

The Adele Analogy for Epistemic Boundaries: Local Models, a Restricted Product of Representations, and a Quantified Validity Radius

The Adele Analogy for Epistemic Boundaries: Local Models, a Restricted Product of Representations, and a Quantified Validity Radius

Abstract

Scientific and engineering practice routinely deploys models outside the domains in which they were validated, and the resulting failures are usually discovered only after the fact. Existing treatments of epistemic uncertainty flag that a prediction is uncertain but do not say where a validated model stops being valid or which parts of its content survive a change of representation. This paper develops a formal framework, modeled on the adele ring, in which each validated model is a "local completion" of an underlying phenomenon, each carries an explicit validity boundary, and the robust content of a model class is exactly what is invariant under changes of representation. Models are formalized as local charts with validation domains, combined in a restricted product over representations, and the framework is tested on the Poisson-to-Gaussian max-entropy bridge. With full arithmetic, the paper shows that at $\lambda = 100$ the Gaussian local model reproduces Poisson probabilities to within 3.3% in log-space at one standard deviation but fails by a factor of 1.30 at three standard deviations, locating the 5% log-tolerance validity radius between 2σ and 2.5σ , and identifies the skewness $1/\sqrt{\lambda} = 0.1$ as the leading representation-dependent quantity excluded from the invariant core. The framework turns "we do not know" into a measured radius, applicable to out-of-distribution risk assessment.

1. Introduction

Deep-learning models and statistical models are now validated on a dataset or a regime and then deployed on inputs that lie outside that regime, and the point at which they stop being trustworthy is not marked anywhere in the deployed artifact. Clinicians who use automated contouring of organs at risk in radiotherapy planning face exactly this problem: deep-learning contouring has improved markedly, yet reliability under out-of-distribution scenarios remains a stated clinical concern, and the consequence of an unnoticed validity failure is a treatment plan that harms a patient [7]. The same structure appears wherever a model certified in one regime is extrapolated: the failure is not random noise but a systematic breakdown at a boundary nobody drew.

The cost of not drawing the boundary is concrete. A radiotherapy team that cannot distinguish "the model is interpolating within its validated domain" from "the model is extrapolating" must either apply costly manual review everywhere or accept unquantified risk [7]. An autonomous reasoning system that aggregates scientific claims needs to know when a well-supported local consensus stops supporting a global conclusion, a need that formal architectures for truth-promoting autonomous scientific reasoning already recognize by making belief a dynamic function over structured claims [1]. In both cases, the missing capability is the same: an explicit, computable statement of where a validated model's guarantees end.

If the framework proposed here holds, that capability becomes available in a specific form. Every model would carry a machine-readable validity boundary, stated as a radius in the space of deviations from the validation regime, calibrated by explicit arithmetic of the kind performed in Section 4. Incompatibilities between well-validated models would be classified not as defects but as structured edges of ignorance that indicate where a deeper global object is missing. And the robust content of a model class would be identified by an invariance test: what survives a change of representation is what one is entitled to extrapolate.

The analogy that organizes the framework is the adele ring of number theory. The adele ring $\mathbb{A}$ combines, for every prime p , a local field $\mathbb{Q}_p$ with the real completion $\mathbb{R}$, in a restricted product: an element of $\mathbb{A}$ is a tuple of local elements that is a p -adic integer for all but finitely many p . No single local field contains the arithmetic of the rationals; the global object is recovered only by imposing coherence across all local completions simultaneously. Arithmetic transfer conjectures in the Langlands program make this coherence principle precise by relating local data attached to p -divisible groups across places, and have been proved in low-dimensional cases [4]. A related research program has proposed adelic consistency conditions across primes as a candidate selection principle for physical constants, formulated as a testable conjecture with a negative control and toy dynamics [11]. The present paper transfers the same structural idea from arithmetic to epistemology: each validated model is a local completion, the unknown is the missing global object, and the boundary of a model is the region where its local coherence with other validated models fails.

The paper's contributions are:

- A formal definitions layer: models as local charts with explicit validation domains, a compatibility relation, and an adele-like restricted product over representations, given in Section 3.
- An invariance criterion for robust knowledge: the invariant core of a model class is what is unchanged under admissible changes of representation, with the first non-invariant term computable as a cumulant ratio, derived in Section 4.
- A quantified validity radius for the Poisson-to-Gaussian max-entropy bridge: log-probability discrepancies of 0.0325 nats at 1σ , 0.0296 nats at 2σ , 0.1197 nats at 2.5σ , and a factor of 1.30 at 3σ , all computed with shown arithmetic in Section 4 and reported in Section 5.

- A mapping of the framework onto existing epistemic-uncertainty practice and its users, with technology-readiness assessment and the smallest falsifying experiment, given in Sections 6 and 7.

2. Background and Related Work

This section reviews the works supplied in the bibliography and states, for each, what its own summary supports. Eight of the twelve entries carry substantive summaries; entries [9], [10] and [12] are supplied with empty or minimal summaries and are discussed only through what is stated.

Formal epistemic architectures. Bayesian Epistemology with Weighted Authority (BEWA) responds to the exponential expansion of scientific literature, which its summary states has surpassed the epistemic processing capabilities of both human experts and current artificial intelligence systems; it operationalises belief as a dynamic, probabilistically coherent function over structured scientific claims [1]. BEWA supplies the aggregation layer that a science of epistemic boundaries would sit on top of: aggregation says how to combine validated local models, while the present framework says where each local model's validity ends. The two are complementary, and Section 7 uses BEWA-style structured claims as the natural carrier format for validity-boundary metadata.

Epistemic intervention in learned systems. Frictive Policy Optimization (FPO) treats clarification, verification, challenge, redirection, and refusal as explicit control actions that a language-model policy can take in order to manage epistemic and normative risk, rather than optimizing only surface-level preference or task utility [6]. FPO's refusal and clarification actions are exactly the behavioral outputs a validity boundary should trigger: when a query falls outside a model's validated chart, the framework of this paper provides the formal condition under which refusal is the correct action. The connection is one of interface, not of proof.

Out-of-distribution risk in a deployed clinical system. The radiotherapy workflow paper states that precision of contouring target structures and organs at risk is crucial for treatment efficacy and patient safety, that deep learning has significantly improved organ-at-risk contouring, and that reliability under out-of-distribution scenarios remains a concern in clinical settings [7]. This is the clearest supplied instance of the real-world problem of Section 1: a validated model deployed near the edge of its domain, with no formal boundary drawn. The summary supplied gives no quantitative OOD results, so no performance figure from that work is used here.

Formal specification and the limits of checking intent. For verification-aware programming languages such as Dafny and F*, the summary states that although checking an implementation against a specification is mechanical, there is no algorithmic way of ensuring the correctness of the user-intent formalization itself, because intent is not mechanically checkable against the informal wish [5]. This is an epistemic boundary in miniature: the validated object is the implementation-relative-to-specification, and the boundary is the unformalized gap between specification and intent. The present framework would represent the specification as a local chart whose validation domain excludes the intent relation.

Model interpretability as boundary probing. The white-box analysis of ColBERT dissects the matching process of a transformer-based ad-hoc information retrieval model through term importance and exact and soft matching patterns, in a setting where the summary notes that transformer models are state of the art but their behavior is far from understood, and that claims have been made that BERT does not satisfy the classical IR axioms, while the traditional axioms are not formally violated in the way claimed [3]. Axiom violations and their precise non-violations are boundary phenomena: the analysis locates where an opaque model's behavior departs from, and merely appears to depart from, a classical axiomatic local model. White-box dissection of this kind is a concrete technique for measuring a validation domain empirically.

Logical validity boundaries. The epistemic ATL work gives a validity-preserving translation from a subset of epistemic Alternating-time Temporal Logic to epistemic Computation Tree Logic; the subset has the finite model property and decidable model-checking, so validity is decidable, but the implied algorithm is unfeasible, and the reduction targets a corresponding CTL validity problem [8]. This supplies a precedent for the framework's central move: a property (validity) is preserved under a change of representation (the translation), and the cost of deciding it in one representation versus another is itself part of the epistemic geometry. The unfeasibility caveat is a boundary on the practical side of the reduction.

Arithmetic coherence across places. The regular formal moduli spaces paper defines formal moduli spaces of p -divisible groups that are regular, with morphisms between them, formulates arithmetic transfer conjectures as variants of the arithmetic fundamental lemma conjecture in the presence of ramification, and proves these conjectures in low-dimensional cases [4]. This is the mathematical source of the restricted-product analogy: local objects at different places, coherence conditions across them, and theorems that hold place by place before they hold globally. The summary supplied gives no further technical detail usable here, so the analogy is drawn only at the structural level stated.

Adelic selection principles for physics. The adelic consistency conditions program addresses the three constants problem, namely why the fine-structure constant takes its observed value, why charged-lepton masses are hierarchically spaced, and what fixes the absolute mass scale, for which there is no accepted resolution within the Standard Model; it develops the adelic idea as a testable conjecture with a negative control and toy dynamics [11]. This is the closest prior statement of the adelic analogy applied to empirical science, and the present paper differs in scope: [11] applies adelic coherence to the selection of physical constants, while this paper applies it to the validity boundaries of models in general and quantifies a worked example.

Corpus-level context. The QNFO entries for Philosophy of Science [9] and Epistemic Dynamics [10] are supplied without summary text; no claim about their content is made here, and they are cited only to mark the corpus context in which the research idea arose. Scale-Invariant Reality [12] is likewise supplied without a usable summary; its title indicates the invariance-under-representation theme that Section 3 formalizes, and nothing more is attributed to it.

3. Methods

3.1 Local charts and validation domains

Fix a phenomenon Φ and a space X of observable outcomes. A local model is a pair

$$M = (f_M, D_M),$$

where $f_M : X \rightarrow Y_M$ is a predictive map into an output space Y_M and $D_M \subseteq X$ is the validation domain: the set of inputs on which f_M has been checked against Φ under a stated protocol and tolerance. The pair (f_M, D_M) is the epistemic analogue of a chart in a manifold atlas: f_M is the coordinate map and D_M is the chart's open set. A model is validated on D_M and silent, in the sense of having no earned guarantee, on $X \setminus D_M$.

The validation protocol is part of the model. Two models with the same f_M but different D_M are different epistemic objects, because the boundary is the object of study here.

3.2 Compatibility and the restricted product

Let $\mathcal{R}$ be a set of representations. A representation is a choice of encoding of X and Y : radix, metric, parametrization, observer frame, or discretization. For each representation $r \in \mathcal{R}$ there is a family of local models $\{M_{r,i}\}_i$ validated in that representation.

Two local models M and M' are compatible on an overlap $U \subseteq D_M \cap D_{M'}$ if their predictions agree within a stated tolerance ϵ :

$$\sup_{x \in U} d_Y(f_M(x), f_{M'}(x)) \leq \epsilon.$$

The adele-like object is the restricted product

$$\mathbb{M}_{\mathcal{R}} = \prod_{r \in \mathcal{R}}' \mathcal{M}_r,$$

the set of tuples $(M_r)_{r \in \mathcal{R}}$ with M_r a validated local model in representation r , such that M_r is trivial (identical to a fixed reference model of that representation) for all but finitely many r . The restriction is what makes the product manageable, exactly as the restricted product over primes makes the adele ring locally compact where the full direct product would not be. A global statement about Φ is admissible if it is a property P such that every compatible tuple in $\mathbb{M}_{\mathcal{R}}$ satisfies P on the intersection of its validation domains.

An incompatibility between two well-validated models with overlapping domains is then a pair of charts that cannot be embedded in one compatible tuple. Under the framework's central conjecture, such incompatibilities mark structured edges of ignorance: places where the missing global object would be required to reconcile the local truths, and therefore places where extrapolation is expected to fail.

3.3 The invariant core

Define an admissible change of representation as a map g between representations that preserves the validation protocol. The invariant core of a model class $\mathcal{C}$ is

$$\text{Inv}(\mathcal{C}) = \bigcap_g \{ \text{quantities } q \text{ of } \mathcal{C} : q(gM) = q(M) \text{ for all } M \in \mathcal{C} \}.$$

The framework's second claim is that the robust known is exactly $\text{Inv}(\mathcal{C})$: quantities invariant under representation change are exactly those one is entitled to extrapolate beyond D_M , while the leading non-invariant quantity measures the distance to the boundary. Section 4 computes both for the worked example.

3.4 Worked example: the Poisson-to-Gaussian max-entropy bridge

The test case is the passage from a Poisson counting model to a Gaussian continuous model. Among distributions on $\mathbb{R}$ with fixed mean μ and variance σ^2 , the Gaussian maximizes differential entropy; among distributions on $\{0, 1, 2, \dots\}$ with fixed mean λ , the Poisson is the corresponding max-entropy object for counting processes. The Gaussian with $\mu = \sigma^2 = \lambda$ is therefore the max-entropy continuous representation of the Poisson counting model, and the bridge between them is a change of representation in the sense of Section 3.1. The question the framework must answer quantitatively is: over what region of outcome space does the Gaussian chart remain a valid representation of the Poisson chart, and what is the first quantity that fails to be invariant?

Throughout Section 4 the Poisson parameter is $\lambda = 100$, chosen so that all quantities are computable by hand with Stirling's approximation

$$\ln n! \approx n \ln n - n + \frac{1}{2} \ln(2\pi n) + \frac{1}{12n},$$

whose relative error in $\ln n!$ is of order $1/(360n^3)$, negligible at the digit level used here.

4. Analysis

Every input number in this section is either a definition (the Poisson parameter $\lambda = 100$, the tolerance levels), a standard mathematical constant, or computed below from the stated formulas. No empirical data are used.

4.1 Input constants

The following constants are used, each to at least seven significant figures:

$$\ln 100 = 4.605170, \quad \ln 110 = 4.700480, \quad \ln 120 = 4.787492,$$

$$\ln 125 = 4.828314, \quad \ln 130 = 4.867534,$$

$$\ln(2\pi) = 1.837877, \quad \ln e = 1.$$

Derived per-case constants: $\ln(2\pi \cdot 100) = \ln(628.3185) = 1.837877 + \ln 100 = 1.837877 + 4.605170 = 6.443047$, and half of this is 3.221524 . This quantity $0.5 \ln(2\pi\lambda) = 3.221524$ enters every Gaussian log-probability below.

4.2 Log-probability of the Poisson model at four deviations

The Poisson probability mass is $P(X = k) = e^{-\lambda} \lambda^k / k!$, so

$$\ln P(X = k) = -\lambda + k \ln \lambda - \ln k!.$$

The Gaussian density approximation at k with $\mu = \sigma^2 = \lambda$ gives

$$\ln p_G(k) = -\frac{1}{2} \ln(2\pi\lambda) - \frac{(k - \lambda)^2}{2\lambda}.$$

Case $k = 110$ (one standard deviation, $x = (k - \lambda)/\sqrt{\lambda} = 1$).

Stirling for $\ln 110!$:

- $110 \ln 110 = 110 \times 4.700480 = 517.0528$
- -110 gives 407.0528
- $\frac{1}{2} \ln(2\pi \cdot 110) = \frac{1}{2} \ln(691.1504) = \frac{1}{2} (1.837877 + \ln 110) = \frac{1}{2} (1.837877 + 4.700480) = \frac{1}{2} (6.538357) = 3.269179$
- $\frac{1}{12 \cdot 110} = 1/1320 = 0.000758$

$$\ln 110! = 517.0528 - 110 + 3.269179 + 0.000758 = 410.3227.$$

Then

$$\ln P_{\text{Pois}}(110) = -100 + 110 \times 4.605170 - 410.3227 = -100 + 506.5687 - 410.3227 = -3.7540.$$

Gaussian:

$$\ln p_G(110) = -3.221524 - \frac{10^2}{200} = -3.221524 - 0.5 = -3.721524.$$

Log-difference:

$$\Delta_{1\sigma} = \ln P_{\text{Pois}} - \ln p_G = -3.7540 - (-3.721524) = -0.0325.$$

The Gaussian overestimates the Poisson log-probability by 0.0325 nats at 1σ , a probability relative error of $e^{0.0325} - 1 = 0.0330$, i.e. 3.3% .

Case $k = 120$ (two standard deviations, $x = 2$).

- $120 \ln 120 = 120 \times 4.787492 = 574.4990$
- -120 gives 454.4990
- $\frac{1}{2} \ln(2\pi \cdot 120) = \frac{1}{2} (1.837877 + 4.787492) = \frac{1}{2} (6.625369) = 3.312685$
- $1/1440 = 0.000694$

$$\ln 120! = 574.4990 - 120 + 3.312685 + 0.000694 = 457.8124.$$

$$\ln P_{\text{Pois}}(120) = -100 + 120 \times 4.605170 - 457.8124 = -100 + 552.6204 - 457.8124 = -5.1920.$$

$$\ln p_G(120) = -3.221524 - \frac{20^2}{200} = -3.221524 - 2 = -5.221524.$$

$$\Delta_{2\sigma} = -5.1920 - (-5.221524) = +0.0296.$$

At 2σ the Gaussian underestimates the Poisson log-probability by 0.0296 nats, a relative error of 3.0% in the opposite direction from the 1σ case.

Case $k = 125$ ($x = 2.5$).

- $125 \ln 125 = 125 \times 4.828314 = 603.5393$
- -125 gives 478.5393
- $\frac{1}{2} \ln(2\pi \cdot 125) = \frac{1}{2}(1.837877 + 4.828314) = \frac{1}{2}(6.666191) = 3.333096$
- $1/1500 = 0.000667$

$$\ln 125! = 603.5393 - 125 + 3.333096 + 0.000667 = 481.8731.$$

$$\ln P_{\text{Pois}}(125) = -100 + 125 \times 4.605170 - 481.8731 = -100 + 575.6463 - 481.8731 = -6.2268.$$

$$\ln p_G(125) = -3.221524 - \frac{25^2}{200} = -3.221524 - 3.125 = -6.346524.$$

$$\Delta_{2.5\sigma} = -6.2268 - (-6.346524) = +0.1197.$$

At 2.5σ the Gaussian underestimates the Poisson log-probability by 0.1197 nats, corresponding to a probability relative error of $e^{0.1197} - 1 = 0.127$, i.e. 12.7%, an order of magnitude larger than the errors at 1σ and 2σ .

Case $k = 130$ (three standard deviations, $x = 3$).

Stirling for $\ln 130!$:

- $130 \ln 130 = 130 \times 4.867534 = 632.7794$
- -130 gives 502.7794
- $\frac{1}{2} \ln(2\pi \cdot 130) = \frac{1}{2}(1.837877 + 4.867534) = \frac{1}{2}(6.705411) = 3.352706$
- $1/(12 \cdot 130) = 1/1560 = 0.000641$

$$\ln 130! = 632.7794 - 130 + 3.352706 + 0.000641 = 506.1328.$$

$$\ln P_{\text{Pois}}(130) = -100 + 130 \times 4.605170 - 506.1328 = -100 + 598.6721 - 506.1328 = -7.4607.$$

$$\ln p_G(130) = -3.221524 - \frac{30^2}{200} = -3.221524 - 4.5 = -7.721524.$$

$$\Delta_{3\sigma} = -7.4607 - (-7.721524) = +0.2608.$$

The probability ratio is $e^{0.2608} = 1.298$: at 3σ the Gaussian underestimates the Poisson probability by a factor of 1.30. This is no longer a small perturbation; it is the kind of discrepancy that would misprice a tail risk by tens of percent while every central-region check passes.

4.3 Leading-order asymptotics and the residual pattern

The four computed points are not scattered; they follow a derivable cubic law. Substituting $k = \lambda + x\sqrt{\lambda}$ into the Stirling-based log-probability and expanding in powers of $x/\sqrt{\lambda}$:

$$\ln P_{\text{Pois}}(k) = k - \lambda + k \ln \frac{\lambda}{k} - \frac{1}{2} \ln(2\pi k) - \frac{1}{12k} + O(k^{-3}).$$

Expanding $k \ln(\lambda/k) = -(\lambda + x\sqrt{\lambda}) \ln(1 + x/\sqrt{\lambda})$ as a power series in $u = x/\sqrt{\lambda}$:

$$(\lambda + x\sqrt{\lambda}) \ln(1 + u) = x\sqrt{\lambda} + \frac{x^2}{2} - \frac{x^3}{6\sqrt{\lambda}} + O(\lambda^{-1}),$$

so that

$$\ln P_{\text{Pois}}(k) = -\frac{x^2}{2} - \frac{1}{2} \ln(2\pi\lambda) + \frac{x^3 - 3x}{6\sqrt{\lambda}} + O(\lambda^{-1}).$$

The first two terms are exactly the Gaussian log-density $\ln p_G(k)$ at x . The leading discrepancy is therefore

$$\Delta(x) = \ln P_{\text{Pois}} - \ln p_G \approx \frac{x^3 - 3x}{6\sqrt{\lambda}} + O(\lambda^{-1}).$$

Checking this against the four hand computations at $\lambda = 100$, so $6\sqrt{\lambda} = 60$:

x	cubic prediction $(x^3 - 3x)/60$	computed Δ	residual
1	$(1 - 3)/60 = -0.0333$	-0.0325	0.0008
2	$(8 - 6)/60 = +0.0333$	+0.0296	0.0037
2.5	$(15.625 - 7.5)/60 = +0.1354$	+0.1197	0.0157
3	$(27 - 9)/60 = +0.3000$	+0.2608	0.0392

The residuals grow with x because the omitted terms are of order x^4/λ and x^6/λ^2 ; at $x = 3$ the term $x^4/(4\lambda) = 81/400 = 0.2025$ enters the raw series before partial cancellation, so a residual of 0.039 at $x = 3$ is consistent with the stated $O(\lambda^{-1})$ remainder. The structural point survives the residuals: the discrepancy is a deterministic cubic in x , scaled by $1/\sqrt{\lambda}$, not random noise. A user who validated the Gaussian chart only on the central region, where $|x| \leq 1$ and $|\Delta| \leq 0.033$, would see small residuals everywhere they looked and would have no warning that the same cubic grows as x^3 and crosses any fixed tolerance at a computable radius.

4.4 The invariant core: cumulants under the representation change

The framework of Section 3.3 requires identifying which quantities of the Poisson model survive the change of representation to the Gaussian chart. Cumulants are the natural probe because the Gaussian family is characterized by the vanishing of all cumulants of order 3 and higher, and because the max-entropy bridge is built by fixing constraints.

For the Poisson distribution with parameter λ , the cumulant generating function is $\ln E[e^{tX}] = \lambda(e^t - 1)$, so every cumulant equals λ :

$$\kappa_1 = \kappa_2 = \kappa_3 = \kappa_4 = \dots = \lambda = 100.$$

For the Gaussian $\mathcal{N}(\lambda, \lambda)$, the cumulant generating function is $\ln E[e^{tY}] = \lambda t + \lambda t^2/2$, so

$$\kappa_1 = \lambda = 100, \quad \kappa_2 = \lambda = 100, \quad \kappa_n = 0 \text{ for } n \geq 3.$$

The representation-invariant quantities are therefore the mean $\kappa_1 = 100$ and the variance $\kappa_2 = 100$: these are exactly the two constraints that define the max-entropy bridge, and they are identical in both charts. The first non-invariant quantity is the third cumulant, $\kappa_3 = 100$ in the Poisson chart versus $\kappa_3 = 0$ in the Gaussian chart. In standardized form, the Poisson skewness is

$$\gamma_1 = \frac{\kappa_3}{\kappa_2^{3/2}} = \frac{\lambda}{\lambda^{3/2}} = \frac{1}{\sqrt{\lambda}} = \frac{1}{10} = 0.1,$$

while the Gaussian skewness is 0. The excess kurtosis is $\kappa_4/\kappa_2^2 = 1/\lambda = 0.01$ in the Poisson chart and 0 in the Gaussian chart.

This identifies the epistemic content of the boundary precisely. The quantity excluded from the invariant core, the standardized skewness $1/\sqrt{\lambda} = 0.1$, is the same quantity whose signature appears in the cubic discrepancy law of Section 4.3: the factor $x^3 - 3x$ in $\Delta(x)$ is (up to convention) the third Hermite polynomial, and its coefficient $1/(6\sqrt{\lambda})$ is exactly $\gamma_1/6$. The leading representation-dependent cumulant and the leading log-probability error are the same object viewed twice. Under the framework's second claim, this is the general pattern: the first non-invariant cumulant, which is computable in the validation domain itself, predicts the shape and scale of the log-probability error outside it.

4.5 The validity radius

Fix the tolerance $\epsilon = 0.05$ in log-space: the Gaussian chart is declared valid at x if $|\Delta(x)| \leq 0.05$. The computed values give:

- $|\Delta(1)| = 0.0325 \leq 0.05$: valid.
- $|\Delta(2)| = 0.0296 \leq 0.05$: valid.
- $|\Delta(2.5)| = 0.1197 > 0.05$: invalid.

The 5% log-tolerance validity radius therefore lies strictly between 2σ and 2.5σ . To locate it more finely, two estimates are available, both labeled projections because they interpolate between computed points rather than evaluating a new one.

Projection 1 (linear interpolation between computed points). Interpolating $|\Delta|$ linearly in x between $(2, 0.0296)$ and $(2.5, 0.1197)$:

$$x_\epsilon = 2 + 0.5 \times \frac{0.05 - 0.0296}{0.1197 - 0.0296} = 2 + 0.5 \times \frac{0.0204}{0.0901} = 2 + 0.5 \times 0.2264 = 2.113.$$

Projection 2 (leading-order cubic). Solving $(x^3 - 3x)/60 = 0.05$, i.e. $x^3 - 3x - 3 = 0$: at $x = 2.10$, $9.261 - 6.3 - 3 = -0.039$; at $x = 2.11$, $9.394 - 6.33 - 3 = +0.064$; linear interpolation between these gives $x \approx 2.10 + 0.01 \times 0.039/0.103 = 2.104$.

Both projections agree that the validity radius is $R \approx 2.1\sigma$, with the honest caveat that the true value requires evaluating Δ at a new point; the two methods differ by about 0.01σ , and the $O(\lambda^{-1})$ corrections of Section 4.3, which push the computed Δ below the cubic prediction at large x , would push the true radius slightly above both estimates. The robust statement, derived without interpolation, is the bracket: $2\sigma < R_{0.05} < 2.5\sigma$.

The same arithmetic gives the radius as a function of λ under the cubic law. Setting $(x^3 - 3x)/(6\sqrt{\lambda}) = \epsilon$ and taking the large- x root $x \approx (6\epsilon\sqrt{\lambda})^{1/3}$: for $\epsilon = 0.05$ and $\lambda = 100$, $(6 \times 0.05 \times 10)^{1/3} = 3^{1/3} = 1.442$, which underestimates the true root 2.1 because the $-3x$ term is not negligible here; the full cubic is the correct instrument at $\lambda = 100$, and the pure-cubic scaling $R \propto \lambda^{1/6}$ emerges only for λ large enough that $x^3 \gg 3x$ at the root. This $\lambda^{1/6}$ scaling is itself a framework prediction: the validity radius of a local continuous chart grows only as the sixth root of the sample-size-like parameter, so a hundredfold increase in λ widens the radius by a factor of only $100^{1/6} = 2.154$.

5. Results

All numbers in this section are those computed with shown arithmetic in Section 4; the two interpolated radii are labeled projections with their assumptions stated in Section 4.5.

R1. Log-probability discrepancies of the Gaussian chart, $\lambda = 100$. At 1σ ($k = 110$), $\Delta = -0.0325$ nats, a 3.25% log-space error, with the Gaussian overestimating the probability by a factor $e^{0.0325} = 1.033$. At 2σ ($k = 120$), $\Delta = +0.0296$ nats, a 3.0% log-space error in the opposite direction. At 2.5σ ($k = 125$), $\Delta = +0.1197$ nats, a probability relative error of $e^{0.1197} - 1 = 12.7\%$. At 3σ ($k = 130$), $\Delta = +0.2608$ nats, and the Gaussian underestimates the Poisson probability by a factor of 1.30.

R2. The discrepancy is a deterministic cubic, not noise. The computed discrepancies follow the derived law $\Delta(x) \approx (x^3 - 3x)/(6\sqrt{\lambda})$ with residuals 0.0008, 0.0037, 0.0157, 0.0392 at $x = 1, 2, 2.5, 3$ respectively, consistent with the derived $O(\lambda^{-1})$ remainder. The error changes sign between 1σ and 2σ , so a one-sided validation check would pass at both points while the underlying cubic grows without bound as x^3 .

R3. The invariant core and the first excluded quantity. The representation-invariant quantities are the mean $\kappa_1 = 100$ and variance $\kappa_2 = 100$, exactly the max-entropy constraints. The first non-invariant quantity is the third cumulant ($\kappa_3 = 100$ versus 0), whose standardized form is the skewness $\gamma_1 = 1/\sqrt{\lambda} = 0.1$; its coefficient appears directly in the discrepancy law as $\gamma_1/6$ multiplying the third Hermite polynomial $x^3 - 3x$. The excess kurtosis $1/\lambda = 0.01$ is the next excluded quantity.

R4. The validity radius. The 5% log-tolerance validity radius is bracketed by the computed points at $2\sigma < R_{0.05} < 2.5\sigma$. Two labeled projections agree on $R_{0.05} \approx 2.1\sigma$: linear interpolation between computed points gives 2.113σ , and the leading-order cubic gives 2.104σ . Under the cubic law, the radius scales as $\lambda^{1/6}$ for large λ , so a hundredfold increase in λ widens it by a factor of 2.154.

R5. The shape of the failure. The failure at the boundary is systematic and signed: the Gaussian chart underestimates right-tail Poisson probabilities beyond roughly 1.7σ (where the cubic $x^3 - 3x$ changes sign, at $x = \sqrt{3} = 1.732$), and the underestimate grows cubically. A model card that reported only central-region agreement would certify this chart with no indication that its tail probabilities are off by 30% at 3σ .

6. Discussion

The worked example supports the framework's two operational claims in a limited and precisely circumscribed way. The invariant-core claim survived its first test in the strong sense that the excluded quantity (skewness 0.1) is not merely correlated with the boundary failure but is its leading coefficient: the cumulant computation done entirely inside the validation domain predicted the functional form of the out-of-domain error. The validity-radius claim survived in the weaker sense that the radius was bracketed rather than computed exactly, and the fine value rests on interpolation.

Limitations. The result rests on the following assumptions, each with its sensitivity:

- **Stirling truncation.** All Poisson log-probabilities use the Stirling series through the $1/(12n)$ term, whose next term is of order $1/(360n^3)$; at $n = 110$ this is $1/(360 \times 1331000) \approx 2.1 \times 10^{-9}$ nats, so it cannot move any reported digit. Sensitivity: negligible.
- **Input log constants.** The logarithms in Section 4.1 are given to six decimal places; an error of 5×10^{-7} in $\ln 110$, multiplied by 110 in $110 \ln 110$, contributes up to 6×10^{-5} nats, an order of magnitude below the smallest reported discrepancy (0.0008 residual at $x = 1$ is itself near this floor). Sensitivity: the $x = 1$ residual is at the precision floor; the $x = 2.5$ and $x = 3$ residuals are far above it.
- **Single parameter value.** All computations use $\lambda = 100$. The cubic law and the $\lambda^{1/6}$ radius scaling are derived asymptotically but tested at only one point. Sensitivity: high; the framework's general claims stand or fall on whether the pattern repeats at other λ , which is the

proposed next experiment (Section 7).

- **Interpolated radius.** The fine value $R_{0.05} \approx 2.1\sigma$ is a projection from two interpolation schemes, not a computed point; the derived, interpolation-free statement is only the bracket $2\sigma < R_{0.05} < 2.5\sigma$. Sensitivity: moderate; the bracket is exact given the four computed points, the fine value could shift by a few hundredths of σ .
- **Tolerance choice.** The radius depends on the chosen $\epsilon = 0.05$; a different tolerance moves the radius along the cubic and changes it by construction, not by error. Sensitivity: definitional.
- **Single worked example.** The identification of the first non-invariant cumulant with the leading error coefficient is demonstrated for one representation change (Poisson to Gaussian). The framework's second claim asserts this is general; that generality is a conjecture supported by one instance, not a theorem. Sensitivity: this is the largest gap between what is derived and what is claimed.

What would falsify the main result. The main quantitative result is that the log-probability discrepancy of the max-entropy bridge is a deterministic cubic $\Delta(x) \approx (x^3 - 3x)/(6\sqrt{\lambda})$ whose coefficient is the excluded skewness $\gamma_1/6$, and that this yields a computable validity radius. Three observations would falsify it. First, evaluating Δ at a new point, for example $x = 2.25$ ($k = 122$ or 123), and obtaining a value inconsistent with the cubic prediction $(11.39 - 6.75)/60 = 0.0773$ beyond the residual trend established at the four computed points would break the law. Second, computing the exact Poisson log-probabilities with full-precision arithmetic and finding discrepancies inconsistent with the hand values by more than the stated precision floor (10^{-4} nats) would invalidate the derivations. Third, repeating the computation at a different λ , for example $\lambda = 400$ where the cubic predicts $\Delta(3) = (27 - 9)/120 = 0.15$, and finding a discrepancy that does not scale as $1/\sqrt{\lambda}$ would falsify the claimed scaling and with it the identification of the skewness as the governing quantity. Each falsification test requires only arithmetic on a known distribution; none requires new data.

Open questions. Whether the restricted-product structure of Section 3.2 supports a locality theorem analogous to local compactness, in a sense meaningful for model classes, is open. Whether incompatibilities between well-validated models always mark recoverable global objects, rather than genuine inconsistency, is open and in the supplied literature unaddressed. Whether the invariant-core criterion extends beyond cumulant-based probes to models without a natural exponential-family structure is open.

7. Practical Implications and Engineering Pathways

(a) Who would use the result and what changes. The immediate users are builders of model-governance and uncertainty tooling, and clinical-ML deployment teams of the kind described in the radiotherapy workflow literature, where the supplied summary states that reliability under out-of-distribution scenarios remains a concern in clinical settings [7]. If the framework holds, the decision that changes is the review policy: instead of a binary "in distribution / out of distribution" flag, the deployed artifact carries a measured radius R_ϵ in a stated deviation coordinate, and manual review is triggered only when the operating point approaches R_ϵ . A second user class is aggregation architectures for autonomous scientific reasoning: BEWA operationalises belief as a dynamic, probabilistically coherent function over structured scientific claims [1], and a validity radius is exactly the kind of structured metadata such a function could consume, weighting a claim by the distance of the query from the validating chart. A third class is language-model policy design: FPO treats refusal and clarification as explicit control actions [6], and the framework supplies the formal trigger condition for those actions. For users that cannot yet be named, the work is foundational: the restricted-product formalism and the invariant-core criterion are definitions and one worked example, not yet a deployable estimator.

(b) Current readiness. Against the NASA Technology Readiness Level definitions, this work is analytical: the framework is defined, one quantitative example is computed by hand, and nothing has been peer reviewed or validated outside the worked example. This places it at TRL 1–2 (basic principles observed and formulated; concept and application stated without proof of feasibility beyond the toy case). Analytical work does not exceed TRL 3, and no claim of implementation, manufacturing, or deployment readiness is made.

(c) The smallest next experiment. The smallest falsifying calculation is a full-precision evaluation of the Poisson and Gaussian log-probabilities at $\lambda = 100$ and one new deviation, $x = 2.25$, plus one new parameter value, $\lambda = 400$ at $x = 3$. What is measured: the log-probability discrepancy Δ at those points, compared against the cubic predictions 0.0773 (at $\lambda = 100$, $x = 2.25$) and 0.15 (at $\lambda = 400$, $x = 3$). The outcome that refutes the main result is a discrepancy inconsistent with the cubic law beyond the residual trend of Section 4.3, or a λ -scaling that is not $1/\sqrt{\lambda}$. This requires only a scripting-language one-liner per point and no data collection.

(d) Binding constraints. For the framework itself, the first binding constraint is generality: the invariant-core criterion is demonstrated only for exponential-family cumulants, and extending it to neural models requires choosing a deviation coordinate and a representation set $\mathcal{R}$, which are design decisions with no canonical answer yet. For deployment in a clinical or governance setting, the binding constraints are, in likely order: (i) defining the deviation coordinate x for real model classes, where the analogue of $(k - \lambda)/\sqrt{\lambda}$ is not given; (ii) calibrating the tolerance ϵ per application, which is a policy input, not a derivable one; (iii) data and compute for empirical calibration of R_ϵ on real validation sets; and (iv) regulatory acceptance of a radius-based review trigger, which is an open question not addressed by the supplied grounding. Cost figures for these steps cannot be derived from the supplied material and are not projected here.

(e) Order-of-magnitude performance bound (projection). One performance projection is derivable from Section 4.5 under its stated assumptions: under the cubic law with the pure-cubic regime $x^3 \gg 3x$ at the root, the validity radius scales as $R \propto \lambda^{1/6}$, so a hundredfold increase in the sample-size-like parameter λ widens the radius by $100^{1/6} = 2.154$. Assumptions: large λ , fixed tolerance $\epsilon = 0.05$,

exponential-family setting. This is a statement about how slowly honest validity boundaries widen with data, and it is the framework's most checkable general prediction.

(f) Risks and misuse paths. A quantified radius can be misused as a certificate: a team could report $R_{0.05} \approx 2.1\sigma$ as if it were an exact, validated boundary when it is an interpolation between four computed points at one parameter value, and could then operate up to the reported radius with less review than before. The radius is also tolerance-relative, so choosing a looser ϵ manufactures a larger radius without any change in the model. Mitigation is disclosure: any deployed radius must carry its tolerance, its deviation coordinate, its parameter regime, and its interpolation status, in the machine-readable form proposed in Section 3.

8. Conclusion

This paper proposed that the adele ring's structure — local completions combined in a restricted product, with a global object recovered only through cross-place coherence — is a productive analogy for epistemic boundaries: each validated model is a local chart, incompatibilities between well-validated charts mark structured edges of ignorance, and the robust content of a model class is its invariant core under changes of representation. The framework was made quantitative on the Poisson-to-Gaussian max-entropy bridge at $\lambda = 100$: the Gaussian chart's log-probability error is a deterministic cubic $\Delta(x) \approx (x^3 - 3x)/(6\sqrt{\lambda})$ with computed values -0.0325 , $+0.0296$, $+0.1197$, and $+0.2608$ nats at $x = 1, 2, 2.5, 3$, the last corresponding to a factor-1.30 probability underestimate; the 5% log-tolerance validity radius is bracketed at $2\sigma < R_{0.05} < 2.5\sigma$ with a projected value near 2.1σ ; and the coefficient of the error law is exactly the first quantity excluded from the invariant core, the skewness $1/\sqrt{\lambda} = 0.1$. The general claims — that invariant cores predict extrapolation failure in general, and that the restricted product is the right global object — remain conjectures supported by one fully worked instance. The framework's contribution at its current readiness is a change of question: from "is the model uncertain?" to "at what measured radius, in which coordinate, under which tolerance does the chart stop being valid?", a question that admits arithmetic answers of the kind computed here and falsification tests that require only a calculator.

Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.16.0-revise-retry). Language models were tools in that process; they are not authors. The named author is accountable for the content.
- Models: drafting by openai/gpt-oss-120b, zai-org/glm-5.3-flash; reconciliation, completion, review, revision and verification by zai-org/glm-5.3-flash, zai-org/glm-5.3, openai/gpt-oss-120b, with the QNFO model router as a fallback.
- Sources: every reference is drawn from a bibliography retrieved from arXiv and the QNFO corpus before drafting; statements about a cited work are restricted to its retrieved title and abstract.
- Checks passed before publication: required sections present and finished, cross-references resolve, reference list rendered from the retrieved bibliography, no assistant residue or placeholders, adversarial model review with its blocking findings resolved, and a citation existence check at publication.
- Independent recomputation: a separately generated script recomputed 16 quantitative claims from the paper's stated inputs; 0 differed from the stated value at the stated precision.
- Not checked: whether each cited source supports the sentence that cites it beyond its abstract; derivation steps other than the recomputed claims. This paper has not been peer reviewed, and this version was published by the pipeline without a prior manual review by the author.

References

- [1] Bayesian Epistemology with Weighted Authority: A Formal Architecture for Truth-Promoting Autonomous Scientific Reasoning. arXiv:2506.16015v1. <https://arxiv.org/abs/2506.16015v1> [3] A White Box Analysis of ColBERT. arXiv:2012.09650v1. <https://arxiv.org/abs/2012.09650v1> [4] Regular formal moduli spaces and arithmetic transfer conjectures. arXiv:1604.02419v2. <https://arxiv.org/abs/1604.02419v2> [5] Evaluating LLM-driven User-Intent Formalization for Verification-Aware Languages. arXiv:2406.09757v2. <https://arxiv.org/abs/2406.09757v2> [6] Frictive Policy Optimization for LLMs: Epistemic Intervention, Risk-Sensitive Control, and Reflective Alignment. arXiv:2604.25136v1. <https://arxiv.org/abs/2604.25136v1> [7] Towards Integrating Epistemic Uncertainty Estimation into the Radiotherapy Workflow. arXiv:2409.18628v1. <https://arxiv.org/abs/2409.18628v1> [8] Reducing Validity in Epistemic ATL to Validity in Epistemic CTL. arXiv:1303.0794v1. <https://arxiv.org/abs/1303.0794v1> [9] QNFO: Philosophy of Science [10] QNFO: Epistemic Dynamics [11] QNFO: Adelic Consistency Conditions Across Primes as a Candidate Selection Principle for Physical Constants: Formulation, Negative Control, and Toy Dynamics [12] QNFO: Scale-Invariant Reality