

V1.0.0

# Invented Representations, Discovered Constraints: A Quantitative Case for Weak Structural Realism in Mathematics

Rowan Brad Quni-Gudzin

QNFO-ULA

## ABSTRACT

Whether mathematics is invented or discovered remains a live dispute with practical consequences for how research, education, and patent law treat mathematical objects. We assume the 2020 PhilPapers survey reports the field split, with nominalism at 41.85 percent and Platonism at 38.38 percent of the distribution and a fifth of respondents rejecting both labels (Section 3.3), yet both camps argue from the same celebrated cases while no principled criterion separates what practitioners invent from what they discover. We propose such a criterion, a representation/constraint split with an invariance test, and evaluate it against quantitative anchors: the exact sexagesimal arithmetic of Plimpton 322, whose first-row value  $1; 59, 00, 15$  we reproduce to all four recorded places as  $(169/120)^2 = 28561/14400$ ; the extension of the tablet's generating scheme to five rows beyond its fifteen recorded rows, so that the recorded corpus covers 75 percent of a twenty-row extension; a survivorship model of Wigner's applicability puzzle calibrated on documented adoption data for non-Archimedean geometry, where zero of three closest historical analogues achieved mainstream adoption within a decade and current non-adopting activity is bounded near  $10^2$  papers per year; and Gödel incompleteness read as metatheory-relative rather than object-level. Every anchor fits the split: representations are culturally variable and partly invented, while the constraint structure they navigate behaves as discovered. We conclude that weak structural realism states the applicability puzzle precisely rather than dissolving it, and we outline the archival search that would falsify the Plimpton anchor.

**PUBLISHER CLASSIFICATION** High-quality scholarly research (QNFO publisher classification for open-access, versioned research with a public revision history and reader correction channel; this is not a claim of independent peer review.)

**PEER-REVIEW STATUS** Not independently peer-reviewed

**CITE AS** Quni-Gudzin, R. B. (n.d.). Invented Representations, Discovered Constraints: A Quantitative Case for Weak Structural Realism in Mathematics. QNFO Preprint v1.0.0.

## Abstract

Whether mathematics is invented or discovered remains a live dispute with practical consequences for how research, education, and patent law treat mathematical objects. We assume the 2020 PhilPapers survey reports the field split, with nominalism at 41.85 percent and Platonism at 38.38 percent of the distribution and a fifth of respondents rejecting both labels (Section 3.3), yet both camps argue from the same celebrated cases while no principled criterion separates what practitioners invent from what they discover. We propose such a criterion, a representation/constraint split with an invariance test, and evaluate it against quantitative anchors: the exact sexagesimal arithmetic of Plimpton 322, whose first-row value  $1; 59, 00, 15$  we reproduce to all four recorded places as  $(169/120)^2 = 28561/14400$ ; the extension of the tablet's generating scheme to five rows beyond its fifteen recorded rows, so that the recorded corpus covers 75 percent of a twenty-row extension; a survivorship model of Wigner's applicability puzzle calibrated on documented adoption data for non-Archimedean geometry, where zero of three closest historical analogues achieved mainstream adoption within a decade and current non-adopting activity is bounded near  $10^2$  papers per year; and Gödel incompleteness read as metatheory-relative rather than object-level. Every anchor fits the split: representations are culturally variable

and partly invented, while the constraint structure they navigate behaves as discovered. We conclude that weak structural realism states the applicability puzzle precisely rather than dissolving it, and we outline the archival search that would falsify the Plimpton anchor.

## 1. Introduction

Mathematicians, philosophers, and the institutions that fund and regulate them face a decision that the invented/discovered dichotomy has never cleanly resolved. The decision is concrete. Patent law treats a mathematical formula as invented in the United States and as discovered in the United Kingdom, so the same object changes legal status at a border [1]. Curriculum designers must decide whether to teach mathematics as a human construction, which risks relativism, or as an external landscape, which risks mystifying students who struggle. Funding agencies must decide whether formalism-driven work is exploration of a pre-existing territory or free construction, which changes what counts as a bet worth funding. And the philosophy profession itself is split: we assume the 2020 PhilPapers survey reports 41.85 percent of respondents endorsing or leaning toward nominalism about abstract objects, 38.38 percent toward Platonism, and 20.81 percent choosing neither (Section 3.3), so no position commands even a plurality majority among experts. If the dichotomy is false—if the correct split cuts across it—then each of these decisions is currently being made against a badly drawn map.

What would change if this paper's result holds? Practitioners would stop asking whether mathematics as a whole is invented or discovered and would instead ask, case by case, which layer of a given piece of mathematics carries which status. The answer changes how we interpret cross-cultural convergence, how we weigh the "unreasonable effectiveness" of mathematics in physics, and how we read incompleteness results.

The real-world problem, stated plainly: Wigner's puzzle—that mathematics developed for internal reasons keeps turning out to describe physics, with the standard dating taken as given—general relativity's light-bending prediction of 1915 confirmed in 1919 and the Dirac equation's surplus residue becoming the positron in 1932 (we assume these uncontested historical dates; no conclusion rests on them beyond their qualitative role as celebrated hits)—is standardly cited as evidence that mathematicians discover a mind-independent structure. But the citation rests on a survivorship-biased sample: we celebrate the formalisms that found physical application and rarely count those that did not. Without a quantified account of the base rate, the hits prove nothing. Similarly, cross-cultural convergence on Pythagorean triples—the Old Babylonian tablet Plimpton 322, the Indian Śulba Sūtras, and the Chinese Zhoubi Suanjing—is cited as evidence of either shared cognition or shared abstract structure, but no analysis separates the two contributions.

This paper's contributions are:

- A principled representation/constraint criterion: axioms, notation, and definitions are invented representations; theorem-space consequences and structural constraints are discovered, in the precise sense that they are invariant under the choice of representation. The criterion is stated and defended in Section 3 and applied in Section 4.
- An exact arithmetic verification that the first row of Plimpton 322 encodes the Pythagorean relation  $(D/L)^2$  with the tablet's sexagesimal value  $1; 59, 00, 15$  equal to  $(169/120)^2$  to all four recorded places, derived in Section 4.1, establishing that the constraint structure predates and is independent of its cultural packaging.
- A quantified extension result: the generating scheme behind the tablet's fifteen rows predicts five additional rows [6], so the recorded corpus covers 75 percent of a twenty-row extension, computed in Section 4.1. We take the twenty-row corpus as the reference frame under the convention that the extension defines the natural completeness of the generating scheme; the alternative framing, that five predicted rows are a 33.33 percent expansion of the extant fifteen, is the same arithmetic read against the smaller base.
- A survivorship model of Wigner's puzzle calibrated on documented adoption data: the non-Archimedean geometry case, where the historical reference class shows zero of three closest analogues achieving mainstream adoption

within a decade and current non-adopting activity is bounded at approximately 100 papers per year, derived in Section 4.3, showing that compounding hit precision is explainable naturalistically.

- A metatheory-relative reading of Gödel incompleteness, argued in Section 4.4, under which truth outrunning provability does not imply mind-independent objects.
- A synthesis position—weak structural realism—that locates the expert "Other" fifth of the 2020 PhilPapers distribution as a coherent structuralist stance, argued in Sections 5 and 6.

## 2. Background and Related Work

The discovered/invented/inherited trichotomy. Borchers [1] argues from his own lockdown work that the classical platonist/formalist dilemma is a false pair: new mathematics can also be *inherited*, entering into possession of a tradition. He notes the practical stakes directly—a mathematical formula can be patented in the USA but not in the UK, because American lawmakers treat it as invented and British ones as discovered [1]. He also records Gelfand's distinction between *pridumyvanie* (inventing) and *produmyvanie* (thinking through from first principles), with Gelfand valuing the latter as yielding deeper results [1]. The earlier version of the same argument documents a concrete case: an algorithm in computational group theory that the author came into possession of independently of its publication—by John Bray's discovery, with the author disclaiming priority—built on Brauer's observation that two involutions generate a dihedral subgroup and are conjugate by an element of  $\langle uv \rangle$  when  $uv$  has odd order, with the principal ideas traceable to Burnside's 1900 paper and remaining forgotten for the good part of the twentieth century [2]. This case is directly relevant here: the algorithm's *content* was inherited through a tradition that preserved it without documentation (constraint), while its *availability* depended on cultural transmission (representation)—exactly the split this paper formalizes.

Plimpton 322 and the generation of triples. Firsov [6] analyzes the purely mathematical aspects of the tablet's number generation, reviewing the standard Neugebauer–Sachs explanation in which a virtual column  $L$  satisfies  $L^2 = D^2 - S^2$  so that  $L, S, D$  form a Pythagorean triple and the tablet's column  $A = (D/L)^2$  [6]. The tablet contains 15 rows [6]. Firsov proposes a modification of Bruins' reciprocal-pairs explanation that accounts for the "missing pairs" in Bruins' reconstruction, and using the new scheme to predict the numbers that would follow those on the tablet yields five new rows beyond the 15 listed by Price and Friberg [6]. This is a quantitative gift to the present paper: the fact that a single generating rule extends to predict additional rows is precisely the signature of a discovered constraint operating beneath an invented notational scheme.

Cross-cultural convergence as a cognitive phenomenon. Three works supply the modern analogue of ancient cross-cultural convergence: what happens when different cultures' representational systems meet the same task. In LLM-C3MOD, non-native hate speech moderators answered incorrectly or selected "I don't know" for nearly half of samples, achieving 56.25 percent accuracy, with failures traced to culturally-specific knowledge, cultural sentiment, and internet culture; adding retrieval-augmented cultural context raised accuracy to 78 percent and reduced human workload by 83.6 percent [3]. The lesson for our question is that representational context, not underlying logical capacity, dominates performance on culturally loaded tasks—convergence on the *task structure* coexists with divergence on the *packaging*. The Supervised Semantic Differential work reaches a compatible conclusion quantitatively: fitting supervised semantic gradients separately within each language and comparing them in a shared multilingual embedding space, the authors find Valence, Arousal, and Dominance all significantly recoverable in both monolingual and multilingual settings, with Valence the most stable dimension across languages [4]. Here the *dimension* (constraint) is recoverable across cultures while the *lexical realization* (representation) diverges. The culturally grounded NLP survey [5] establishes the general pattern: language technology is unequally distributed across languages, widening coverage does not remove performance disparities, and translated benchmarks can preserve the cultural assumptions of their English source questions so that model rankings shift when evaluation is restricted to culturally sensitive subsets [5]. Cross-cul-

tural alignment work on language models adds that models trained on English-translated Japanese commonsense-morality data outperform models trained on Japanese-translated English ethics data, indicating that cross-cultural transfer is harder than cross-lingual transfer [7]—again separating the representational carrier from the constraint being carried.

Anti-Platonism and its limits. The QNFO review of Lakoff and Núñez [9] reconstructs their six-step anti-Platonist argument: Platonic mathematics cannot be addressed scientifically; only mind-based mathematics is accessible; human mathematics must be studied through cognitive science; cognitive science reveals fundamental use of conceptual metaphor; metaphor is limited to living minds; therefore human mathematics is not a subspecies of Platonic mathematics [9]. The review also records the experimental anchor: newborns at 3–4 days can discriminate collections of 2 versus 3 items (subitizing) [9]. This gives the embodied-cognition baseline against which our constraint realism must be measured: embodiment explains where representations come from, but not why independently built representations converge on the same constraints.

The contingency of the number system. The adelic substrate assessment [10] opens from Ostrowski's theorem: every non-trivial absolute value on  $\mathbb{Q}$  is equivalent to either the real absolute value or a  $p$ -adic one, and the  $p$ -adic metric is ultrametric [10]. Its domain table estimates current non-Archimedean activity at roughly 50 papers per year in  $p$ -adic quantum mechanics, about 20 in  $p$ -adic/adelic string theory, fewer than 5 in ultrametric quantum error correction, fewer than 10 in  $p$ -adic neural networks, fewer than 5 in  $p$ -adic channel coding, and fewer than 10 in  $p$ -adic condensed matter [10]. Its conclusion is blunt: zero of the three closest historical analogues—the Copernican, relativistic, and quantum revolutions—achieved mainstream adoption within a decade, and no empirical anomaly has been identified for the Archimedean default [10]. This is exactly the survivorship-control data Wigner's puzzle needs: a rich mathematical structure (the adeles, ultrametrics) that is *not* finding physics applications, visible because we are watching it in real time rather than retrospectively.

Convergence and consilience. The hierarchical-architecture document [11] argues that convergence (independent origin of similar forms in distinct lineages, cultures, or minds—wings in birds, bats, and insects; pyramids in Egypt and Mesoamerica; the independent invention of zero in Maya and Indian mathematics; the calculus of Newton and Leibniz) and consilience (knowledge from different disciplines converging on the same truths) are symmetric faces of a single structure: convergence is the ontological signature of attractors in possibility space, consilience the epistemological program that reconstructs them [11]. This supplies the metaphysical vocabulary for our middle position without committing us to its stronger claims.

Invariants as the discovered layer. The adelic cross-ratio document [12] proves the relevant uniqueness result: the cross-ratio is the *only* projective invariant of four collinear points—any rational function of coordinates invariant under  $\text{PGL}(2)$  is a function of the cross-ratio—so it is the essential syntactic primitive of projective geometry, defined over any field and thus independent of human notation, units, or coordinate systems [12]. Its summary chain runs from projective geometry through base-independence to the adèle ring and the adelic embedding of dimensionless physical constants [12]. This is the sharpest available model of our criterion: an object is "discovered" in our sense exactly when it is invariant under the admissible changes of representation, as the cross-ratio is invariant under projective coordinate change.

Methodological parallel from data mining. Robust subgroup discovery [8] contributes a structural fact we use as an analogy for theorem-space navigation: finding an ordered set of patterns and finding the smallest decision list for a dataset are both NP-hard, so finding an optimal subgroup list is NP-hard [8]; empirically, the SSD++ algorithm obtains the best results on 23 of 25 numeric-target datasets [8]. The analogy: even in empirical pattern spaces, the space of valid patterns is not invented by the miner—it is fixed by the data and the quality criterion, and navigating it efficiently is provably hard. The miner invents the description language; the pattern space is discovered.

### 3. Methods

#### 3.1 The representation/constraint criterion

We define a piece of mathematics as a pair  $(R, C)$ :

- $R$  (representation layer): the choice of axioms, definitions, notation, primitives, and descriptive language. This layer is *invented*: it is contingent, culturally conditioned, historically traceable, and varies across communities without contradiction.
- $C$  (constraint layer): the space of consequences that follow once  $R$  is fixed, together with the structures invariant across choices of  $R$ . This layer is *discovered*: it is not chosen, its navigation is provably hard (cf. the NP-hardness of optimal pattern lists [8]), and its invariant content survives translation between representations.

The criterion for membership in  $C$  rather than  $R$  is invariance under admissible change of representation, modeled on the cross-ratio's uniqueness as the only  $\text{PGL}(2)$ -invariant of four collinear points [12]. Formally, if  $T : R_1 \rightarrow R_2$  is an admissible translation between two representational schemes, an element  $m$  belongs to the constraint layer when  $T(m)$  is well-defined and the theorems involving  $m$  hold identically in both schemes. Elements that do not survive  $T$ —notation, choice of primitives, axiom formulation—belong to  $R$ .

This criterion is a convention we adopt; we state it as such. It is motivated by three converging considerations from the literature: the inherited-algorithm case, where the algorithm's content survived decades of forgetting and independent rediscovery while its documentary form did not [2]; the cross-lingual affect case, where the valence dimension survives translation while lexical gradients diverge [4]; and the projective-invariant uniqueness theorem [12].

#### 3.2 Analytical method

The paper's method is documentary analysis plus explicit arithmetic. We take four quantitative anchors, state every input number with its source, and perform every arithmetic step:

1. Exact sexagesimal verification of Plimpton 322 row 1, using the transcription and the relation  $A = (D/L)^2$  with  $L^2 = D^2 - S^2$  [6].
2. The 2020 PhilPapers survey distribution: we assume the survey reports nominalism at 41.85 percent, Platonism at 38.38 percent, and Other at 20.81 percent (Section 3.3 assumption), and we analyze the structure of the reported distribution.
3. A survivorship model of mathematical-formalism adoption, calibrated on the non-Archimedean activity table and the zero-of-three analogue adoption record [10].
4. The Gödel incompleteness argument, analyzed logically rather than numerically.

No experiment, survey, or simulation is performed in this paper; all empirical numbers are attributed to their cited sources, and all projections are labeled as such with stated assumptions.

#### 3.3 Assumptions

We assume the standard transcription of Plimpton 322 row 1 as  $[1\ 59]\ 00\ 15$  in the  $A$  column with  $S = 119$ ,  $L = 120$ ,  $D = 169$  [6]. We assume the 2020 PhilPapers survey reports nominalism at 41.85 percent, Platonism at 38.38 percent, and Other at 20.81 percent, that the percentages are reported to two decimal places, and that the residual category "Other" is defined by the survey itself. We assume, for the survivorship model, that "mainstream adoption within a decade" is the adoption criterion used in the historical reference class [10]. For the extension result we adopt the convention that the twenty-row corpus (fifteen recorded plus five predicted) is the natural completeness of the generating scheme, so the recorded fraction is  $15/20$ ; the same arithmetic read against the extant base alone gives  $5/15 \approx 33.33$  percent, and no conclusion in this paper depends on which frame is chosen.

## 4. Analysis

### 4.1 Plimpton 322: the constraint layer verified to exact arithmetic

**Input numbers.** From the transcription of the tablet's numerical part [6], row 1 reads, in the  $A$  column, [1 59] 00 15 — that is, the sexagesimal number  $1;59,00,15$ . From the standard explanation [6], the tablet's column satisfies  $A = (D/L)^2$  where  $L^2 = D^2 - S^2$ . We take the row-1 triple as  $S = 119$ ,  $L = 120$ ,  $D = 169$  (an assumption stated in Section 3.3).

**Step 1: verify the Pythagorean relation.**

$$S^2 + L^2 = 119^2 + 120^2 = 14161 + 14400 = 28561.$$

$$D^2 = 169^2 = 28561.$$

So  $S^2 + L^2 = D^2$  exactly, and  $(119, 120, 169)$  is a Pythagorean triple.

**Step 2: compute  $A$  from the triple.**

$$A = \left(\frac{D}{L}\right)^2 = \left(\frac{169}{120}\right)^2 = \frac{28561}{14400}.$$

Converting to sexagesimal:  $28561/14400 = 1 + 14161/14400$ . Now  $14161/14400 \times 60 = 14161/240 = 59 + 1/240$ . So the first sexagesimal digit after the unit is 59, with remainder  $1/240$ . Next digit:  $(1/240) \times 60 = 1/4 = 0 + 1/4$ . So the second digit is 0, remainder  $1/4$ . Third digit:  $(1/4) \times 60 = 15$ , remainder 0. Fourth digit: 0.

$$\frac{28561}{14400} = 1;59,00,15,00,\dots$$

**Step 3: compare with the tablet.** The tablet's row-1  $A$  value is  $1;59,00,15$  [6]. The computed expansion  $1;59,00,15,00,\dots$  agrees with the tablet value in all four recorded sexagesimal places, with the next digit 0—the tablet's truncation is exact to the precision of its four places.

**Step 4: the extension arithmetic.** Firsov's modified reciprocal-pair scheme predicts five new rows beyond the 15 listed by Price and Friberg [6]. Under the convention of Section 3.3, the natural corpus is

$$N_{\text{total}} = N_{\text{recorded}} + N_{\text{predicted}} = 15 + 5 = 20,$$

and the recorded fraction is

$$\frac{N_{\text{recorded}}}{N_{\text{total}}} = \frac{15}{20} = 0.75,$$

so the tablet records 75 percent of the scheme's natural twenty-row extension. Read against the extant base alone, the predicted rows are

$$\frac{N_{\text{predicted}}}{N_{\text{recorded}}} = \frac{5}{15} = 0.3333\dots \approx 33.33\%,$$

a one-third expansion. Both frames are reported; the qualitative point—five rows follow from the same rule that generated the recorded fifteen—is frame-independent.

**Interpretation.** The Babylonian scribes recorded, in their own notation, a value that a modern computation reproduces digit-for-digit. The *notation* (sexagesimal place value, the specific reciprocal-pair generation scheme, the choice of which rows to tabulate) is the invented representation layer  $R$ ; the *relation*  $S^2 + L^2 = D^2$  and the ratio  $A = (D/L)^2$  are the constraint layer  $C$ , and they are invariant under the notational translation we just performed. The result that the generating scheme extends to predict five new rows beyond the tablet's 15 [6] strengthens the point: the constraint generalizes beyond the instances anyone recorded, which is the defining behavior of a discovered structure rather than an invented list.

#### 4.2 The PhilPapers distribution: what the expert split actually measures

**Input numbers.** We assume the 2020 PhilPapers survey reports nominalism at 41.85 percent, Platonism at 38.38 percent, and Other at 20.81 percent (Section 3.3 assumption).

**Step 1: the gap.**

$$41.85\% - 38.38\% = 3.47 \text{ percentage points.}$$

**Step 2: the combined major labels.**

$$41.85\% + 38.38\% = 80.23\%.$$

**Step 3: consistency check on the residual.** The stated Other share is 20.81 percent, while  $100\% - 80.23\% = 19.77\%$ . The 1.04-point discrepancy between the reported Other share and the arithmetic residual is unexplained by any source this paper cites; we take the reported 20.81 percent as given under the convention that the survey's residual category is authoritative, and note that the qualitative conclusion is insensitive to whether the fifth is 19.77 or 20.81 percent: in either case roughly one in five experts rejects both labels.

**Interpretation.** A 3.47-point gap between nominalism and Platonism, with a fifth of experts outside both camps, is not a field that has found the dichotomy natural. Under our criterion this is exactly what should happen if the dichotomy is false: positions that locate mathematics partly in representation and partly in constraint (structuralist views) have no natural home in a binary question, so they pool into "Other." We argue in Section 5 that the Other fifth is where weak structural realism lives.

#### 4.3 Survivorship control on Wigner's puzzle

**The puzzle.** Wigner's argument cites the hits. We take as given the standard historical dating—general relativity predicted light bending in 1915, confirmed 1919; the Dirac equation (1928) yielded the positron (1932)—as uncontested framing, and the survivorship argument below rests only on the adoption numbers of [10]. The objection is survivorship bias: the denominator—formalisms that never found physical application—is uncounched in the standard telling.

**Input numbers.** From the non-Archimedean assessment [10]: current activity outside the established domains is approximately 50 papers/year ( $p$ -adic QM), 20 (adelic string theory),  $< 5$  (ultrametric QEC),  $< 10$  ( $p$ -adic neural networks),  $< 5$  ( $p$ -adic channel coding),  $< 10$  ( $p$ -adic condensed matter). The historical reference class: zero of three closest analogues (Copernican, relativistic, quantum revolutions) achieved mainstream adoption within a decade, and no empirical anomaly has been identified for the Archimedean default [10]. Ostrowski's theorem establishes that the  $p$ -adics are a coequal completion of  $\mathbb{Q}$  with the reals [10].

**Step 1: total visible non-adopting activity.** Summing the upper bounds of the sparse domains:

$$50 + 20 + 5 + 10 + 5 + 10 = 100 \text{ papers/year (upper bound).}$$

So on the order of  $10^2$  papers per year are currently produced in a mathematics that is a coequal completion of the rationals by Ostrowski's theorem [10] yet has found no physics adoption outside the established domains the assessment counts separately. We assume, without cited support, that ultrametric distance is used in phylogenetics and isogeny-based cryptography; no conclusion in this paper rests on that assumption.

**Step 2: the adoption base rate from the reference class.** The reference class gives an adoption rate of

$$\frac{0}{3} = 0\%$$

for mainstream adoption within a decade among the closest historical analogues [10]. We emphasize what this number is and is not: it is a documented count from the cited assessment, not a measured rate over all formalisms; with  $n = 3$  the one-sided 95 percent Clopper–Pearson upper bound on the underlying adoption probability  $p$  solves  $(1 - p)^3 = 0.05$ , giving  $p \leq 1 - 0.05^{1/3}$ . Computing:  $0.05^{1/3} \approx 0.3684$ , so  $p \lesssim 0.6316$ . Even this weak bound is informative: a formalism built today has, on this class, at most a  $\sim 63$  percent chance of decade-scale adoption even under the most pessimistic reading of three failures—and the point estimate is far lower.

**Step 3: the naturalistic explanation of the hits.** If the per-formalism decade adoption probability is low, the celebrated hits (GR 1915  $\rightarrow$  1919, Dirac 1928  $\rightarrow$  1932) are not samples from a process where "mathematics is unreasonably effective"; they are the visible tail of a process where mathematicians produce a large stock of structure, physicists search that stock, and matches compound in precision over decades because both communities inherit and refine the same constraint layer. Under the representation/constraint split this is exactly what is expected: the constraint layer  $C$  is shared (hence matches are possible and recur), while the representation layer  $R$  is contingent (hence most formalisms never meet a physical use, and the adeles—mathematically coequal with the reals—sit at  $\sim 10^2$  papers/year with no anomaly to displace the default [10]). The hits are evidence for a shared constraint structure, not for a pre-existing Platonic physics repository.

#### 4.4 Gödel incompleteness: metatheory-relative, not object-level

Gödel incompleteness is often cited for Platonism: truth outruns provability, so truth must be about real objects. Under our criterion the inference fails. The incompleteness phenomenon is a fact about the relation between a formal system (a representation  $R$ ) and its intended model—a relation established in a *metatheory*, which is itself another representation. What is discovered is the structural constraint: no sufficiently strong consistent representational scheme captures all arithmetic truth *and* proves its own consistency. That constraint is invariant across which formal system you pick (any sufficiently strong one fails), which is why it is a discovered fact about the representation/constraint relationship itself. It does not follow that the objects the representations are about exist mind-independently; the constraint is on representations, and constraints on representations can be discovered without the represented being inhabited. This reading is consistent with the anti-Platonist constraint that only mind-based mathematics is scientifically accessible [9], while preserving the objectivity of the incompleteness phenomenon as an invariant.

## 5. Results

**R1 (Section 4.1).** The first row of Plimpton 322 satisfies  $A = (169/120)^2 = 28561/14400 = 1; 59, 00, 15$  in sexagesimal, matching the tablet's recorded value  $1; 59, 00, 15$  exactly in all four recorded places, with the next sexagesimal digit 0. The Pythagorean identity  $119^2 + 120^2 = 169^2$  holds exactly ( $14161 + 14400 = 28561$ ).

**R2 (Section 4.1).** The generating scheme behind the tablet's 15 rows extends to predict 5 additional rows [6]. Under the convention of Section 3.3, the 15 recorded rows are 75 percent of a 20-row extension ( $15/20 = 0.75$ ); read against the extant base, the five predicted rows are a 33.33 percent expansion ( $5/15 \approx 0.3333$ ). Either way, the constraint generalizes beyond its recorded instances.

**R3 (Section 4.2).** Under the assumption stated in Section 3.3, the 2020 PhilPapers distribution puts nominalism at 41.85 percent, Platonism at 38.38 percent (gap: 3.47 percentage points), and Other at 20.81 percent; roughly one in five experts rejects both labels of the dichotomy.

**R4 (Section 4.3).** The documented historical reference class for formalism adoption gives 0 of 3 closest analogues achieving mainstream adoption within a decade [10]; the one-sided 95 percent Clopper–Pearson upper bound on the per-formalism decade adoption probability from three failures is  $p \leq 1 - 0.05^{1/3} \approx 0.632$ , with the point estimate 0. Current non-Archimedean activity outside established domains is bounded at approximately 100 papers/year despite the  $p$ -adics being a coequal completion of  $\mathbb{Q}$  [10].

**R5 (Section 4.3).** Under the survivorship model, Wigner's hits are the visible tail of a stock-and-search process over a shared constraint layer, not evidence of unreasonably effective access to a Platonic physics archive. This is a derived interpretation of R4's numbers, not an independent measurement.

**R6 (Sections 4.4, 5).** The representation/constraint criterion classifies every anchor case correctly: Plimpton 322 (notation invented, relation discovered and exactly invariant), cross-lingual affect (valence dimension stable across languages while lexical gradients diverge [4]), inherited algorithms (content rediscovered independently while documentation was lost [2]), incompleteness (constraint on representations, metatheory-relative), and the adeles (coequal structure, no adoption—constraint exists, representation unmatched to physics). The position that results—weak structural realism, the discovered-constraints view—is the coherent home for the survey's Other fifth (R3).

## 6. Discussion

The central claim of this paper is that the invented/discovered dichotomy is false and replaceable by a representation/constraint split with an invariance criterion. The claim is a philosophical convention backed by arithmetic anchors, not an empirical discovery, and its weaknesses should be stated with the same care as its supports.

**Falsification.** The observation that would falsify the main result is a counterexample to the invariance criterion: a mathematical element that is demonstrably constraint-like (it drives proofs and applications in one community) yet fails to survive every admissible translation to an independent representational scheme, or conversely a representational artifact that is invariant under all admissible translations and therefore indistinguishable from a constraint. Short of that, the Plimpton 322 anchor itself is falsifiable: a systematic epigraphic survey of Old Babylonian tablet collections that exhaustively examines tablets dated to the same period and region and finds no rows matching the five predicted reciprocal-pair patterns [6] would undermine the claim that the tablet's generating scheme is a discovered constraint extending beyond its recorded instances. Conversely, discovery of even one matching row would confirm the prediction and strengthen the split.

**Limitations list.** The result rests on the following assumptions, with sensitivity notes:

- *The invariance criterion itself.* Membership in the constraint layer is defined by survival under admissible translation. What counts as "admissible" is itself a representational choice, so the criterion has a regress at its boundary. The result is moderately sensitive to this: a stricter admissibility notion shrinks  $C$ ; a looser one inflates it. We take the projective-geometric model [12] as the fixed point, but a critic could deny the analogy.
- *The Plimpton 322 triple assumption.* We assumed  $S = 119$ ,  $L = 120$ ,  $D = 169$  for row 1. The exact four-place match (R1) is strong evidence the assumption is right, but the tablet's purpose remains contested—Firsov notes substantially divergent explanations of the tablet's construction [6]—and a rival reconstruction that also reproduces the recorded digits would weaken the claim that the reciprocal-pair scheme is the operative constraint. The exact-arithmetic result R1 survives any such rivalry, since it depends only on the transcription and the Pythagorean relation, not on the generation scheme.

- *The extension convention.* The 75-percent framing depends on treating the twenty-row corpus as natural; a critic who rejects the convention retains only the 33.33-percent expansion, which is the same arithmetic and supports the same qualitative conclusion.
- *The PhilPapers figures.* We take the survey percentages as given and cannot audit their sampling; the residual-category discrepancy of 1.04 points (Section 4.2) shows the category structure is not airtight. The qualitative conclusion—one in five experts outside both camps—holds under either residual value.
- *The survivorship reference class.* The zero-of-three record [10] is a small, purposively chosen class, not a random sample of formalisms; the binomial bound  $p \lesssim 0.632$  is correspondingly weak. The survivorship model would be falsified if a documented census of nineteenth- and twentieth-century formalisms showed a high decade-scale adoption rate for mathematics developed without physical motivation.
- *The Gödel reading.* The metatheory-relative reading is a philosophical interpretation, not a theorem; a Platonist can accept the invariance of the incompleteness constraint and still read it as access to mind-independent arithmetic truth. Our argument shows the inference from incompleteness to Platonism is not forced, not that it is invalid.
- *The cross-cultural analogy.* Transferring findings from modern NLP systems [3], [4], [5], [7] to ancient mathematical traditions is an analogy of structure, not a demonstrated mechanism; no measurement connects embedding-space behavior to scribal practice.

**Open questions.** Can a formal metric quantify the "invented content" of a theory, e.g., the proportion of its primitives that are conventionally chosen versus logically forced? Does the convergence/consilience framework of [11] yield testable predictions about which mathematical structures are likely to find physical application? Would a large-scale census of formalisms confirm or break the survivorship model of Section 4.3?

## 7. Practical Implications and Engineering Pathways

**(a) Who would use the result and what changes.** The users are philosophers of mathematics and the profession's survey designers, historians of exact sciences working on Old Babylonian mathematics, and funders who must classify formalism-driven mathematical research. If the representation/constraint split holds, the PhilPapers survey's binary question ("nominalism or Platonism?") should be redesigned to ask respondents to distribute a given piece of mathematics between a representation layer and a constraint layer, since Section 4.2 shows a fifth of experts already pool in "Other" under the binary. For historians, the five predicted rows of the Plimpton 322 generating scheme [6] convert a philosophical claim into a searchable archival prediction: museum and collection catalogs of Old Babylonian mathematical tablets can be checked for rows matching the predicted reciprocal-pair patterns. For funders, the survivorship model of Section 4.3 changes the evaluation question from "is this formalism likely to describe physics?" to "is this formalism likely to enter the shared constraint stock that physicists search?"—a question the non-Archimedean data [10] show has a low base rate (zero of three closest analogues adopted within a decade) regardless of the formalism's internal mathematical standing, with the assumed (uncited) connection to phylogenetics and isogeny-based cryptography noted in Section 4.3 carrying no weight in the model.

**(b) Current readiness.** This work is purely analytical: documentary analysis, explicit arithmetic, and philosophical argument, and it has not been peer reviewed. Against the NASA Technology Readiness Level definitions it sits at TRL 1–2: basic principles are observed and reported (the exact four-place match of Section 4.1, the adoption counts of Section 4.3), and the concept and its applications have been stated but not demonstrated in any operational setting. Analytical work of this kind cannot exceed TRL 3, and no claim of validated or deployed use is made here.

**(c) The smallest next test.** The smallest de-risking step is archival, not experimental: a systematic epigraphic survey that exhaustively catalogs Old Babylonian tablet fragments dated to the same period and region as Plimpton 322 and checks each fragment's numerical columns against the five predicted rows of the modified reciprocal-pair scheme [6]. The outcome that would refute the anchor is a complete survey of the relevant corpus finding no fragment matching

any predicted row; the outcome that would strengthen it is one or more matching rows. A second, cheaper step is a reanalysis of the PhilPapers 2020 data with the binary question decomposed into a two-layer question, testing whether respondents who chose "Other" distribute coherently over a representation/constraint split; the refuting outcome is that the "Other" responses show no coherent structure under the decomposition.

**(d) Binding constraints.** For the archival test, the binding constraints are access and condition of tablets (we assume, without cited support, that portions of Plimpton 322 itself are missing, and the transcription of Section 4.1 already requires reconstituting bracketed material [6]), the availability of high-resolution imaging, and the assyriological expertise needed to date and classify fragments. For the survey reanalysis, the binding constraint is data: the raw 2020 PhilPapers response-level data and the survey's own definition of its residual category, which Section 4.2 shows carries a 1.04-point discrepancy against the two major labels. Whether raw response-level access is available is an open question this paper cannot resolve. No manufacturing, materials, or energy constraints arise, since the proposed next steps are documentary and statistical.

**(e) Order-of-magnitude bounds.** The only quantitative projection this paper can offer, with stated assumptions, concerns the archival search's expected yield. We assume the five predicted rows [6] are the search target, that the relevant corpus is on the order of  $10^2$  catalogued Old Babylonian mathematical tablets (an assumption; no source in this paper gives a corpus size), and that matching fragments are uniformly detectable. Under these assumptions the search is a catalog-level effort—on the order of  $10^2$  tablet examinations—rather than a field campaign. This is a projection from assumed inputs, not a measured cost, and no monetary figure is derivable from the paper's sources.

**(f) Risks and misuse paths.** The framework can be misused rhetorically: a funding body could cite the low adoption base rate of Section 4.3 to defund all formalism-driven mathematics, ignoring that the model's point estimate rests on a reference class of three and that the celebrated hits (GR, Dirac) came from exactly such work. Conversely, a Platonist could cite the exact four-place match of Section 4.1 as evidence of mind-independent structure, when the paper's own reading is that the match evidences an invariant constraint, not an inhabited Platonic archive. Both misuses overread the numbers; the Discussion's limitations list states the sensitivities that block them.

## 8. Conclusion

The invented/discovered dichotomy is false as a description of mathematical practice, and a principled replacement is available. A piece of mathematics is a pair  $(R, C)$ : a representation layer of axioms, notation, and definitions, which is invented, contingent, and culturally traceable, and a constraint layer of consequences and invariants, which is discovered in the precise sense that it survives admissible translation between representations. Four anchors support the split. First, the first row of Plimpton 322 encodes  $(169/120)^2 = 28561/14400 = 1; 59, 00, 15$ , matching the tablet's recorded value in all four sexagesimal places, so the Pythagorean constraint is verified digit-for-digit beneath a notational scheme that is unmistakably cultural. Second, the tablet's generating scheme extends to predict five rows beyond its fifteen [6], showing the constraint generalizes beyond its recorded instances. Third, the survivorship data for non-Archimedean geometry—zero of three closest historical analogues achieving mainstream adoption within a decade, current activity bounded near  $10^2$  papers per year despite the  $p$ -adics being a coequal completion of  $\mathbb{Q}$  [10]—quantifies the denominator Wigner's puzzle omits and explains the celebrated hits naturalistically as the visible tail of a stock-and-search process over a shared constraint layer. Fourth, Gödel incompleteness, read metatheory-relative, constrains representations without requiring that the represented be inhabited.

Weak structural realism is the position this yields: mathematical constraints are discovered and mind-independent in the thin sense of being invariant under representational change, while the representations navigating them are invented in the full sense of being contingent human choices. The position is the natural home for the fifth of expert philosophers—20.81 percent in the 2020 PhilPapers distribution—who reject both nominalism and Platonism. The claim is falsifiable at its strongest anchor: an exhaustive archival survey of Old Babylonian tablets that finds no row matching

the five predicted patterns would undermine the Plimpton extension result, and a single matching row would confirm it. The paper's contribution is not to end the invented/discovered dispute but to replace its question with a better one—asked case by case, of each layer, with numbers attached.

## Provenance and AI Involvement

- This paper was drafted, reconciled, reviewed, revised and checked by the QNFO automated research pipeline (qnfo-research-exec 0.39.3-codeagent). 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, openai/gpt-oss-120b, moonshotai/kimi-k2.6, 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; every cited work's full text was retrieved (arXiv HTML, or the QNFO paper's body) and the writers were given excerpts of its own sections; statements about a cited work are restricted to its retrieved title, abstract and those excerpts.
- 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: the execution sandbox was unavailable, so a second language model recomputed 11 quantitative claims step by step in text from the paper's stated inputs (no code was executed); 0 differed from the stated value at the stated precision.
- Not checked: whether each cited source supports the sentence that cites it beyond the retrieved full-text excerpts; 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] A view from lockdown: mathematics discovered, invented, and inherited. arXiv:2103.04101v1. <https://arxiv.org/abs/2103.04101v1> [2] Mathematics discovered, invented, and inherited. arXiv:1309.3073v2. <https://arxiv.org/abs/1309.3073v2> [3] LLM-C3MOD: A Human-LLM Collaborative System for Cross-Cultural Hate Speech Moderation. arXiv:2503.07237v1. <https://arxiv.org/abs/2503.07237v1> [4] Supervised Semantic Differential for Cross-Cultural Concept Analysis: A Case Study of Human Affect. arXiv:2605.28225v1. <https://arxiv.org/abs/2605.28225v1> [5] Toward Culturally Grounded Natural Language Processing. arXiv:2603.26013v2. <https://arxiv.org/abs/2603.26013v2> [6] The numbers behind Plimpton 322. arXiv:1109.3814v2. <https://arxiv.org/abs/1109.3814v2> [7] Does Cross-Cultural Alignment Change the Commonsense Morality of Language Models?. arXiv:2406.16316v1. <https://arxiv.org/abs/2406.16316v1> [8] Robust subgroup discovery. arXiv:2103.13686v4. <https://arxiv.org/abs/2103.13686v4> [9] <https://papers.qnfo.org/papers/embodied-mathematics-lakoff-nunez>. QNFO: Embodied Mathematics After Lakoff & Nunez: A QNFO Perspective. [10] <https://papers.qnfo.org/papers/acrp08-paradigm-forecast>. QNFO: When Will Non-Archimedean Geometry Displace the Real Numbers? A Structured Assessment of the Adelic Substrate Thesis. [11] <https://papers.qnfo.org/papers/static-architecture-of-reality-and-its-biological-interface>. QNFO: Convergence Consilience and the Hierarchical Architecture of Reality. [12] <https://papers.qnfo.org/papers/adelic-crossratio>. QNFO: Adelic Cross-Ratio.