What Would Prove It?

Definitional Indeterminacy, Methodological Undecidability, and a Decision Framework for the Indus Sign System

Authors to be determined

Corresponding author: Boris Kriger
Institute of Integrative and Interdisciplinary Research
Department of Computational Humanities and Cultural Systems
boriskriger@interdisciplinary-institute.org

Abstract

The Indus sign system, attested on approximately 4,000 inscribed objects from the Harappan civilization (c. 2600–1900 BCE), remains undeciphered after a century of scholarship. This paper argues that the impasse reflects two compounding failures. The first is definitional: the term "script" encompasses both linguistic and non-linguistic systems, but the debate has failed to formalize the distinction that actually matters—whether the system encodes natural language. The second failure is methodological: the question is undecidable within the framework of statistical analysis alone. The exchange between Rao et al. [2009] and Sproat [2010] demonstrated that entropic methods can neither confirm nor exclude the linguistic hypothesis—a symmetry that signals the framework's insufficiency. This paper addresses both failures. We formalize three precisely defined hypotheses—linguistic writing, administrative registration code, and metrological hybrid—and construct a Decision Framework comprising five testable predictions supplemented by an algebraic substitution test applied to individual inscriptions. The existing statistical evidence—including sequence uniqueness rates of 98.3%, edge-peaked positional entropy, numeral clustering, and near-duplicate analysis—is evaluated against all three hypotheses. The framework is generalizable to other undeciphered or contested sign systems.

Keywords: Indus script, undecidability, definition-dependent provability, decision framework, Harappan civilization, administrative codes, undeciphered writing systems, computational linguistics

1 Introduction

For nearly a century, scholars have debated whether the signs inscribed on Harappan seals, tablets, and other objects constitute a writing system encoding spoken language. The debate has produced a substantial literature but no consensus. Proponents of the linguistic hypothesis point to the system's combinatorial structure, its conditional entropy profile, and parallels with known scripts [Parpola, 1994, Rao et al., 2009, Yadav et al., 2010]. Proponents of the non-linguistic hypothesis point to the extreme brevity of inscriptions, the absence of longer texts, and the failure of all phonetic decipherment attempts [Farmer et al., 2004, Sproat, 2010]. Bonta [2010] reached the non-linguistic conclusion independently through a different method—systematic substitution testing on individual texts. As Sproat [2023] has argued, nonlinguistic systems can possess rich, complex structure; the boundary between categories is less categorical than typically assumed. A third position—that the signs encode structured administrative information without representing spoken language—has been advanced by Mukhopadhyay [2019], Mukhopadhyay [2023a], Pollard [2025], and Kriger and Hunt [2026].

The standard diagnosis of the impasse is that we lack sufficient evidence. This paper proposes a different diagnosis: the impasse reflects two compounding epistemological failures that have not been distinguished from each other.

The first failure is **definitional**. The term "script" is broad enough to encompass both linguistic and non-linguistic systems—a programming language such as JavaScript is called a "script" but does not encode natural language. The question that actually divides the field is not whether the Indus system is a "script" but whether it encodes natural language. This narrower question has never been formalized with sufficient precision.

The second failure is **methodological**. Statistical methods have been deployed to argue both for and against the linguistic hypothesis. The exchange between Rao et al. [2010] and Sproat [2010] exposed the limits of entropy-based approaches from both sides: entropic similarity to linguistic systems does not prove linguistic status, and entropic dissimilarity from random sequences does not exclude non-linguistic structure. The symmetry is exact, and it signals that the methodological system of statistical analysis alone lacks the resources to resolve the question.

This paper addresses both failures. Section 2 analyzes the definitional indeterminacy. Section 3 replaces the undefined binary with three precisely specified hypotheses. Section 4 evaluates existing statistical evidence. Section 5 introduces an algebraic substitution test. Section 6 constructs a Decision Framework. Sections 7 and 8 address specific objections. Section 9 formalizes the argument. Section 10 discusses generalizability.

2 Definitional Indeterminacy

2.1 The Real Divide: Linguistic vs. Non-Linguistic

The term "script" is not the source of the disagreement. A script, broadly understood, is any conventional system of visual signs with internal structure. Under this definition, the Indus system trivially qualifies—as do mathematical notation, heraldic systems, and programming languages. The question that actually divides the field is narrower: does the system encode natural language?

This is the distinction that matters, and it has never been formalized. Scholars who assert that the Indus system is "a script" and scholars who deny it often talk past each other because they are answering different questions. The former operate with a broad semiotic definition; the latter operate with a phonological or morphological definition requiring the encoding of spoken language. Both are correct relative to their definitions. The debate is underdefined, not undecided.

2.2 Formalizing the Divide

We propose that the productive question is not "Is this a script?" but "Does this system encode natural language?" This question admits three structurally distinct answers, corresponding to three hypotheses developed in Section 3. The binary—"linguistic" versus "non-linguistic"—is itself too coarse, because it obscures the possibility of a hybrid system with limited linguistic content (Section 3.3).

2.3 Dissolving the Binary

As Sproat [2023] has argued from a comparative perspective, sign systems occupy a spectrum from fully linguistic to fully non-linguistic, and the boundaries between categories are not always sharp. Damodaram Pillai [2023] has demonstrated that even Brāhmī—a system universally accepted as linguistic writing—originated as a hybrid of multiple source scripts (Aramaic, Phoenician, and Greek), suggesting that hybridity is not exotic but historically commonplace in South Asian sign systems. The productive question is not "Is this writing?" but "What is the structure of this system, and what function did it serve?"

3 Three Hypotheses

To escape the definitional impasse, we replace the binary question with three precisely defined structural hypotheses.

3.1 H_1: Linguistic Writing

The signs encode natural language—logographically, syllabically, or through some combination—permitting the free expression of spoken utterances. Structural predictions: (a) inscription length correlates with informational complexity of context; (b) morphological or syntactic regularities reflect the grammar of the encoded language; (c) signs combine in word-level or phrase-level patterns; (d) weak or no correlation between animal motif and sign sequences; (e) individual signs repeat within inscriptions at rates comparable to known scripts; (f) substituting putative linguistic values for signs yields coherent utterances across multiple texts.

3.2 H_2: Administrative Registration Code

The signs encode structured identifying information—personal or institutional identity, clan affiliation, administrative status—without representing spoken language. Structural predictions: (a) the vast majority of sign sequences are unique; (b) strong positional structure with fixed elements at edges and variable elements in the middle; (c) numeral signs cluster at specific positions; (d) numeral distribution deviates from Benford's law; (e) near-duplicate pairs exist, differing in the variable middle zone; (f) inscription length is independent of contextual complexity; (g) animal motif correlates strongly with sign prefixes; (h) individual signs rarely repeat within an inscription.

3.3 H_3: Metrological Hybrid

Some signs function as logographs for commodities or measures while the overall sequence structure serves an administrative rather than narrative function. Under this hypothesis, an inscription might encode something like "3 wheat measures" or "4 rice bags"—structured information with limited lexical content, but not a system capable of freely expressing spoken language. This is a critical distinction: H_3 does not represent "proper linguistic writing" in the sense of encoding arbitrary utterances; it represents a metrological notation with constrained lexical elements.

Roman imperial inscriptions provide a partial precedent: the Roman alphabet is used formulaically (*IMP CAES DIVI F AVGVSTVS*), though with far greater linguistic flexibility than H_3 would predict for the Indus system. A closer parallel may be the hybrid origin of Brāhmī script itself, which Damodaram Pillai [2023] has shown to be a composite of multiple source traditions—demonstrating that sign systems in South Asia have a documented history of structural hybridity.

Structural predictions: (a) signs at fixed positions show semantic clustering with the animal motif; signs in the variable middle do not; (b) some signs function as commodity or unit markers with consistent positional behaviour; (c) conditional entropy is intermediate between linguistic and random baselines—which is what Rao et al. [2009] demonstrated; (d) substitution testing yields metrological readings (quantities, commodities) in some positions but incoherent results in others.

4 Existing Statistical Evidence and Its Limits

4.1 Sequence Uniqueness

In a corpus of 179 Mohenjo-Daro seals, 98.3% of sign sequences are unique [Kriger and Hunt, 2026]. Strongly predicted by H_2; consistent with H_3; compatible with but not naturally explained by H_1.

4.2 Positional Entropy Profile

The entropy profile shows low entropy at positions 1 and 5, high entropy at positions 2–4. Strongly predicted by H_2 (fixed categorical markers at edges, variable code in the middle). H_1 would predict a flatter profile.

4.3 Numeral Distribution

Stroke-numeral signs constitute 15–17% of all signs and cluster at the penultimate position. The digit "2" accounts for 58–67% of all numeral tokens, drastically violating Benford's law [Kriger and Hunt, 2026]. Strongly predicted by H_2; inconsistent with H_1 if the numerals represent quantities. Under H_3, the numerals may represent metrological quantities, but their positional rigidity and distributional skew remain anomalous for a freely expressive system.

4.4 Near-Duplicate Pairs

Eight pairs of seals differ by exactly one sign, with the difference never at the first or last position [Kriger and Hunt, 2026]. Predicted by H_2; under H_1, the restriction to the middle zone would require explanation.

4.5 Intra-Inscription Non-Repetition

Individual signs very rarely repeat within a single inscription—a robust finding first noted by Farmer et al. [2004] and consistently emphasized in subsequent work [Sproat, 2010, 2023]. Under H_2, this follows from combinatorial selection without replacement. Under H_1, it is anomalous: most writing systems exhibit substantial within-text sign repetition.

4.6 Seal Wear

Gokhale and Ameri [2026] present agent-based modelling evidence that many seals show minimal wear, inconsistent with repeated stamping. Compatible with carried identification objects. However, the physical function of the object is logically independent of the sign system's nature: Mesopotamian seals included true writing and were used as seals.

4.7 Summary

The evidence is consistently compatible with H_2 and partially with H_3, while creating multiple anomalies for H_1. But no finding definitively excludes H_1, and no finding uniquely supports H_2 over H_3. The methodological system of statistical analysis alone cannot resolve the question.

5 The Algebraic Substitution Test

Statistical evidence operates at the corpus level. A complementary approach operates at the level of individual inscriptions: assign a putative linguistic value to a frequently occurring sign and test whether the resulting readings are coherent across multiple texts.

Table 1: Evidential status of existing findings.

Finding	H_1 Ling.	H_2 Code	H_3 Hybrid
98.3% unique sequences	Compatible	Str. predicted	Compatible
Edge-peaked entropy	Weakly comp.	Str. predicted	Compatible
Numeral clustering	Anomalous	Str. predicted	Compatible
Benford violation	Anomalous	Str. predicted	Constraining
Near-duplicates in middle	Req. expl.	Predicted	Compatible
Intra-text non-repetition	Anomalous	Predicted	Constraining
Minimal seal wear	Neutral	Compatible	Neutral

Bonta [2010] applied this method systematically and concluded that the Indus signs do not encode prose language. Damodaram Pillai (personal communication, April 2026) reports independently reaching the same conclusion through substitution testing: assigning the value *fish/mīn* to the fish sign (following Parpola's well-known proposal) and testing the resulting readings across multiple inscriptions rapidly produces incoherent results incompatible with natural language.

This test is particularly powerful because it directly engages with the strongest version of the linguistic hypothesis. If no consistent linguistic reading can be sustained across even a small set of inscriptions, the burden shifts to proponents of H_1 to explain why. The failure is not merely negative evidence; it is a constructive result: it tells us what kind of system would produce such failures (one in which signs do not map to words or morphemes in a spoken language).

We propose formalizing this as **Test 6**: for each of the 10 most frequent signs, assign the most plausible linguistic value (drawing on Parpola's, Mahadevan's, and other proposed readings) and test whether the resulting translations are semantically coherent across a random sample of 20 inscriptions. Under H_1, at least some assignments should yield coherent readings. Under H_2, none will. Under H_3, assignments may yield coherent readings for metrological terms (commodity names, units) but not for other positions.

6 The Decision Framework: Five Tests

If the question is undecidable within the current methodological system, and if definitional indeterminacy has been resolved by specifying three precise hypotheses, the remaining task is to extend the system. The Decision Framework specifies five tests whose outcomes differentially favour one hypothesis. Two can be performed with existing digital corpora; three require archaeological collaboration. Test 6 (Section 5) supplements the framework with a text-level diagnostic.

6.1 Test 1: Inscription–Context Correlation

Question: Does inscription length correlate with the informational complexity of the archaeological context?

Predictions: H_1: positive correlation. H_2: no correlation. H_3: weak correlation (longer inscriptions may list more commodities).

Status: Requires compilation of provenance data from excavation reports. Archaeological collaboration is indispensable.

6.2 Test 2: Animal Motif–Sign Prefix Correlation

Question: Do specific animal motifs (unicorn, zebu, elephant, rhinoceros) correlate with specific sign prefixes?

Predictions: H_2: strong correlation (both encode the same categorical variable). H_1: weak or no correlation. H_3: moderate, position-dependent.
Preliminary observation: A cursory examination of the available corpus (Damodaram Pillai, personal communication, April 2026) does not reveal an obvious animal–prefix correlation, though the unicorn motif appears modal—possibly reflecting its higher frequency in the corpus rather than a systematic prefix association. A full statistical analysis using mutual information measures across the complete ICIT database [Wells and Fuls, 2010–2023] is required before drawing conclusions. A negative result on this test would weaken pure H_2 and strengthen H_1 or H_3.

6.3 Test 3: Cross-Site Format Consistency

Question: Does the positional structure remain consistent across geographically distant sites?
Predictions: H_2: high consistency (uniform format). H_3: high consistency in formulaic elements, variation in commodity terms. H_1: the prediction is ambiguous—a dialectally diverse language would show regional variation, but an elite standardized language (comparable to Classical Sanskrit or literary Tamil) might show equally high consistency. This ambiguity limits the discriminating power of this test for H_1 and should be noted explicitly.
Status: Feasible with existing data. Yadav et al. [2010] noted cross-corpus bigram consistency but did not test positional entropy by site.

6.4 Test 4: Commodity–Inscription Independence

Question: When inscribed objects co-occur with specific commodities, do different inscriptions appear with the same commodity?
Predictions: H_2: independent (different owners, same commodity). H_1: correlated (inscriptions describe contents). H_3: partially correlated (some sign positions name commodities, others identify owners).
Status: Requires systematic compilation of provenance data. Mukhopadhyay [2023b] provides unpublished analyses of commodity-related sign distributions that may bear on this test. This is where archaeological collaboration is most critical.

6.5 Test 5: Prefix Distribution Across Sites

Question: Is there evidence for coordinated prefix allocation across sites?
Predictions: H_2: site-specific or coordinated allocation. H_1: distribution reflecting linguistic categories. H_3: dual structure.
Implications: The Harappan civilization shows remarkable standardization—weights, measures, brick ratios, urban planning—across vast distances, consistent with distributed administrative protocols.

6.6 The Decision Matrix

No two hypotheses generate identical predictions across all six tests. A clear result on even three would substantially narrow the hypothesis space.

7 The Roman Inscription Objection

A natural objection: Roman imperial inscriptions are formulaic and code-like (*IMP CAES DIVI F AVGVSTVS COS XII...*), yet unambiguously linguistic. What prevents the Indus system from being identical?

Table 2: Predictions of each hypothesis for the six tests.

Test	H_1 Prediction	H_2 Prediction	H_3 Prediction
1. Length–context	Pos. correlation	No correlation	Weak corr.
2. Animal–prefix	Weak/none	Strong corr.	Moderate
3. Cross-site format	Ambiguous*	High consistency	High/mixed
4. Commodity indep.	Correlated	Independent	Partial corr.
5. Prefix distribution	Ling. distributed	Admin. structured	Dual structure
6. Algebraic subst.	Coherent readings	No coherent readings	Partial (metrol.)

*An elite standardized language may also show high consistency; see Section 5.3.

Roman inscriptions differ from the Indus corpus in measurable ways. Individual letters repeat freely within inscriptions, reflecting alphabetic combinatorics. Abbreviations derive transparently from known words. The corpus contains inscriptions of widely varying length. Syntactic connectives reveal grammatical structure. The intra-text non-repetition finding is particularly significant: Roman inscriptions show abundant letter repetition because alphabetic writing requires it. The Indus system does not. If it includes linguistic elements, those elements must operate under combinatorial rules unlike any known writing system. This constrains H_3 without excluding it.

8 The Linear B Comparison and Archaeological Evidence

Linear B provides an instructive contrast. Its inscriptions include pictographic representations of recorded commodities, providing independent evidence that the texts describe inventory transactions. This correlation was crucial to Ventris's decipherment. The Indus corpus lacks this transparent correlation. No credible claim has established a systematic link between animal motifs and sign sequences—a point reinforced by preliminary examination of the corpus (Section 5.2). This is why Test 2 is critical: a statistically robust animal–prefix correlation would provide the kind of independent evidence that has been missing.

The absence of longer texts is often cited against the linguistic hypothesis. But the Harappans may have written on perishable materials. What survives is the subset carved into durable materials—precisely the category most likely to serve identity-marking or administrative functions. The survival bias of the archaeological record is itself evidentially relevant.

A note on the corpus: some objects cited in the secondary literature as evidence for longer texts require critical evaluation. Provenance and authenticity must be established independently for each artefact before it can serve as evidence in this debate.

9 The Structure of the Argument

The argument can be stated concisely. Let Q be: "Which functional architecture—H_1, H_2, or H_3—does the Indus sign system instantiate?" Let S be the methodological system of statistical analysis of sign sequences. Let S' be the extended system incorporating the six tests of Sections 5–6.

Claim 1 (Definitional Resolution): Q is well-defined. It replaces the ambiguous question "Is this a script?" with a question about functional architecture that admits three structurally distinct answers.

Claim 2 (Undecidability in S): Q is undecidable in S. The existing statistical evidence is consistent with multiple hypotheses. The Rao–Sproat exchange demonstrated this symmetry.

Claim 3 (Signal, Not Endpoint): The undecidability of Q in S signals S's explanatory inadequacy, not Q's inherent unanswerability.

Claim 4 (Adequate Extension): S' extends S by principled methods—archaeologically grounded predictions and text-level substitution testing—without introducing the desired answer as a bare assumption.
Claim 5 (Decidability in S'): If the six tests yield clear results, the hypothesis space narrows to at most one candidate.

10 Generalizability

The dual framework—resolving definitional indeterminacy, then extending the methodological system—applies to any contested sign system: Proto-Elamite, Vinča symbols, Rongorongo. In each case, progress requires specifying what precisely is being asked, then specifying what evidence would answer it. The algebraic substitution test (Section 5) is particularly portable: it can be applied to any system for which putative linguistic readings have been proposed.

The broader lesson is that "Is this writing?" is not a question with a fixed answer awaiting discovery. It is a question whose *answerability* depends on both the definitional precision with which it is posed and the methodological system in which it is evaluated. Constructing the right definitions and the right system is not a preliminary to the investigation. It *is* the investigation.

11 Conclusion

The century-long debate over the Indus sign system has been impeded by two compounding failures. The first is definitional: the question was framed around the ambiguous term "script" when the real divide—linguistic versus non-linguistic—was never formalized. The second is methodological: the statistical tools deployed lack the resources to answer even the correctly formulated question. This paper has addressed both: three precisely defined hypotheses replace the undefined binary, and six testable predictions extend the methodological system from corpus-level statistics to text-level substitution testing and archaeological evidence.

Two tests can be performed immediately with existing corpora. One (the algebraic substitution test) can be performed on individual inscriptions. Three require archaeological collaboration. Together, they form a decision tree in which no two hypotheses generate identical predictions.

The existing evidence is consistently explained by H_2, the administrative-code hypothesis. But consistency is not proof. The path forward requires specifying what would constitute proof—which is what this paper has done.

Acknowledgments

The author is deeply grateful to Richard Sproat for extensive and detailed comments that directly shaped every section of this paper. His published critique of entropy-based methods [Sproat, 2010], his analysis of nonlinguistic symbol systems in *Symbols* [Sproat, 2023], and his unpublished observations on the limits of statistical inference, the significance of archaeological context, the Roman inscription parallel, the intra-text non-repetition constraint, and the seal-function question (personal communication, March 2026) collectively defined the methodological challenges that the Decision Framework was constructed to address.

The author also thanks Karan Damodaram Pillai for detailed comments on the manuscript, including the observation that the linguistic/non-linguistic divide—not the term "script"—is the true axis of disagreement; the refinement of H_3 as a metrological rather than freely expressive system; the caveat regarding elite standardized languages in Test 3; the recommendation of the algebraic substitution method; and the identification of a fraudulent artefact in the secondary literature. His published work on the hybrid origin of Brāhmī

[Damodaram Pillai, 2023] provided the historical precedent for sign-system hybridity in South Asia.

References

Bonta, Steven. The Indus Valley script: A new interpretation. Academia.edu, 2010. https://www.academia.edu/8691385/

Damodaram Pillai, Karan. The hybrid origin of Brāhmī script from Aramaic, Phoenician and Greek letters. *Indialogs*, 10:93–122, 2023. https://doi.org/10.5565/rev/indialogs.213

Farmer, Steve, Richard Sproat, and Michael Witzel. The collapse of the Indus-script thesis: The myth of a literate Harappan civilization. *Electronic Journal of Vedic Studies*, 11(2):19–57, 2004.

Gokhale, Pallavee and Marta Ameri. Modelling the possible archaeological past(s): Agent-based modelling of Harappan seal use and survival. *Computer Applications and Quantitative Methods in Archaeology Proceedings*, 51(1), Article 3, 2026.

Kriger, Boris and Treasure A. Hunt. Positional constraints, sequence uniqueness, and stroke numerals in Indus seal inscriptions from Mohenjo-Daro: A statistical analysis. IIIR Computational Humanities and Cultural Systems, 2026. https://doi.org/10.5281/zenodo.19103880

Mahadevan, Iravatham. *The Indus Script: Texts, Concordance and Tables*. Memoirs of the Archaeological Survey of India, 1977.

Mukhopadhyay, Bahata Ansumali. Interrogating Indus inscriptions to unravel their mechanisms of meaning conveyance. *Humanities and Social Sciences Communications*, 6:49, 2019.

Mukhopadhyay, Bahata Ansumali. Semantic scope of Indus inscriptions comprising taxation, trade and craft licensing, commodity control and access control. *Humanities and Social Sciences Communications*, 10:912, 2023.

Mukhopadhyay, Bahata Ansumali. Unpublished working papers on Indus inscriptions. SSRN, 2019–2023. https://papers.ssrn.com/sol3/cf_dev/AbsByAuth.cfm?per_id=3011592

Parpola, Asko. *Deciphering the Indus Script*. Cambridge University Press, 1994.

Pollard, Matthew F. The Indus Valley inscriptions as merchant marks: A non-linguistic interpretation. Cambridge Open Engage, 2025.

Rao, Rajesh P. N., Nisha Yadav, Mayank N. Vahia, Hrishikesh Joglekar, Ronojoy Adhikari, and Iravatham Mahadevan. Entropic evidence for linguistic structure in the Indus script. *Science*, 324(5931):1165, 2009.

Rao, Rajesh P. N., Nisha Yadav, Mayank N. Vahia, Hrishikesh Joglekar, Ronojoy Adhikari, and Iravatham Mahadevan. Entropy, the Indus script, and language: A reply to R. Sproat. *Computational Linguistics*, 36(4):795–805, 2010.

Sproat, Richard. Ancient symbols, computational linguistics, and the reviewing practices of the general science journals. *Computational Linguistics*, 36(3):585–594, 2010.

Sproat, Richard. *Symbols: An Evolutionary History from the Stone Age to the Future*. SpringerNature, 2023.

Wells, Bryan K. *The Archaeology and Epigraphy of Indus Writing*. Archaeopress, 2015.

Wells, Bryan K. and Andreas Fuls. Indus Corpus of Inscribed and Iconographic Texts (ICIT). Database, 2010–2023.

Yadav, Nisha, Hrishikesh Joglekar, Rajesh P. N. Rao, Mayank N. Vahia, Iravatham Mahadevan, and Ronojoy Adhikari. Statistical analysis of the Indus script using n-grams. *PLoS ONE*, 5(3):e9506, 2010.