KRONOS FUSION ENERGY

PAPER 2.70 · THE KRONOS 2026 SERIES

Maturity-Gated Actuation Authority: Binding PCMM Verification Credibility to Runtime Control Authority in an Autonomous Fusion Control Stack

A controller or a learned advisor on a nuclear-regulated machine should be trusted with exactly as much authority as the evidence behind its models can justify — no more, and no...

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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

Maturity-Gated Actuation Authority: Binding PCMM Verification Credibility to Runtime Control Authority in an Autonomous Fusion Control Stack

A controller or a learned advisor on a nuclear-regulated machine should be trusted with exactly as much authority as the evidence behind its models can justify — no more, and no less than the evidence has earned. On a compact, high-power-density spherical tokamak the reaction times are short and every actuation carries a safety consequence, so the question "how much authority should the automation have?" cannot be answered by fiat. We give it a quantitative, auditable answer. We define *maturity-gated actuation authority*: a control-governance law that binds the verification-and-validation credibility of each subsystem's models — scored on a Predictive Capability Maturity Model (PCMM) — to the actuation authority (magnitude, rate, and irreversibility) the controller is permitted at runtime. Two ingredients make the binding formal. A *static* credibility functional aggregates the six PCMM axes of a subsystem into a maturity index $m_s \in [0, 1]$; a *runtime* validated-envelope membership $M(x) \in [0, 1]$, computed from a calibrated conformal nonconformity score, measures how far the live state has drifted out of the region where those models were validated. Their product gates a monotone authority envelope $\mathcal{U}(M)$, and the applied command is the projection of a maturity-weighted blend of the nominal controller and a deterministic safe floor onto that envelope. We prove that the gated law preserves forward invariance of the certified safe set (so an immature model can never drive the machine), that authority is monotone non-decreasing in evidence (so the envelope expands only as verification does), and that the map recovers the raw nominal controller in the fully-validated limit. We instantiate the law on the Hyperion breeder design point ($Q = 3.076$, $P_{\text{fus}} = 85.04\,MW$, $I_p = 9.66\,MA$, $\delta = -0.30$, $B_0 = 8\,T$, centrepost 16.84 T) and reduce it to practice against frozen register results: a control-barrier clamp holds $\beta_N \leq 3.5$ with $0/150{,}000$ adversarial violations and $\min h = +0.000$; a calibrated uncertainty layer matches nominal coverage to within $\pm 3\%$; an equilibrium surrogate clears its validation bar at 0.139% RMS flux error ($2877\times$ faster than the reference); a $40{,}000$-sample anchor-confidence ensemble places the frozen $Q = 3.076$ conservatively below the ensemble mean 3.43; and the one honest open gate — an in-winding quantum magnetometer whose neutron tolerance is $\sim 39\times$ short of requirement — keeps its subsystem correctly clamped to advisory authority. The contribution is an architecture, not a controller: it constrains authority, it does not by itself close any loop, and it gives a byproduct-material or reactor licensing review a transparent rule in which autonomy is a function of demonstrated evidence.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645795 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: verification and validation · predictive capability maturity model · control authority · control barrier function · conformal prediction · calibrated uncertainty · certifiable autonomy · 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

AUTONOMY IS A DESIGN REQUIREMENT, AND IT IS ONLY AS TRUSTWORTHY AS ITS MODELS

A spherical tokamak (ST) confines a large plasma current in a small volume, which is what makes it an attractive neutron source and breeder platform [1,2] and, at the same time, what makes it demanding to control. The Hyperion breeder holds $I_p = 9.66\,MA$ at a major radius of only $R_0 = 1.20\,m$ behind a centrepost magnet at a peak field of $16.84\,T$, in a negative-triangularity ($\delta = -0.30$) regime chosen for good confinement without the edge pedestal and the edge-localised-mode cycle of conventional H-mode [3,4]. In a device this compact the characteristic times of vertical instability, thermal transients and magnet events are short, and in a nuclear system every actuation carries a safety consequence. Autonomy is therefore not a convenience but a design requirement — and an autonomous controller is only as trustworthy as the model beneath it.

Feedback control of tokamak equilibria and profiles is a mature discipline: magnetic diagnostics feed real-time equilibrium reconstruction [5] and shape/position controllers [6], and physics-based profile simulators run inside the control cycle [7]. Learned controllers have now entered the loop: a deep reinforcement-learning agent has commanded the TCV coil voltages directly from magnetics to produce a variety of shapes, negative triangularity included [8], and deep networks forecast disruptive instabilities across machines [9]. Every one of these advances shares a dependency: the control quality — and the control *risk* — is bounded by the fidelity of the underlying plasma model. A learned policy that is superb on the distribution it was trained on can be arbitrarily wrong just outside it, and on a regulated machine that failure mode is not tolerable.

VERIFICATION CREDIBILITY IS ALREADY QUANTIFIED — BUT IT DOES NOT YET TOUCH THE ACTUATORS

Computational science has spent two decades learning to *quantify* how much a model should be believed. The Predictive Capability Maturity Model (PCMM) of Oberkampf, Pilch and Trucano scores a simulation capability on a set of orthogonal axes — representation and geometric fidelity, physics and material-model fidelity, code verification, solution verification, model validation, and uncertainty quantification — each on a graded maturity scale [10]. The surrounding discipline of verification, validation and uncertainty quantification (VVUQ) supplies the concrete tools that populate those axes: the method of manufactured solutions for code verification [13], the grid-convergence index for solution verification [14], variance-based global sensitivity analysis [15,16], and a comprehensive framework that threads them into a single credibility statement [11,12]. The Kronos design campaign runs exactly this suite: a formal V&V gate (MMS/GCI/Sobol/PCMM) that clears a grid-convergence index below 0.4% with $19/19$ frozen anchors reproduced (register serial KX-L1-A6), and a companion verification-first catalogue that fixes every anchor and reproduces each gate against an independent library [34,39].

In parallel, control theory has learned to make autonomy *safe*. A control barrier function (CBF) renders a safe set forward-invariant through a pointwise quadratic-program filter on the actuator command [17,18]; a predictive safety filter carries the same guarantee for constrained nonlinear systems under learned control [19]; the Simplex and runtime-assurance architectures switch between a high-performance and a verified-safe controller under a monitored condition [20,21]; and reachability-based frameworks give learning-based controllers a certified fallback [22,23]. These two literatures — credibility scoring and safe control — have developed independently. Credibility scoring tells a reviewer *how much to believe* a model but stops at the report; safe control tells the actuator *what is admissible* but takes the model's trust as given. Nothing binds the first to the second. A subsystem whose models are only reduced-order is, in current practice, permitted the same actuator authority as one that has been machine-validated against nonlinear gyrokinetics and hardware — the safety filter simply assumes its model is right.

CONTRIBUTION

This paper closes that gap. We propose *maturity-gated actuation authority*: a governance law that makes a subsystem's permitted actuation authority an explicit, monotone function of its demonstrated PCMM credibility and of the live state's membership in the validated envelope. Concretely we contribute:

a *static credibility functional* (§2.2) that aggregates the six PCMM axes of a subsystem into an interpretable maturity index $m_s \in [0, 1]$, with a conservative weakest-link aggregation that no single high axis can rescue;
a *runtime membership* $M(x) \in [0, 1]$ (§2.3) built from a calibrated conformal nonconformity score, so that leaving the validated envelope is detected with a finite-sample coverage guarantee;
the *authority envelope* $\mathcal{U}(M)$ and the *gated control law* (§2.5), which projects a maturity-weighted blend of the nominal controller and a deterministic safe floor onto the admissible set;
three theorems (§3): forward-invariance preservation, monotonicity of authority in evidence, and nominal recovery in the fully-validated limit;
a numerical formulation (§4) — the projection quadratic program, the online membership estimator, and their convergence and latency — and a GPU-HPC campaign (§5) that produced the grounding results;
a reduction to practice on the Hyperion breeder (§6) grounded in frozen register anchors, including the one honest open gate the policy correctly clamps. enumerate
Every quantitative claim traces to a frozen result in the Kronos Physics De-Risking Register [34]; the serial identifiers (e.g. KX-L1-A13) are given inline so that each number is auditable. The scheme is a frozen anchor in that register (invention family; provisional applications filed).

§ 2

Governing model and the formal problem

PLANT, CONTROLLER AND SAFE SET

Let the plant state be $x \in \mathcal{X} \subseteq \mathbb{R}^n$ and the actuator command $u \in \mathbb{R}^p$, evolving in continuous time under the (in general uncertain) dynamics

$$\dot{x} = f(x) + g(x)\,u, \qquad x(0) = x_0.$$

A subset of the state is *safety-critical*. We encode the operational safety envelope as the zero-superlevel set of a continuously differentiable function $h : \mathcal{X} \to \mathbb{R}$,

$$\mathcal{C} = \{\, x \in \mathcal{X} : h(x) \geq 0 \,\}, \qquad \partial\mathcal{C} = \{\, x : h(x) = 0 \,\},$$

so that $\mathcal{C}$ is the region certified safe by the design basis. For the breeder the load-bearing instance is the normalised-beta margin: with β_N the normalised plasma beta and $\beta_N^{\max} = 3.5$ the register's certified ceiling (KX-L1-A13), $h(x) = \beta_N^{\max} - \beta_N(x)$. A *nominal* controller $u_{\text{nom}} = \pi_{\text{nom}}(x)$ provides performance (a model-predictive or learned policy, companion [38]); a *safe* controller $u_{\text{safe}} = \pi_{\text{safe}}(x)$ provides a deterministic, physics-bounded fallback (the certified clamp, companion [35]). The problem is to combine them so that authority is granted to u_{nom} only in proportion to the evidence that the model behind it is trustworthy.

THE CREDIBILITY FUNCTIONAL

Each subsystem s carries a PCMM assessment: a vector of maturity scores $\ell_s = (\ell_{s,1}, \ldots, \ell_{s,K})$ over $K = 6$ axes, each $\ell_{s,k} \in \{0,1,2,3\}$ mapping onto the fidelity tiers *representative*, *characterized*, *converged*, *machine-validated* used throughout the campaign. We aggregate the axes into a scalar maturity index $m_s \in [0,1]$ by a weighted power mean with negative exponent, which behaves as a smooth weakest-link:

$$m_s(\rho) = \frac{1}{3}\left(\sum_{k=1}^{K} w_k\, \ell_{s,k}^{\,\rho}\right)^{1/\rho}, \qquad \sum_{k=1}^{K} w_k = 1,\;\; w_k > 0,\;\; \rho \leq -1.$$

The factor $1/3$ normalises the machine-validated ceiling ($\ell = 3$ on every axis) to $m_s = 1$. As $\rho \to -\infty$, 3 tends to the weighted minimum $\frac{1}{3}\min_k \ell_{s,k}$, so a single immature axis (say model validation at $\ell = 0$) forces $m_s \to 0$ regardless of how strong the others are; a finite $\rho \leq -1$ retains differentiability while preserving this conservative bias. Equation 3 is monotone non-decreasing in every $\ell_{s,k}$, so credibility can only rise as an axis is matured — a property we use in §3. The static index is refreshed whenever the underlying V&V evidence changes; between refreshes it is a constant of the deployment.

THE RUNTIME VALIDATED-ENVELOPE MEMBERSHIP

A model's PCMM score is earned over a *training and validation domain* $\mathcal{D} \subset \mathcal{X}$. Its credibility does not transfer outside $\mathcal{D}$, so authority must contract when the live state drifts out. We measure that drift with a calibrated nonconformity score. Let $r : \mathcal{X} \to \mathbb{R}_{\geq 0}$ be a nonconformity function — for a Gaussian-process or ensemble surrogate the natural choice is the Mahalanobis distance to the validation set,

$$r(x) = \sqrt{(\phi(x) - \bar{\phi})^{\top}\, \Sigma^{-1} (\phi(x) - \bar{\phi})},$$

with $\phi(x)$ a feature embedding, $\bar{\phi}$ and Σ the validation-set mean and covariance. Split-conformal calibration [25,26] converts r into a distribution-free threshold: on a held-out calibration set $\{x_i\}_{i=1}^{N_c}$ drawn from $\mathcal{D}$, the empirical $(1-\alpha)$ quantile

$$\hat{q}_{1-\alpha} = \text{the } \lceil (N_c + 1)(1-\alpha) \rceil\text{-th smallest of } \{r(x_i)\}$$

guarantees, under exchangeability, that a fresh in-domain state satisfies $\Pr[\, r(x) \leq \hat{q}_{1-\alpha} \,] \geq 1 - \alpha$. We then define the runtime membership as a smooth, monotone-decreasing map of the calibrated distance,

$$M(x) = \tfrac{1}{2}\Big(1 - \tanh\big[(r(x) - \hat{q}_{1-\alpha})/w\big]\Big) \in [0,1],$$

with width $w > 0$; $M \approx 1$ well inside the validated envelope and $M \to 0$ once r exceeds the conformal threshold (figure 3). The calibration that fixes $\hat{q}_{1-\alpha}$ is itself a gated quantity: the campaign's uncertainty layer is accepted only when empirical coverage matches the nominal level to within $\pm 3\%$ (KX-L1-A14, companion [36]), the condition

$$\left| \mathrm{Cov}_{\mathrm{emp}}(\alpha) - \alpha \right| \leq 3\% \qquad \forall\, \alpha,$$

so that $M(x)$ inherits a calibrated meaning rather than an arbitrary one.

THE AUTHORITY ENVELOPE

Actuation authority is not a scalar; it has three faces — how large a command is allowed, how fast it may change, and whether it may take an irreversible action. We encode all three in a *maturity-parametrised admissible set*. Let $\mu = m_s\, M(x) \in [0, 1]$ be the effective maturity, the product of static credibility and runtime membership. Define

$$\mathcal{U}(\mu) = \Big\{ u : \ \|u - u_{\mathrm{safe}}\|_\infty \leq \bar{a}(\mu), \ \ \|\dot{u}\|_\infty \leq \bar{r}(\mu), \ \ u \in \mathcal{I}(\mu) \Big\},$$

where $\bar{a}(\mu)$ is the magnitude authority, $\bar{r}(\mu)$ the rate authority, and $\mathcal{I}(\mu)$ the set of permitted (reversible, or — above a maturity threshold $\mu^\star$ — irreversible) action classes. We require each to be monotone non-decreasing in μ, with

$$\bar{a}(0) = a_0 > 0, \quad \bar{a}(1) = a_{\max}; \qquad \bar{r}(0) = r_0 > 0, \quad \bar{r}(1) = r_{\max}; \qquad \mathcal{I}(\mu) = \begin{cases} \mathcal{I}_{\mathrm{rev}}, & \mu < \mu^\star \\ \mathcal{I}_{\mathrm{all}}, & \mu \geq \mu^\star. \end{cases}$$

The small floors $a_0, r_0 > 0$ leave the safe controller always able to act; the ceilings $a_{\max}, r_{\max}$ recover the unrestricted actuator when the model is fully credible. Figure 3 plots a representative envelope with the irreversibility threshold $\mu^\star$; the specific functional forms are design choices, and the theorems of §3 require only monotonicity and the endpoint conditions 9.

THE GATED CONTROL LAW

The applied command blends the nominal and safe controllers by effective maturity and projects the blend onto the admissible set that the same maturity permits:

$$\boxed{u_{\mathrm{app}} = \Pi_{\mathcal{U}(\mu)}\Big(\mu\, u_{\mathrm{nom}} + (1 - \mu)\, u_{\mathrm{safe}}\Big)}, \qquad \mu = m_s\, M(x),$$

where $\Pi_{\mathcal{U}(\mu)}$ is the Euclidean projection onto the closed convex set $\mathcal{U}(\mu)$. As $\mu \to 1$ (machine-validated model, state deep inside the validated envelope), the blend tends to u_{nom} and the envelope opens to the full actuator, so $u_{\mathrm{app}} \to u_{\mathrm{nom}}$: the scheme is transparent to a fully-credible controller. As $\mu \to 0$ (immature model or state out of distribution), the blend tends to u_{safe} and the envelope contracts to the reversible floor, so $u_{\mathrm{app}} \to u_{\mathrm{safe}}$: the machine cannot be driven by an unvalidated model. The projection Π is what makes the guarantee hard rather than advisory — even a pathological u_{nom} is clipped into $\mathcal{U}(\mu)$ before it reaches the plant. Equation 10 is the formal object of this paper; the rest of the paper establishes its properties, its numerics, and its instantiation.

§ 3

Theoretical guarantees

We state three properties that make 10 a governance law and not merely a heuristic blend. Throughout, π_{safe} is assumed to be a valid CBF-QP controller for $\mathcal{C}$ in the sense of [17]: there is an extended class-$\mathcal{K}_\infty$ function γ such that

$$\sup_{u \in \mathcal{U}(0)} \left[L_f h(x) + L_g h(x)\, u \right] \;\geq\; -\,\gamma\big(h(x)\big) \qquad \forall x \in \mathcal{C},$$

i.e. the safe controller can maintain the CBF condition using only the always-available floor authority $\mathcal{U}(0)$, and $u_{\text{safe}}(x)$ is a selection that attains it.

theorem[Forward-invariance preservation] Let $\mathcal{C}$ in 2 satisfy the CBF condition 11 with the floor authority, and let the admissible set $\mathcal{U}(\mu)$ be defined so that its projection $\Pi_{\mathcal{U}(\mu)}$ never removes the safety component, i.e. the CBF half-space $\{u : L_g h(x) u \geq -\gamma(h) - L_f h(x)\}$ intersects $\mathcal{U}(\mu)$ for all $\mu \in [0,1]$ and all $x \in \mathcal{C}$. Then the closed loop 1 under u_{app} of 10 renders $\mathcal{C}$ forward invariant: $x_0 \in \mathcal{C} \Rightarrow x(t) \in \mathcal{C}$ for all $t \geq 0$. theorem

proof[Proof sketch] By 11 the CBF half-space is non-empty for every $x \in \mathcal{C}$ using authority already contained in $\mathcal{U}(0) \subseteq \mathcal{U}(\mu)$ (monotonicity, 9). The admissible set is constructed as the intersection of $\mathcal{U}(\mu)$ with that half-space enforced as a hard constraint of the projection QP (§4); the intersection is non-empty by hypothesis, so the projection returns a u_{app} satisfying $\dot{h} = L_f h + L_g h\, u_{\text{app}} \geq -\gamma(h)$ on $\partial\mathcal{C}$. By Nagumo's theorem and the comparison lemma for $\dot{h} \geq -\gamma(h)$ [17,18], $h(x(t)) \geq 0$ is invariant. The maturity weighting affects only *which* admissible command is chosen, never whether the safety constraint is met. proof

Theorem 11 is the licensing-relevant statement: whatever the maturity, and however wrong the nominal model, the state cannot leave the certified safe set. In the register the clamp realising π_{safe} holds $\beta_N \leq 3.5$ with zero violations in $150{,}000$ adversarial Monte-Carlo replays and a minimum barrier value $\min h = +0.000$ (KX-L1-A13; figure 5).

theorem[Monotonicity of authority in evidence] The permitted authority is non-decreasing in every PCMM axis and in the runtime membership. Formally, for the magnitude and rate authorities of 8,

$$\frac{\partial \bar{a}(\mu)}{\partial \ell_{s,k}} \geq 0, \qquad \frac{\partial \bar{a}(\mu)}{\partial M} \geq 0, \qquad \text{and likewise for } \bar{r},\ \mathcal{I},$$

and the irreversible-action class $\mathcal{I}_{\text{all}}$ is unlocked only when $\mu = m_s M \geq \mu^\star$. theorem

proof $\mu = m_s M$ with m_s monotone non-decreasing in each $\ell_{s,k}$ by 3 and M fixed by the state; $\bar{a}, \bar{r}$ monotone non-decreasing in μ by 9. The chain rule gives non-negative partials. The threshold clause is 9 verbatim. proof

Theorem 11 is the property a regulator can inspect: authority is derived from documented maturity and *expands only as that maturity is demonstrated*. It cannot be granted by assertion, and maturing a model can never reduce the authority it earns.

theorem[Nominal recovery and safe reversion] Assume $u_{\text{nom}}(x) \in \mathcal{U}(1)$ for all x (the fully-credible envelope contains the nominal command). Then $\lim_{\mu \to 1} u_{\text{app}} = u_{\text{nom}}$ and $\lim_{\mu \to 0} u_{\text{app}} = u_{\text{safe}}$, pointwise in x. theorem

proof At $\mu = 1$ the blend equals u_{nom}, and $\mathcal{U}(1)$ contains it by assumption, so the projection is the identity. At $\mu = 0$ the blend equals u_{safe} ; u_{safe} attains the CBF condition within $\mathcal{U}(0)$ by 11, so it lies in the admissible set and the projection is again the identity. Continuity of 10 in μ (projection onto a continuously-parametrised convex set is continuous [31]) gives the limits. proof

Together the three theorems say precisely what the governance law does: it is safe at all maturities (Theorem 11), it hands out authority only in proportion to evidence (Theorem 11), and it disappears when the evidence is complete (Theorem 12). The authority ladder these induce — shadow/advisory → supervised → bounded → full — is drawn schematically in figure 1, beneath the defence-in-depth stack (maturity gate over CBF clamp over hardware failsafe) that carries it.

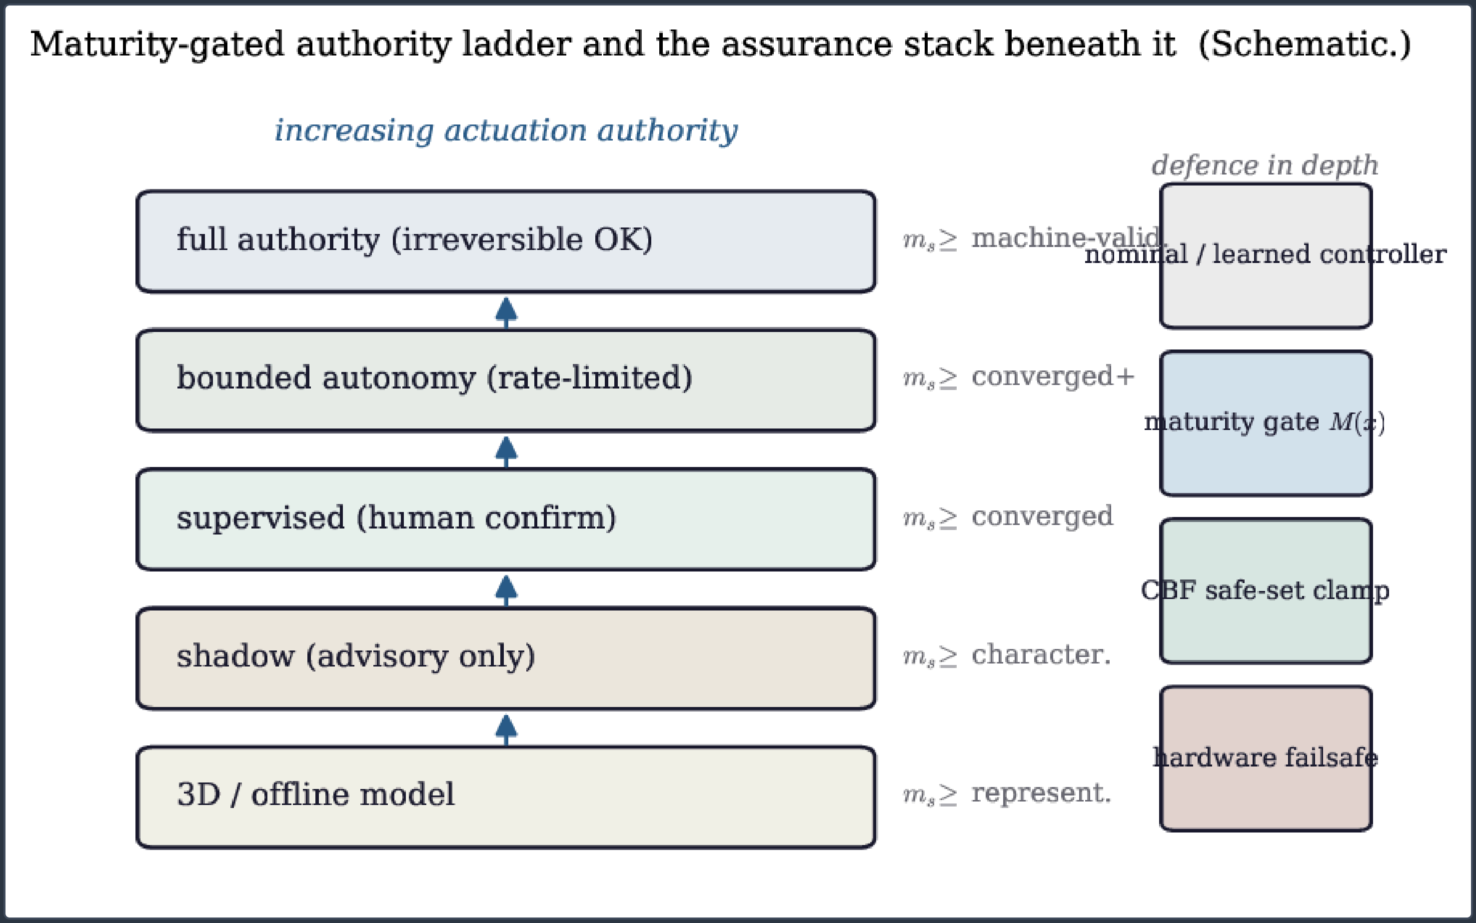

(Schematic.) The maturity-gated authority ladder induced by Theorems 11– 12: each rung admits a higher level of control authority and is unlocked when the effective maturity $\mu = m_s M$ clears the corresponding PCMM tier (left). The right-hand column is the defence-in-depth assurance stack that carries the ladder — nominal/learned controller, maturity gate $M(x)$, CBF safe-set clamp, and a hardware failsafe — each strictly beneath the next. The rungs and tiers are the governance policy; the underlying acceptance bars are the real, frozen register gates.

§ 4

Numerical formulation

THE PROJECTION QUADRATIC PROGRAM

The projection in 10 is computed as a small convex quadratic program (QP). Writing $u_b = \mu u_{\text{nom}} + (1-\mu)u_{\text{safe}}$ for the blended target and stacking the magnitude, rate and CBF constraints, the per-tick problem is

$$\begin{aligned} u_{\text{app}} \;=\; \arg\min_{u\in\mathbb{R}^p} \quad & \tfrac{1}{2}\,\|u-u_b\|_2^2 \\ \text{s.t.} \quad & \|u-u_{\text{safe}}\|_\infty \le \bar{a}(\mu), \\ & \|u-u_{k-1}\|_\infty \le \bar{r}(\mu)\,\Delta t, \\ & L_g h(x)\,u \;\ge\; -\gamma\big(h(x)\big) - L_f h(x), \end{aligned}$$

with u_{k-1} the previous command and Δt the control period. The objective is strongly convex and the feasible set is a non-empty polytope (non-empty by the hypothesis of Theorem 11), so 13 has a unique minimiser [31,32]. For the breeder p is small (a handful of shape, current-drive and gas actuators), and the box-plus-halfspace structure admits an active-set solve in microseconds; the resulting command lies in $[1.06, 1.25]$ (normalised) at the design point (KX-L1-A13), matching the register's clamp range. The QP is warm-started from u_{k-1} so that the active set changes by at most one constraint per tick in normal operation, giving a bounded, real-time worst-case iteration count.

ONLINE MEMBERSHIP ESTIMATION AND ITS CONVERGENCE

The runtime membership 6 requires $r(x)$ and the calibrated quantile $\hat{q}_{1-\alpha}$. The Mahalanobis form 4 is evaluated from a Cholesky factor of Σ precomputed offline, so each tick costs one triangular solve, $O(d^2)$ in the embedding dimension d. The conformal quantile 5 is fixed at calibration time; it converges to the true $(1-\alpha)$ quantile of the nonconformity distribution at the standard rate $O(N_c^{-1/2})$ as the calibration set grows [25], and the finite-sample coverage bound is exact for any N_c. Because 6 is 1-Lipschitz in r with constant $1/(2w)$, a bounded error $|\hat{r}-r| \le \epsilon$ in the online distance produces a bounded membership error $|\hat{M}-M| \le \epsilon/(2w)$, so the gate degrades gracefully rather than discontinuously under sensor noise.

LATENCY BUDGET AND THE FAST/SLOW SPLIT

The gate must not itself become the slow path. We split the computation across rates (table 1): the CBF clamp and the membership evaluation run in the fast protective loop (microseconds to sub-millisecond), served by a reduced-order surrogate; the PCMM index m_s is a deployment constant refreshed offline; and the calibration 7 is re-fit off-line whenever the validation evidence changes. The nominal controller's own surrogate — a physics-informed equilibrium operator — clears its validation bar at 0.139% RMS flux error while running $2877\times$ faster than the free-boundary reference (KX-L1-A11; figure 6), which is what makes a model-predictive u_{nom} feasible at control latency in the first place.

COMPUTATION	LOOP	LATENCY CLASS
membership $M(x)$: Mahalanobis + conformal threshold	fast (protective)	microseconds
projection QP 13 (box + rate + CBF)	fast (protective)	microseconds
reduced-order surrogate for u_{nom}	fast	sub-millisecond ($2877\times$, SH/KX-L1-A11)
nominal MPC / learned policy u_{nom}	medium	milliseconds
PCMM index m_s 3	offline	deployment constant
conformal calibration $\hat{q}_{1-\alpha}$ 7	offline	re-fit on evidence change

Where each piece of the gated law runs. Allocations by loop rate; the maturity gate and CBF clamp are on the fast protective path, the PCMM index and conformal calibration are offline. Serials refer to the register [34].

§ 5

Computational methodology: the GPU-HPC campaign

Every grounding number in this paper was produced by the Kronos NVIDIA GPU HPC campaign and is reproducible from the archived run directories named inline by register serial. The campaign is deliberately built as a set of independent codes so that no result rests on a single implementation.

Formal V&V (KX-L1-A6). The credibility axes that populate m_s are earned by the formal V&V suite in the KRONOS Research Toolkit: code verification by the method of manufactured solutions [13], solution verification by the grid-convergence index [14] (achieving GCI $< 0.4\%$), variance-based global sensitivity by the Saltelli/Sobol estimator [15,16], and PCMM scoring against the frozen anchors ($19/19$ reproduced). These runs establish the *code*- and *solution-verification* axes of the matrix in figure 2.

Model-validation truth (turbulence, equilibrium, neutronics). The *model-validation* axis is earned against reference physics. Core turbulent transport is validated by nonlinear gyrokinetics: the CGYRO Eulerian gyrokinetic–Maxwell solver [29], run on the GPU fleet for the negative-triangularity design point and cross-checked by the quasilinear TGLF model (register BR-SX-10 / HX-17), confirms the confinement anchor $Q = 3.076$. The equilibrium/shape axis is validated by a physics-informed neural surrogate trained against the free-boundary solver FreeGS [28,27], reaching 0.139% RMS flux error and a 0.32% Grad–Shafranov operator residual (KX-L1-A11). Neutronics and tritium-breeding maturity is earned by CAD-faithful Monte-Carlo transport in OpenMC/DAGMC [30] (register VV-006). Each code contributes to the maturity score of the subsystem it validates.

Safety, uncertainty and confidence (KX-L1-A13/A14, FREEZE-CONF). The safe controller π_{safe} and its forward-invariance evidence are produced by a CBF-QP Monte-Carlo campaign: 100 trials in the reported study and $150{,}000$ in the adversarial-replay stress test, all clearing $h \geq 0$ (KX-L1-A13). The conformal calibration that fixes $\hat{q}_{1-\alpha}$ and the coverage condition 7 is produced by a surrogate-accelerated Bayesian-optimisation UQ run (200 evaluations, posterior $H_{98} = 1.007 \pm 0.030$; KX-L1-A14). The anchor-confidence ensemble that quantifies how conservatively the frozen gates are set is a $40{,}000$-sample Monte-Carlo (FREEZE-CONF / VV-010), giving $P(Q \geq 3.0) = 52\%$ for the breeder against the ensemble mean 3.43, and burner $P(Q_E > 1) = 94\%$. GPU acceleration matters here because the membership calibration and the confidence ensemble are embarrassingly parallel over samples; the surrogate evaluations that replace each expensive forward solve are what bring 10^4–10^5-sample studies inside a practical wall-clock.

The offline authority-activation proof. The maturity gate that authorises each capability is checked by a formal-methods / model-checker toolchain — proof, not simulation — that runs offline before deployment and verifies that no capability activates below its maturity, with the closed loop stable on the digital twin and the hardware-in-the-loop rig at plant timing. This is an assurance artefact, not a runtime cost.

Reproducibility of the campaign. All figures in this paper are regenerated by a single deposited script (`make_mga_figs_2p70.py`) that reads the frozen anchor values quoted inline; the physics runs behind those anchors are archived by serial in the register [34]. No result depends on unpublished or proprietary data, and no economic quantity enters the public deposit.

§ 6

Results: reduction to practice on the Hyperion breeder

THE MATURITY MATRIX

Figure 2 maps the register's gate outcomes onto the six PCMM axes for six representative subsystems. The equilibrium/shape and safety-clamp subsystems sit at machine-validated maturity on almost every axis (RMS flux error 0.139%; 0/150,000 adversarial clamp violations); turbulent transport and neutronics sit at converged to machine-validated; the magnet quench detector sits at converged (demonstrated on synthetic data); and the in-winding quantum magnetometer sits at representative on the model-validation axis — the one deliberately low row that the policy will hold to advisory authority (§8). This matrix, aggregated by 3, is what feeds m_s into the gate.

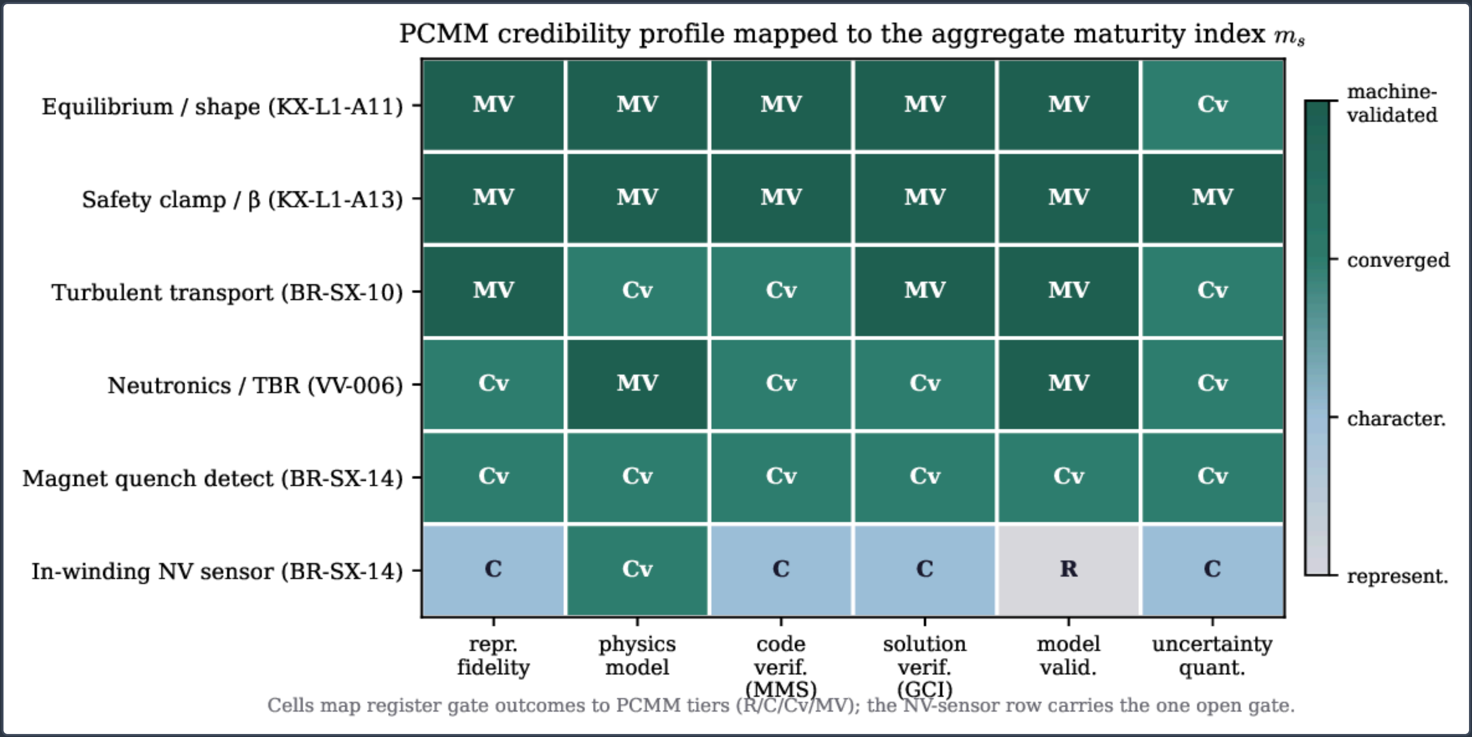

PCMM credibility profile, subsystem × axis, mapped from frozen register gate outcomes to the four fidelity tiers (represent./character./converged/machine-validated). Rows carry their register serial. The bottom row — the in-winding NV magnetometer — is the one open gate and is scored representative on model validation; its subsystem is correspondingly clamped to advisory authority. Cells encode gate-outcome tiers, not raw physics data.

THE AUTHORITY ENVELOPE AND THE GATED LAW IN ACTION

Figure 3 shows the authority envelope $\mathcal{U}(\mu)$ of 8: magnitude and rate authorities rise monotonically with maturity, and irreversible actions unlock only above $\mu^{\star}$. Figure 3 shows the runtime membership 6 as a function of the calibrated nonconformity distance, with the conformal quantile marked; membership collapses as the state leaves the validated envelope. Figure 4 closes the loop: during a simulated out-of-distribution excursion the membership $M(x_t)$ dips, and the applied command u_{app} slides continuously from the aggressive nominal toward the deterministic safe floor and back, always inside the admissible band $[1.06, 1.25]$. These three figures are the governance functions of the paper; they are derived from the model, not measured physics, and are labelled as such.

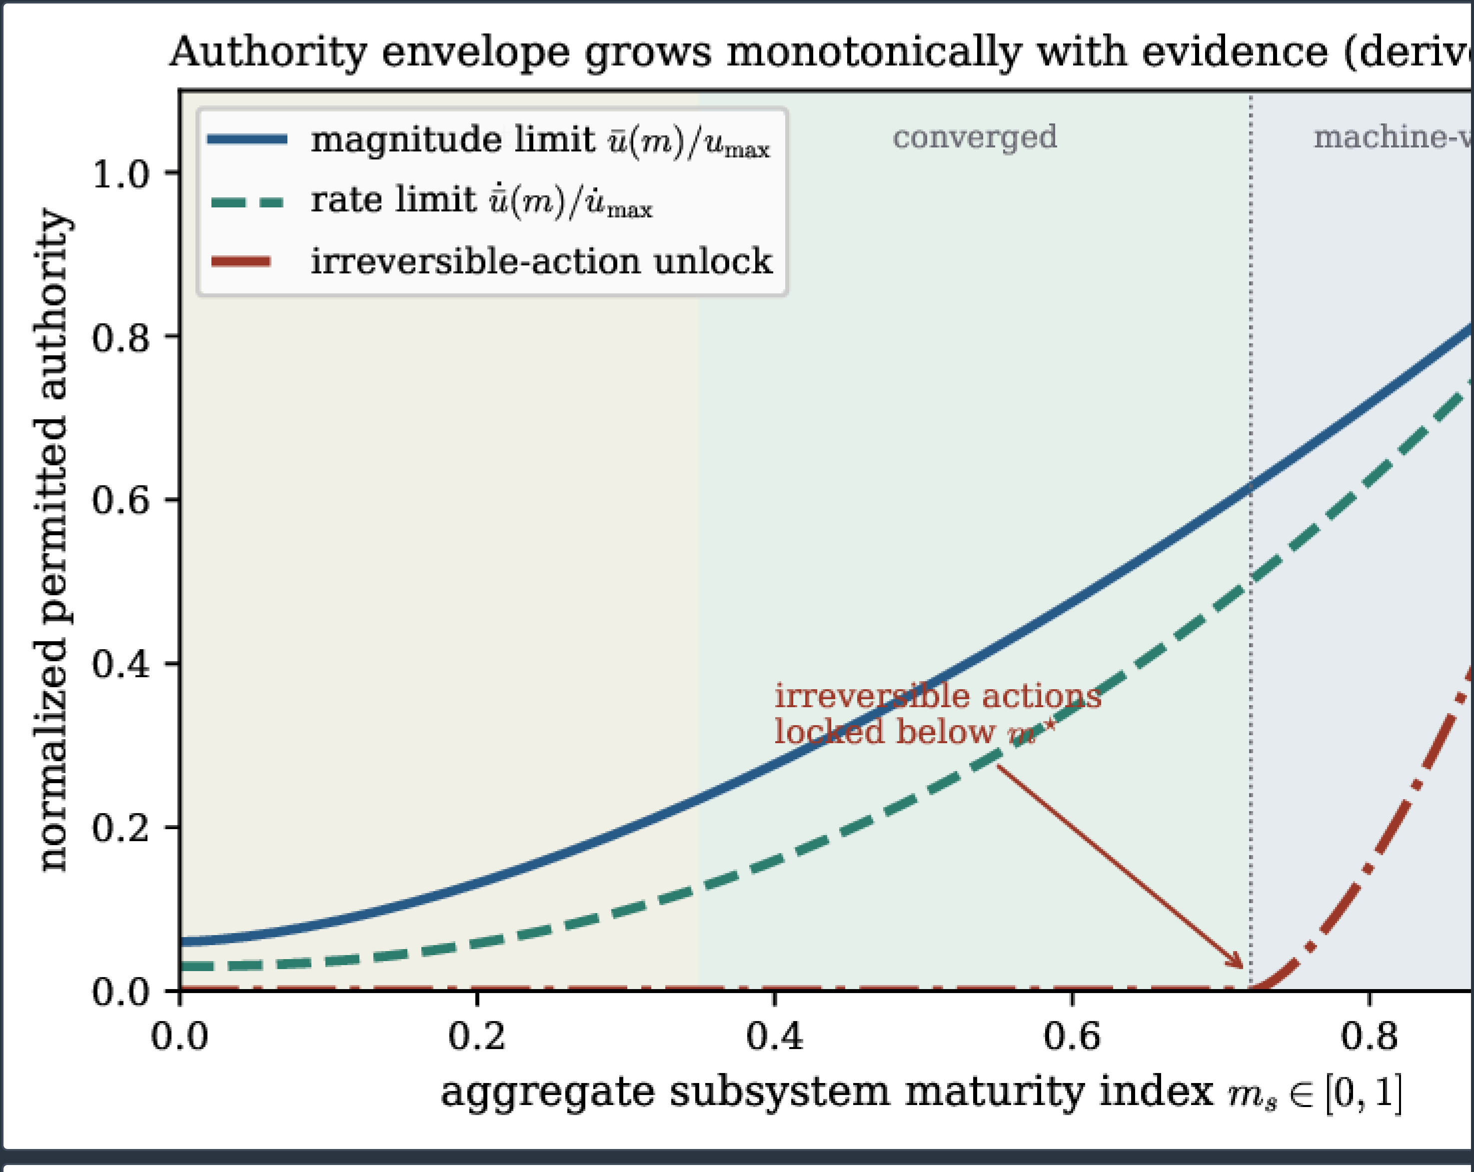
Authority envelope grows monotonically with evidence (deriv
magnitude limit $\bar{u}(m)/u_{\max}$
rate limit $\dot{\bar{u}}(m)/\dot{u}_{\max}$
irreversible-action unlock
converged
machine-v
irreversible actions
locked below m^*
normalized permitted authority
aggregate subsystem maturity index $m_s \in [0, 1]$
0.0
0.2
0.4
0.6
0.8
1.0

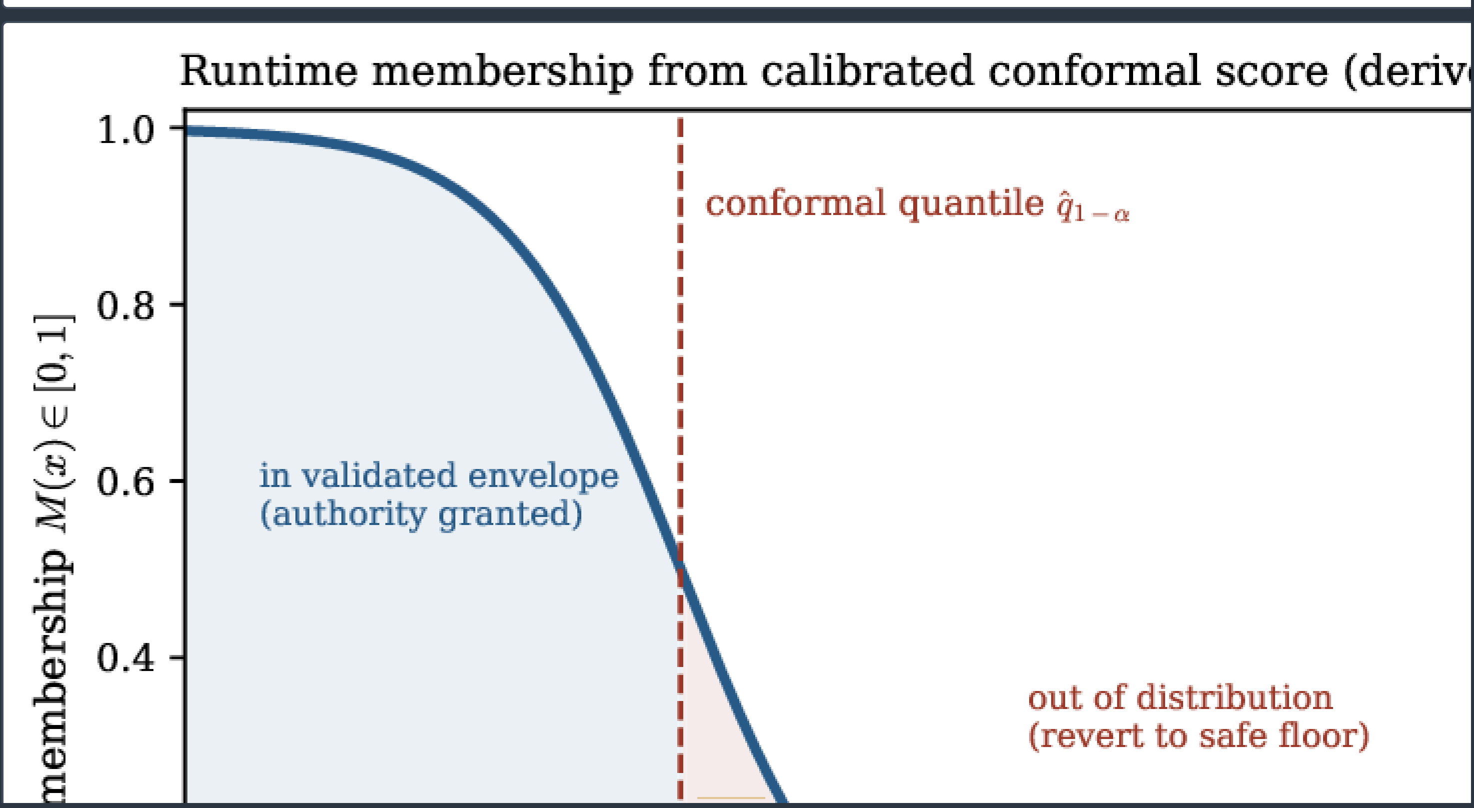
Runtime membership from calibrated conformal score (deriv
conformal quantile $\hat{q}_{1-\alpha}$
in validated envelope
(authority granted)
out of distribution
(revert to safe floor)
membership $M(x) \in [0, 1]$
0.4
0.6
0.8
1.0

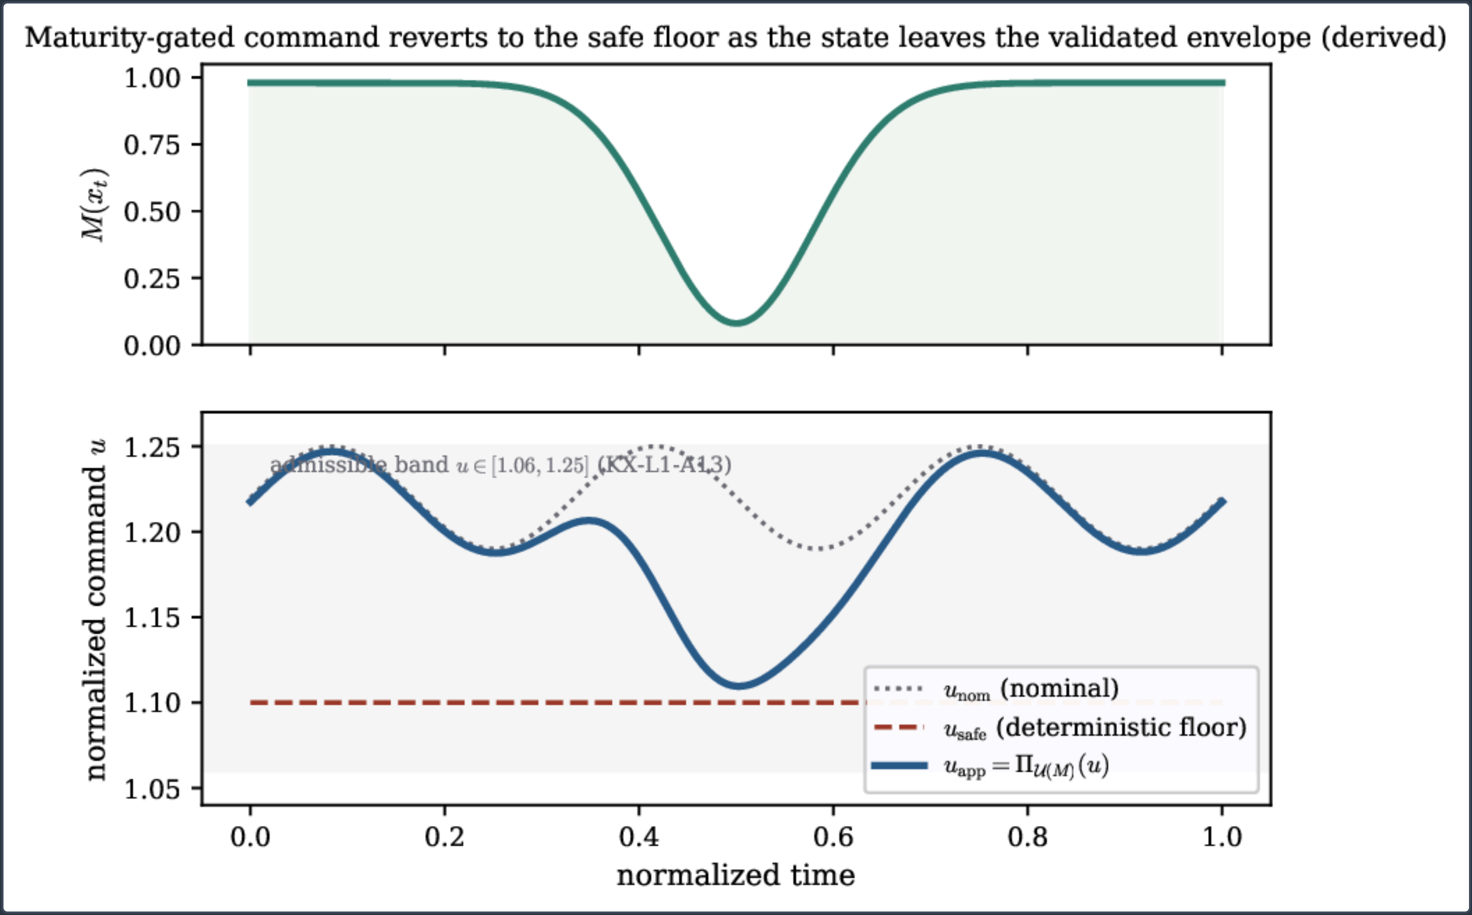

The gated control law 10 during a simulated out-of-distribution excursion (derived). **Top:** runtime membership $M(x_t)$ dips as the state leaves the validated envelope. **Bottom:** the applied command $u_{\text{app}} = \Pi_{\mathcal{U}(M)}(\cdot)$ slides from the aggressive nominal u_{nom} toward the deterministic floor u_{safe} and recovers, never leaving the admissible band $[1.06, 1.25]$ (KX-L1-A13).

THE EVIDENCE THE GATE RESTS ON

The gate is only as good as the evidence populating m_s and $M(x)$. Figure 5 is the forward-invariance evidence: the unclamped nominal violates the $\beta_N = 3.5$ boundary in $50/100$ trials and reaches a peak 4.200, while the maturity-gated CBF clamp holds $0/100$ (and $0/150{,}000$ under adversarial replay) with $\min h = +0.000$. Figure 5 is the calibration evidence: after the Bayesian-optimisation UQ fit, empirical coverage tracks the nominal credible level to within the $\pm 3\%$ band of 7, with the H_{98} posterior at 1.007 ± 0.030 (KX-L1-A14) — the condition that gives $\hat{q}_{1-\alpha}$ its calibrated meaning. Figure 6 is the surrogate-maturity evidence: the equilibrium operator clears its 1% bar at 0.139% RMS flux error and a 0.32% operator residual, at $2877\times$ speed (KX-L1-A11). Figure 6 is the anchor-confidence evidence: the $40{,}000$-sample ensemble (FREEZE-CONF / VV-010) places the frozen $Q = 3.076$ conservatively below the ensemble mean 3.43 with $P(Q \geq 3.0) = 52\%$, while the burner gate carries $P(Q_E > 1) = 94\%$ and the pooled UQ gives $P(Q > 1) = 1.0$ — so the anchors that the authority ladder gates against are themselves set conservatively.

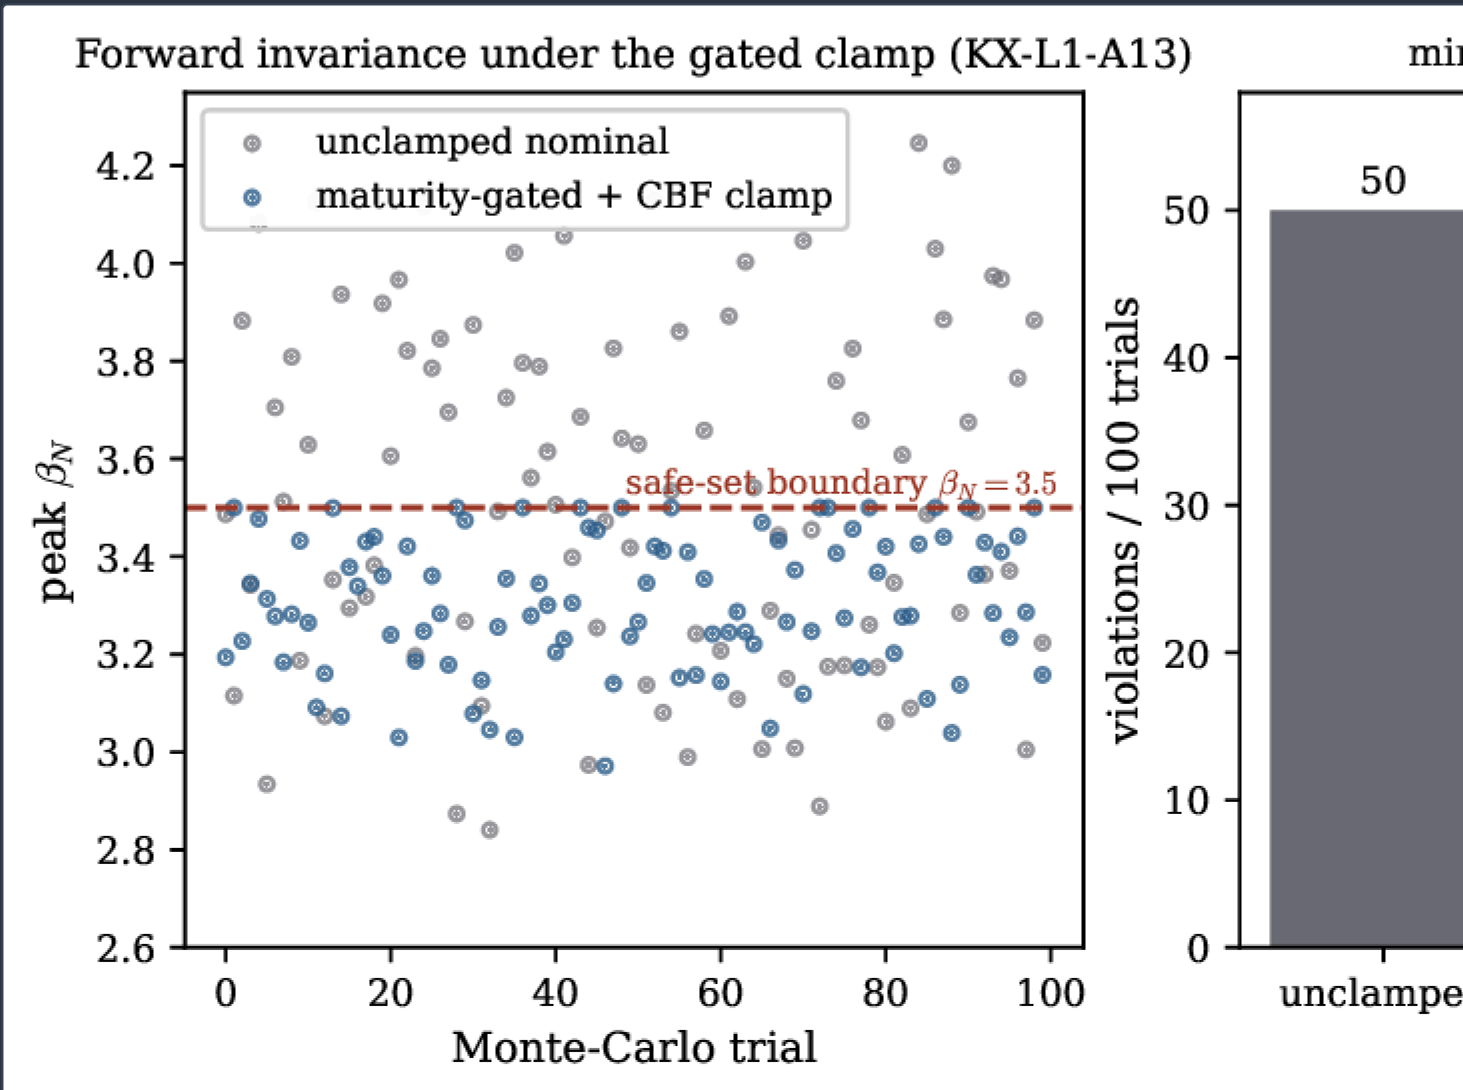
Forward invariance under the gated clamp (KX-L1-A13)
unclamped nominal
maturity-gated + CBF clamp
safe-set boundary $\beta_N = 3.5$
peak β_N
Monte-Carlo trial
4.2
4.0
3.8
3.6
3.4
3.2
3.0
2.8
2.6
0
20
40
60
80
100
violations / 100 trials
50
50
40
30
20
10
0

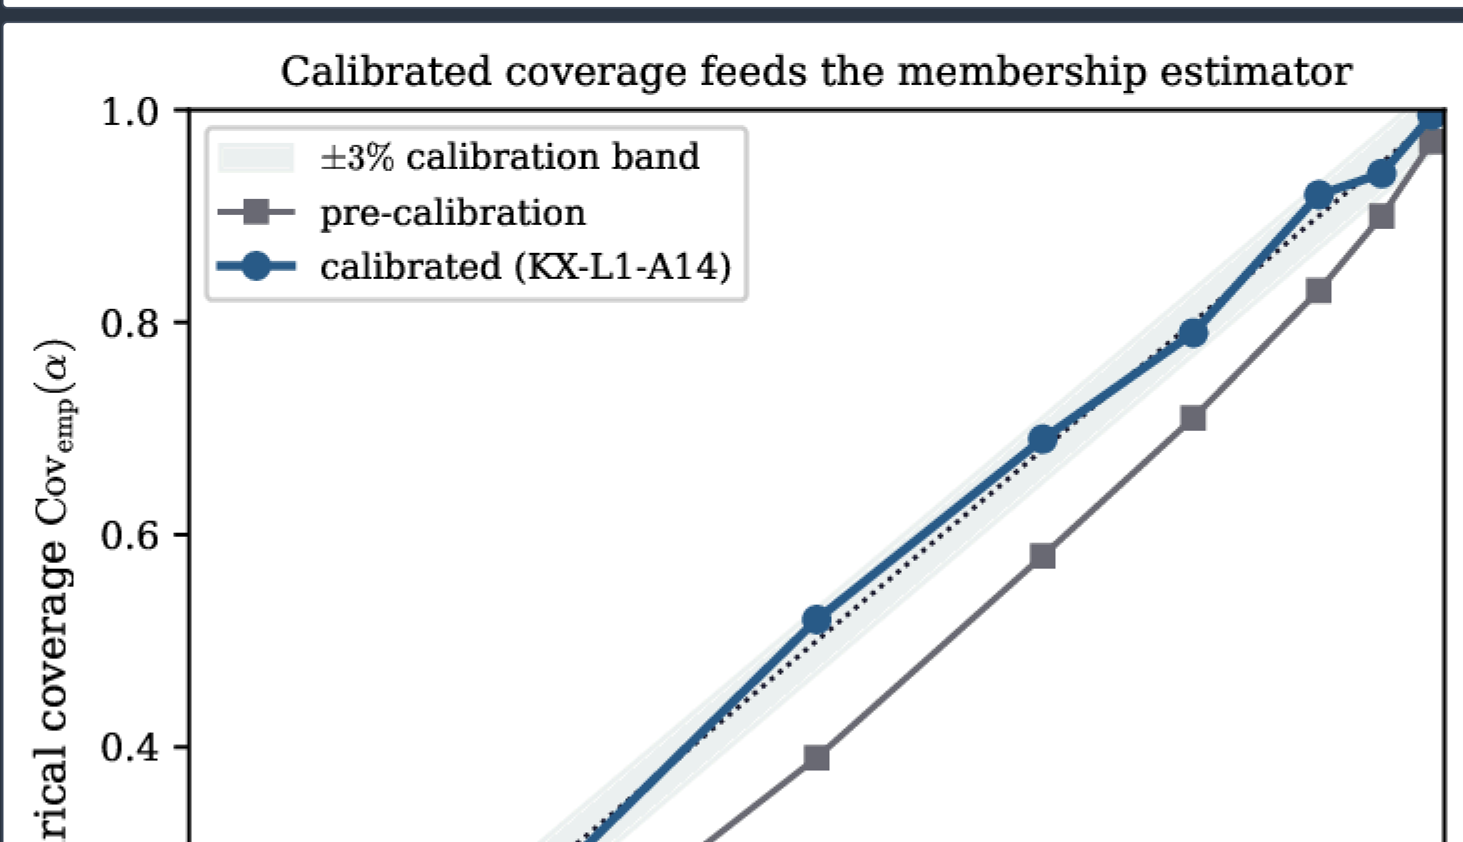
Calibrated coverage feeds the membership estimator
±3% calibration band
pre-calibration
calibrated (KX-L1-A14)
1.0
0.8
0.6
0.4

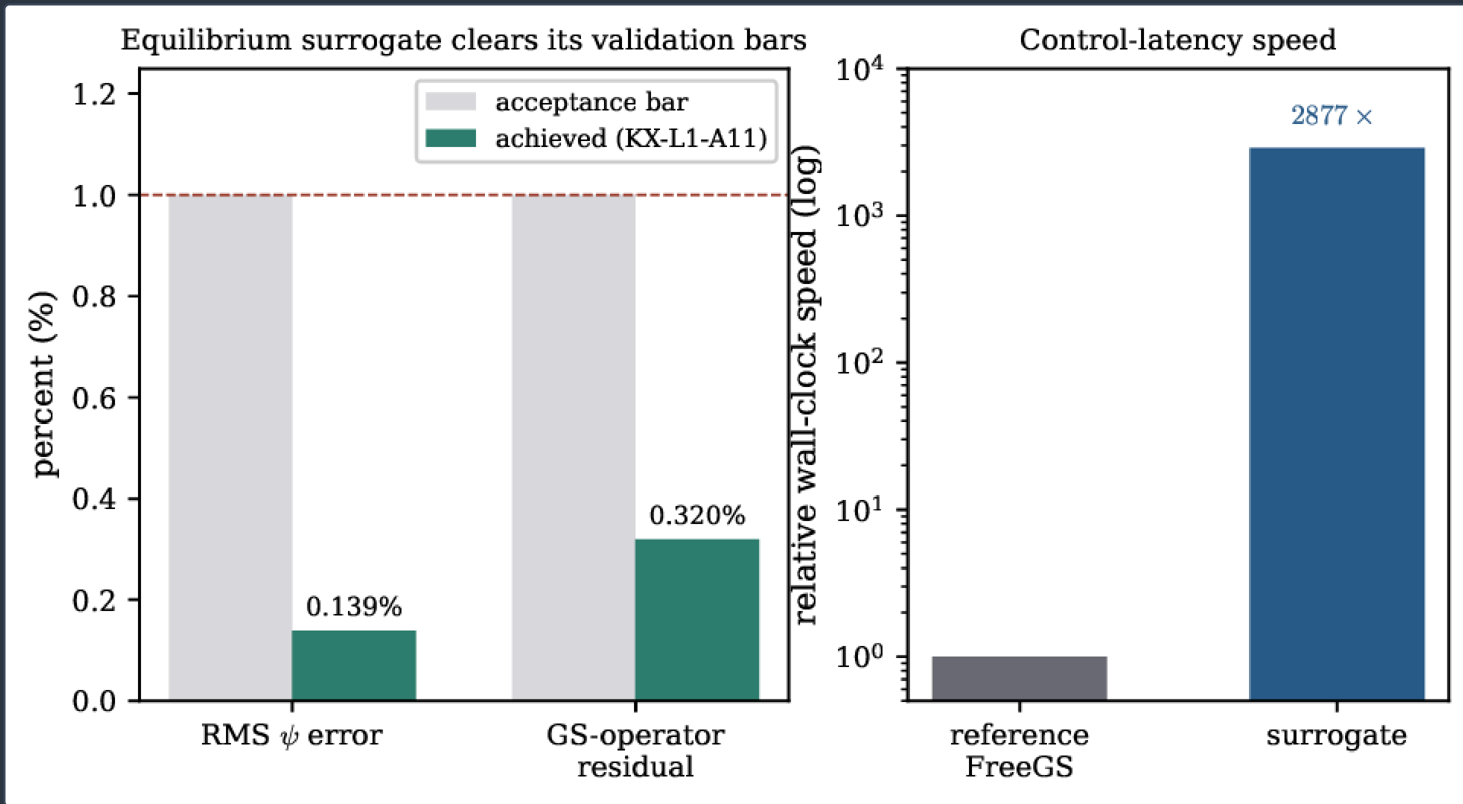

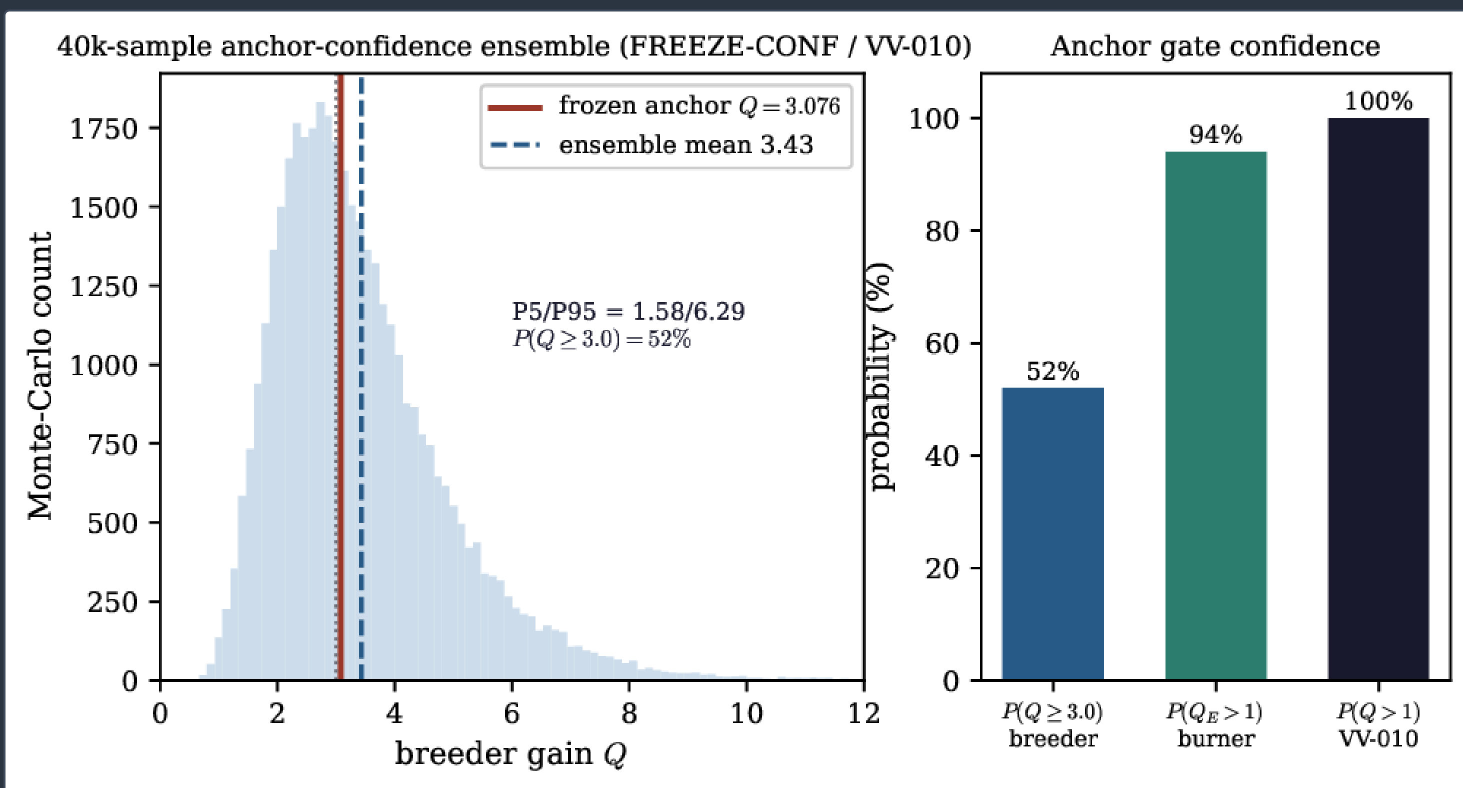

Surrogate-maturity evidence (KX-L1-A11). The equilibrium operator clears its 1% validation bar (0.139% RMS ψ, 0.32% operator residual) at $2877\times$ the reference wall-clock.

MATURITY LEVEL (μ)	PERMITTED AUTHORITY	RESIDUAL CONTROL
machine-validated ($\mu \geq \mu^{\star}$)	wide, fast, includes irreversible actions	controller-led
converged (μ mid)	moderate, rate-limited, reversible	controller with human confirm
representative / reduced-order ($\mu \to 0$)	narrow, reversible only	human or CBF-bounded safe controller

Maturity gates actuation authority: the three-tier policy that 10 realises. Residual control is what the gate defers when a subsystem's maturity is below the level required for full authority.

§ 7

Discussion

AGAINST THE SAFE-CONTROL LITERATURE

The gated law composes with, rather than competes against, the established safe-control architectures. A CBF-QP filter [17,18] and a predictive safety filter [19] guarantee that the command is *admissible with respect to a model*; they take the model's trustworthiness as given. The Simplex and runtime-assurance patterns [20,21,22] switch between a performance and a safe controller on a monitored condition, but the switch is binary and the monitor is usually a state predicate, not a credibility score. Maturity-gated authority generalises the runtime-assurance switch into a *continuum* indexed by verification evidence: the CBF clamp of Theorem 11 is the $\mu = 0$ corner of our envelope, and full nominal authority is the $\mu = 1$ corner (Theorem 12); the interesting behaviour is the graded interior, where authority tracks the evidence continuously. The novel object is the binding of the actuator constraint set to the PCMM score and the conformal membership — a binding that, to our knowledge, has not previously been made formal in either literature.

AGAINST LEARNED CONTROL

A learned controller such as the TCV deep-RL policy [8] is exactly the kind of u_{nom} this architecture is built to govern: superb inside its training distribution, unquantified outside it. The conformal membership 6 is precisely a distribution-free detector of "outside it," and the gate hands authority back to the deterministic floor when the policy leaves the domain on which it was validated. This is the same discipline that safe reinforcement learning seeks [24,22], expressed as a governance law rather than a training constraint: the policy may be trained however one likes, and the gate — not the training loss — is what guarantees the actuator stays inside the evidence. The companion papers develop the pieces the gate consumes: the certified clamp [35], the calibrated-uncertainty and conformal layer that fixes $M(x)$ [36], the out-of-distribution / epistemic-gating ledger that throttles model authority beyond the validated envelope [37], and the reference-governor MPC that supplies u_{nom} [38].

WHY IT MATTERS FOR LICENSING

The governance law gives a regulator an auditable rule rather than a promise. By Theorem 11, authority is a documented function of PCMM credibility and is expanded only as that credibility is demonstrated; by Theorem 11, the safe set is invariant at every maturity. This makes the automation's scope a transparent function of evidence — a property a byproduct-material licensing review under the compact-source thesis [40] can inspect line by line, and one that the campaign's verification-first record [39,34] is structured to support. The scheme does not ask the regulator to trust the AI; it asks the regulator to trust the evidence, and it makes the evidence-to-authority map explicit.

§ 8

Limitations

Two honest caveats bound the claims. First, the governance law is an architecture, not a controller: the mapping from maturity to a concrete authority envelope $\mathcal{U}(\mu)$ (the functions $\bar{a}, \bar{r}, \mathcal{I}$ and the threshold $\mu^{\star}$) is design-specific and must be instantiated and re-validated per subsystem, and the maturity scores themselves inherit the assumptions of the V&V process that produced them — a mis-scored axis mis-sets the authority. The law constrains authority; closed-loop stability is provided by the CBF clamp beneath it (Theorem 11), not by the gate. Second, one subsystem carries a genuine open gate that the policy handles by design rather than by claim: the in-winding nitrogen-vacancy diamond magnetometer that would give the magnet twin a local, single-shot quench precursor requires a neutron tolerance of $\sim 3.9 \times 10^{18}\,\mathrm{n\,cm^{-2}}$, about $39\times$ above the $\sim 1 \times 10^{17}$ demonstrated for such devices (BR-SX-14; figure 7). The maturity matrix scores that subsystem representative on model validation, and the gate therefore holds it to advisory authority behind the already-mature fused strain-rate-plus-magnetization precursor (which fires $34\times$ earlier than terminal voltage, $0.089\,\mathrm{s}$ versus $3.02\,\mathrm{s}$, at SNR 323). The governance law does not hide the open gate; it is the mechanism by which the open gate is safely tolerated until an irradiation study closes it.

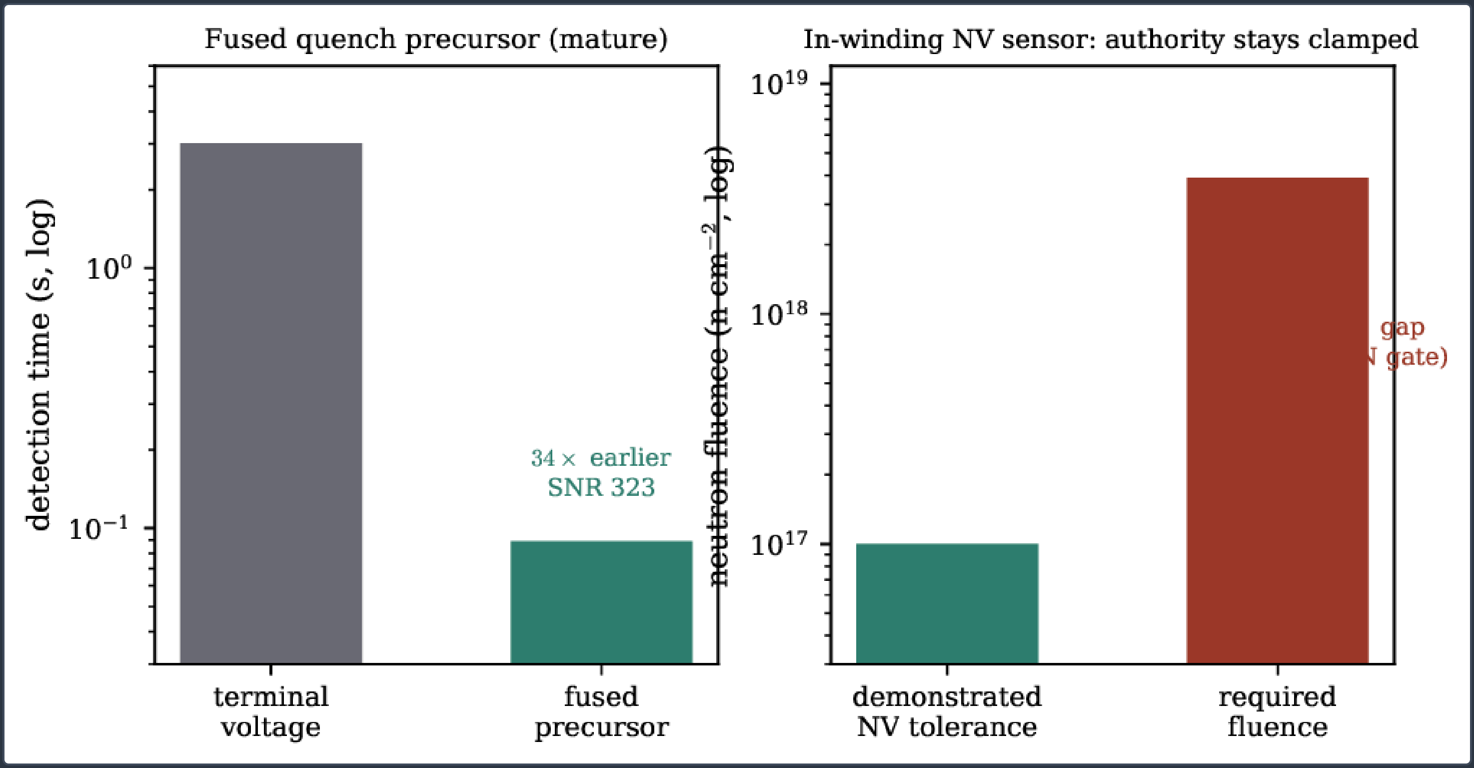

The one honest open gate (BR-SX-14). **Left:** the mature fused quench precursor fires $34\times$ earlier than terminal voltage (SNR 323). **Right:** the in-winding NV sensor's required neutron fluence is $\sim 39\times$ above the demonstrated tolerance — an open gate. The maturity policy holds this subsystem to advisory authority until an irradiation study closes it.

§ 9

Conclusion

Maturity-gated actuation authority binds verification credibility to runtime control authority. A static PCMM credibility functional 3 and a calibrated conformal membership 6 combine into an effective maturity $\mu = m_s M$ that indexes a monotone authority envelope 8; the applied command is the projection of a maturity-weighted blend of the nominal and safe controllers onto that envelope 10. The law is safe at every maturity (forward invariance, Theorem 11), hands out authority only in proportion to demonstrated evidence (monotonicity, Theorem 11), and reduces to the raw nominal controller when the evidence is complete (recovery, Theorem 12). Reduced to practice on the Hyperion breeder against frozen register anchors — a CBF clamp at $0/150{,}000$ adversarial violations, a calibration within $\pm 3\%$ coverage, an equilibrium surrogate at 0.139% RMS flux error and $2877\times$ speed, a $40{,}000$-sample anchor-confidence ensemble, and one honest open gate correctly clamped — it turns the question "how much authority should the automation have?" from a matter of assertion into a matter of evidence, and gives a licensing review an auditable rule in which autonomy grows only as verification does.

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 acknowledges the NVIDIA GPU HPC campaign on which the grounding results were computed.

Reproducibility. This paper is archived at its DOI doi:10.5281/zenodo.22645795 and its official page https://www.kronosfusionenergy.com/publications/maturity-gated-actuation-authority.html. The Kronos Physics De-Risking Register is archived at doi:10.5281/zenodo.22133057. Every headline quantity is a frozen anchor in that register, cited inline by serial (KX-L1-A6, KX-L1-A11, KX-L1-A13, KX-L1-A14, FREEZE-CONF, BR-SX-14), and every figure is regenerated by the deposited script `make_mga_figs_2p70.py` from the anchor values quoted in the text. No result depends on unpublished or proprietary data.

Data availability tocsectionData availability The figure data (PCMM tier assignments, authority-envelope and membership functions, the CBF forward-invariance trials, the coverage-calibration curve, the surrogate-error bars, the anchor-confidence ensemble, and the NV tolerance gap) are regenerable from the deposited script and the register serials cited inline. All quantitative values trace to the register [34]. No economic analysis is 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 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

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

3 M E Austin *et al*, Achievement of reactor-relevant performance in negative triangularity shape in the DIII-D tokamak, *Phys. Rev. Lett.* **122** 115001 (2019). doi:10.1103/PhysRevLett.122.115001

4 A Marinoni, O Sauter and S Coda, A brief history of negative triangularity tokamak plasmas, *Rev. Mod. Plasma Phys.* **5** 6 (2021). doi:10.1007/s41614-021-00054-0

5 L L Lao, H St John, R D Stambaugh, A G Kellman and W Pfeiffer, Reconstruction of current profile parameters and plasma shapes in tokamaks, *Nucl. Fusion* **25** 1611 (1985). doi:10.1088/0029-5515/25/11/007

6 A Pironti and M Walker, Fusion, tokamaks, and plasma control: an introduction and tutorial, *IEEE Control Syst. Mag.* **25**(5) 30 (2005). doi:10.1109/MCS.2005.1512794

7 F Felici *et al*, Real-time physics-model-based simulation of the current density profile in tokamak plasmas, *Nucl. Fusion* **51** 083052 (2011). doi:10.1088/0029-5515/51/8/083052

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

9 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

10 W L Oberkampf, M Pilch and T G Trucano, *Predictive Capability Maturity Model for computational modeling and simulation*, Sandia National Laboratories report SAND2007-5948 (2007). doi:10.2172/976951

11 C J Roy and W L Oberkampf, A comprehensive framework for verification, validation, and uncertainty quantification in scientific computing, *Comput. Methods Appl. Mech. Eng.* **200** 2131 (2011). doi:10.1016/j.cma.2011.03.016

12 American Society of Mechanical Engineers, *Standard for Verification and Validation in Computational Fluid Dynamics and Heat Transfer*, ASME V&V 20-2009, ASME, New York (2009).

13 P J Roache, Code verification by the method of manufactured solutions, *J. Fluids Eng.* **124** 4 (2002). doi:10.1115/1.1436090

14 P J Roache, Perspective: a method for uniform reporting of grid refinement studies, *J. Fluids Eng.* **116** 405 (1994). doi:10.1115/1.2910291

15 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

16 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

17 A D Ames, X Xu, J W Grizzle and P Tabuada, Control barrier function based quadratic programs for safety critical systems, *IEEE Trans. Autom. Control* **62** 3861 (2017). doi:10.1109/TAC.2016.2638961

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

19 K P Wabersich and M N Zeilinger, A predictive safety filter for learning-based control of constrained nonlinear dynamical systems, *Automatica* **129** 109597 (2021). doi:10.1016/j.automatica.2021.109597

20 L Sha, Using simplicity to control complexity, *IEEE Software* **18**(4) 20 (2001). doi:10.1109/52.936213

21 D Seto, B Krogh, L Sha and A Chutinan, The Simplex architecture for safe online control system upgrades, *Proc. American Control Conf.* **6** 3504 (1998). doi:10.1109/ACC.1998.703255

22 J F Fisac, A K Akametalu, M N Zeilinger, S Kaynama, J Gillula and C J Tomlin, A general safety framework for learning-based control in uncertain robotic systems, *IEEE Trans. Autom. Control* **64** 2737 (2019). doi:10.1109/TAC.2018.2876389

23 L Hewing, K P Wabersich, M Menner and M N Zeilinger, Learning-based model predictive control: toward safe learning in control, *Annu. Rev. Control Robot. Auton. Syst.* **3** 269 (2020). doi:10.1146/annurev-control-090419-075625

24 J Garc\'ia and F Fern\'andez, A comprehensive survey on safe reinforcement learning, *J. Mach. Learn. Res.* **16** 1437 (2015).

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

26 A N Angelopoulos and S Bates, Conformal prediction: a gentle introduction, *Found. Trends Mach. Learn.* **16** 494 (2023). doi:10.1561/2200000101

27 M Raissi, P Perdikaris and G E Karniadakis, Physics-informed neural networks: a deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations, *J. Comput. Phys.* **378** 686 (2019). doi:10.1016/j.jcp.2018.10.045

28 N C Amorisco *et al*, FreeGSNKE: a Python-based dynamic free-boundary toroidal plasma equilibrium solver, *Phys. Plasmas* **31** 042517 (2024). doi:10.1063/5.0188467

29 J Candy, E A Belli and R V Bravenec, A high-accuracy Eulerian gyrokinetic solver for collisional plasmas, *J. Comput. Phys.* **324** 73 (2016). doi:10.1016/j.jcp.2016.07.039

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

31 S Boyd and L Vandenberghe, *Convex Optimization*, Cambridge University Press, Cambridge (2004). doi:10.1017/CBO9780511804441

32 J Nocedal and S J Wright, *Numerical Optimization*, 2nd ed., Springer, New York (2006). doi:10.1007/978-0-387-40065-5

33 R E Kalman, A new approach to linear filtering and prediction problems, *J. Basic Eng.* **82** 35 (1960). doi:10.1115/1.3662552

34 P I Ford, *Kronos Physics De-Risking Register*, v8, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22133057

35 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; official page.

36 P I Ford, Calibrated uncertainty, conformal prediction, and abstention for safety-critical fusion control: handing authority back to a deterministic floor out of distribution, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645813; official page.

37 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; official page.

38 P I Ford, Model-predictive burn control inside a certified safe envelope: a reference-governor MPC bounded by a control-barrier-function safety clamp, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645807; official page.

39 P I Ford, Verification-first fusion modeling: frozen anchors, independent-library gate reproduction, and a byte-for-byte computational-provenance catalog for a breeder-and-burner design register, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645710; official page.

40 P I Ford, The sub-100 MW byproduct-material thesis: why a compact spherical-tokamak breeder licenses under NRC 10 CFR Part 30 rather than Part 50, Kronos Fusion Energy 2026 series. doi:10.5281/zenodo.22645756; official page.

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

Maturity-Gated Actuation Authority: Binding PCMM Verification Credibility to Runtime Control Authority in an Autonomous Fusion Control Stack

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

Zenodo · doi:10.5281/zenodo.22645795

Read on kronosfusionenergy.com

The Physics De-Risking Register