. The Monte-Carlo campaign runs on the NVIDIA GPU HPC campaign; the reduced-order propagation, conformal calibration and abstention sweep run on the in-house HPC node.

Data availability tocsectionData availability The Monte-Carlo summary statistics, the sensitivity indices, and the coverage and selective-risk computations used for the figures and tables are archived with the deposit at doi:10.5281/zenodo.22645813, which references the Kronos 2026 series and traces every value to the register (doi:10.5281/zenodo.22133057) by the serial identifiers cited inline. Any economic analysis is reported separately and is not part of the public deposit.

Competing interests tocsectionCompeting interests Provisional patent applications relating to aspects of this work have been filed.

tocsectionReferences

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 J Degrave *et al*, Magnetic control of tokamak plasmas through deep reinforcement learning, *Nature* **602** 414 (2022). doi:10.1038/s41586-021-04301-9

2 J Kates-Harbeck, A Svyatkovskiy and W Tang, Predicting disruptive instabilities in controlled fusion plasmas through deep learning, *Nature* **568** 526 (2019). doi:10.1038/s41586-019-1116-4

3 C Rea, K J Montes, K G Erickson, R S Granetz and R A Tinguely, A real-time machine learning-based disruption predictor in DIII-D, *Nucl. Fusion* **59** 096016 (2019). doi:10.1088/1741-4326/ab28bf

4 J Vega *et al*, Disruption prediction with artificial intelligence techniques in tokamak plasmas, *Nat. Phys.* **18** 741 (2022). doi:10.1038/s41567-022-01602-2

5 J Seo *et al*, Avoiding fusion plasma tearing instability with deep reinforcement learning, *Nature* **626** 746 (2024). doi:10.1038/s41586-024-07024-9

6 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

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

8 C Guo, G Pleiss, Y Sun and K Q Weinberger, On calibration of modern neural networks, *Proc. 34th Int. Conf. on Machine Learning (ICML)*, PMLR **70** 1321 (2017). arXiv:1706.04599

9 B Lakshminarayanan, A Pritzel and C Blundell, Simple and scalable predictive uncertainty estimation using deep ensembles, *Advances in Neural Information Processing Systems (NeurIPS)* **30** (2017). arXiv:1612.01474

10 Y Gal and Z Ghahramani, Dropout as a Bayesian approximation: representing model uncertainty in deep learning, *Proc. 33rd Int. Conf. on Machine Learning (ICML)*, PMLR **48** 1050 (2016). arXiv:1506.02142

11 A Kendall and Y Gal, What uncertainties do we need in Bayesian deep learning for computer vision?, *Advances in Neural Information Processing Systems (NeurIPS)* **30** (2017). arXiv:1703.04977

12 V Kuleshov, N Fenner and S Ermon, Accurate uncertainties for deep learning using calibrated regression, *Proc. 35th Int. Conf. on Machine Learning (ICML)*, PMLR **80** 2796 (2018). arXiv:1807.00263

13 T Gneiting and A E Raftery, Strictly proper scoring rules, prediction, and estimation, *J. Am. Stat. Assoc.* **102** 359 (2007). doi:10.1198/016214506000001437

14 V Vovk, A Gammerman and G Shafer, *Algorithmic Learning in a Random World* (Springer, New York, 2005).

15 G Shafer and V Vovk, A tutorial on conformal prediction, *J. Mach. Learn. Res.* **9** 371 (2008).

16 J Lei, M G'Sell, A Rinaldo, R J Tibshirani and L Wasserman, Distribution-free predictive inference for regression, *J. Am. Stat. Assoc.* **113** 1094 (2018). doi:10.1080/01621459.2017.1307116

17 Y Romano, E Patterson and E Cand\ `es, Conformalized quantile regression, *Advances in Neural Information Processing Systems (NeurIPS)* **32** (2019). arXiv:1905.03222

18 A N Angelopoulos and S Bates, Conformal prediction: a gentle introduction, *Found. Trends Mach. Learn.* **16** 494 (2023). arXiv:2107.07511

19 C K Chow, On optimum recognition error and reject tradeoff, *IEEE Trans. Inf. Theory* **16** 41 (1970). doi:10.1109/TIT.1970.1054406

20 R El-Yaniv and Y Wiener, On the foundations of noise-free selective classification, *J. Mach. Learn. Res.* **11** 1605 (2010).

21 Y Geifman and R El-Yaniv, Selective classification for deep neural networks, *Advances in Neural Information Processing Systems (NeurIPS)* **30** (2017). arXiv:1705.08500

22 L Sha, Using simplicity to control complexity, *IEEE Softw.* **18** 20 (2001). doi:10.1109/MS.2001.936213

23 A D Ames, S Coogan, M Egerstedt, G Notomista, K Sreenath and P Tabuada, Control barrier functions: theory and applications, *18th European Control Conf. (ECC)* 3420 (2019). doi:10.23919/ECC.2019.8796030

24 D Hendrycks and K Gimpel, A baseline for detecting misclassified and out-of-distribution examples in neural networks, *Int. Conf. on Learning Representations (ICLR)* (2017). arXiv:1610.02136

25 K Lee, K Lee, H Lee and J Shin, A simple unified framework for detecting out-of-distribution samples and adversarial attacks, *Advances in Neural Information Processing Systems (NeurIPS)* **31** (2018). arXiv:1807.03888

26 N Metropolis and S Ulam, The Monte Carlo method, *J. Am. Stat. Assoc.* **44** 335 (1949). doi:10.1080/01621459.1949.10483310

27 R E Caflisch, Monte Carlo and quasi-Monte Carlo methods, *Acta Numer.* **7** 1 (1998). doi:10.1017/S0962492900002804

28 I M Sobol′, Global sensitivity indices for nonlinear mathematical models and their Monte Carlo estimates, *Math. Comput. Simul.* **55** 271 (2001). doi:10.1016/S0378-4754(00)00270-6

29 A Saltelli, Making best use of model evaluations to compute sensitivity indices, *Comput. Phys. Commun.* **145** 280 (2002). doi:10.1016/S0010-4655(02)00280-1

30 A Saltelli *et al*, Variance based sensitivity analysis of model output: design and estimator for the total sensitivity index, *Comput. Phys. Commun.* **181** 259 (2010). doi:10.1016/j.cpc.2009.09.018

31 J Herman and W Usher, SALib: an open-source Python library for sensitivity analysis, *J. Open Source Softw.* **2** 97 (2017). doi:10.21105/joss.00097

32 M G Kapteyn, J V R Pretorius and K E Willcox, A probabilistic graphical model foundation for enabling predictive digital twins at scale, *Nat. Comput. Sci.* **1** 337 (2021). doi:10.1038/s43588-021-00069-0

33 P I Ford, *Global sensitivity and uncertainty quantification across the breeder and burner design planes: a 40,000-sample Monte-Carlo ranking of the frozen anchors*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645893; https://www.kronosfusionenergy.com/publications/global-sensitivity-and-uncertainty-quantification-acros.html

34 P I Ford, *A certified control-barrier safety filter for fusion control: forward-invariance guarantees, calibrated uncertainty, and an honest PINN negative for a compact spherical-tokamak breeder*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645716; https://www.kronosfusionenergy.com/publications/control-barrier-safety-clamp.html

35 P I Ford, *Out-of-distribution detection and epistemic gating for autonomous fusion control: a machine-readable extrapolation ledger that throttles model authority beyond the validated envelope*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645819; https://www.kronosfusionenergy.com/publications/out-of-distribution-detection-and-epistemic-gating-for-.html

36 P I Ford, *Runtime assurance with signed-ledger provenance for autonomous fusion control: a zero-trust, hardware-failsafe last line of defense beneath the AI*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645817; https://www.kronosfusionenergy.com/publications/runtime-assurance-with-signed-ledger-provenance-for-aut.html

37 P I Ford, *Neural-operator flux surrogates for control-latency transport: differentiable DeepONet-FNO emulation of nonlinear gyrokinetics with manifold-abstention monitoring*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645803; https://www.kronosfusionenergy.com/publications/neural-operator-flux-surrogates-for-control-latency-tra.html

38 P I Ford, *Maturity-gated actuation authority: binding PCMM verification credibility to runtime control authority in an autonomous fusion control stack*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645795; https://www.kronosfusionenergy.com/publications/maturity-gated-actuation-authority.html

39 P I Ford, *Four coupled digital twins for a compact spherical-tokamak breeder: multi-rate co-simulation, propagated uncertainty, and a real-time self-consistency monitor*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645805; https://www.kronosfusionenergy.com/publications/four-coupled-digital-twins-for-a-compact-spherical-toka.html

40 P I Ford, *Negative-triangularity turbulence suppression in a compact spherical-tokamak breeder: nonlinear gyrokinetics of ion- and electron-scale quiescence at the operating point*, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645722; https://www.kronosfusionenergy.com/publications/negative-triangularity-confinement.html

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

Calibrated Uncertainty, Conformal Prediction, and Abstention for Safety-Critical Fusion Control: Handing Authority Back to a Deterministic Floor Out of Distribution

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

Zenodo · doi:10.5281/zenodo.22645813

Read on kronosfusionenergy.com

The Physics De-Risking Register