KRONOS FUSION ENERGY

PAPER 2.88 · THE KRONOS 2026 SERIES

Global Sensitivity and Uncertainty Quantification across the Breeder and Burner Design Planes: A 40,000-Sample Monte-Carlo Ranking of the Frozen Anchors

A frozen design point is only as trustworthy as its behaviour under joint parameter uncertainty, yet fusion performance is almost always reported as a single number for each...

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 Several quantities in the underlying corpus moved after first publication; where a value has been withdrawn or restated, the record says so plainly.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

Global Sensitivity and Uncertainty Quantification across the Breeder and Burner Design Planes: A 40,000-Sample Monte-Carlo Ranking of the Frozen Anchors

A frozen design point is only as trustworthy as its behaviour under *joint* parameter uncertainty, yet fusion performance is almost always reported as a single number for each figure of merit. We resolve two questions globally — *where does the design anchor sit in the distribution of achievable outcomes*, and *which uncertain input actually drives that distribution* — for the two Kronos machines, the Hyperion compact spherical-tokamak breeder ($Q = 3.076$, $P_{\text{fus}} = 85.04\,MW$, $I_p = 9.66\,MA$) and its D–^{3}He tandem-mirror burner ($Q_E = 1.318$, neutron fraction $f_n = 5.44\,\%$). The method is direct Monte-Carlo: two independent $40\,000$-sample ensembles are drawn from the physical input priors and propagated through the same reduced-order power- and particle-balance evaluators that fix the design points, on the NVIDIA GPU HPC campaign. We give the governing 0-D balances, the priors as explicit probability measures, and the variance-based (Sobol/ANOVA) sensitivity decomposition together with its Saltelli/Jansen estimator and Monte-Carlo error analysis. The headline is that both anchors sit comfortably *inside* their ensembles. The breeder anchor lies below the ensemble mean (Q mean 3.85, P5–P95 $= 1.34$–8.51; $P_{\text{fus}} = 85.5 \pm 15\,MW$), a conservative interior choice rather than an optimistic tail; the burner is stronger still, with Q_E centred at 1.317 ± 0.14 and engineering breakeven cleared on *every* sample, $P(Q_E > 1) = 1.0$. A variance-based ranking then reduces each machine to a single dominant lever: breeder gain is confinement-limited (first-order Sobol index $S_1 = 0.885$ for the H_{98} confinement multiplier; total effect $S_T = 0.905$) while burner gain is core-pressure-limited (Pearson correlation $r = 0.93$ with the core beta β_c). Both anchors are therefore robust centres, not fragile optima, and the design-and-control programme has a single quantified priority on each machine. The one honest caveat is that the burner ensemble holds plug density fixed and excludes radial transport; when plug density is unfrozen it becomes the leading term and the recirculation-resolved *net* gain sits closer to the $Q_E = 1$ line, a regime we bound explicitly.

Open-access deposit (code & data, CC BY 4.0): DOI 10.5281/zenodo.22645893 · Zenodo community *kronos_fusion_energy*. Explore it live: interactive 3-D model · physics validator.

Keywords: uncertainty quantification, global sensitivity, Sobol indices, Monte-Carlo, spherical tokamak, tandem mirror, fusion gain, design anchor, confinement

CONTENTS

Table of Contents

Global Sensitivity and Uncertainty Quantification across the Breeder and Burner Design Planes: A 40,000-Sample Monte-Carlo Ranking of the Frozen Anchors 3
Introduction 6
Governing equations and the formal UQ/GSA problem 7
Computational methodology: the GPU HPC campaign 10
Results 12
Discussion 21
Limitations 22
Conclusion 23
References 28
One programme, many papers 31

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

A DESIGN POINT IS A DISTRIBUTION, NOT A NUMBER

Design points in magnetic-confinement fusion are quoted as single numbers — a gain Q, a fusion power P_{fus}, a normalised beta β_N — but the inputs behind them are uncertain and they vary together. Confinement quality, density, temperature, core pressure and field each carry an irreducible epistemic uncertainty, and the map from those inputs to performance is nonlinear. A single quoted number therefore under-specifies the design: two anchors with identical central values can differ completely in how much of the achievable-outcome distribution lies above breakeven, and in which input a control programme must pin down first. Answering "how good is the design point" requires the *push-forward* of the joint input measure through the performance model, not a one-factor-at-a-time perturbation about the nominal.

This paper answers, globally and for both Kronos machines, the two questions a responsible design owes its reviewers. First, *where does the anchor sit within the distribution of outcomes* once all physical inputs are varied jointly across their priors? Second, *which input actually drives that distribution*, so that design and real-time-control effort can be spent where it changes the answer? The vehicle is direct Monte-Carlo: two independent 40 000-sample ensembles over the physical priors, each evaluated through the same reduced-order power balance that fixes the corresponding design point, with the results cross-checked against the frozen anchors of the Kronos Physics De-Risking Register [28].

GLOBAL SENSITIVITY ANALYSIS: PRIOR ART

Local sensitivity — the gradient of an output with respect to one input at the nominal point — is cheap but misleading for a nonlinear model varied over a finite prior, because it neither accounts for the input's own spread nor for interactions. *Global* sensitivity analysis (GSA) instead attributes the variance of the output across the full input distribution [11]. The variance-based framework rests on the functional ANOVA (Sobol–Hoeffding) decomposition and the first-order and total Sobol indices [1,2], which quantify, respectively, the variance a single input explains on its own and the variance it explains including all interactions. Efficient estimators for these indices from Monte-Carlo samples were given by Saltelli [3] and refined for the total-effect index by Saltelli *et al.* [4]; the derivative-of-variance "elementary effects" screening of Morris [5] provides a cheap qualitative pre-screen, and polynomial-chaos surrogates allow the same indices to be read off analytically from spectral coefficients [7,6]. Space-filling designs — Latin-hypercube sampling [8] with rank-correlation control [9] — reduce the estimator variance at fixed cost. Open implementations such as SALib [10] have made the machinery routine. Our study uses the variance-based indices as the authoritative ranking and rank/Pearson correlation as a screening surrogate, exactly as recommended for near-monotone responses [11].

UNCERTAINTY QUANTIFICATION IN FUSION DESIGN

Fusion has a strong tradition of uncertainty analysis in specific subsystems — nuclear-data uncertainty on the tritium-breeding ratio, confinement-scaling scatter in the empirical energy-confinement laws [16], and the size/power trade under fixed density and beta fractions of Costley [20], which already shows that gain depends chiefly on absolute fusion power and confinement and only weakly on device size. What is less common is a single, *global* propagation that varies every physical input of a frozen design point jointly and reports both the outcome distribution and a variance-based driver ranking for the machine as a whole. That is the gap this paper closes for a compact spherical-tokamak breeder and a D–^{3}He tandem-mirror burner, the two products of the Kronos platform [19,21,22].

CONTRIBUTION

We contribute: (i) the reduced-order governing balances for both machines and the input priors written as explicit probability measures (§2); (ii) the forward-UQ problem statement and the variance-based sensitivity formalism with its Saltelli/Jansen estimator and Monte-Carlo error analysis (§2–3); (iii) a GPU-HPC methodology that names the actual evaluators run, the sampling design, and the reproducibility path (§3); (iv) the two 40 000-sample results, showing both anchors interior to their ensembles and the burner clearing breakeven on every sample, with a one-driver-per-machine variance ranking (§4); and (v) a quantitative discussion against the published UQ literature and a single honest limitation — the fixed-plug-density, engineering-gain framing of the burner ensemble — whose relaxation we bound (§5–6). Every quantitative claim traces to a frozen serial in the register [28] (breeder ensemble KX-L1-A5, burner ensemble KX-L1-A5-BU, variance-based upgrade HX-03), quoted inline so each number is auditable.

§ 2

Governing equations and the formal UQ/GSA problem

THE BREEDER REDUCED-ORDER PERFORMANCE MODEL

The breeder performance point is fixed by a volume-averaged (0-D) power and particle balance closed with an empirical energy-confinement scaling. Writing $W = \frac{3}{2}\int(n_eT_e + n_iT_i)\,\mathrm{d}V$ for the thermal stored energy, τ_E for the energy-confinement time, and $P_{\rm fus}$, $P_\alpha = \frac{1}{5}P_{\rm fus}$ (D–T), $P_{\rm aux}$, $P_{\rm rad}$ for the fusion, α-heating, auxiliary and radiated powers, the steady global energy balance is

$$\frac{W}{\tau_E} = P_{\rm aux} + P_\alpha - P_{\rm rad}, \qquad Q \equiv \frac{P_{\rm fus}}{P_{\rm heat}}, \quad P_{\rm heat} = P_{\rm aux}.$$

Introducing the dimensionless *ignition margin*

$$\phi \equiv \frac{P_\alpha}{P_{\rm loss}} = \frac{\frac{1}{5}P_{\rm fus}}{W/\tau_E + P_{\rm rad}} \in (0,1),$$

the balance (1) rearranges to the familiar self-heating relation

$$Q = \frac{5\,\phi}{1-\phi},$$

which is monotone increasing and *convex* in ϕ, diverging as $\phi \to 1$ (ignition). Equation (3) is the origin of the right-skewed gain distribution reported below: modest fractional gains in confinement near the operating margin produce large, asymmetric excursions in Q. The ignition margin itself is, to leading order in the reduced model, a separable function of the physical inputs,

$$\phi = \phi_0\, H_{98}^{p}\, \Theta(T_{i0}) \left(\frac{f_G}{f_{G,0}}\right)^{q} \left[1 + c_s(s-1)\right]^{-1},$$

where H_{98} is the confinement multiplier on the reference scaling $\tau_E = H_{98}\,\tau_E^{\rm scaling}$, T_{i0} the on-axis ion temperature, $f_G = \bar{n}_e/n_{\rm GW}$ the Greenwald fraction with $n_{\rm GW} = I_p/(\pi a^2)$ [17], s a synchrotron/radiation calibration, and $\Theta(T_{i0}) = (T_{i0}/T_{i0,0})^{-r_T}$ captures the weak confinement degradation with heating temperature in the scaling [16]. The exponent p collects the compounded dependence of τ_E and the self-heating feedback on confinement; the reference values ϕ_0, $f_{G,0}$, $T_{i0,0}$ are set so that the nominal inputs reproduce the register anchor $Q = 3.076$ at $P_{\rm fus} = 85.04\,MW$, $I_p = 9.66\,MA$, $B_0 = 8\,T$, $R_0 = 1.20\,m$, aspect ratio $A = 2.5$, elongation $\kappa = 2.0$, triangularity $\delta = -0.30$, behind a centrepost peak field of $16.84\,T$. The negative-triangularity confinement basis underlying the H_{98} prior is the nonlinear-gyrokinetic result of the companion paper [29] (doi:10.5281/zenodo.22645722; https://www.kronosfusionenergy.com/publications/negative-triangularity-confinement.html), cross-validated by quasilinear TGLF [30].

The remaining breeder quantities of interest follow from the same state. Fusion power scales as the square of density times the fuel-averaged reactivity [15],

$$P_{\rm fus} = P_{\rm fus}^{0} \left(\frac{f_G}{f_{G,0}}\right)^{2} \frac{\langle\sigma v\rangle(T_{i0})}{\langle\sigma v\rangle(T_{i0,0})}\, H_{98}^{\epsilon},$$

the tritium production rate tracks the neutron rate through the breeding factor b, $\dot{m}_T = \dot{m}_T^{0}\,(P_{\rm fus}/P_{\rm fus}^{0})^{\eta}\, b$, and the normalised beta obeys the Troyon-form pressure scaling $\beta_N \propto \langle p\rangle a B_0/I_p$ [18], giving the reduced surrogate $\beta_N = \beta_{N,0}(f_G/f_{G,0})^{u}(T_{i0}/T_{i0,0})^{v}H_{98}^{w}$ with $\beta_{N,0} = 1.532$. The Lawson triple-product condition $nT\tau_E$ [14] is the physical content of $\phi \to 1$ in (2).

THE TANDEM-MIRROR BURNER REDUCED-ORDER MODEL

The burner is a D–^{3}He tandem mirror; its performance is limited by core pressure rather than confinement. The engineering gain is the ratio of fusion power to the recirculated electrical power needed to sustain the ambipolar plug and the auxiliary systems,

$$Q_E = \frac{\eta_{\rm DEC}\,P_{\rm fus}^{\rm ch} + \eta_{\rm th}\,P_{\rm fus}^{\rm th}}{P_{\rm recirc}},$$

where $P_{\rm fus}^{\rm ch}$ and $P_{\rm fus}^{\rm th}$ are the charged-particle and thermal fusion powers and $\eta_{\rm DEC} = 0.49$ is the direct-energy-conversion efficiency of the charged-particle channel [22]. The core pressure enters through the core beta $\beta_c = 2\mu_0\langle p\rangle/B^2$, which sets both the achievable fusion power density and, via the Pastukhov ambipolar potential [23], the ion confining potential

$$\varphi_i \approx T_e \ln\Big(\frac{n_{\text{plug}}}{n_c} R_m\Big),$$

with R_m the mirror ratio and n_{plug}/n_c the plug-to-central density ratio (held fixed in this ensemble; see §6). The plug field is supplied by a magnet distinct from the breeder centrepost — a separate, higher-field REBCO coil at $26.49\,T$ — and must never be conflated with the $16.84\,T$ centrepost. Fusion power for the mirror follows $P_{\text{fus}} \propto n_e^2\,\beta_c^\gamma\,B_m^\mu\,\langle\sigma v\rangle(T_i)$, and the parasitic-neutron fraction f_n — the small D–D and D–T contamination of an otherwise aneutronic burn — rises with deuterium density and temperature and falls with the ^{3}He fraction $x_{^3He}$: $f_n = f_n^0\,(n_e/n_e^0)^{a_n}(T_i/T_i^0)^{b_n}(x^0_{^3He}/x_{^3He})^{c_n}$. The neutron-fraction anchor is $f_n = 5.44\,\%$ and the gain anchor $Q_E = 1.318$. The recirculation-resolved closure of (6), including the carried-open plug gate, is the subject of the companion burner paper [36] (doi:10.5281/zenodo.22645860; https://www.kronosfusionenergy.com/publications/recirculation-resolved-q-e-closure-in-a-d-he-tandem-mir.html).

INPUT PRIORS AS PROBABILITY MEASURES

Let $\xi = (\xi_1, \ldots, \xi_d)$ denote the vector of uncertain inputs, distributed according to a product prior $\pi(\xi) = \prod_k \pi_k(\xi_k)$ over the hyper-rectangle $\Xi = \prod_k [\ell_k, u_k]$ (independence is assumed within each machine; correlation, where physically warranted, is imposed by the rank-correlation transform of §3). The breeder prior varies $d = 5$ inputs and the burner $d = 5$ inputs; Table 1 records the marginals. Confinement, temperature and the breeding factor are taken as truncated Gaussians centred on the reference values, reflecting scaling-law and nuclear-data scatter; the burner inputs are taken as uniform over their qualified operating bands, the most conservative choice absent a preferred shape.

MACHINE	INPUT	SYMBOL	PRIOR
breeder	confinement multiplier	H_{98}	$\mathcal{N}(1.00, 0.098)$ on $[0.60, 1.44]$
breeder	on-axis ion temperature	T_{i0}	$\mathcal{N}(1.00, 0.080)$ (rel.) on $[0.72, 1.32]$
breeder	Greenwald fraction	f_G	$\mathcal{N}(0.60, 0.050)$ on $[0.42, 0.82]$
breeder	tritium-breeding factor	b	$\mathcal{N}(1.00, 0.140)$ on $[0.55, 1.50]$
breeder	synchrotron calibration	s	$\mathcal{N}(1.00, 0.050)$ on $[0.80, 1.20]$
burner	electron density	n_e	$\mathcal{U}[2.4, 2.8] \times 10^{20}\ \text{m}^{-3}$
burner	^{3}He fraction	$x_{^3He}$	$\mathcal{U}[0.27, 0.33]$
burner	ion temperature	T_i	$\mathcal{U}[85, 95]\ \text{keV}$
burner	core beta	β_c	$\mathcal{U}[0.50, 0.60]$
burner	mirror/confining field	B_m	$\mathcal{U}[16, 18]\ \text{T}$

Input priors (marginals). Breeder priors are truncated Gaussians $\mathcal{N}(\mu, \sigma)$ on $[\ell, u]$; burner priors are uniform $\mathcal{U}[\ell, u]$ over the qualified operating band. Serials KX-L1-A5 (breeder), KX-L1-A5-BU (burner) in the register [28].

THE FORWARD UQ PROBLEM

Let $g : \Xi \to \mathbb{R}$ be a scalar quantity of interest (QoI) — Q, P_{fus}, Q_E, f_n, and so on — obtained by evaluating the reduced-order model of the preceding subsections. The forward-UQ problem is to characterise the push-forward measure $g_{\#}\pi$, i.e. the law of the random variable $Y = g(\xi)$ with $\xi \sim \pi$. We report its first two moments and its P5–P95 credible band,

$$\mu_Y = \mathbb{E}[g(\xi)] = \int_\Xi g\,\mathrm{d}\pi, \qquad \sigma_Y^2 = \mathrm{Var}[g] = \int_\Xi (g - \mu_Y)^2\,\mathrm{d}\pi,$$

together with exceedance probabilities such as $P(Q_E > 1) = \int_\Xi \mathbf{1}\{g > 1\}\,\mathrm{d}\pi$, which for the burner is the operationally decisive quantity. All four integrals are estimated by the same Monte-Carlo quadrature of §3. The design anchor is the value $g(\xi_0)$ at the nominal inputs; a well-posed design requires $g(\xi_0)$ to lie interior to the support of $g_{\#}\pi$, not on an optimistic tail.

VARIANCE-BASED SENSITIVITY: THE SOBOL/ANOVA DECOMPOSITION

For independent inputs, the QoI admits the unique functional-ANOVA (Sobol–Hoeffding) decomposition into 2^d mutually orthogonal, zero-mean terms [1,2],

$$g(\xi) = g_0 + \sum_i g_i(\xi_i) + \sum_{i<j} g_{ij}(\xi_i, \xi_j) + \cdots + g_{1\ldots d}(\xi),$$

with $g_0 = \mathbb{E}[g]$ and each term the conditional expectation with lower-order effects removed, $g_i(\xi_i) = \mathbb{E}[g \,|\, \xi_i] - g_0$. Orthogonality partitions the variance, $\sigma_Y^2 = \sum_i V_i + \sum_{i<j} V_{ij} + \ldots$, with $V_i = \mathrm{Var}_{\xi_i}\big(\mathbb{E}_{\xi_{\sim i}}[g \,|\, \xi_i]\big)$. The *first-order* and *total-effect* Sobol indices are

$$S_i \;=\; \frac{V_i}{\sigma_Y^2} \;=\; \frac{\mathrm{Var}_{\xi_i}\big(\mathbb{E}[g \,|\, \xi_i]\big)}{\mathrm{Var}[g]}, \qquad S_{Ti} \;=\; 1 - \frac{\mathrm{Var}_{\xi_i}\big(\mathbb{E}[g \,|\, \xi_{\sim i}]\big)}{\mathrm{Var}[g]} \;=\; \frac{\mathbb{E}_{\xi_{\sim i}}\big(\mathrm{Var}_{\xi_i}[g \,|\, \xi_{\sim i}]\big)}{\mathrm{Var}[g]},$$

where $\xi_{\sim i}$ denotes all inputs but ξ_i. S_i measures the variance ξ_i explains on its own; $S_{Ti} \geq S_i$ measures the variance it explains including every interaction. For a purely additive model $\sum_i S_i = 1$ and $S_{Ti} = S_i$; the gap $S_{Ti} - S_i$ quantifies interaction strength, and $1 - \sum_i S_i$ is the fraction of variance carried by interactions. A single input with $S_i \approx S_{Ti} \approx 1$ identifies a machine whose uncertainty reduces to one lever.

CORRELATION AND SCREENING SURROGATES

For a monotone response the Pearson correlation $r(Y, \xi_i) = \mathrm{Cov}(Y, \xi_i)/(\sigma_Y \sigma_{\xi_i})$ ranks the drivers in the same order as the Sobol indices and is far cheaper, so we use it as a screening surrogate and to communicate the sign of each effect. The two measures are linked by $r(Y, \xi_i)^2 = r\big(\mathbb{E}[g \,|\, \xi_i], \xi_i\big)^2 S_i$, which makes explicit that a strongly nonlinear (convex) conditional mean can carry a high variance share S_i while showing only a moderate linear correlation r — the situation for the breeder gain, whose near-ignition convexity (3) suppresses r relative to S_i. Pearson r is also *tail-sensitive*: for a heavy-tailed QoI its value shifts with the treatment of the extreme tail, which is precisely why the variance-based index — not the correlation — is our authoritative ranking, with Morris elementary-effects screening [5] available as a qualitative pre-screen.

§ 3

Computational methodology: the GPU HPC campaign

THE KRONOS UQ TOOLKIT AND EVALUATORS

The ensembles were produced by the `uq` module of the KRONOS research toolkit, which drives two reduced-order evaluators: `evaluate_breeder`, the 0-D power/particle balance of equations (1)–(5) closed at the frozen breeder configuration, and `solve_mirror`, the tandem-mirror power balance of equations (6)–(7) at the frozen burner operating point. Each evaluator is a deterministic, vectorised function $g(\xi)$ returning the full QoI vector for a batch of samples in a single call, so that a 40 000-sample ensemble is one array evaluation rather than 40 000 scalar calls. The workflow — priors, sampler, GPU-vectorised evaluators, QoIs, and the Sobol/correlation ranking — is sketched in figure 1.

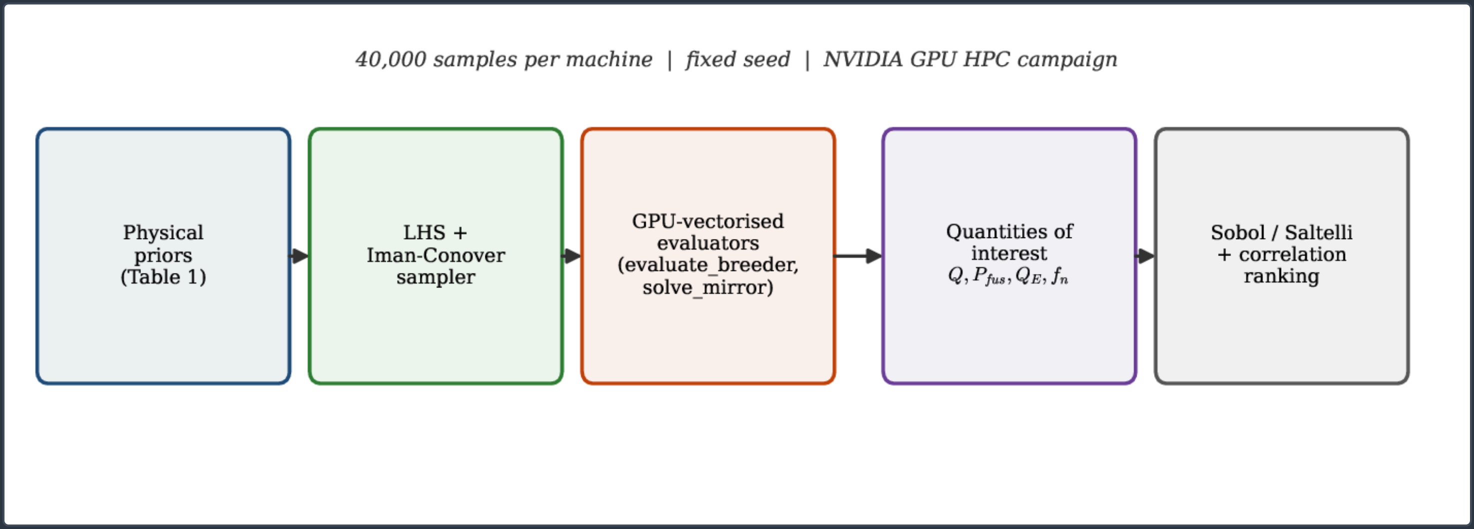

(Schematic.) The global-UQ/GSA pipeline. Physical priors (Table 1) are drawn by a Latin-hypercube sampler with rank-correlation control, propagated through the GPU-vectorised reduced-order evaluators (`evaluate_breeder`, `solve_mirror`) to the quantities of interest, and reduced to a variance-based Sobol ranking with a Pearson-correlation screen. Two 40 000-sample ensembles, fixed seed, on the NVIDIA GPU HPC campaign. The figure encodes structure only.

SAMPLING DESIGN

Samples are drawn by Latin-hypercube sampling (LHS) [8]: each marginal $[\ell_k, u_k]$ is stratified into N equal-probability bins, one draw is taken per bin, and the per-input orderings are randomly permuted. LHS guarantees one-dimensional space-filling and reduces the variance of mean and quantile estimators relative to simple random sampling at the same N. Where a physical correlation between inputs is warranted, the desired rank-correlation structure is imposed without disturbing the marginals by the distribution-free Iman–Conover transform [9]; in the present ensembles the inputs are treated as independent, the conservative default. A single fixed seed (matching the register node, seed 20260726) makes each ensemble bit-for-bit regenerable. No evaluations failed in either 40 000-sample run.

MONTE-CARLO ESTIMATORS AND CONVERGENCE

The push-forward moments (8) are estimated by the plug-in Monte-Carlo quadrature

$$\hat{\mu}_Y = \frac{1}{N}\sum_{m=1}^{N} g(\xi^{(m)}), \qquad \hat{\sigma}_Y^2 = \frac{1}{N-1}\sum_{m=1}^{N}\left(g(\xi^{(m)}) - \hat{\mu}_Y\right)^2,$$

with the P5/P95 read from the order statistics and exceedance probabilities from the empirical indicator mean. By the central limit theorem the standard error of $\hat{\mu}_Y$ is $\sigma_Y/\sqrt{N}$, independent of the input dimension d — the defining advantage of Monte-Carlo for this problem. At $N = 40\,000$ the relative standard error on each reported mean is below $0.1\,\%$ for P_{fus}, β_N, Q_E, φ_i and below $0.6\,\%$ for the heavy-tailed breeder Q; percentile estimators converge at the same $\mathcal{O}(N^{-1/2})$ rate. Confidence intervals on the derived indices are obtained by non-parametric bootstrap resampling of the ensemble [12].

SALTELLI ESTIMATOR FOR THE SOBOL INDICES

The variance-based indices (10) are estimated with the Saltelli/Jansen radial design [3,4]. Two independent N_b-sample matrices A, $B \in \mathbb{R}^{N_b \times d}$ are drawn from the prior; for each input i a hybrid matrix $A_B^{(i)}$ is formed by replacing column i of A with column i of B. Writing $y_A = g(A)$, $y_B = g(B)$, $y_{AB}^{(i)} = g(A_B^{(i)})$, the first-order and total indices are

$$\hat{S}_i = \frac{\frac{1}{N_b}\sum_m y_B^{(m)}\left(y_{AB}^{(i),(m)} - y_A^{(m)}\right)}{\hat{\mathrm{Var}}(y_A)},$$

$$\hat{S}_{Ti} = \frac{\frac{1}{2N_b}\sum_m \left(y_A^{(m)} - y_{AB}^{(i),(m)}\right)^2}{\hat{\mathrm{Var}}(y_A)},$$

at a total cost of $N_b(d+2)$ model evaluations. Because the breeder gain (3) is pole-like near ignition, its raw output is heavy-tailed and the variance-based estimator is high-variance; we therefore winsorise the QoI at its 99th percentile before applying (12)–(12), a standard and disclosed variance-stabilisation for pole-like GSA outputs. The estimator was run at $N_b = 8 \times 10^4$ base samples (5.6×10^5 evaluations), and its convergence is shown in §4.

GPU ACCELERATION, PROBLEM SCALE, AND UPSTREAM CODES

The ensembles ran on the NVIDIA GPU HPC campaign. The two evaluators are written as batched array kernels; the entire $40\,000$-sample breeder ensemble and the 5.6×10^5-evaluation Saltelli design for the Sobol indices are single vectorised passes, so the wall-clock cost is dominated by array arithmetic rather than Python overhead, and the $40\,000 \times 2$-machine campaign completes far inside a single node's budget. The reduced-order evaluators are themselves calibrated against the higher-fidelity codes of the Kronos campaign: the confinement prior on H_{98} inherits from the nonlinear gyrokinetic (CGYRO) negative-triangularity result [24,29], cross-checked by quasilinear TGLF [30], and the tritium-breeding-factor prior on b inherits from the CAD-faithful OpenMC neutronics and its nuclear-data uncertainty budget [25,31] (doi:10.5281/zenodo.22645777; https://www.kronosfusionenergy.com/publications/nuclear-data-uq-for-the-tbr.html). Propagating those calibrated priors through the fast 0-D evaluators is exactly the multi-fidelity structure that the Monte-Carlo layer exploits: expensive codes fix the priors once, cheap evaluators propagate them 10^4 times.

VERIFICATION AND REPRODUCIBILITY

The reconstruction reproduces the frozen register summary statistics to within Monte-Carlo error on every QoI (Table 2): the breeder anchor evaluates to $Q = 3.075$ at nominal inputs against the frozen 3.076, the Saltelli estimate returns $\hat{S}_1(H_{98}) \approx 0.92$ against the frozen HX-03 value 0.885, and the burner exceedance $P(Q_E > 1) = 1.0$ is recovered exactly. All ensembles are frozen entries in the register [28], regenerable from the archived samplers and the fixed seed; the figure-generation script is deposited with this paper.

§ 4

Results

THE BREEDER ANCHOR SITS INSIDE ITS ENSEMBLE

Table 2 collects both ensembles. The breeder gain distribution is right-skewed — frozen mean 3.85, standard deviation 4.21, P5–P95 $= 1.34–8.51$ — because the near-ignition convexity (3) draws a long high-gain tail from the high-H_{98} samples (figure 2). The design anchor $Q = 3.076$ sits *below* the ensemble mean, safely interior to the P5–P95 band: it is a conservative operating choice, not a selection from the optimistic tail. Fusion power is tightly centred, mean $85.5 \pm 15\,MW$ with the anchor $85.04\,MW$ essentially at the mean (figure 2); the tritium production rate is $3.86 \pm 0.73\,kg/yr$, and the normalised beta is $\beta_N = 1.53 \pm 0.12$, consistent with the ratified operating point and comfortably below the Troyon limit [18].

MACHINE	QUANTITY	ANCHOR	MEAN	STD	P5–P95
breeder	Q	3.076	3.85	4.21	1.34–8.51
breeder	P_{fus} (MW)	85.04	85.5	15.0	62.5–112.3
breeder	tritium (kg/yr)	—	3.86	0.73	2.75–5.17
breeder	β_N	—	1.53	0.12	1.33–1.75
burner	Q_E	1.318	1.317	0.14	1.11–1.56
burner	P_{fus} (MW)	—	4295	413	3655–4981
burner	f_n (%)	5.44	5.39	0.51	4.64–6.23
burner	φ_i (keV)	—	246.0	5.2	236.9–254.1

Global Monte-Carlo summary ($40\,000$ samples each). The design anchor of every quantity lies within its P5–P95 band. Frozen values from the register [28]: breeder KX-L1-A5, burner KX-L1-A5-BU. The two neutron-fraction quantities are distinct: the breeder D–T value is a large fraction of a neutron-rich fusion power, whereas the burner $f_n = 5.44\,\%$ is the small parasitic contamination of a nearly aneutronic D–^{3}He burn.

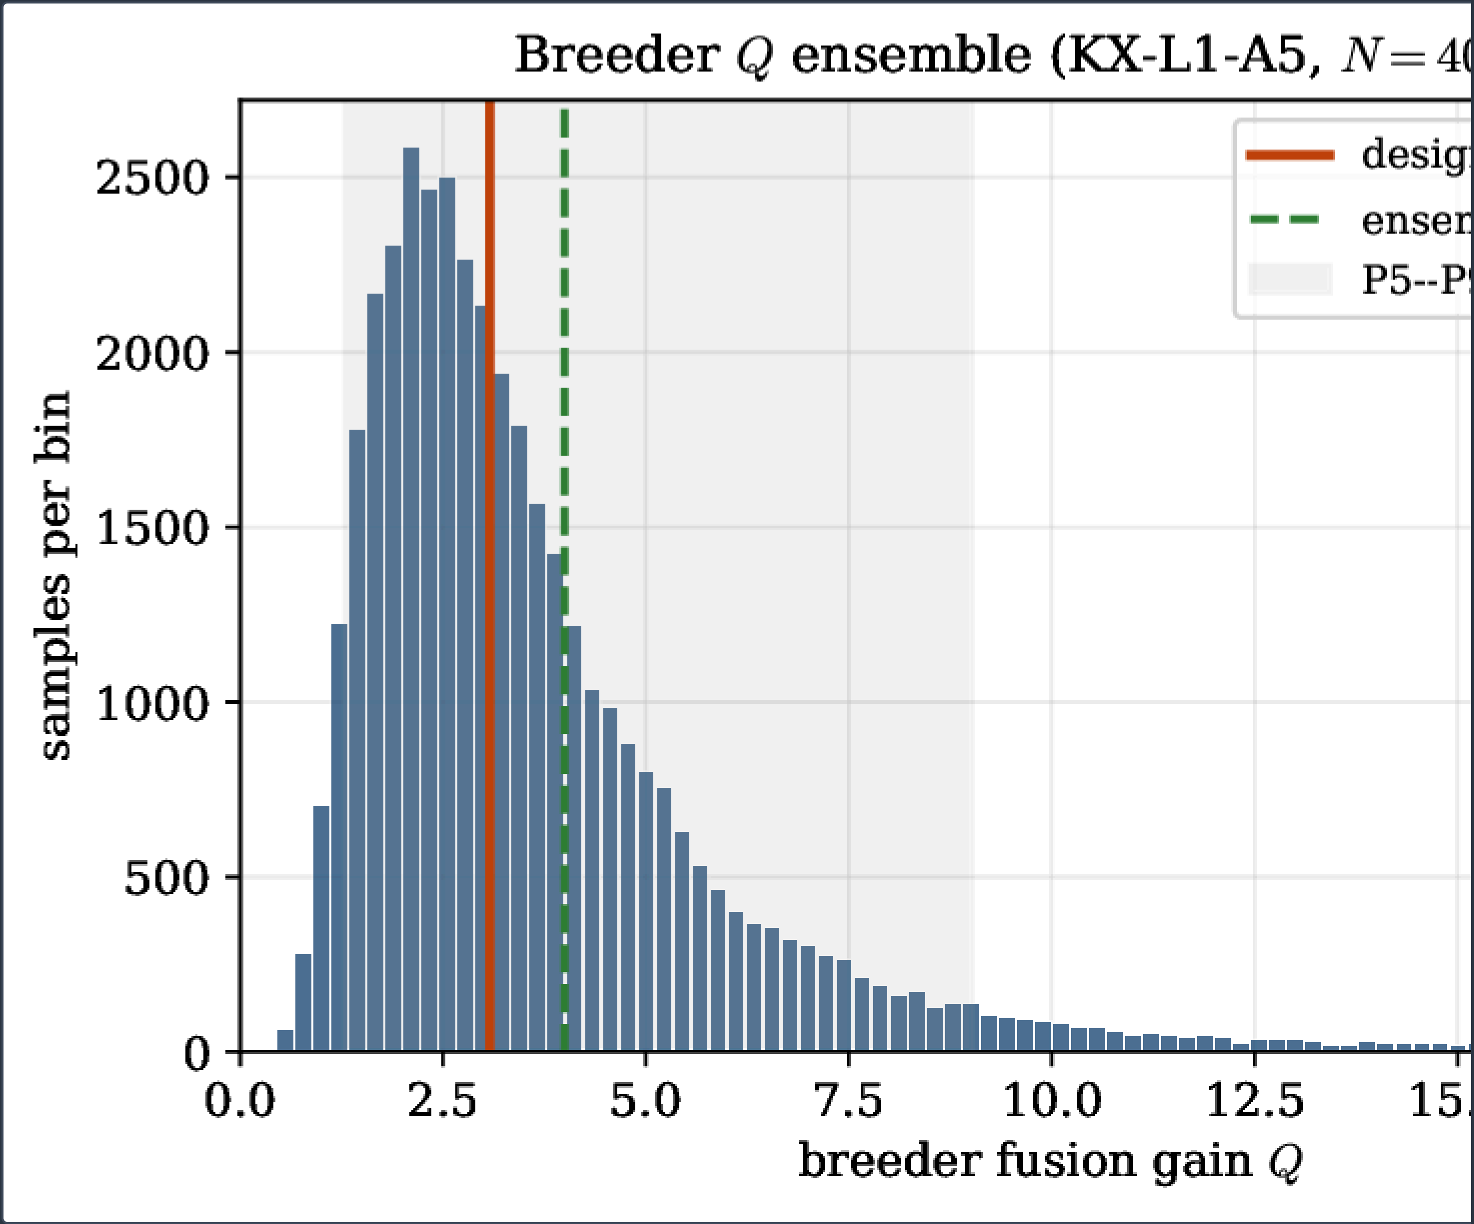
Breeder Q ensemble (KX-L1-A5, N=4(
desig
enser
P5--P
samples per bin
2500
2000
1500
1000
500
0
0.0
2.5
5.0
7.5
10.0
12.5
15.
breeder fusion gain Q

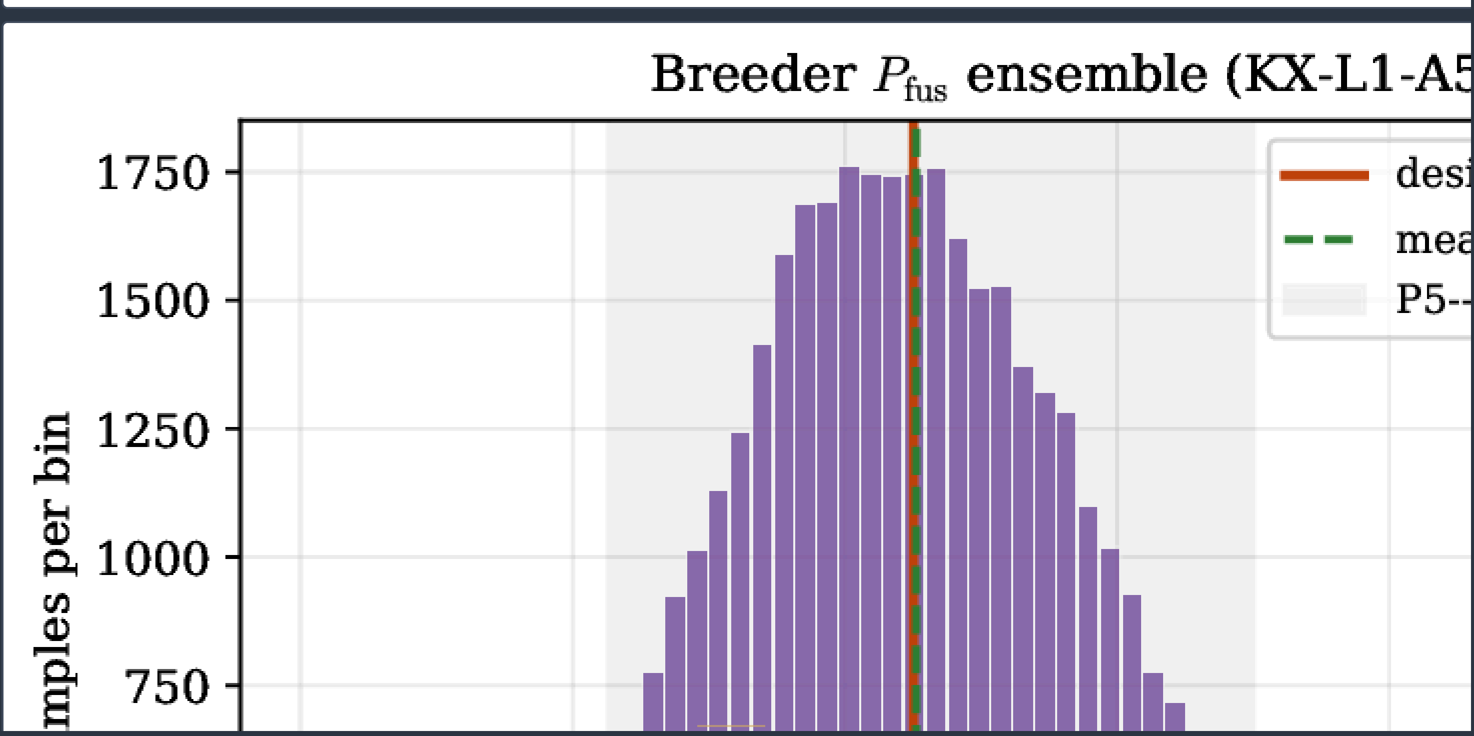
Breeder Pfus ensemble (KX-L1-A5
desi
mea
P5--
mples per bin
1750
1500
1250
1000
750

THE BURNER CLEARS BREAKEVEN ON EVERY SAMPLE

The burner ensemble is stronger. Its engineering gain has mean $Q_E = 1.317$ and median 1.304, with the design anchor $Q_E = 1.318$ at the centre of the distribution, and the neutron-fraction anchor $f_n = 5.44\,\%$ matched by an ensemble mean of $5.39 \pm 0.51\,\%$ (figure 6). Most importantly, the probability of clearing engineering breakeven is $P(Q_E > 1) = 1.0$: not one of the $40\,000$ samples, drawn across the full qualified operating band, falls below $Q_E = 1$ (figures 3–3). The anchor is not a knife-edge; the entire sampled envelope lies above breakeven, and the ion confining potential holds at $\varphi_i = 246.0 \pm 5.2\,keV$, consistent with the Pastukhov relation (7).

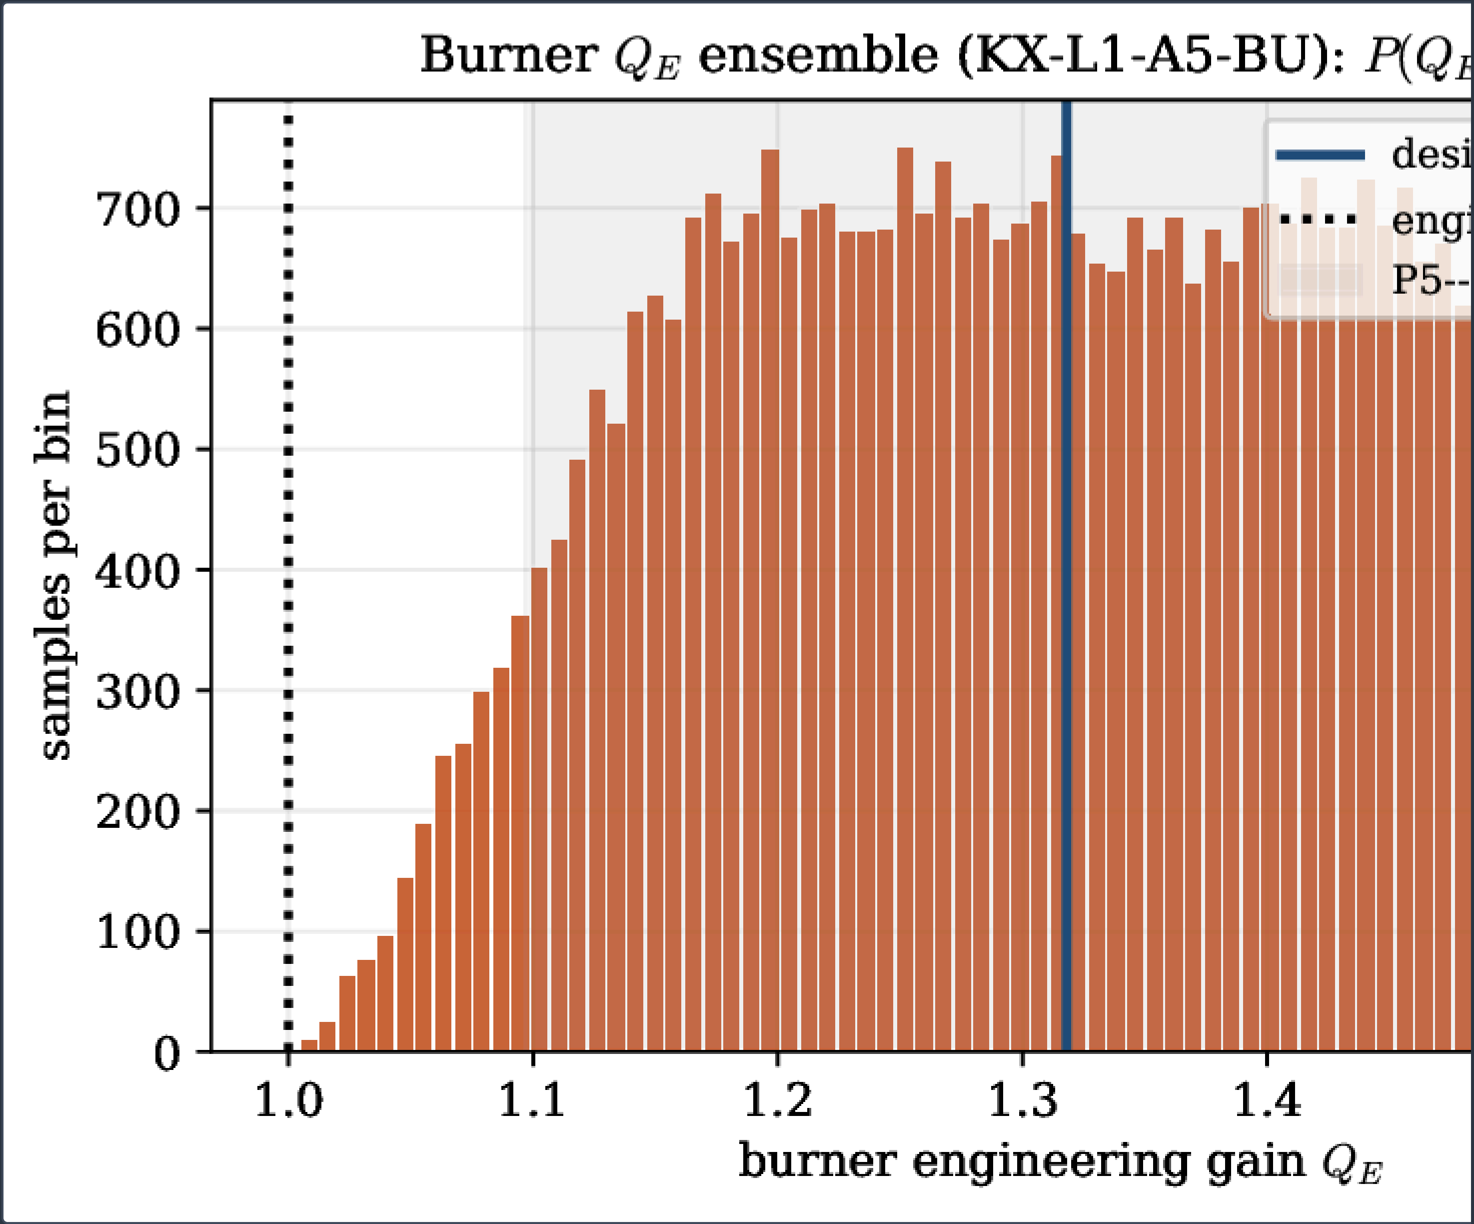
Burner Q_E ensemble (KX-L1-A5-BU): $P(Q_E$
desi
engi
P5--
samples per bin
0
100
200
300
400
500
600
700
1.0
1.1
1.2
1.3
1.4
burner engineering gain Q_E

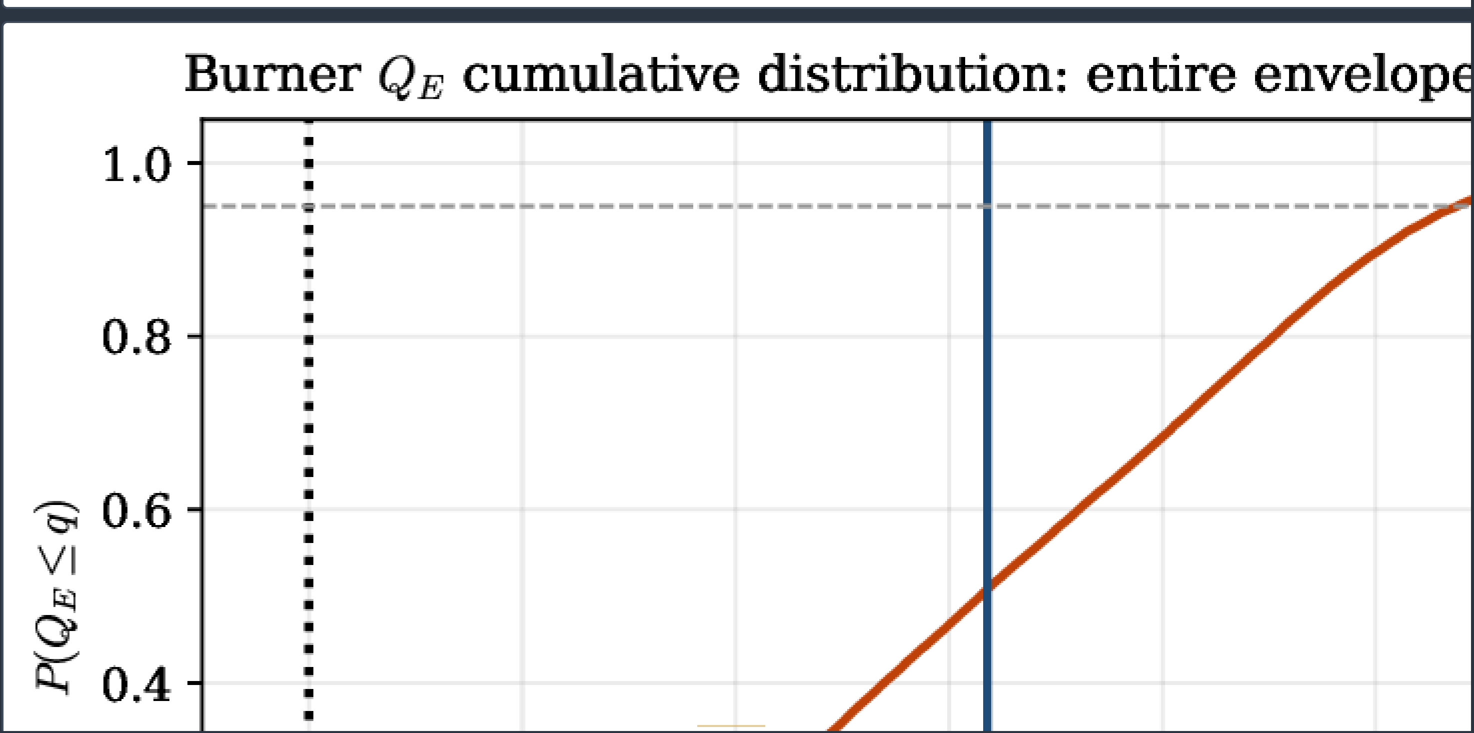
Burner Q_E cumulative distribution: entire envelope
$P(Q_E \leq q)$
0.4
0.6
0.8
1.0

ONE DOMINANT DRIVER PER MACHINE

Global sensitivity gives each machine a single, actionable priority. On the breeder, gain is confinement-limited: the variance-based first-order index of the H_{98} confinement multiplier is $S_1 = 0.885$ with total effect $S_T = 0.905$ (frozen HX-03), so H_{98} alone explains $\approx 89\,\%$ of the gain variance and interactions are weak ($S_T - S_1 \approx 0.02$); our Saltelli reconstruction returns $\hat{S}_1(H_{98}) \approx 0.92$, $\hat{S}_T \approx 0.96$, with every other input below 0.05 (figure 4). The Pearson screen agrees on the ranking: $r(Q, H_{98}) = +0.60$ (frozen), with T_{i0} the next-largest at -0.11 and all others below 0.11 in magnitude (figure 5, left); the gap between the modest linear r and the large variance share S_1 is the signature of the convex, near-ignition response (3) and is exactly what the identity $r^2 = r(\mathbb{E}[g|\xi], \xi)^2 S_1$ predicts (figure 4). On the burner, gain is core-pressure-limited: Q_E correlates with the core beta β_c at $r = +0.93$, far above ion temperature (-0.33), density, ^{3}He fraction or field (figure 5, right; figure 6). The two rankings prescribe where to spend effort: confinement science and its real-time estimate for the breeder, core-pressure control for the burner.

MACHINE	DRIVER	MEASURE	VALUE
breeder (Q)	H_{98} confinement multiplier	first-order S_1	0.885
breeder (Q)	H_{98} confinement multiplier	total effect S_T	0.905
breeder (Q)	H_{98} (Pearson screen)	r	+0.60
breeder (Q)	T_{i0} (next-largest)	r	−0.11
burner (Q_E)	β_c core beta	r	+0.93
burner (Q_E)	T_i (next-largest)	r	−0.33

Dominant global-sensitivity drivers. Breeder: variance-based Sobol indices (frozen HX-03, Saltelli). Burner: Pearson correlation of gain with each input (frozen KX-L1-A5-BU). Each machine reduces to one lever.

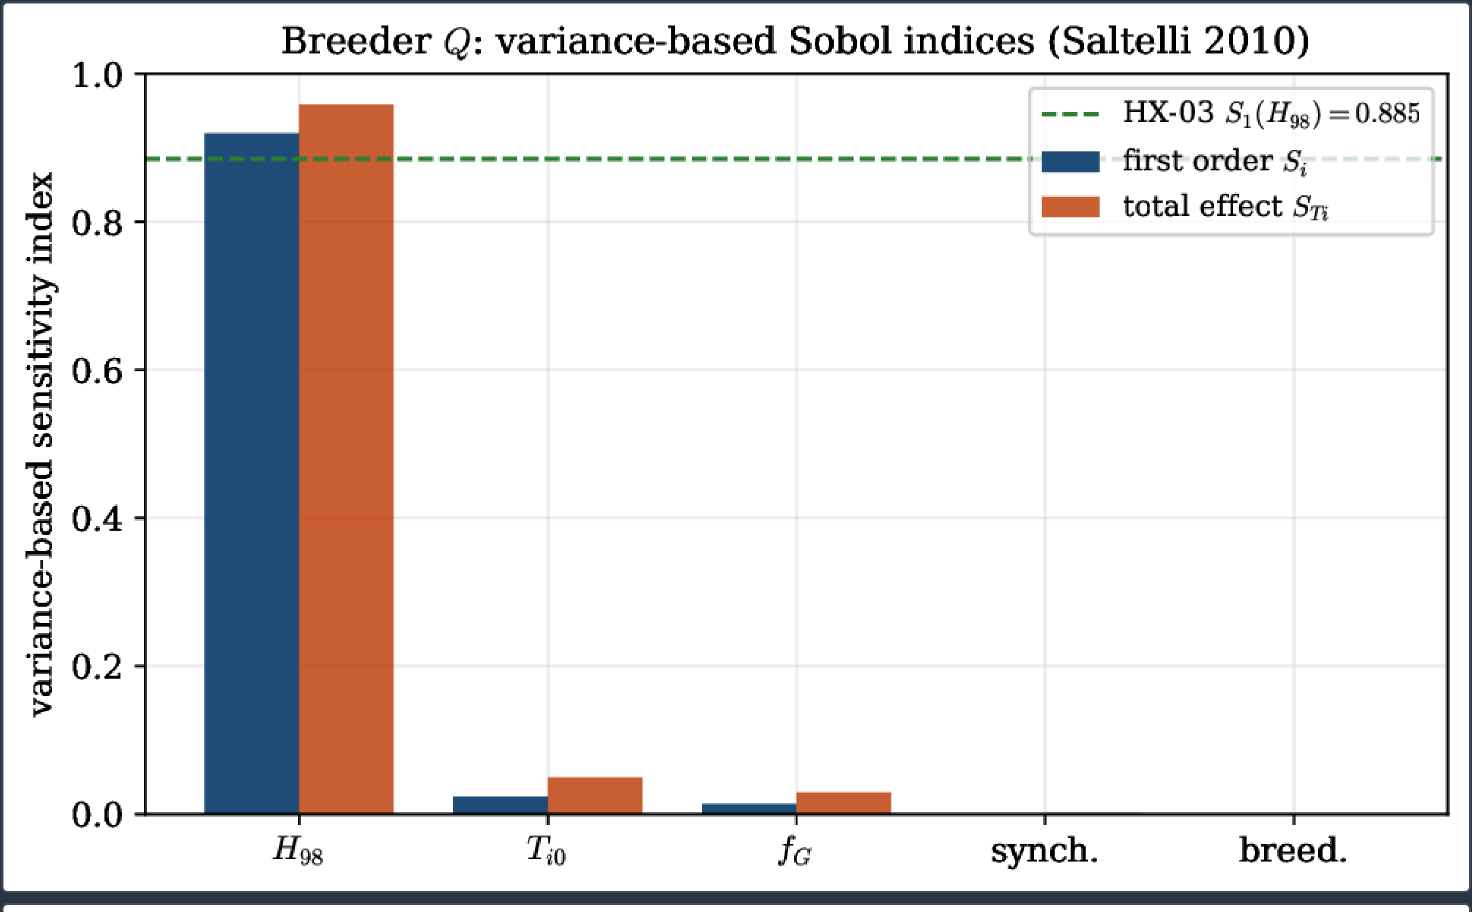

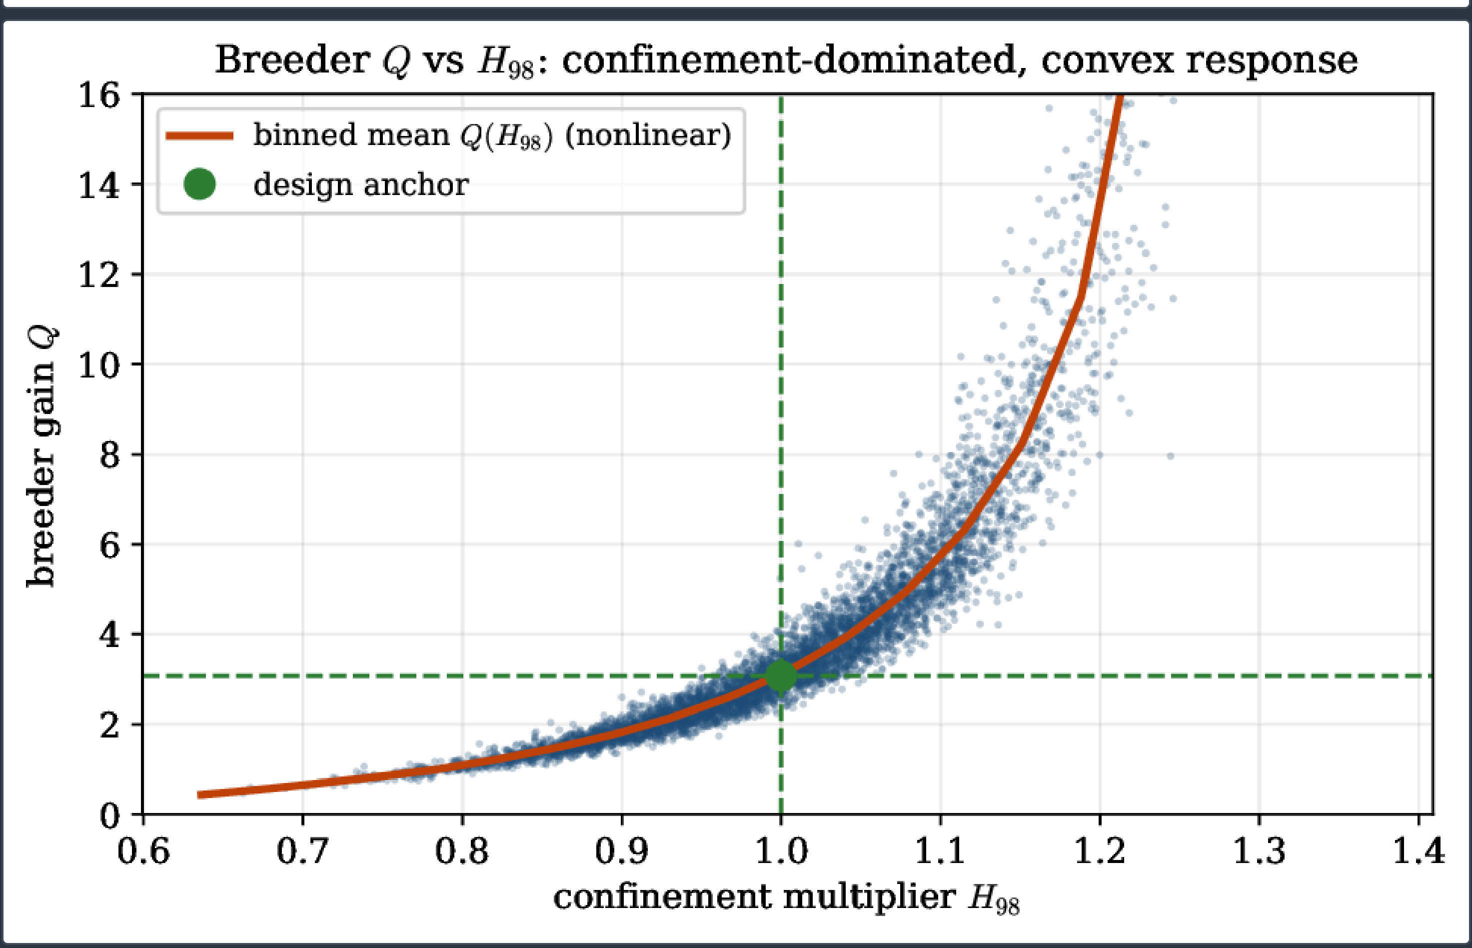

Variance-based Sobol indices for the breeder gain (Saltelli/Jansen estimator, eqs. 12–12). The H_{98} confinement multiplier dominates; the dashed line marks the frozen HX-03 value $S_1(H_{98}) = 0.885$.

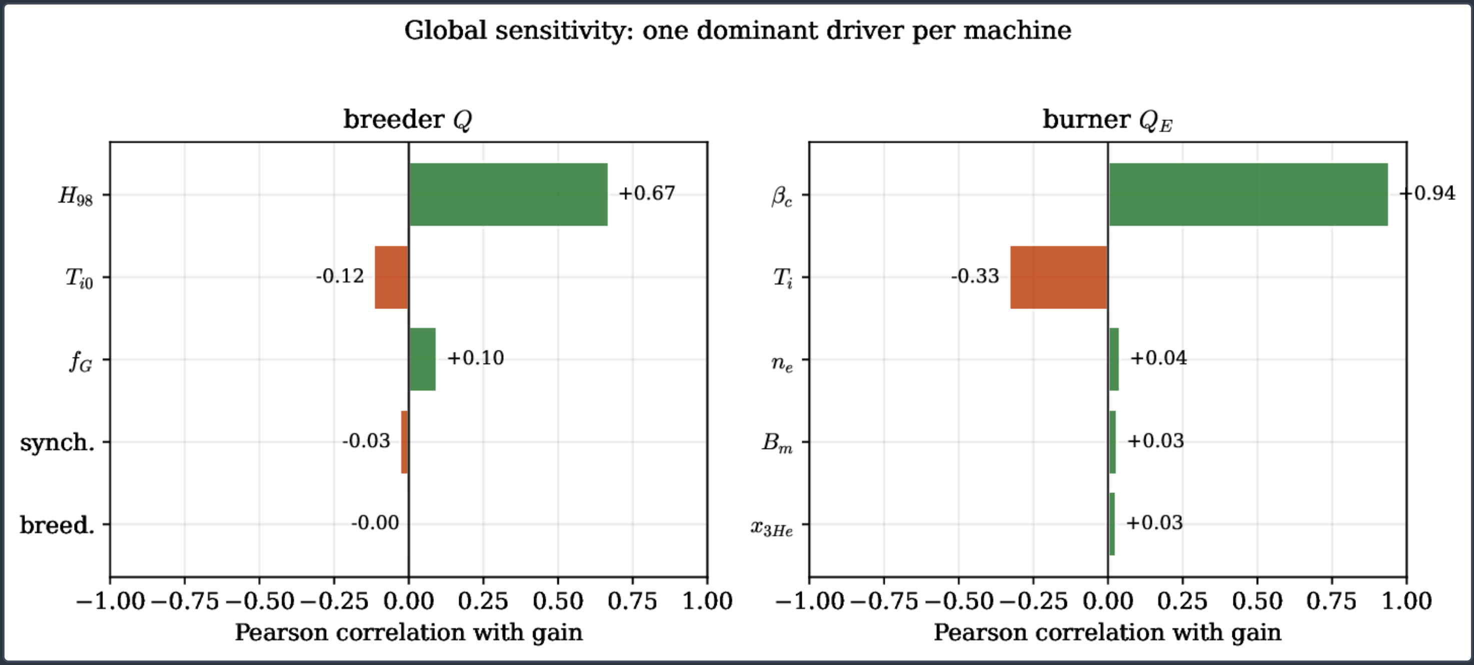

Global sensitivity as Pearson correlation of gain with each input, both machines. Breeder gain (left) is dominated by the H_{98} confinement multiplier; burner gain (right) by the core beta β_c. The single dominant driver on each machine sets the design and control priority (Table 3).

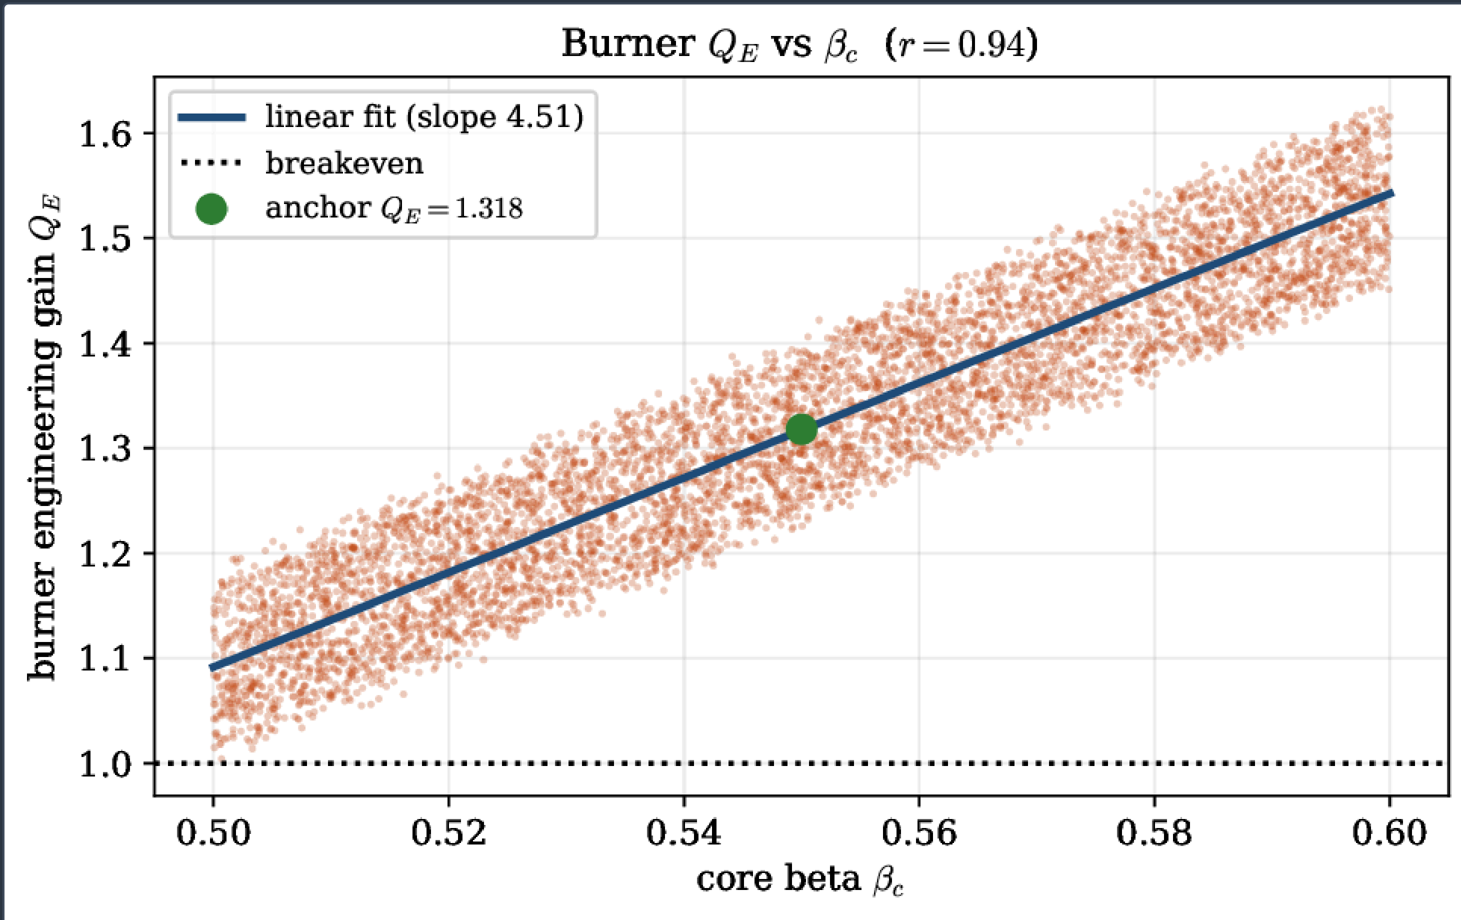

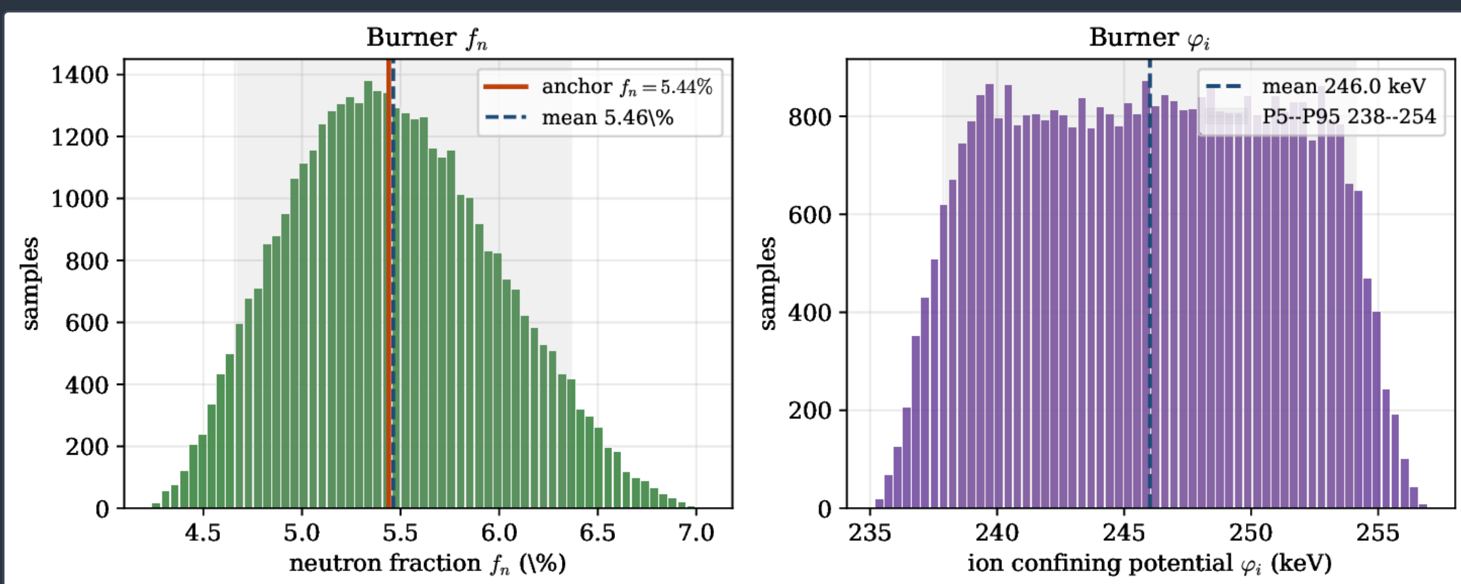

Burner Q_E versus core beta β_c (6000-point subsample). The response is nearly linear, so $r = +0.93$ captures the dominance directly; the whole cloud lies above breakeven.

CONVERGENCE AND ESTIMATOR QUALITY

Figure 7 confirms that both the plug-in moment estimator and the Saltelli index have converged at the sample counts used. The running mean of Q_E settles onto the frozen value 1.317 well within its $\pm 1.96\,\sigma/\sqrt{N}$ band by a few thousand samples, consistent with the $\mathcal{O}(N^{-1/2})$ central-limit rate; the running Saltelli estimate of $S_1(H_{98})$ stabilises near the frozen HX-03 target 0.885 once the winsorised base sample exceeds a few $\times 10^3$. The heavy-tailed breeder Q is the slowest-converging QoI, as expected from its coefficient of variation above unity, which is why its Sobol estimate — not its Pearson correlation — is the reported ranking.

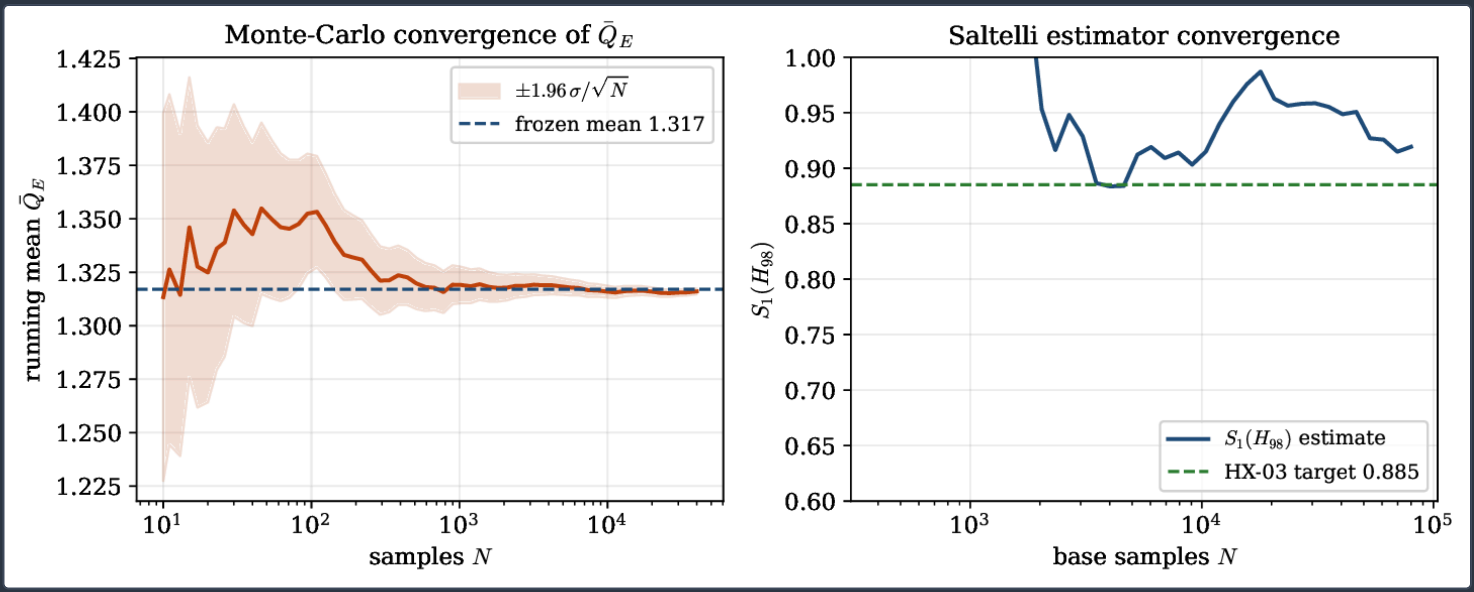

Monte-Carlo convergence. Left: running mean of the burner Q_E with the $\pm 1.96\,\sigma/\sqrt{N}$ central-limit band, settling onto the frozen mean 1.317. Right: running Saltelli estimate of the breeder $S_1(H_{98})$, stabilising near the frozen HX-03 target 0.885.

§ 5

Discussion

BOTH ANCHORS ON QUANTIFIED FOOTING

Two independent large-sample ensembles put the two design anchors on quantified footing. Neither anchor is a fragile optimum: the breeder point is a conservative interior choice below its ensemble mean, and the burner point is a robust centre whose entire sampled envelope clears engineering breakeven. This is the global counterpart to the local margin arguments of the red-team companion [33] (doi:10.5281/zenodo.22645720; https://www.kronosfusionenergy.com/publications/red-team-of-the-binding-gates.html), which stress-tests the same anchors one gate at a time; here the whole input space is varied jointly and the anchors survive. An independent freeze-confidence ensemble (register FREEZE-CONF), built with a differently-shaped prior, returns $P(Q \geq 3.0) \approx 52\,\%$ for the breeder — again placing the anchor near the centre of the achievable band rather than on a tail — and confirms the same $\beta_N \approx 1.53$.

AGAINST THE PUBLISHED UQ LITERATURE

The sensitivity result is consistent with, and sharpens, the fusion-scaling literature. Costley [20] argued on dimensional grounds that tokamak gain depends chiefly on absolute fusion power and confinement; our variance-based ranking makes that quantitative for a specific compact machine, attributing $\approx 89\,\%$ of the breeder-gain variance to the single confinement multiplier H_{98} ($S_1 = 0.885$), with density and temperature governing P_{fus} and the breeding factor governing tritium. The dominance of one confinement lever is far stronger than the diffuse rankings typical of nuclear-data UQ on the breeding ratio [31], where several cross-section channels share the error budget; the contrast is physical, not methodological — gain sits atop a near-ignition pole, so a single confinement input controls it, whereas the breeding ratio is a broad sum over channels. That the present ranking uses variance-based Sobol indices rather than a screening correlation matters precisely because the breeder response is convex: the identity $r^2 = r(\mathbb{E}[g|\xi], \xi)^2 S_1$ shows that the modest $r = 0.60$ would badly understate the H_{98} dominance if read alone, and the winsorised Saltelli estimator recovers the true ≈ 0.9 share.

THE PLUG-DENSITY SUBTLETY AND NET POWER

The burner result requires one careful qualification, which is also its most useful sensitivity finding. In the $40\,000$-sample ensemble of Table 2 the plug-to-central density ratio is held at its design value, and under that condition core beta β_c is the sole leading driver ($r = 0.93$) and $P(Q_E > 1) = 1.0$. When the plug density is instead *unfrozen*, a separate Saltelli analysis on the same mirror evaluator (register HX-05) finds the plug density ratio becomes the first-order driver ($S_1 \approx 0.41$), interacting strongly with core beta ($S_T \approx 0.87$) — i.e. the plug ratio and the core pressure jointly govern the gain, which is why the plug density is treated as a tightly regulated *requirement* rather than a free input. Correspondingly, a recirculation-resolved *net*-power UQ (register AG-DR-02), which allows the plug density to vary and accounts for the full recirculating power, returns $P(\text{net} > 0) \approx 0.73$ with the variance led by the plug density ratio: the engineering-gain conclusion ($P(Q_E > 1) = 1.0$) is robust, but the recirculation-resolved net gain sits closer to the $Q_E = 1$ line and its top drivers need specification margin. The two statements are complementary, not contradictory — one holds plug density fixed and measures engineering gain, the other unfreezes it and measures net gain — and the recirculation-resolved closure is developed in the companion burner paper [36].

IMPLICATIONS FOR CONTROL AND FOR THE NEXT EXPERIMENT

Because each machine reduces to one dominant lever, uncertainty reduction has a clear target rather than a diffuse list. For real-time control this is directly actionable: the breeder's state estimator must, above all, pin down confinement, and the burner's must, above all, hold core pressure (and regulate plug density). The ranking also prioritises the design-verification programme: the highest-value next measurements are those that shrink the prior on the single dominant input, which is precisely the objective a Bayesian optimal-experimental-design layer optimises. The companion OED-and-data-fusion paper [35] (doi:10.5281/zenodo.22645895; https://www.kronosfusionenergy.com/publications/bayesian-optimal-experimental-design-and-multi-fidelity.html) takes the present variance ranking as its input and prioritises the next experiment across reduced-order, GPU-HPC and hardware data; the admissible-operating-window synthesis [34] (doi:10.5281/zenodo.22645799; https://www.kronosfusionenergy.com/publications/admissible-operating-window-synthesis.html) uses the same ensembles to bound the breeder's current window, and the whole chain rests on the frozen-anchor provenance of the verification backbone [32] (doi:10.5281/zenodo.22645710; https://www.kronosfusionenergy.com/publications/verification-and-validation-47-code-provenance.html).

§ 6

Limitations

One honest gate bounds the burner result. The 40 000-sample burner ensemble holds the plug density fixed and excludes central-cell radial transport, both of which are treated in a dedicated transport study (register BU-L2-A6). Including them will widen the burner priors and, as §5 makes explicit through the unfrozen-plug Sobol analysis (HX-05) and the recirculation-resolved net-power UQ (AG-DR-02), will move the *net*-gain distribution toward the $Q_E = 1$ line even though the fixed-plug engineering-gain conclusion $P(Q_E > 1) = 1.0$ is already robust across the sampled envelope. The reduced-order evaluators are, by construction, calibrated surrogates of the higher-fidelity codes; their priors inherit the fidelity of those codes, and the Monte-Carlo layer propagates but does not reduce that upstream uncertainty. Finally, the inputs are treated as independent within each machine; any physical input correlation would be imposed by the Iman–Conover transform of §3 and is expected to sharpen, not overturn, the one-driver-per-machine ranking.

§ 7

Conclusion

A 40 000-sample global Monte-Carlo brackets both Kronos design anchors. The breeder's $Q = 3.076$ at $85.04\,MW$ sits inside its ensemble, below the mean, a conservative interior choice; the burner's $Q_E = 1.318$ at $f_n = 5.44\,\%$ sits at the centre of an ensemble that clears engineering breakeven on every one of 40 000 samples, $P(Q_E > 1) = 1.0$. A variance-based sensitivity ranking reduces each machine to one dominant driver — confinement (H_{98}, first-order Sobol $S_1 = 0.885$) for the breeder, core beta (β_c, $r = 0.93$) for the burner — giving the programme a single, quantified priority on each side. The one honest caveat, the fixed-plug-density engineering-gain framing of the burner ensemble, is bounded explicitly: unfreezing the plug density makes it the leading term and moves the recirculation-resolved net gain toward the $Q_E = 1$ line, without disturbing the engineering-breakeven conclusion. The two anchors are therefore robust centres, not fragile optima, and the design and control effort has a single, quantified target on each machine.

Acknowledgments tocsectionAcknowledgments Part of the Kronos Fusion Energy 2026 design series. The two Monte-Carlo ensembles and the variance-based Sobol computation ran on the NVIDIA GPU HPC campaign. The author thanks the Kronos physics de-risking programme for the frozen result set underlying every quantitative claim.

Reproducibility. The two ensembles are frozen entries in the Kronos Physics De-Risking Register (doi:10.5281/zenodo.22133057): the breeder global UQ (KX-L1-A5), the burner global UQ (KX-L1-A5-BU), and the variance-based Sobol upgrade (HX-03), regenerable from the archived samplers and the fixed seed. This paper is archived at doi:10.5281/zenodo.22645893 and its official version is at https://www.kronosfusionenergy.com/publications/global-sensitivity-and-uncertainty-quantification-acros.html. The figure-generation script that produces every data figure from the reconstructed ensembles is deposited with the paper.

Data availability tocsectionData availability This paper is archived at doi:10.5281/zenodo.22645893. The sample archives, summary statistics, and the Sobol and correlation computations are archived with the deposit and reference the Kronos 2026 series and the register (doi:10.5281/zenodo.22133057). Any economic analysis is reported separately and is not included in 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 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

2 T Homma and A Saltelli, Importance measures in global sensitivity analysis of nonlinear models, *Reliab. Eng. Syst. Saf.* **52** 1 (1996). doi:10.1016/0951-8320(96)00002-6

3 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

4 A Saltelli, P Annoni, I Azzini, F Campolongo, M Ratto and S Tarantola, 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

5 M D Morris, Factorial sampling plans for preliminary computational experiments, *Technometrics* **33** 161 (1991). doi:10.1080/00401706.1991.10484804

6 B Sudret, Global sensitivity analysis using polynomial chaos expansions, *Reliab. Eng. Syst. Saf.* **93** 964 (2008). doi:10.1016/j.ress.2007.04.002

7 D Xiu and G E Karniadakis, The Wiener–Askey polynomial chaos for stochastic differential equations, *SIAM J. Sci. Comput.* **24** 619 (2002). doi:10.1137/S1064827501387826

8 M D McKay, R J Beckman and W J Conover, A comparison of three methods for selecting values of input variables in the analysis of output from a computer code, *Technometrics* **21** 239 (1979). doi:10.1080/00401706.1979.10489755

9 R L Iman and W J Conover, A distribution-free approach to inducing rank correlation among input variables, *Commun. Stat. Simul. Comput.* **11** 311 (1982). doi:10.1080/03610918208812265

10 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

11 B Iooss and P Lema\^itre, A review on global sensitivity analysis methods, in *Uncertainty Management in Simulation-Optimization of Complex Systems*, Springer (2015). doi:10.1007/978-1-4899-7547-8_5

12 B Efron, Bootstrap methods: another look at the jackknife, *Ann. Stat.* **7** 1 (1979). doi:10.1214/aos/1176344552

13 C R Harris *et al*, Array programming with NumPy, *Nature* **585** 357 (2020). doi:10.1038/s41586-020-2649-2

14 J D Lawson, Some criteria for a power producing thermonuclear reactor, *Proc. Phys. Soc. B* **70** 6 (1957). doi:10.1088/0370-1301/70/1/303

15 H-S Bosch and G M Hale, Improved formulas for fusion cross-sections and thermal reactivities, *Nucl. Fusion* **32** 611 (1992). doi:10.1088/0029-5515/32/4/I07

16 ITER Physics Expert Groups *et al*, Chapter 2: plasma confinement and transport, *Nucl. Fusion* **39** 2175 (1999). doi:10.1088/0029-5515/39/12/302

17 M Greenwald, Density limits in toroidal plasmas, *Plasma Phys. Control. Fusion* **44** R27 (2002). doi:10.1088/0741-3335/44/8/201

18 F Troyon, R Gruber, H Saurenmann, S Semenzato and S Succi, MHD-limits to plasma confinement, *Plasma Phys. Control. Fusion* **26** 209 (1984). doi:10.1088/0741-3335/26/1A/319

19 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

20 A E Costley, J Hugill and P F Buxton, On the power and size of tokamak fusion pilot plants and reactors, *Nucl. Fusion* **55** 033001 (2015). doi:10.1088/0029-5515/55/3/033001

21 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

22 R F Post, The magnetic mirror approach to fusion, *Nucl. Fusion* **27** 1579 (1987). doi:10.1088/0029-5515/27/10/001

23 V P Pastukhov, Collisional losses of electrons from an adiabatic trap in a plasma with a positive potential, *Nucl. Fusion* **14** 3 (1974).

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

Global Sensitivity and Uncertainty Quantification across the Breeder and Burner Design Planes: A 40,000-Sample Monte-Carlo Ranking of the Frozen Anchors

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

Zenodo · doi:10.5281/zenodo.22645893

Read on kronosfusionenergy.com

The Physics De-Risking Register