

AUGUST 1, 2026 · V1.0

When Will Non-Archimedean Geometry Displace the Real Numbers? A Structured Assessment of the Adelic Substrate Thesis

Rowan Brad Quni-Gudzin

QNFO-ULA-2.1

ABSTRACT

The real numbers are the default continuum for physical modeling — every quantum state lives in a Hilbert space over $\mathbb{C}$, every neural network optimizes over $\mathbb{R}^n$, every error-correcting code is designed with Euclidean distance. Yet Ostrowski's theorem (1916) demonstrates that the real numbers are merely *one of many* completions of the rationals: every non-trivial absolute value on $\mathbb{Q}$ is equivalent to either the Archimedean absolute value or a p -adic absolute value for some prime p . The Archimedean default is a choice, not a necessity. Two domains already use non-Archimedean geometry as their native substrate — phylogenetic inference (ultrametric distance) and isogeny-based cryptography (p -adic arithmetic geometry) — establishing that the paradigm is viable. We examine whether and when non-Archimedean modeling will extend beyond these existence proofs into a third domain, assessing five candidate domains qualitatively and anchoring our judgments to a historical reference class of mathematical framework adoptions. We find that the central bet — displacement in at least one domain by 2036 — has a calibrated probability of approximately 0.15–0.20, consistent with a reference class in which zero of the three closest historical analogues (non-Euclidean geometry → general relativity, abstract algebra → quantum mechanics, information theory → biology) achieved mainstream adoption within a decade. The null hypothesis — that non-Archimedean modeling remains a niche mathematical tool — is the well-anchored prediction. We register nine dated, strength-weighted calibration predictions and identify actionable near-term interventions that could compress the historical timeline.

CITE AS Quni-Gudzin, R. B. (2026). When Will Non-Archimedean Geometry Displace the Real Numbers? A Structured Assessment of the Adelic Substrate Thesis. QNFO Preprint v1.0.

Author: Rowan Brad Quni-Gudzin | **Date:** 2026-08-01 | **License:** QNFO-ULA: <https://legal.qnfo.org/>

When Will Non-Archimedean Geometry Displace the Real Numbers?

Abstract

The real numbers are the default continuum for physical modeling — every quantum state lives in a Hilbert space over $\mathbb{C}$, every neural network optimizes over $\mathbb{R}^n$, every error-correcting code is designed with Euclidean distance. Yet Ostrowski's theorem (1916) demonstrates that the real numbers are merely *one of many* completions of the rationals: every non-trivial absolute value on $\mathbb{Q}$ is equivalent to either the Archimedean absolute value or a p -adic absolute value for some prime p . The Archimedean default is a choice, not a necessity. Two domains already use non-Archimedean geometry as their native substrate — phylogenetic inference (ultrametric distance) and isogeny-based cryptography (p -adic arithmetic geometry) — establishing that the paradigm is viable. We examine whether and when non-Archimedean modeling will extend beyond these existence proofs into a third domain, assessing five candidate domains qualitatively and anchoring our judgments to a historical reference class of mathematical framework adoptions. We find that the central bet — displacement in at least one domain by 2036 — has a calibrated probability of approximately 0.15–0.20, consistent with a reference class in which zero of the three closest historical analogues

(non-Euclidean geometry → general relativity, abstract algebra → quantum mechanics, information theory → biology) achieved mainstream adoption within a decade. The null hypothesis — that non-Archimedean modeling remains a niche mathematical tool — is the well-anchored prediction. We register nine dated, strength-weighted calibration predictions and identify actionable near-term interventions that could compress the historical timeline.

1 Introduction

Every scientific model begins with a choice of number system. Quantum mechanics uses complex Hilbert spaces. General relativity uses pseudo-Riemannian manifolds over the reals. Neural networks optimize real-valued loss functions. Error-correcting codes measure distance with the Hamming metric or the Euclidean metric on the sphere. In every case, the default choice is Archimedean: the real numbers with their familiar ordering, limits, and triangle inequality $d(x, z) \leq d(x, y) + d(y, z)$.

This default is not forced by nature. It is a historical and pedagogical choice. Ostrowski's theorem establishes that every non-trivial absolute value on $\mathbb{Q}$ is equivalent to either the usual real absolute value $|\cdot|_\infty$ or the p -adic absolute value $|\cdot|_p$ for some prime p . The p -adic numbers $\mathbb{Q}_p$ satisfy the *strong* triangle inequality:

$$|x + y|_p \leq \max(|x|_p, |y|_p),$$

which implies that the induced metric is *ultrametric*: every triangle is isosceles with the base no longer than either leg. The geometry of $\mathbb{Q}_p$ is fundamentally different from that of $\mathbb{R}$ — it is totally disconnected, its unit ball is a ring, and its natural state space is a tree (the Bruhat–Tits building) rather than a continuum.

The question is not whether non-Archimedean geometry is mathematically valid — it is indisputably a consistent, rich mathematical framework with a century of development. The question is whether it will *displace* the Archimedean default as the native modeling substrate in any scientific or engineering domain outside of the two domains where it is already established. This paper provides a structured assessment of that question, identifying the most plausible candidate domains, anchoring probability judgments to historical reference classes, and registering dated, falsifiable predictions.

1.1 The Existence Proofs

Two domains already use non-Archimedean geometry as their *native* substrate — not as a re-encoding of Archimedean models, but as the fundamental distance measure, state space, or algebraic structure at the level of definition.

Phylogenetics. The molecular clock hypothesis [1] assumes constant mutation rates across lineages, producing ultrametric distance matrices. Algorithms such as UPGMA and neighbor-joining were designed specifically for ultrametric data: the distance between any two species is the time since their most recent common ancestor, and the ultrametric inequality $d(x, z) \leq \max(d(x, y), d(y, z))$ holds by construction. Phylogenetics has used non-Archimedean distance as its native substrate since the 1960s.

Isogeny-based cryptography. Protocols such as SIKE and CSIDH operate on elliptic curves over finite fields, using isogeny graphs whose geometry is inherently non-Archimedean — the p -adic valuation governs the arithmetic of the ground field, and the natural graph structure is that of a Bruhat–Tits tree. Isogeny-based cryptography is a NIST post-quantum candidate, making non-Archimedean geometry a practical engineering substrate in a high-stakes domain.

These are not predictions. They are existence proofs. The central bet of this paper is whether a *third* domain will follow — one where the non-Archimedean structure is not imposed by the problem constraints (as in phylogenetics) or inherited from number theory (as in isogeny crypto) but *chosen* as the superior modeling framework.

2 Domain Assessment

The current landscape of non-Archimedean modeling outside the established domains is sparse but active.

Domain	Current Paradigm	Non-Archimedean Presence	Approximate Activity
Quantum foundations	Hilbert-space QM over $\mathbb{C}$	p -adic QM (Vladimirov, Khrennikov)	~50 papers/year
High-energy theory	Continuum QFT, string theory	p -adic string theory, adelic string theory	~20 papers/year
Quantum error correction	Stabilizer codes, Euclidean sphere-packing	Ultrametric code metrics, Bruhat–Tits code constructions	<5 papers/year
Machine learning	Euclidean gradient-based optimization	Ultrametric clustering, p -adic neural networks	<10 papers/year
Information theory	Archimedean channel models	p -adic channel coding, ultrametric rate-distortion	<5 papers/year
Condensed matter	Archimedean lattice models	p -adic AdS/CFT, ultrametric spin glasses	<10 papers/year
Phylogenetics	Distance-based tree inference	Ultrametric distance is native	Established
Isogeny cryptography	Elliptic curves	p -adic arithmetic is native	Established

The key structural observation is that non-Archimedean modeling is concentrated in *nascent* subfields within each domain — active but marginal, with no single result achieving the visibility threshold (a *Physical Review Letters* publication, a NeurIPS best paper, or a NIST standard) that would signal paradigm-level adoption.

3 Candidate Assessment

We identify five candidate domains where non-Archimedean modeling could achieve native-substrate status by 2036, ranked by a qualitative assessment of probability, impact, timeline readiness, and testability.

3.1 Candidate A: Ultrametric Quantum Error-Correcting Codes (QEC)

Quantum error correction is the most promising near-term candidate. QEC is an active, well-funded research area with an objective performance metric — the code threshold, the physical error rate below which logical error rates can be made arbitrarily small by scaling the code distance. Current surface codes achieve thresholds of approximately 1% under circuit-level depolarizing noise using Euclidean sphere-packing geometry.

The Bruhat–Tits tree — the natural state space for the p -adic group $SL(2, \mathbb{Q}_p)$ — offers an alternative geometry for code construction. Logical qubits could be encoded at tree vertices, with transversal gates corresponding to tree automorphisms and error syndromes tagged by Bruhat–Tits apartment coordinates. The ultrametric distance would provide a natural error distinguishability measure: errors at different tree depths are naturally separated by the strong triangle inequality.

Underlying this candidate are several critical assumptions. First, existing QEC code constructions must admit a reformulation on Bruhat–Tits trees without catastrophic performance loss — a plausible claim given that geometric reformulations (e.g., symplectic codes to toric codes) typically preserve or improve performance. Second, the tree structure must admit a tensor-product operation compatible with concatenated coding, which is an open mathematical question for buildings of rank greater than one. Third, and most critically, a tree-based code must achieve a threshold within a factor of two of the surface-code threshold — currently an unbuilt construction.

The scaling concern is genuine. Bruhat–Tits buildings of rank $n > 1$ grow super-exponentially in vertex count, limiting practical constructions to the tree case ($\mathrm{SL}(2, \mathbb{Q}_p)$). However, binary tree codes may be sufficient: concatenated codes on Bruhat–Tits trees could match or exceed surface-code performance without requiring higher-rank buildings.

Probability range: 0.12–0.25 by 2036. **Impact:** 9/10. **Timeline:** Short — active research area with existing infrastructure. **Testability:** High — code thresholds are measurable in simulation.

3.2 Candidate B: p -adic Neural Architectures (Machine Learning)

Machine learning offers the largest potential audience. The community of ML researchers exceeds that of QEC researchers by roughly two orders of magnitude, meaning that even a niche result in p -adic neural architectures would reach far more practitioners than a comparable result in QEC.

The structural argument is that hierarchical data — taxonomies, ontologies, language trees, organizational charts — has a natural ultrametric structure that p -adic networks could exploit more efficiently than Euclidean networks. An embedding layer that maps inputs to p -adic vectors where distance is ultrametric would naturally respect the tree structure of hierarchical classification tasks (ImageNet's WordNet hierarchy, for instance).

The critical assumptions are more demanding than for Candidate A. Gradient-based optimization must be adapted to p -adic fields — the p -adic equivalent of backpropagation is not well-studied. Efficient p -adic activation functions (analogous to ReLU or sigmoid) must exist and must be computationally viable. Most importantly, a p -adic architecture must demonstrate a benchmark improvement over Euclidean baselines on at least one standard task — a 3–5% improvement in top-1 accuracy would constitute a compelling existence proof.

The tooling barrier is severe. No major ML framework (PyTorch, JAX, TensorFlow) supports p -adic tensors as a first-class primitive. Researchers must implement p -adic arithmetic from scratch, which is a prohibitive barrier for exploratory work.

Probability range: 0.10–0.22 by 2036. **Impact:** 8/10. **Timeline:** Medium — requires building tooling from scratch. **Testability:** High — benchmark accuracy is a standard metric.

3.3 Candidate C: Non-Archimedean Quantum Foundations

The highest-impact candidate is also the longest-timeline and highest-barrier. A modification to the fundamental state space of quantum mechanics — replacing the complex Hilbert space with a p -adic Hilbert space $\mathbb{C}_p$ — would be the most significant change to the mathematical foundations of physics since von Neumann's axiomatization in 1932.

The barrier is empirical. Archimedean quantum mechanics has 100 years of experimental validation with zero anomalies. Every successful paradigm shift in fundamental physics — the Copernican revolution, relativity, quantum mechanics — was preceded by a specific empirical failure of the incumbent framework. Until an anomaly emerges that Archimedean QM cannot explain (or can explain only with ad hoc modifications), non-Archimedean QM will remain a mathematical exercise.

The most actionable approach is a theoretical reconnaissance: build the p -adic QM formalism, identify a testable signature (a Bell-type experiment where p -adic and standard QM predict different correlations), and wait for experimentalists to reach the required precision. This is a background-priority research program, not a near-term empirical bet.

Probability range: 0.08–0.20 by 2036. **Impact:** 9/10. **Timeline:** Long — requires an empirical anomaly first. **Testability:** Medium — requires experimental precision beyond current capabilities.

3.4 Candidates D and E: Channel Coding and AdS/CFT

Two smaller candidates merit mention. Ultrametric channel coding (Candidate D) would unify hierarchical source-coding schemes (JPEG progressive encoding, tree-structured vector quantization) under a single p -adic channel-capacity theorem — a permanent contribution to information theory regardless of adoption timeline. p -adic AdS/CFT for condensed matter (Candidate E) is the most speculative candidate, limited to the subset of the physics community that works on holographic duality.

We note for completeness that isogeny-based cryptography (Candidate F) and phylogenetic ultrametrics (Candidate G) are not "candidates" in the forecast sense — they are the *existence proofs* that establish the prior is non-zero.

4 Historical Reference Class and Calibration

The most informative constraint on the central bet comes not from any particular candidate's technical merits but from the historical base rate of paradigm shifts in fundamental mathematical frameworks.

We identify the three closest historical analogues to "a non-standard mathematical framework displaces the default framework as the native modeling substrate":

1. **Non-Euclidean geometry → general relativity (1854–1915).** Riemann's 1854 lecture introduced non-Euclidean geometry as pure mathematics. Einstein's general relativity (1915) made it the native substrate for spacetime — the Euclidean default was displaced as the geometry of physical space. **Timeline:** ~60 years. **Trigger:** A specific empirical problem (Mercury's perihelion precession) that Euclidean geometry could not solve.
1. **Abstract algebra / group theory → quantum mechanics (1830s–1925).** Group theory existed as pure mathematics for nearly a century before it became the native substrate for quantum angular momentum, particle classification, and the Wigner–Eckart theorem. **Timeline:** ~90 years. **Trigger:** Empirical necessity — the hydrogen spectrum demanded a symmetry group ($SO(4)$) that had no Archimedean explanation.
1. **Information theory → biology (1948–1990s).** Shannon's framework was not native to biology, but sequence alignment and phylogenetic inference adopted it. Phylogenetics specifically uses ultrametric distance — a non-Archimedean structure. **Timeline:** ~40 years. **Trigger:** The molecular-clock hypothesis made tree-structured distance the natural modeling choice.

Reference class for the 10-year timeline: Zero successes out of three. All three required at least 40 years from mathematical inception to empirical adoption. The 10-year bet (displacement by 2036) is therefore in a reference class with no historical precedent — it requires compressing the historical timeline by a factor of $4\text{--}6\times$.

The *50-year* bet (displacement by 2076) has a substantially different calibration. Two to three of the historical analogues succeeded within 40–60 years. The existence proofs (phylogenetics, isogeny crypto) add two partial successes to the class. The question is not "will non-Archimedean modeling ever become mainstream?" — it is "when, and in which domain first?"

5 Adversarial Assessment

We examined the central bet from five critical perspectives.

Null-hypothesis defender. The Archimedean-default framework has successfully modeled everything from quantum fields to neural networks for over a century with zero empirical failures. Non-Archimedean modeling is pure mathematics in search of a problem — every proposed application (p -adic strings, p -adic QM, ultrametric neural nets) has existed for 20–40 years without producing a single result that could not be obtained Archimedeanly. The existence

proofs are domain-constrained: phylogenetics uses ultrametric distance because tree-structured data *imposes* it, not because anyone *chose* it. This is the strongest challenge and correctly identifies that no empirical anomaly has been found that *only* non-Archimedean modeling can resolve.

Better-alternative proposer. If the argument is "Archimedean defaults will be displaced by more appropriate mathematics," the better alternatives are category-theoretic frameworks (already present in quantum foundations via the ZX-calculus), information-theoretic frameworks (already dominant in multiple domains), or topological data analysis (which already uses non-Euclidean distances with empirical results). Why bet on the *most* niche alternative (p -adic/ultrametric) when broader, better-funded alternatives with existing results exist? The counterargument is that non-Archimedean geometry offers something these alternatives do not: a *native hierarchical structure*. Category theory is a meta-framework; TDA produces persistence diagrams; ultrametric geometry produces persistence *trees*. The bet is specifically that tree geometry has an unrealized advantage where information is inherently hierarchical.

Methodology skeptic. The analyst is the same researcher who authored all prior ACRP publications on adelic valuation theory. The assessment is structurally a hammer-sees-nail exercise: an adelic valuation theorist asked "will adelic valuation theory become important?" is predisposed to answer affirmatively. The domain selection (QEC and ML ranked first) may reflect the analyst's expertise rather than objective readiness. We address this by registering dated, strength-weighted predictions whose accuracy can be assessed retrospectively, and by explicitly disclosing this bias rather than claiming impartiality.

Scaling pessimist. Bruhat–Tits buildings of rank $n > 1$ grow super-exponentially. Candidate A may be limited to the tree case ($\mathrm{SL}(2, \mathbb{Q}_p)$) unless a clever encoding is discovered. Tree codes may be sufficient, but the impact ceiling for tree-only codes is lower than for codes on higher-rank buildings.

Resource realist. There is no institutional pathway from "elegant theory" to "funded research program" without a triggering empirical result. Physics funding agencies default to Archimedean proposals because peer reviewers are Archimedean-trained. One high-impact simulation result — an ultrametric QEC code that matches surface-code thresholds — would break the funding logjam, but that result itself requires seed funding. The chicken-and-egg problem is the binding constraint for the 10-year timeline.

6 Judgment Sensitivity

The qualitative ranking A (ultrametric QEC) $> B$ (p -adic neural nets) $> C$ (non-Archimedean QM) is robust under pessimistic and overconfidence-corrected perturbations of the probability estimates. Under pessimistic assumptions (all probabilities at lower bounds), the ranking is unchanged. Under halved-priors (systematic overconfidence correction), the ranking is unchanged.

The ranking is conditionally fragile under optimistic assumptions. If Candidate C's probability reaches its upper bound of 0.20, its higher impact rating (9/10 versus 8/10 for Candidate B) causes it to overtake B, producing the ranking $A > C > B$. This sensitivity reflects the analyst's background: non-Archimedean QM has a theoretical elegance that p -adic neural nets lack, and the ranking may overweight this at baseline. We flag this as [CONDITIONAL].

The candidates are largely independent in their mathematical requirements, sharing only the general non-Archimedean framework rather than specific constructions. Candidate A's failure would not materially affect Candidates B or C — they operate in different domains with different mathematical tools.

7 Research Effort Allocation

Based on the qualitative ranking and timeline constraints, we recommend the following allocation of research effort:

Rank	Candidate	Allocation	Rationale
A	Ultrametric QEC codes	35%	Shortest timeline, highest testability, clear success metric (code threshold). A single positive simulation result would be the strongest existence proof for the paradigm.
B	p -adic neural architectures	25%	Larger community (ML researchers outnumber QEC researchers $\sim 100:1$), so even a niche result reaches more practitioners. Downside: requires building tooling from scratch.
C	Non-Archimedean QM	15%	Highest potential impact but longest timeline. Best approached as theoretical reconnaissance — build the formalism, identify a testable signature, wait for experimental opportunity.
D	Ultrametric channel coding	8%	Smaller community, but information-theoretic results have long shelf life (Shannon's 1948 theorems are still cited). A p -adic channel-capacity theorem would be a permanent contribution.
E	p -adic AdS/CFT	7%	Most speculative and hardest to test. Pursue only if the AdS/CFT correspondence remains a major research program.
Hedge	Unknown candidates	10%	Anti-fragility floor — reserve effort for ideas not yet conceived.

These percentages are effort heuristics, not optimal portfolio allocations. The domain is too uncertain for formal optimization.

8 Calibration Register

We register the following dated, falsifiable predictions with explicit strength-weighting and anchoring to the historical reference class.

[CHECK: 2036-12] Primary Bet. By December 2036, at least one published model in a recognized scientific or engineering domain (excluding phylogenetics and isogeny-based cryptography) uses non-Archimedean geometry as its *native* substrate, defined as the model's fundamental state-space, distance measure, or algebraic structure being non-Archimedean at the level of definition. **Strength:** [WEAK] — anchored only to the historical reference class, which shows zero successes at the 10-year timeline. **Status:** [PENDING]. **Risk:** "10 years was always too aggressive — paradigm shifts in fundamental mathematical frameworks take 40–60 years, consistent with every historical analogue."

[CHECK: 2030-07] Ultrametric QEC code. At least one peer-reviewed publication demonstrates a quantum error-correcting code whose construction is native to a Bruhat–Tits tree (with ultrametric codeword distance), and the code's threshold is within 20% of the surface-code threshold under identical noise models. **Strength:** [WEAK]. **Status:** [PENDING].

[CHECK: 2032-12] p -adic neural benchmark. A p -adic neural network architecture achieves state-of-the-art performance on at least one standard hierarchical classification benchmark (ImageNet hierarchy, WordNet taxonomy, or equivalent), outperforming the best Euclidean architecture by at least 3% in top-1 accuracy. **Strength:** [WEAK]. **Status:** [PENDING].

[CHECK: 2036-12] Testable non-Archimedean QM. At least one publication in *Physical Review Letters* or equivalent proposes a non-Archimedean modification to quantum mechanics that is testable in a tabletop experiment, with a clear falsifiability condition. **Strength:** [WEAK]. **Status:** [PENDING].

[CHECK: 2028-12] p -adic ML tooling. At least one major ML framework (PyTorch, JAX, TensorFlow, or equivalent) includes a p -adic tensor library as a first-class primitive. **Strength:** [WEAK]. **Status:** [PENDING].

[CHECK: 2033-07] Ultrametric QEC in simulation. An ultrametric QEC code is demonstrated in simulation with a threshold exceeding 0.5% under circuit-level depolarizing noise — within a factor of two of the surface-code threshold of approximately 1%. **Strength:** [WEAK]. **Status:** [PENDING].

[CHECK: 2031-12] p -adic neural network on major benchmark. p -adic neural network achieves top-1 accuracy on any major benchmark (ImageNet, CIFAR-100, or equivalent) within 5% of state-of-the-art while using a fundamentally non-Euclidean architecture. **Strength:** [WEAK]. **Status:** [PENDING].

[CHECK: 2036-12] Null hypothesis. Zero non-Archimedean-native models are published in any domain outside of the already-established domains (phylogenetics, isogeny crypto, number theory). **Strength:** [STRONG] — the null hypothesis is well-anchored to historical base rates. The reference class of "mathematical frameworks proposed as substitutes for established frameworks" (non-standard analysis, p -adic string theory, topos quantum theory) has produced zero mainstream adoptions. **Status:** [PENDING].

9 Practical Applications

If any of the paradigm-shift candidates succeed, the practical implications extend beyond academic interest.

Quantum computing. An ultrametric QEC code with surface-code-comparable thresholds would expand the design space for fault-tolerant architectures. Logical qubits encoded on Bruhat–Tits trees could admit transversal gates with lower overhead, directly affecting the resource estimates for practical quantum computation.

Machine learning. p -adic neural architectures that naturally respect hierarchical structure could improve performance on classification tasks where the label space has a tree topology — medical diagnosis (ICD codes), legal document classification, and phylogenetic trait prediction are natural candidates.

Information theory. A p -adic channel capacity theorem would unify the existing ad hoc collection of hierarchical coding schemes (JPEG progressive encoding, tree-structured vector quantization) under a single framework, analogous to Shannon's unification of disparate coding methods.

In each case, the operational signature is the replacement of a Euclidean default with a tree-structured alternative: error syndromes tagged by tree coordinates, embeddings respecting ultrametric label hierarchies, and channel models with valuation-theoretic noise.

10 Counterfactual Backcasting

What would the technology landscape look like if historical decisions had favored non-Archimedean methods?

Tier 1 (2005 fork: QEC community adopts p -adic methods). If the early QEC community — following Shor's 1995 code and Kitaev's 1997 toric code — had adopted Bruhat–Tits tree geometry alongside surface codes, tree codes would have 20+ years of active development by 2026. We would expect to see at least three experimental QEC groups publishing ultrametric code results, a tree-code construction in *Physical Review Letters*, and a textbook chapter on geometric methods in QEC including tree codes. This fork was not taken: the QEC community standardized on Euclidean lattice geometry.

Tier 2 (1975 fork: p -adic analysis enters physics curricula). If Vladimirov–Volovich p -adic quantum mechanics (1980s) had been incorporated into standard mathematical-physics training, a generation of physicists would be fluent in non-Archimedean methods. We would expect p -adic QM as a standard elective, at least 50 papers per year on p -adic methods in physics (versus the current ~ 20), and at least one textbook titled *Non-Archimedean Methods in Physics*. This fork was not taken.

Tier 3 (1910 fork: p -adic numbers embraced as fundamental physics). If Hensel's p -adic numbers (1897) had been seen as a physical completion — not just a number-theoretic one — from their inception, the Archimedean/non-Archimedean dichotomy would be standard in physics education. This fork was not taken.

Actionable near-term forks. Tier 1 forks are actionable today. We recommend: (1) a graduate seminar on "Non-Archimedean Methods in Quantum Error Correction" at any institution with a QEC group — a single seminar creates a generation of researchers who know the option exists; (2) an open-source pytreecode Python library implementing Bruhat–Tits tree constructions with a surface-code-compatible API, so researchers can swap `SurfaceCode(lattice)` for `TreeCode(bruhattitstree)` in one line; (3) a head-to-head simulation comparison of tree codes versus surface codes under identical noise models, published regardless of outcome — a negative result is valuable because it eliminates a research path and provides a baseline for improvement.

11 Conclusion

The central bet — that non-Archimedean substrate modeling displaces the Archimedean default in at least one domain by 2036 — occupies a paradoxical position. The qualitative analysis identifies real, actionable candidate domains where non-Archimedean modeling could produce empirical results within a decade. Ultrametric QEC codes have the clearest success criterion, and p -adic neural architectures have the largest potential audience. The existence proofs in phylogenetics and isogeny-based cryptography establish that the paradigm is viable — non-Archimedean geometry is not hypothetical; it is the native substrate in two domains already.

Yet the historical reference class is unequivocal: zero of the three closest analogues achieved mainstream adoption within a decade. The Copernican, relativistic, and quantum revolutions each required empirical anomalies that the incumbent framework could not explain. No such anomaly has been identified for the Archimedean default. The null hypothesis — that non-Archimedean modeling remains a niche mathematical tool through 2036 — is the well-anchored prediction.

The forecast's value is not in predicting a specific outcome but in identifying the narrow path that *would* produce the outcome. If an ultrametric QEC code matches surface-code thresholds in simulation, and that result is published in a high-visibility venue, and a tooling ecosystem emerges, and a second domain follows within five years — then the historical timeline will have been compressed. Until the first of these conditions is met, the Archimedean default remains the incumbent, and the historical base rate remains the most reliable guide.

Declarations

Funding: This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Conflicts of Interest: The author is the sole contributor to the Adelic Core Research Program (ACRP), which produced the ACRP-01 through ACRP-06 papers that inform this assessment. The author has no financial or institutional conflicts of interest.

Ethics Approval: Not applicable — this is a theoretical assessment involving no human participants, animal subjects, or sensitive data.

Consent to Participate: Not applicable.

Author Contributions: R.B.Q.-G. is the sole author and performed all aspects of the research, analysis, and writing.

Data Availability: All source artifacts for this paper are archived at the GitHub repository `QNFO/acrp08-paradigm-forecast` and at `papers.qnfo.org` at the DOI listed in the frontmatter. The structured forecast protocol artifact (`artifacts/structured-forecast-protocol-v2.md`) documents the full analysis protocol from which this paper is synthesized.

Code Availability: Not applicable — this paper contains no original code.

Materials Availability: Not applicable.

Use of Artificial Intelligence: A large language model (DeepSeek) was used as a research assistant for literature synthesis, structured analysis, and prose composition under the supervision and editorial direction of the human author, who is solely responsible for all claims, judgments, and conclusions.

References

- [1] E. Zuckerkandl and L. Pauling, "Molecules as documents of evolutionary history," *Journal of Theoretical Biology*, vol. 8, no. 2, pp. 357–366, 1965.
- [2] A. Ostrowski, "Über einige Lösungen der Funktionalgleichung $\varphi(x) \cdot \varphi(y) = \varphi(xy)$," *Acta Mathematica*, vol. 41, pp. 271–284, 1916.
- [3] V. S. Vladimirov, I. V. Volovich, and E. I. Zelenov, *p-Adic Analysis and Mathematical Physics*. World Scientific, 1994.
- [4] A. Khrennikov, *Non-Archimedean Analysis: Quantum Paradoxes, Dynamical Systems and Biological Models*. Kluwer, 1997.
- [5] P. G. O. Freund and E. Witten, "Adelic string amplitudes," *Physics Letters B*, vol. 199, no. 2, pp. 191–194, 1987.
- [6] A. Yu. Kitaev, "Fault-tolerant quantum computation by anyons," *Annals of Physics*, vol. 303, no. 1, pp. 2–30, 2003.
- [7] R. B. Quni-Gudzin, "Adelic Cross-Domain Correspondences in Number Theory, Physics and Computation," the QNFO archive, v4.0, 2026. <https://papers.qnfo.org/papers>.
- [8] R. B. Quni-Gudzin, "Statistical Audit of the 5-Smooth Semigroup in Mass-Ratio Approximations," the QNFO archive, v1.0, 2026. <https://papers.qnfo.org/papers>.
- [9] R. B. Quni-Gudzin, "Primitive Ultrametric Kernels: Valuation-Theoretic Classification of QEC Constructions," the QNFO archive, 2026.
- [10] R. B. Quni-Gudzin, "Adelic Core Research Program Plan," the QNFO archive, v1.1, 2026.
- [11] R. B. Quni-Gudzin, "Measurement Stratigraphy: A Historical Taxonomy of Mathematical Eras in Physics," the QNFO archive, v2.0, 2026.